

Keith J. Dreyer  
David S. Hirschorn  
James H. Thrall  
Amit Mehta *Editors*

# PACS



A Guide to the Digital Revolution  
Second Edition

 Springer

# **PACS**

## **A GUIDE TO THE DIGITAL REVOLUTION**

Second Edition

# PACS

## A GUIDE TO THE DIGITAL REVOLUTION

SECOND EDITION

KEITH J. DREYER, DO, PHD

DAVID S. HIRSCHORN, MD

JAMES H. THRALL, MD

AMIT MEHTA, MD

EDITORS

With 97 Illustrations

 Springer

Keith J. Dreyer, DO, PhD  
Assistant Professor of Radiology  
Harvard Medical School  
Vice Chairman of Radiology Informatics  
Massachusetts General Hospital  
Boston, MA 02114  
USA

James H. Thrall, MD  
Professor of Radiology  
Harvard Medical School  
Chairman of Radiology  
Massachusetts General Hospital  
Boston, MA 02114  
USA

David S. Hirschorn, MD  
Research Fellow in Radiology Informatics  
Harvard Medical School  
Massachusetts General Hospital  
Boston, MA 02114  
*and*  
Director of Radiology Informatics  
Staten Island University Hospital  
Staten Island, NY 10305  
USA

Amit Mehta, MD  
Director of Interventional Radiology  
St. Josephs Health Center  
Toronto, Ontario  
Canada M6R 1B5

Library of Congress Control Number: 2005925976

ISBN 10: 0-387-26010-2

ISBN 13: 978-0387-26010-5

Printed on acid-free paper.

© 2006 Springer Science+Business Media, Inc.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, Inc., 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden. The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of going to press, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed in the United States of America. (BS/EB)

9 8 7 6 5 4 3 2 1

springeronline.com

*To my parents  
and to my family and loved ones  
whose genuine enthusiasm  
inspires my passion for new ideas  
KJD*

*To my wife Elisbeva,  
for your love and warmth  
DSH*

*To my wife Jean, who manages information technology in our household  
and my children, Trevor and Keely, who have joined me in working  
in the digital world  
JHT*

*To my parents, MHM and NM for their guidance and support,  
and my family, SAM, CKM, SKM, HKM, and RM (on her way)  
for their constant humor,  
and my teachers and colleagues for their help along the way  
AM*

# PREFACE

---

---

The digital transformation of radiology marches on. Slow and inefficient film- and paper-based methods are giving way to quicker and simpler computer-based ones. Report turnaround times are being measured in minutes or hours instead of days. There are more than a hundred information technology vendors in the radiology market; they all claim that they have best system to suit your needs. How can you tell which ones really are best for you? How can you distinguish between truly important features and ones that are just marketing ploys? How can you know the right questions to ask to make sure you are getting all that you need and avoiding hidden costs? This book covers the full spectrum of radiology information technology in the digital department. It brings together the expertise of many of the respected leaders in PACS, RIS, and speech recognition systems from academic centers such as Harvard and the University of Maryland, community hospitals, and even international teleradiology practices. Recent changes in image display technologies are explored, as well as the maturation of digital mammography, three-dimensional imaging, the electronic medical record, and teleradiology.

The process of assessing the needs of the institution and developing a request for proposal that matches those unique requirements is covered in

depth. This includes information on writing the primary evaluation criteria, evaluating proposals from different vendors, and choosing appropriate vendors. To justify the considerable investment of a PACS, financial concepts and tools are included that are useful in the financial evaluation. Legal issues that arise with teleradiology and formal policies that address these issues are also discussed.

This book is intended for radiologists, technologists, administrators, and IT professionals who want to better understand these technologies and their impact. It is also useful for industry vendors, consultants, and health-care leaders who have an interest and modest knowledge of IT management issues.

What's new in the second edition? It presents some of the latest research on reading room design and radiologist workflow. Recent developments in CR and digital mammography are also included. Major changes in display and storage technologies which can have a huge impact on the cost of PACS are discussed. Experience gained from maturing teleradiology practices is shared. The role of decision support tools for order entry and digital teaching files are also explored. These updates and additions will provide you with the most current information about the digital transformation of radiology.

We would like to acknowledge our developmental editor, Merry Post, for keeping track of the myriad of details needed to make this second edition a reality. She also deserves credit for dealing with all of our crazy schedules. Her persistence is what kept this book on track. Kudos to the chapter authors for sharing their expertise with all of us, and thanks to my fellow section editors for organizing and keeping track of progress of the chapters in their sections. I also can not thank my wife, Elisheva, enough for supporting me in editing this book. The warm and loving home that she provides allows me to focus on the task at hand; write, edit, and review chapters; and send out countless e-mails.

*David S. Hirschorn, MD*

# CONTENTS

---

---

PREFACE	vii
CONTRIBUTORS	xiii

SECTION I	INTRODUCTION	1
	<i>Keith J. Dreyer</i>	

1	INTRODUCTION	3
	<i>David S. Hirschorn</i>	

SECTION II	ADMINISTRATIVE ISSUES	7
	<i>Gordon Smith</i>	

2	INTRODUCTION TO RIS AND PACS	9
	<i>Gordon Smith</i>	

3	PACS STRATEGIC PLAN AND NEEDS ASSESSMENT	27
	<i>Leonard A. Levine</i>	



---

4	CREATING THE PACS REQUEST FOR PROPOSAL AND SELECTING A VENDOR	45	<i>Alan L. Schweitzer and Gordon Smith</i>
5	REENGINEERING WORKFLOW: A FOCUS ON PERSONNEL AND PROCESS	73	<i>Bruce I. Reiner and Eliot L. Siegel</i>
6	REENGINEERING WORKFLOW: THE RADIOLOGIST'S PERSPECTIVE	97	<i>Eliot L. Siegel, Bruce I. Reiner, and Nancy Knight</i>
7	FINANCIAL MODELING	125	<i>Syrene R. Reilly and David Avrin</i>
8	LEGAL ISSUES AND FORMAL POLICIES	145	<i>Gordon Smith and David S. Hirschorn</i>
SECTION III TECHNICAL ISSUES 171			
<i>Thomas J. Schultz</i>			
9	COMPUTER FUNDAMENTALS	173	<i>Keith J. Dreyer</i>
10	DIGITAL IMAGING FUNDAMENTALS	183	<i>Keith J. Dreyer and Mannudeep K. Kalra</i>
11	IMAGE ACQUISITION	189	<i>Katherine P. Andriole</i>
12	IMAGE COMPRESSION	229	<i>Bradley J. Erickson</i>
13	PACS ARCHITECTURE	249	<i>Kenneth Heckman and Thomas J. Schultz</i>
14	NETWORKING FUNDAMENTALS	269	<i>Scott M. Rogala</i>
15	SERVERS AND OPERATING SYSTEMS	303	<i>Kenneth M. Nesbitt, Thomas J. Schultz, and Roberto Dasilva</i>

- 16 STORAGE AND ENTERPRISE  
ARCHIVING 319  
*Paul G. Nagy and Thomas J. Schultz*
- 17 IMAGE DISPLAYS 347  
*David S. Hirschorn*
- 18 DIGITAL MAMMOGRAPHY 363  
*Martin J. Yaffe*
- 19 WEB DISTRIBUTION 373  
*Keith J. Dreyer*
- SECTION IV CLINICAL ISSUES 383  
*David S. Hirschorn*
- 20 PACS WORKSTATION SOFTWARE 385  
*Steven C. Horii*
- 21 BREAST IMAGING, COMPUTER-AIDED  
DETECTION, AND COMPUTER-ASSISTED  
CLASSIFICATION 433  
*Shalom S. Buchbinder*
- 22 THREE-DIMENSIONAL IMAGING IN  
RADIOLOGY 447  
*Gordon J. Harris*
- 23 VOICE RECOGNITION 467  
*Michael J. Mardini and Amit Mehta*
- 24 ORDER ENTRY IN RADIOLOGY 483  
*Daniel I. Rosenthal*
- 25 DIGITAL TEACHING FILES AND  
EDUCATION 495  
*Khan M. Siddiqui and Barton F. Branstetter IV*
- 26 TELERADIOLOGY 523  
*Giles Boland, Jonathan T. Schlakman, and  
James H. Thrall*
- INDEX 561

# CONTRIBUTORS

---

---

*Katherine P. Andriole, PhD.*

Associate Professor of Radiology, Harvard Medical School, Assistant Medical Director, Imaging IT, Director of Imaging Informatics, Center for Evidence-Based Imaging, Brigham and Women's Hospital, Boston, MA 02120, USA

*David Avrin, MD, PhD*

Professor of Radiology, Adjunct Professor of Medical Informatics, University of Utah, University of Utah Hospital and Clinics, Salt Lake City, UT 84132, USA

*Giles Boland, MD*

Associate Professor of Radiology, Harvard Medical School, Vice Chairman, Business Development, Massachusetts General Hospital, Boston, MA 02114, USA

*Barton F. Branstetter IV, MD*

Assistant Professor of Radiology and Otolaryngology, Director of Head and Neck Imaging, Associate Director of Informatics, University of Pittsburgh Medical Center, Pittsburgh, PA 15213, USA

*Shalom S. Buchbinder, MD*

Clinical Associate Professor, Albert Einstein College of Medicine, Chairman of Radiology, Clinical Associate Professor of Radiology, Obstetrics, Gynecology and Womens' Health, Staten Island University Hospital, Staten Island, NY 10305, USA

*Roberto Dasilva, MCSE*

Data Center Manager, Department of Radiology, Massachusetts General Hospital, Boston, MA 02114, USA

*Keith J. Dreyer, DO, PhD*

Assistant Professor of Radiology, Harvard Medical School, Vice Chairman of Radiology Informatics, Massachusetts General Hospital, Boston, MA 02114, USA

*Bradley J. Erickson, MD, PhD*

Associate Professor of Radiology and Medical Informatics, Director, Radiology Informatics Laboratory, Department of Radiology (E-2), Mayo Clinic, Rochester, MN 55905, USA

*Gordon J. Harris, PhD*

Director, 3D Imaging Service, Massachusetts General Hospital, Boston, MA 02114, USA

*Kenneth Heckman, BSN*

Information Systems Analyst, Partners HealthCare System, Inc., Boston, MA 02115, USA

*David S. Hirschorn, MD*

Research Fellow in Radiology Informatics, Harvard Medical School, Massachusetts General Hospital, Boston, MA 02114; Director of Radiology Informatics, Staten Island University Hospital, Staten Island, NY 10305, USA

*Steven C. Horii, MD*

Professor of Radiology, University of Pennsylvania, University of Pennsylvania Medical Center, Philadelphia, PA 19104, USA

*Mannudeep K. Kalra*

Director of CT Research, Assistant Professor of Radiology, Emory University Hospital, Atlanta, GA 30322, USA

*Nancy Knight, PhD*

Coordinator, Research Publications and Grants, Veterans Affairs Maryland Healthcare System, Baltimore, MD 21201, USA

*Leonard A. Levine, BS, MSIE*

Manager, Radiology Data Services, Department of Radiology, Massachusetts General Hospital, Boston, MA 02114, USA

*Michael J. Mardini, MBA*

Chief Executive Officer, Commissure, Inc., New York, NY 10011, USA;  
Founder and Former CEO, Talk Technology, Inc.

*Amit Mehta, MD*

Director of Interventional Radiology, St. Josephs Health Center, Toronto, Ontario, Canada M6R 1B5

*Paul G. Nagy, PhD*

Assistant Professor of Radiology, University of Maryland, University of Maryland Medical Center, Baltimore, MD 21201, USA

*Kenneth M. Nesbitt*

Systems Engineer, Partners IS—Enterprise Medical Imaging, Boston, MA 02114, USA

*Syrene R. Reilly, MBA*

Director of Quality Management Services, Partners HealthCare System, Inc., Boston, MA 02199, USA

*Bruce I. Reiner, MD*

Director of Radiology Research, Veterans Affairs Maryland Healthcare System, Baltimore, MD 21201, USA

*Scott M. Rogala*

Corporate Manager Network Engineering, Partners Healthcare, Charlestown, MA 02129, USA

*Daniel I. Rosenthal, MD*

Professor of Radiology, Harvard Medical School, Vice Chairman for Administration, Department of Radiology, Massachusetts General Hospital, Boston, MA 02114, USA

*Jonathan T. Schlakman, MD*

Radiologist, Remote Radiology International, Efrat, Israel 90435

*Thomas J. Schultz, BSE*

Chief Engineer, Partners HealthCare System, Inc., Boston, MA 02114, USA

*Alan L. Schweitzer, MEE*

Chief Technology Officer, Radiology Consulting Group, Boston, MA 02114, USA

*Khan M. Siddiqui, MD*

Chief, Imaging Informatics and Body MR Imaging, Veterans Affairs Maryland Healthcare System, Baltimore, MD 21201, USA

*Eliot L. Siegel, MD*

Professor of Diagnostic Radiology, University of Maryland School of Medicine, Director of Imaging, Veterans Affairs Maryland Healthcare System, Baltimore, MD 21201, USA

*Gordon Smith, MBA*


Director of Radiology Informatics, Department of Radiology, Massachusetts General Hospital, Boston, MA 02114, USA

*James H. Thrall, MD*

Professor of Radiology, Harvard Medical School, Chairman of Radiology, Massachusetts General Hospital, Boston, MA 02114, USA

*Martin J. Yaffe, PhD*

Professor, Department of Medical Imaging and Medical Biophysics, University of Toronto, Senior Scientist, Imaging/Bioengineering Research, Sunnybrook and Women's College Health Sciences Centre, Toronto, Ontario, Canada M4N 3M5



# **INTRODUCTION**

KEITH J. DREYER

# INTRODUCTION

---

---

DAVID S. HIRSCHORN

**T**he first edition of this book made the point that picture archiving and communication systems (PACS) were no longer just a possibility but a reality. The second edition takes this statement one step further: PACS is not just a reality but a necessity. Most larger radiology departments have gone digital, and smaller departments and imaging centers are not far behind. Printing CT and MRI exams for interpretation is like printing your e-mail in order to read it. Ten years ago this analogy would be lost on most radiologists because they didn't know what e-mail was, but now virtually all radiologists know what it is and use it on a regular basis. Digital cameras are ubiquitous, and millions of consumers, radiologists among them, are filling up hard drives instead of shoeboxes with their family photos. By now most radiologists have viewed images on the Internet and have begun to recognize the benefits of managing images on a computer.

The benefits of PAC systems are clear. Within seconds after an image is acquired, it can be viewed by the radiologist and any number of referring and treating physicians simultaneously. There is no film to be lost or stolen. CT exams with a thousand images are becoming common and simply cannot be managed effectively on film. PACS viewing software can be used



to dissect, analyze, magnify, or reformat image data in an infinite number of ways. Virtual private networks can transmit whole exams across the globe within seconds for remote consultation, perhaps in the middle of the night to a radiologist who is just starting her day. Today's archives can keep decades of studies online in a cost-effective manner and in a much more organized and accessible manner than ever possible in a traditional fileroom.

The PAC system is the most visible component of a digital radiology department but is by no means the only one. A successful PACS requires a strong radiology information system (RIS) to feed it patient and exam information and to keep track of the life cycle of all exams from order placement to final result. The RIS ties together all the computer systems within the department and is typically the sole point of communication to the world outside the department, such as the hospital information system and the billing system. As such, it is perhaps the most complex system in the department.

The third key component of a digital radiology department is the speech recognition system for report transcription. Speech recognition systems have been available for several years and are slowly becoming a necessity. As hospitals begin to realize that it is not unreasonable to expect a final report within hours instead of days, the pace of adoption will begin to pick up. Speech interfaces to computers are still uncommon in the general computing environment, but this will inevitably change. Many commercial telephone answering services routinely use speech recognition, and the trend is increasing. As will be explained later in detail, speech recognition systems do far more than convert speech to text. They yield numerous other benefits derived from using a computer-based dictation system that just weren't worth pursuing until speech recognition technology came along.

These 3 systems—the RIS, the PACS, and the speech recognition system—form the backbone of a digital radiology department and are discussed in detail in this book. The RIS directs information flow of exams from the ordering process, scheduling, and image acquisition through interpretation, communication of results, and billing. The PACS serves to receive and store the images from the modalities and to distribute them to radiologists for primary interpretation and throughout the healthcare enterprise for clinical review. The speech recognition system is a sophisticated and powerful tool to help the radiologist generate a clear and accurate report in a timely fashion. When implemented correctly with maximum system integration, the result is better, faster, and more cost-effective patient care.

## ORGANIZATION

This second edition is organized differently from the first. The first edition mirrored the topics required for the process of developing a request for proposal (RFP) for PACS. In this edition, we chose to organize the topics around the 3 main perspectives from which most people approach digital radiology: administrative, technical, and clinical. In this way the reader can more quickly focus on topics of personal interest. Some may choose to focus on only a few chapters of one section; others who play multiple roles will need to draw on 2 or all 3 of the sections.

The administrative section begins with an introduction to RIS and PACS and proceeds to explore the issues involved in obtaining these systems. The effects that these systems have on the technical staff and the radiologists are then discussed. Different financing options are presented next. The section concludes with a discussion of the legal issues surrounding the transition to a digital department.

The technical section starts with some basic computing and image-processing information and then focuses on digital imaging. Image acquisition and compression raise issues that were not experienced with film; these are covered in separate chapters. Various PACS architectures are presented, along with their practical differences. Basic topics on hardware and software choices, such as networking and servers and operating systems, come next. Clinical storage techniques merit a chapter of their own, as storage can be a large part of the operational cost of running a PACS. Next explored are image displays, as they are the main component of the PACS viewing station that differentiates it from a regular desktop PC. Digital mammography, the most challenging modality to bring into the digital world, also merits a chapter of its own, which discusses the technical requirements for this special modality. The section ends with the topic of web distribution, which may or may not be built into a PAC system.

The clinical section deals with some of the same topics mentioned above, such as PACS viewing stations and digital mammography, but from a clinical perspective. Also discussed in greater depth are three-dimensional imaging, speech recognition, and physician order entry systems. Teaching files and education are covered as well. Teleradiology, what it means today and what it might mean tomorrow, concludes the last section.

It is our hope that you will find this book rich with ideas and information that you can use as you enter the digital transformation of radiology.



# **ADMINISTRATIVE ISSUES**

GORDON SMITH

# INTRODUCTION TO RIS AND PACS

---

---

GORDON SMITH

In the current marketplace, forces acting upon the radiology practice are mandating the conversion from the analog paper- and film-based systems to a purely digital department. These forces range from market competition to demands from the referring base and, most prominent, to the need to become more efficient to balance the losses from the steady decline in reimbursement rates.

The efficiency driver has the greatest direct impact on the practice due to the reduced reimbursement rates for procedures, which is driving practices to increase productivity just to break even. However, if a practice is already at maximum capacity and costs are not being covered, the practice is in for a tough decision regarding increasing efficiency. Does the practice make the investment in technology to help increase efficiency, or does the practice add another radiologist with the hope that the increased overhead will be offset by the increased volume? These are decisions that practice management often faces. However, in today's market the problem is compounded further by the lack of available human resources (radiologists) to correct the problem, thus driving practices to the technological solution. The administrative end (billing) has been forced into being digital by

the Health Insurance Portability and Accountability Act (HIPAA), which requires by law that all submissions be in digital format; those that are not receive an automatic penalty. The efficiency driver, along with the needs for increased quality, clinical effectiveness, and meeting the pressures of market competition, should not be perceived as forcing a new way of practicing radiology. The new digital world should be seen as an opportunity to take a practice into the 21st century and to provide the patient with clinical services that could never be provided in the analog world, such as three-dimensional (3-D) reconstruction.

This decision will be one of the most important ones that will be made by the practice. It will have an overwhelming impact upon the way you do work, where you do work, and the culture of the department in which you work. The decision to go digital is the decision to take your current practice apart and define what is good and what is bad. This is an opportunity to leverage what you do well and correct what you do poorly.

What exactly is meant by being digital, and what is needed to accomplish becoming digital? Becoming digital simply means that where there is currently a physical element that is used to perform the management of information to run the practice, that element is changed into an electronic format. Schedules, tracking forms, film jackets, and reports are all produced digitally. How does technology address the pressures defined previously?

Imaging technology improves efficiency through the use of:

- ▶ Information management
  - ▶ Radiology information system (RIS) deployment
  - ▶ Digital modality deployment
  - ▶ PACS deployment
- ▶ Computer-aided diagnosis
- ▶ Remote access

Quality issues are addressed by:

- ▶ Instantaneous access to priors
  - ▶ Deep online clinical archive
- ▶ Online diagnostic information
- ▶ Subspecialty collaboration

Clinical effectiveness is enhanced by:

- ▶ Better data
  - ▶ Direct and computed radiography
- ▶ More data
  - ▶ Multidetector computed tomography
- ▶ New data
  - ▶ MRI
- ▶ Same data but more information
  - ▶ Image fusion (PET/CT)
  - ▶ 3-D rendering

In addition to the demands of managing the practice and providing modern clinical care, you must meet the demands of the users of the information, your referring base. Most practicing physicians are aware of the advances in imaging technology and the advantages it can provide. Clinicians need imaging as a screening tool, and they are demanding almost instantaneous access to imaging information. Five service areas should be addressed to meet the demands of the practice's customers: accessibility, urgency, security, simplification, and service.

Accessibility is addressed through the use of Web access, which provides access to the images independent of location. The image data can be incorporated into the enterprise medical record (EMR), which provides the clinician with a single point of access to information for relevant clinical data from multiple departments. There is an emerging technology that involves the use of online collaboration.

Urgency has always been an issue; with the advent of the digital world, this need has increased substantially. The increased perception of the urgent need for data, right or wrong, is a demand that still needs to be met. The urgency for imaging information to meet the increased demand for "quick reads" and to provide instantaneous access to imaging data and interpretation can be addressed through the use of the image distribution process to flag and distribute the images digitally to the appropriate radiologist for interpretation. This process can be provided outside of the common hours of operation, as well, through the use of teleradiology or "nighthawk" services. These are interpretation services that are provided digitally by an off-site radiologist with a report provided to the clinician in a matter of a few hours instead of the next day.

The need for security in a digital environment is paramount. The ability to secure information in the digital environment is substantially better than it is in the analog world. This is accomplished through the use of computer-level and application-level security along with the implementation of tiered access to data.

In the digital environment, access to data is provided on a need-to-know basis. In contrast, the processes of the analog world expose patient information to many individuals who should not have access to it. Tiered access to patient data is not just the best practice for a department; it also meets the information security regulations established by HIPAA.

Simplification is providing more information in a manner that does not overwhelm the clinician with too much data. Important clinical information can be communicated to the clinician by supplying just annotated key images instead of a complete study without annotations. Developing technologies that are making their presence known are the use of 3-D rendering of images and the use of multimedia reports.

The final key point that needs to be addressed is providing the referring clinician with services that improve the ability and ease of scheduling exams, increase access to the radiologists during the exam process to facilitate changes in scheduling based on the urgency of the exam, and expedite the distribution of results. Digitally based scheduling provides the clinician with quicker access to available appointment times, and electronically submitted requests are less prone to being lost. The service that has the greatest impact is the ability to present a Web-based self-scheduling interface. This type of interface benefits both the clinician by offering more control over the ordering process and the radiology department by allowing structured input that can increase the capture of the correct CPT and ICD-9 codes for each exam.

We have discussed the areas within and outside the department driving the move toward or the expansion of a digital department. Next are the components of the digital department. At the heart of the digital radiology department exist two main computer systems: the radiology information system (RIS) and the picture archiving and communication system (PACS). The RIS encompasses many text-based computing functions including transcription, reporting, ordering, scheduling, tracking, and billing. PACS deals with image-based computing functions such as acquisition, interpretation, storage, and local image distribution.

## THE RADIOLOGY INFORMATION SYSTEM

The RIS is the nervous system of the digital department (Figure 2.1). Every aspect of the digital department relies in some manner on the RIS. The RIS drives the workflow of the information of the department. It is responsible for scheduling orders, capturing relevant clinical information about an exam and providing this clinical information only to areas of the department that require it, preparing prior exams if needed, and providing the PACS with the information it needs to perform its role. Once an image is captured, the RIS and PACS work together to provide the radiologist with the necessary information to interpret the exam and to deliver the report to the clinicians. In addition to the clinical functions of the RIS, the system manages billing for the exams and provides the necessary data to support management reporting for the department.

Scheduling is where the process begins. The scheduling step kicks off a number of events within the RIS to prepare for an exam to be performed. The process of scheduling an exam captures the appropriate clinical information to determine the exam to be performed. It is also the point in the process at which the patient demographics are captured. Accurate patient information is required for proper acquisition of relevant prior exam information and to ensure that billing can be performed correctly.

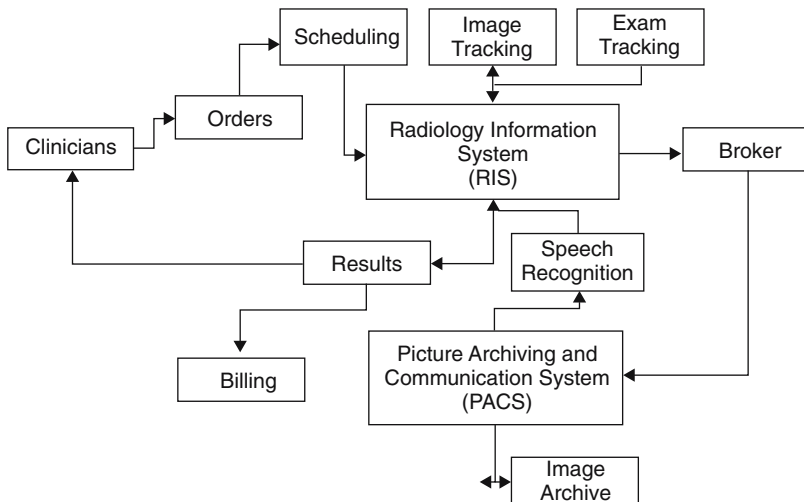


FIGURE 2.1

Traditional radiology digital infrastructure.



The scheduling process is where a majority of the data errors occur within the system. Input data errors at this point will for the most part eliminate any operational efficiency gained by moving to the digital department. The three traditional interfaces of scheduling provide various levels of control over the integrity of the data being put into the system. In the hospital information system (HIS) scheduling method, traditionally the lowest level of accuracy exists due to the lack of control over managing the sources of data for the HIS. The second most inaccurate is the manual scheduling method. This is the scheduling of exams within the department. In this area you do have control through programs that increase accuracy such as competency-based training which is discussed in Chapter 25. The Web-based scheduling method is the most accurate because there is more control over the incoming data, assuming there is a structured method of gathering the required information.

Once the exam is scheduled, what happens with that data? The most beneficial processes are the acquisition of relevant prior exam information and the validation of patient information. This information is used in the pre-fetching of prior films, either by moving studies in the PACS from long-term storage to near-line cache or by the creation of pick lists for the film library. This pre-fetching process reduces the time needed to gather the appropriate prior studies, which will allow for the finding of lost films before the time of interpretation, thereby increasing the quality of care. Additionally, it improves the process of protocolling exams by allowing for the process to occur well before the exam. This further increases efficiency by reducing the number of interruptions in the workflow to protocol the exam at the time of the exam.

The RIS provides the technologist and the radiologist with relevant information for performing the exam. The technologist interacts with the RIS either by receiving a paper request or, in the digital environment, by checking an electronic worklist that provides the details of the exam, including the protocol assigned by the radiologist. During this process the RIS tracks the exam status and the patient. This information is used to manage the rest of the exam transaction.

When the exam is complete and the images are ready for interpretation, the RIS and PACS interact to validate that the images acquired match the order information. Once the images are determined to be valid, the exam data are routed to populate worklists for the appropriate radiology specialty for interpretation. This routing can be driven by either the RIS or the PACS, and there are different schools of thought regarding which is preferable. Either way, the relevant exam information is provided to the radiologist to interpret the study.

The report in the digital department is captured by speech recognition, and, after it is signed by the radiologist, it is delivered to the appropriate destinations. These are primarily the requesting clinician and the billing office; delivery methods may include fax, secure e-mail, and, of course, regular mail. The RIS also serves as an archive for all the exam data, including the report. Thus, the RIS is the backbone for almost all the clinical operations of the department.

Beyond controlling the exam management process for the department, the RIS can also provide a wealth of information for improving the operations management for the department. The exam mixes, volume, turnaround times, and billing data can facilitate the measurement of key department metrics. Some examples of these measurements are throughput for the department overall and by area or device. Report turnaround times and the changes in efficiency due to changes made in department processes are also key operations measurements.

The next section introduces key digital imaging technologies and the fundamentals of PACS operations, followed by a discussion of how the RIS and PACS are tied together to form a single powerful core for “being digital.”

## **THE PICTURE ARCHIVING AND COMMUNICATION SYSTEM**

The PACS by far is the portion of the digital department that gets the most attention, and rightly so, for this is the where the bulk of the work is performed. This is also the area of the department where the greatest change occurs. It is important to understand the fundamentals of the functions and basic technology of PACS. Many other chapters in this book explore the details of each area discussed in this introduction. Figure 2.2 illustrates a general overview of the basic functions and relationships of the PACS core elements.

### **ELEMENTS OF A PACS**

Following are the basic elements of a PACS:

- ▶ Image acquisition
- ▶ PACS core
- ▶ Interpretation workstations

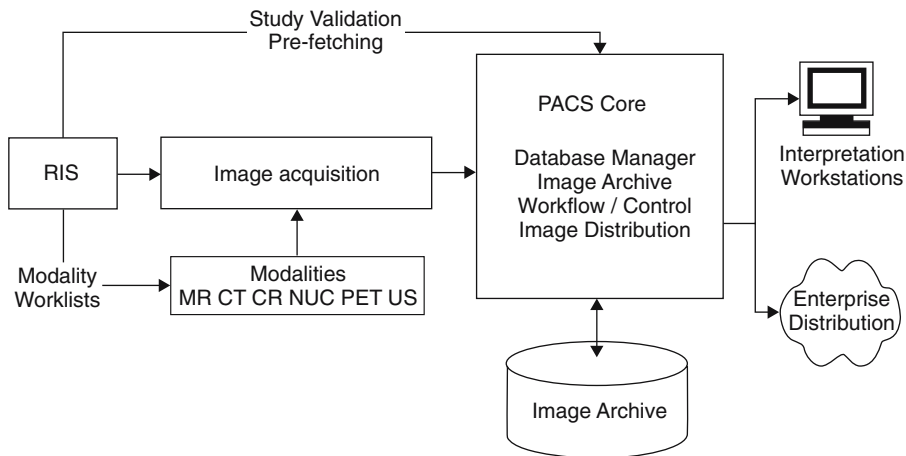


FIGURE 2.2

Database management.

## IMAGE ACQUISITION

Image acquisition is the first point of image data entry into a PACS, and, as a result, errors generated here can propagate throughout the system, adversely affecting clinical operations. General predictors for successful incorporation of image acquisition devices into a digital imaging department include ease of device integration into the established daily workflow routine of the clinical environment, high reliability and fault tolerance of the device, simplicity and intuitiveness of the user interface, and device speed.

Digital image acquisition from the inherently digital modalities such as computed tomography (CT) and magnetic resonance imaging (MRI) makes sense. There are two methods for accomplishing this: direct capture and frame grabbing. Direct digital interfaces allow capture and transmission of image data from the modality at the full spatial resolution and bit depth or gray scale inherent in the modality, while analog (video) frame grabbers digitize the video signal voltage output going to an image display, such as a scanner console monitor. In the frame-grabbing method, as in printing an image to film, the image quality is limited by the process to only 8 bits (or 256 gray values). This may not allow viewing in all the appropriate clinical windows and levels or contrast and brightness settings.

For example, when viewing a CT of the chest, one may wish to view in lung window and level settings and in mediastinal and bone windows and levels. Direct capture of the digital data will allow the viewer to dynamically window and level through each of these settings on the fly (in real time) at the softcopy display station. To view all appropriate window and level settings on film, several copies of the study would have to be printed, one at each window and level setting. If one performs the analog acquisition or frame grabbing of the digital data, the viewer can only window and level through the 8 bits captured, which will not be sufficient. Thus, direct capture of digital data from the inherently digital modalities is the preferred method of acquisition. Methods for digital image acquisition of the conventional projection x-ray include devices such as computed radiography (CR) or imaging with photostimulable or storage phosphors and digitization of existing analog film, as well as direct digital detectors falling under the general heading of digital radiography (DR). Digital acquisition of images already on film can be accomplished using a variety of image digitization devices or film scanners. These include the infrequently used analog video cameras with analog-to-digital converters (ADCs), digital cameras, charge-coupled devices (CCDs), and laser scanners.

**FILM DIGITIZERS** Film digitizers will still be necessary even in the all-digital or filmless imaging department, so that film images from outside referrals without digital capabilities can be input into the system and viewed digitally. Film digitizers convert the continuous optical density values on film into a digital image by sampling at discrete, evenly spaced locations and quantizing the transmitted light from a scan of the film into digital numbers. Several types of film digitizers exist today, with some used more frequently than others in PACS and teleradiology applications.

A commonly used film scanner for PACS is the CCD or flat-bed scanner, which uses a row of photocells and uniformly bright light illumination to capture the image. A lens focuses the transmitted light from the collimated, diffuse light source onto a linear CCD detector, and the signal is collected and converted to a digital electronic signal by an ADC.

The laser scanner or laser film digitizer uses either a helium-neon (HeNe) gas laser or a solid-state diode laser source. The laser beam is focused by lenses and directed by mirror deflection components, and the light transmitted through the film is collected by a light guide, and its intensity detected by a photomultiplier tube, converted to a proportional electronic signal, and digitized in an ADC. They are semi- or fully automatic in operation and are currently the scanner of choice for PACS applications.

**COMPUTED RADIOGRAPHY** Computed radiography refers to projection x-ray imaging using photostimulable or storage phosphors. In this modality, x-rays incident on a photostimulable phosphor-based image sensor or imaging plate produce a latent image that is stored in the imaging plate until stimulated by laser light. This released light energy can be captured and converted to a digital electronic signal for transmission of images to display and archival devices. Unlike conventional screen-film radiography in which the film functions as the imaging sensor, or recording medium, as well as the display and storage media, CR eliminates film from the image-recording step, resulting in a separation of image capture from image display and image storage. This separation of functions allows optimization of each of these steps individually. In addition, CR can capitalize on features common to all digital images, namely, electronic transmission, manipulation, display, and storage of radiographs.

Computed radiography can be used for the digital image acquisition of projection radiography examinations into a PACS. As a result of its wide exposure latitude and relative forgiveness of exposure technique, CR can improve the quality of images acquired in difficult imaging situations, as in portable or bedside examinations of critically ill or hospitalized patients. As such, CR systems have been successfully used in the intensive care unit (ICU), in the emergency room (ER) or trauma center, as well as in the operating room (OR). CR can also be cost-effective for a high-volume clinic setting or for a low-volume setting for input to a teleradiology service, and it has successfully reduced retake rates for portable and other examinations.

Technologic advances in CR hardware and software have contributed to the increased acceptance of CR as the current counterpart to conventional screen-film projection radiography, making its use for clinical purposes more widespread. CR is compatible with existing x-ray equipment, yet separates the functions of image acquisition or capture, image display, and image archiving, as opposed to traditional screen-film radiography, in which film serves for image capture, display, and archival medium. This separation of functions by CR enables optimization of each of these steps individually. Potential benefits are improved diagnostic capability via the wide dynamic range of CR and the ability to manipulate the exam through image processing as well as enhanced radiology department productivity via networking capabilities for transmission of images to remote digital softcopy displays and for storage and retrieval of the digital data.

**DIGITAL RADIOGRAPHY** In addition to the current clinical devices for digital image acquisition of projection x-rays, such as CR or imaging with

photostimulable or storage phosphors, are the direct digital detectors, which fall under the general heading of DR.

Digital radiography refers to devices in which the digitization of the x-ray signal occurs within the detector itself, providing an immediate full fidelity image on a softcopy display monitor. Compare this with CR, which uses a photostimulable phosphor imaging plate detector in a cassette design that must be processed in a CR reader following x-ray exposure for conversion to a digital image. Digital radiography devices may be classified as direct or indirect based on their detector design and conversion of absorbed x-rays into an image. Note that the acronym DR may be used by some to refer to direct radiography, also called direct digital radiography (DDR), a subset of digital radiography in which x-ray absorption within the detector is converted into a proportional electric charge without an intermediate light conversion step.

Recent technologic advances in CR and DR have made digital projection radiography more prevalent in the clinical arena; CR currently has a greater clinical installation base. Hardware and software improvements in detector devices, in image reading-scanning devices, in image-processing algorithms, and in the cost and utility of image-display devices have contributed to the increased acceptance of these digital counterparts to conventional screen-film radiography.

**COMPARISON OF COMPUTED RADIOGRAPHY AND DIGITAL RADIOGRAPHY** Digital radiography devices have more efficient detectors, offering direct energy conversion of x-ray for immediate readout. These detectors have all the benefits of digital or filmless imaging. But cost is still high because detector production is difficult and expensive, and DR is a one-room-at-a-time detector. DR may be cost-effective in high-volume sites and for imaging examinations requiring higher spatial resolution, such as upright chest exams and bone work.

The ease of use, straightforward integration, and proven reliability of CR systems over DR systems adds to the attractiveness of CR as a replacement for screen-film systems in general radiography in a PACS digital-imaging network. Digital radiography, however, has potential for excellent image quality available immediately at the time of exposure. It is likely that CR and DR devices will coexist for some time.

While CR and DR have been used for general radiography for many years, it is only recently that they have been successfully applied to mammography. Furthermore, meeting the cost competitiveness of screen-film systems is difficult unless film printing is eliminated from the cost equation. Future improvements in image-processing algorithms, with a better

understanding of optimum display settings for softcopy viewing, have the potential to greatly facilitate and standardize softcopy reading of digital projection radiographs and further the acceptance of CR and DR in the clinical arena. There is more detailed explanations of these technologies in later chapters.

## PACS CORE

Once the images have been acquired, they need to be managed appropriately to ensure that storage, retrieval, and delivery all occur without error. The PACS should also guarantee that the images are stored using long-term methods that meet the minimum legal obligations for the retention of images for the given state. Additionally, they need to be delivered for interpretation in a timely manner. These requirements are satisfied by the PACS core.

The PACS core consists of the following:

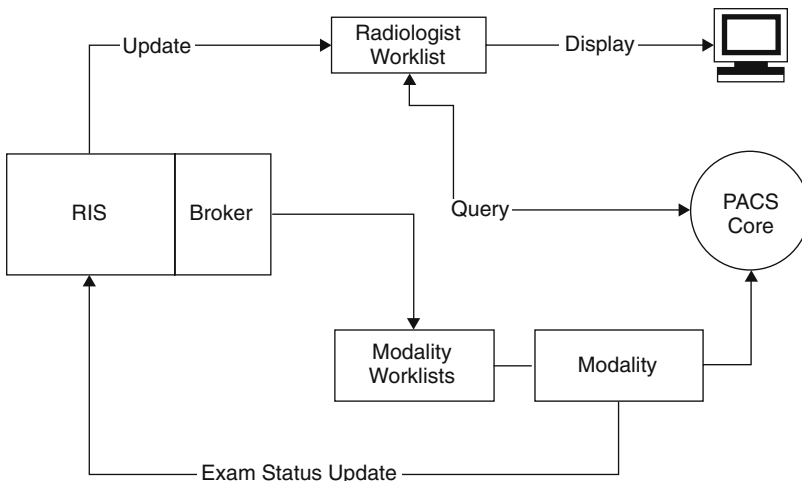
- ▶ Database manager (e.g., Oracle, MS-SQL, Sybase)
- ▶ Image archive (e.g., RAID, Jukebox)
- ▶ Workflow/control software (image manager)
- ▶ RIS interface

The database manager is the heart of the PACS. The relationship between the image and the storage location is stored and managed within the database along with all the relevant data required to retrieve the image (see Figure 2.3). The database manager must also be able to retrieve images for a given patient's current or prior exams when queried by the RIS or other outside systems. The types of queries that the database responds to are defined by the Digital Imaging and Communications in Medicine (DICOM) standards. DICOM and these associated properties will be discussed later in this section. The database architecture is typically relational, utilizing Oracle or Microsoft SQL Server.

The image archive works in conjunction with the database manager by storing the images in a highly available system to provide online images for nearly instant retrieval and long-term storage to meet retention regulations and disaster recovery. The images available for nearly instantaneous access consist of the recently acquired exams and those that were pre-fetched (requested by the RIS) from the previously scheduled exams and pulled from long-term storage. This storage is often referred to as online or near-line. It

consists of redundant array of inexpensive disks (RAID), where the images are stored on hard disk and are readily available when the database manager makes a request for the images to be distributed. The second tier of storage is referred to as long-term storage; this is intended to be the location for studies that need to be kept on hand but are not needed for immediate access. The platform for this type of storage ranges from tape and optical jukeboxes to storage-area networks. As the cost of RAID continues to plunge, it is unclear how much longer there will be a need to utilize this secondary storage for clinical image retrieval at all; it may evolve into a disaster recovery system only.

Image management (workflow control) is the role of the core that is the most visible and drives the functionality of the PACS. The image management process is where the data from the RIS and the data from the core meet and are managed in a number of different ways. Image management/workflow of the PACS determines where and how images are routed throughout the system to ensure they are stored appropriately once received from the imaging devices. Image management is also responsible for the routing of exams to the appropriate location, responding either to the PACS database or to the RIS. In addition to managing the storage and distribution of images, the image manager is also the area within the PACS where the system administrator has tools to correct for system and data errors to ensure data integrity.

**FIGURE 2.3**

RIS-driven PACS workflow.



The RIS interface is where the two principal computing systems within the digital department come together. This interface is responsible for passing the appropriate scheduling and exam information to the core to facilitate the pre-fetching of prior exams, the validation of the demographic/exam information stored within the image prior to storage in the core and subsequent distribution. Depending on the configuration and architecture of the PACS-RIS relationship, this interface is managed with or without a broker.

The role of the broker is to negotiate between the PACS and the RIS the data required and perform any data format conversions that may be required. The use of a broker is not preferred since it is another point for management and for failure, in addition to the limited functionality it may impose on the system.

PACS-RIS architectures are moving away from the use of brokers and are either combining the RIS into the PACS or vice versa. Essentially where is the functionality of the broker embedded? The combination of the RIS-PACS is based on which system actually controls the workflow. In an RIS-driven environment, the modality and the radiologists' worklists are controlled directly by the RIS, with the PACS acting in the passive role of serving the images at the request of the RIS (Figure 2.3).

In this method, the RIS is driving workflow by drawing on its database to populate the modality worklists, therefore driving the work performed at each modality. The schedule, status of each exam, and changes in status (canceled or completed) are communicated directly with the RIS, enabling the RIS to make direct updates to its database. Once the completed status event is received from the modality and the study is validated (RIS and PACS compare information to ensure accuracy), the exam is sent from the RIS to the radiologist workstation worklist. The radiologist is then presented the exam for interpretation simply by using the data provided by the RIS to query the PACS and having the image displayed on the workstation. The RIS is driving the work, and the PACS is simply an image repository that is queried on an as-needed basis. This is the model that PACS architectures are moving toward since the RIS is the primary repository for all departmental data and it does not make sense to duplicate data sources.

The second type of architecture is declining in use but is still very common in legacy systems and some current PACS. This model is the design in which the broker functionality is embedded within the PACS, thereby having workflow controlled by the PACS, with the RIS acting as a passive source of data (Figure 2.4).

The workflow in this mode shows that when an exam is ordered, the request travels over a link to the RIS and is then sent to the modality to

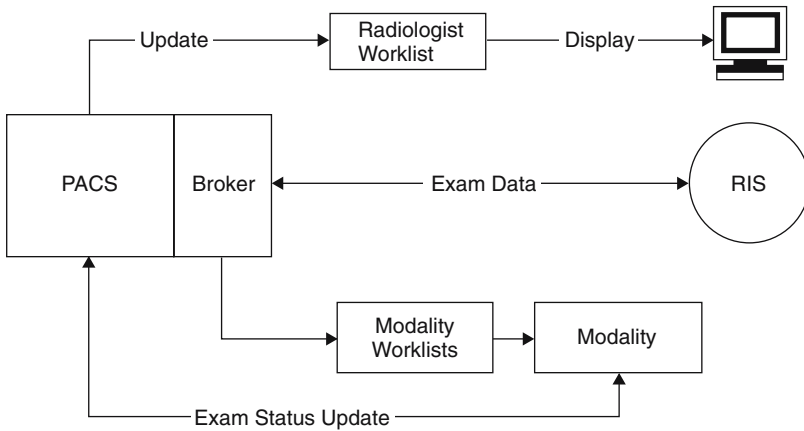


FIGURE 2.4

PACS workflow.

generate the worklist. Once the exam status changes (completed or canceled), the change is communicated to the PACS, then to the RIS. Once the PACS and RIS have exchanged information to validate the exam, the worklist for the radiologist is updated. The workstation then queries the PACS for the related images, and the process then continues as normal. As can be seen in this model, the data needed to perform the exam have several more paths to travel outside the RIS before the exam information returns to the central data repository.

The details of the PACS core designs and operations are discussed in more detail later in the book, but as demonstrated here the PACS core is the heart of the PACS. The future will bring changes in architecture (i.e., Web-based technology) and ongoing discussion of the merits of each. In many cases, when the discussion is over, there will be a new technology to discuss.

## INTERPRETATION WORKSTATIONS

Now that what goes on behind the scenes has been introduced, the tool that allows the radiologists and clinicians to interact with the data contained within the PACS needs to be introduced. The workstation is where the physician and clinician see the results of the capture of the relevant exam information within the RIS and the images acquired and stored within the PACS.

There are two general classifications of workstations: diagnostic and review. The distinguishing characteristics between them are resolution and functionality. The diagnostic workstation is the type that is used by the radiologist to perform primary interpretation of the exam. These workstations are the highest in resolution and brightness and contain the highest level of functionality. Historically, they have been dedicated to the task with the application loaded locally, in some cases, on hardware and operating systems other than the Windows/Intel (Wintel) platforms. These systems are quite expensive and require support skills that are not found in the typical hospital or clinic setting. Also, due to the expense, the number of workstations that could be deployed was limited by available capital, which, in many cases, after the purchase of the core of the PACS was quite limited. As technology has moved forward, many of these workstations have moved to the Wintel platform, which has somewhat reduced the capital cost of the workstations and enabled a greater number of workstations to be deployed. Workstation availability is an issue that needs to be taken into consideration when making the PACS purchase decision. It may be that due to limited capital resources, the potential efficiencies gained by the deployment of PACS are outweighed by the limited number of available workstations. However, there is a new trend in the industry that is quickly becoming the standard, Web-based PACS. In this case the primary interpretation workstation is any computer that meets the performance and video resolution requirements to support the interface and has network connectivity. This new technology facilitates increased usage and acceptance of the workstation in the department, thereby increasing availability of workstations and the desired efficiencies they bring. This new trend will drive many more practices into the PACS world. Also, this availability of Web-based workstations is facilitating the increased level of teleradiology solutions due to the Web-based clients that allow access from almost anywhere.

The next type of workstation is the clinical review workstation. This workstation is not as powerful as the diagnostic workstation. The difference can be in hardware (resolution), available software functionality, or both. In the past, the sheer cost of deploying diagnostic workstations made it difficult for referring clinicians to benefit from the advantages of PACS. This drove the need for a step down in the type of services provided by the workstation. In the past, many PACS vendors either created a scaled-down version of their workstation or leveraged the rising technology of the Web. The clinical review workstation allows referring clinicians to have direct access to the images. The quality of images is sufficient for the interpretation of clinicians, allowing them to review the images along with the radiology report and possibly to share those results. This area of workstations is

benefiting the most from the advent of Web-based workstations. Web-based clients allow access to the images to be distributed more widely within and outside the practice.

It is important to note that the technology behind the workstation is increasing at an incredible rate and that the hardware needed to support it is becoming ever-more accessible. As this trend continues, the penetration of PACS technology within the institution will increase, allowing the practice to realize more and more of the efficiencies that may have driven the PACS decision.

## CONCLUSION

As the practice of radiology is faced with the challenges of reduced reimbursements and the lack of both financial and human capital, some critical decisions need to be made to try to stave off the potential failure of the practice. Practices must attempt to drive their operations in the most efficient manner possible to better defend against these forces. Not only are there financial pressures; there is significant pressure from the referring base for advanced imaging services. These services can be provided in an analog practice; however, as technology continues to advance, the data will continue to increase, and this will eventually overrun the analog practice. The forces acting on the modern radiology practice dictate that a practice will have to become digital sometime to survive. The following chapters provide you with the detailed knowledge you need to start on the path of both understanding and becoming digital.

# PACS STRATEGIC PLAN AND NEEDS ASSESSMENT

---

---

LEONARD A. LEVINE

## WHY DO YOU NEED A PACS STRATEGIC PLAN?

For more than 100 years, the efficiency of radiology practices has been limited by film and film-handling activities. Picture archiving and communication systems (PACS) completely reengineer radiology practices by enabling images to be electronically viewed virtually anywhere on a clinical workstation or an ordinary PC. Film is printed on demand instead of after each and every exam by the technologists. Prior examination films do not need to be retrieved, matched to current films, distributed to radiologists, and retrieved again for refiling and storage. Radiologists' reports are not delayed due to missing films or because of inadequate hanging space for film on an alternator or view box.

PACS enables ubiquitous availability of images, resulting in improved clinical care and productivity throughout the healthcare enterprise. Patient

care is improved due to image availability and faster report turnaround and because the image dataset can be manipulated to yield more clinical information (e.g., three-dimensional [3-D] reconstruction and computer-aided diagnosis [CAD]).

Because a PACS is expensive technology that impacts the entire health-care enterprise, a strategic business plan is essential to define the costs, benefits, technical changes, and operational changes that will need to occur to make the PACS a success. In many cases, a principal goal of the PACS strategic business plan is to help secure funding for the PACS and to provide a roadmap for its implementation. Building a strategic plan often provides the first opportunity in the process to educate end users about how the PACS will impact their operations. The issues and functional requirements identified in the strategic planning process will become input to the request for proposal (RFP). PACS vendors will be required to respond in detail to the RFP about how they will meet the unique needs of your institution.

## WHAT IS IN THE PACS STRATEGIC PLAN?

A PACS strategic business plan is typically comprised of operational, technical, and financial sections. The following documentation is required to develop the plan:

- ▶ A list of all sites where images are acquired, clinically reviewed, or interpreted, which will help to define the scope of the project.
- ▶ A modality list for each site, which will identify any upgrades and associated costs that will be required to successfully integrate the imaging equipment with the PACS.
- ▶ An organizational chart, which will reveal the contacts responsible for each area that may be affected by the PACS. These stakeholders will need to be educated about how a PACS can affect their operations. In turn, they may become some of the key decision makers that will help define how the PACS should be implemented in their areas.
- ▶ Mission and vision statements to illustrate how a PACS is aligned with the organization's other strategic goals.
- ▶ Technical staffing data, throughput statistics, and hours of operation by modality, which will project productivity and capacity improvements that may result from a PACS implementation.

- ▶ Professional staffing data, including reading room locations and the degree of subspecialty interpretation, which will help estimate the number and type of primary interpretation workstations that will be needed.
- ▶ Exam volume and commensurate film expenses, which will help project film savings and estimate digital image archive needs.
- ▶ Network diagrams, including wide area networks to offsite imaging sites, to determine the adequacy of imaging dataset transmission.
- ▶ Number and type of different radiology information systems (RIS) and hospital information systems (HIS) deployed across the enterprise and whether patients have a unique medical record number across the various sites in the enterprise. Vendors should be able to explain how their PACS solution will seamlessly and effectively facilitate the flow of images across disparate HIS and RIS systems.

## PACS STRATEGIC PLANNING: OPERATIONS

There are six recommended components to the operational section of the plan: (1) alignment with other strategic initiatives, (2) a PACS readiness assessment, (3) a basic phased implementation plan, (4) a PACS operational impact analysis, (5) a market assessment, and (6) a concluding section that illustrates how PACS can leverage existing human and capital resources to meet future demand for radiological services. Each component is discussed below.

### ALIGNMENT TO STRATEGIC GOALS AND OBJECTIVES

The first phase of the strategic planning process is to outline the strategic goals and objectives for the project. Strategic goals may be logistical in nature, for example, moving into a new “digitally ready” department. There may also be strategic business objectives for the PACS, for example, aligning the PACS plans for a private radiologist’s practice with the PACS plans for the hospital served. Other objectives may include protection or expansion of market share and professional recruitment. Information systems (IS) initiatives will need to be integrated with the PACS; for example, deployment of an electronic medical record or a new RIS will need to be aligned with the PACS strategy.

## PACS READINESS ASSESSMENT

A PACS readiness assessment consists of an assessment of the organizational behavior, technical infrastructure, and existing operations. Assessment of organizational readiness includes a candid evaluation of whether the current leadership, departmental culture, and available support personnel are ready to implement and manage the change processes associated with a PACS implementation. A leader with prior PACS experience and project management skills is recommended in order to develop the institutional “vision” for a PACS. Operational and technical aspects of the organization’s readiness to implement PACS may be gleaned from interviews with key stakeholders and decision makers.

Since a strategic business plan is often the vehicle that is used to secure funding for a PACS, interviews conducted during the strategic planning process represent an opportunity to both educate and build support for a PACS throughout all levels of the institution. Consultants with direct PACS experience can provide the education that is necessary for building support for the strategic business plan. The results of the interviews will drive the development of the PACS implementation plan. The specific objectives of each interview are discussed below.

Radiology administration, including the administrative director and the chief of radiology, can identify the drivers for PACS from a radiology perspective, including service and productivity issues. Radiology administrators can assist with computed radiography (CR) or digital radiography (DR) deployment strategies. With PACS, on demand printing often replaces printing each study after acquisition. Ironically, this may result in the need for more printing capacity at certain locations (e.g., the film library). Radiology administrators can assist with the deployment strategy for Digital Imaging and Communications in Medicine (DICOM) compliant printers.

Hospital executives, including the chief executive officer, chief operating officer, chief information officer, chief financial officer, and vice presidents, can identify PACS drivers from an institutional perspective. The chief financial officer should discuss competing capital-intensive initiatives and describe the requirements for the capital decision-making process. These executives can also assist in setting the financial goals for reducing film printing. PACS does not completely eliminate the need to print film, and it is important to plan for the ability both to print on demand and to print film as a backup strategy in the event the PACS is down.

Referring physicians can define their service level expectations from radiology and provide insight into how a PACS can be optimized throughout the enterprise to improve patient throughput and care. Referring



physician support is critical to the success of a PACS project. High-volume clinical areas, such as the intensive care unit (ICU) and emergency room, may have special requirements for electronic viewing of images. These areas are accustomed to having direct access to film and may require workstations comparable to those deployed in radiology, including dual monitors for side-by-side comparison of current and prior exams, with high-brightness, high-resolution, and grayscale monitors. The plan should include estimates of the associated costs of these workstations.

Radiologists should describe how the current analog environment negatively affects their ability to provide optimal service and how the PACS will mitigate or remedy the situation. A privately owned radiology practice may have PACS goals that are inconsistent with the hospital's. Outside consultants may be able to offer solutions that are more closely aligned with both the hospital's and the radiologists' goals. Professional staffing and interpretation practices will help to determine the location and type of diagnostic workstations and monitors that should be deployed.

The role of film librarians will be significantly different in a PACS environment. Film librarians will be managing electronic images, printing on demand, and managing the unread case list. They will need to be PC proficient and able to understand complex decision rules, particularly during the transition period from analog to digital images, to ensure that all relevant prior images, regardless of medium, are available at the time of interpretation. Therefore, the skill sets of this group need to be assessed. The film librarians can also provide estimates of the impact of lost film on patient care and throughput.

Technical managers and technologists should describe workflow, particularly as it relates to film-handling activities. The interviews of technical managers and technologists along with observations should be used to quantify the expected positive impact of a PACS on technologists' productivity and equipment utilization.

The interviews conducted during the PACS readiness assessment should be synthesized to highlight common themes about the PACS strategic goals and objectives.

## IMPLEMENTATION PLAN

The third component of the operational strategic plan is a high-level implementation plan based on a phased approach to the PACS deployment. The phases of the implementation plan are typically based on strategic financial and operational PACS drivers that were identified in the interview process.

The plan should include potential schedules for CR or DR deployment, digital archive building, transitioning radiology operations to PACS, and enterprise-wide PACS deployment of electronic images as well as plans to reduce film expenses.

The implementation plan should include a description of each major phase of the project, including its primary objectives and benefits, estimated time frame, dependencies, and costs (Table 3.1). The implementation plan sets the parameters on when PACS-related expenses will be incurred and when film and film-related savings will be realized. Therefore, the implementation plan is a prerequisite to developing the PACS return on investment (ROI).

## PACS IMPACT ANALYSIS

The fourth component of the operations section is the PACS impact analysis. The impact analysis includes estimates of the expected PACS benefits that may be realized through improved workflow and clinical care at each phase of the PACS deployment. Much of the emphasis will be on the workflow and productivity improvements that occur within radiology. Enterprise-wide PACS benefits should be discussed, although they are generally more difficult to quantify.

### PACS IMPACT ON FILM AND FILM-RELATED EXPENSE

Film savings are usually the most quantifiable cost savings used to justify investments in PACS. Film savings include not only the cost of the actual film used but also all film-related costs, such as chemistry, film handling, storage, transportation, reprints and retakes, and film processor maintenance.

Film cost avoidance projections should include estimates of additional film that would otherwise be used to meet expected exam volume growth. Estimates should also include projected film increases that would otherwise occur with the use of newly deployed multislice computed tomography (CT) scanners and other image-intensive modalities.

Cost savings can also include an anticipated reduction in the film library workforce, but many plans attempt to estimate this more aggressively than is reasonable. It is important to remember that the need to manage and retrieve prior films may persist for several years after the deployment of PACS, and the need to print some films on demand may persist indefinitely.

**TABLE 3.1**  
PACS Implementation Plan

<b>Phase</b>	<b>Key Objectives</b>	<b>Primary Benefits</b>	<b>Time Frame</b>	<b>Dependencies</b>	<b>Estimated Cost</b>
1	Develop and communicate PACS strategic plan Install telecommunications links Upgrade CT, MRI, US, and NM modalities to DICOM, worklist-enabled Prepare sites for viewing locations Install networking (wiring and switches) Prepare computer room Install PACS core system, including Web server Implement RIS interface Install speech recognition server and RIS interface Train radiologists	Set stage for remaining phases	3–6 months	Vendor selected Contract signed Funding approved RIS stable	\$1,600,000 (PACS) \$220,000 (SR)
2	Bring CT, MRI, US, and NM modalities online Integrate modalities with PACS Train technologists Acquire images QA image quality	Improve radiologist productivity Reduce lost film	6–12 months	Phase 1 completed	\$600,000 (PACS)

*(Continued)*

**TABLE 3.1**  
PACS Implementation Plan (Continued)

<b>Phase</b>	<b>Key Objectives</b>	<b>Primary Benefits</b>	<b>Time Frame</b>	<b>Dependencies</b>	<b>Estimated Cost</b>
	Implement softcopy interpretation of CT and MRI modalities Install diagnostic workstations Train radiologists Implement limited filming/print on demand Integrate laser printers with PACS Develop and implement limited filming program	Improve report turnaround Improve access to digital modality images in radiology			
3	Install CR Install CR equipment Integrate CR with PACS Train technologists Acquire images Perform QA on image quality	Reduce lost film Improve access to plain film images in radiology	3–9 months	Phase 1 completed CR funding approved PACS support plan	\$230,000 (PACS) \$1,000,000 (CR)

4	<p>Implement softcopy interpretation of CR                  Install clinical review workstations (ICUs and EDs)                  Train clinicians                  Limited filming/print on demand                  Develop and implement limited filming program</p>	<p>Improve radiologist productivity                  Improve report turnaround                  Improve access to plain film images in ICU and ED</p>	6–12 months	Phases 1, 2, and 3 completed	\$250,000
	<p>Integrate imaging Web user interface with electronic medical record                  Image distribution to clinicians                  Market capability                  Inform clinicians of availability and hardware requirements                  Assist with installation                  Train users</p>	<p>Provide access to images to referring clinicians                  Improve service                  Reduce film cost</p>			
	<p>Implement limited filming/print on demand                  Expand limited filming program</p>				

CR indicates computed radiography; CT, computed tomography; DICOM, Digital Imaging and Communications in Medicine; ED, emergency department; ICU, intensive care unit; MRI, magnetic resonance imaging; NM, nuclear medicine; PACS, picture archiving and communication systems; QA, quality assurance; SR, speech recognition; US, ultrasound.

In addition, films from modalities that cannot easily be converted to digital format (e.g., mammography) must continue to be managed by the film librarians.

Film-related cost savings must be related to the implementation plan. For example, if deployment of enterprise-wide electronic distribution of images is delayed until the latter phases of the project, you will not be able to stop printing film until after this phase of the project is complete. If you are phasing in PACS one modality at a time, film savings will be realized only for the modalities implemented and only if electronic distribution is available for those modalities.

Some areas are particularly challenging for electronic distribution and may need to continue to print film until the operational issues unique to these areas are addressed. For example, orthopedists may need to be able to perform measurements on the images for comparison with prosthetic devices. If the PACS does not provide this capability, the orthopedists will need to continue to receive film. Deploying workstations in operating rooms (ORs) may be a challenge due to space constraints, the bright lighting conditions, and requirement for a sterile operating field, resulting in continued film printing for this area. Digital mammography places some extraordinary demands on PACS due to the large image sizes and requirement for high-resolution workstations. As a result, digital mammography is not widely deployed at the current time.

### PACS OPERATIONAL IMPACT ANALYSIS

The PACS impact analysis should include an analysis of technologist workflow in order to quantify the productivity and capacity opportunities that may result from a PACS implementation within radiology. Productivity and capacity improvements are commensurate with the amount of film-handling activities that are eliminated through a PACS implementation. The total time required for these activities can be estimated based on a percentage of total work time spent on performing them or on a per-exam basis. The time required for film-handling activities can then be annualized to hours based on the number of studies performed. The film-handling time estimates may be obtained through any combination of interviews, heuristics, or observations. Examples of film-handling activities include but are not limited to sneaker-nets, film printing, matching priors, refilming, searching for lost films, and slow teleradiology systems. Increases in equipment capacity due to faster throughput through elimination of film-handling activities should also be estimated.

The recouped equipment capacity may not be commensurate with the technologist productivity gains, since adjustments for multi-technologist staffing and other factors may be required. For example, if two technologists are staffing a CT scanner, patient throughput may be reduced if one technologist is processing film while the second technologist is focusing on patient throughput.

Capacity opportunities and productivity improvements during the off shifts should be excluded if the demand is exceptionally low during those periods. Always include benefits and shift differentials when translating the hours saved to salary avoidance.

## PACS IMPACT ON MARKET SHARE ANALYSIS

The competitor analysis includes a definition of the immediate radiology market, a description of competitors, and the status of their PACS implementations. The strategic plan should note any PACS advertising done by local competitors in order to highlight any potential threat to market share. Complete the threat assessment with a description of the market share at risk if PACS is not implemented at your institution. Typically, estimates for the at-risk market share are provided by senior hospital executives and the hospital's marketing department. The goal of projecting the market-share-at-risk assessment is to estimate the amount of additional revenue that may be shifted to or from a competitor that has respectively superior or inferior radiology services.

The market share analysis is important because enhanced revenue from PACS has a far greater impact on ROI analysis than do film and film-handling savings alone. However, revenue assessments are difficult to make and are based on soft assumptions about current and future market share. Therefore, it is imperative to include the most senior levels of institutional and radiology administration in generating the assumptions that lead to expected PACS-related revenue changes. It is ultimately up to this same team to decide whether projected revenue enhancements should be included directly in the ROI. Projected revenue enhancements can often turn a negative PACS ROI positive. Even if potential revenues are not included in the ROI, they should be discussed in the strategic plan to illustrate the point that without a PACS, current revenue sources may be at risk to PACS-enabled competitors who may be capable of delivering electronic images to the clinicians' desktop computers.

The example in Table 3.2 illustrates the point. Community Hospital knows its current annual volume and estimates that it currently captures 60%

**TABLE 3.2**  
Potential Shift in Market Share Exam Volume

<b>Modality</b>	<b>Community Hospital Volume (est. 60% market share)</b>	<b>Estimated Total Outpatient Market Volume</b>	<b>Outstanding Market Share (40%)</b>	<b>5% Potential Market Shift in Volume to Community Hospital</b>
CT	13,470	22,450	8,980	449
Diagnostic	51,174	85,290	34,116	1,706
MRI	4,718	7,863	3,145	157
US	5,964	9,940	3,976	199
<b>Total</b>	<b>75,326</b>	<b>125,543</b>	<b>50,217</b>	<b>2,511</b>

CT indicated computed tomography; MRI, magnetic resonance imaging; US, ultrasound.

of the outpatient radiology business in its target market area. Based on its known volume and its estimate of current market share, Community Hospital estimates the total market volume and outstanding market share. Senior administrators then discuss the referring physician community, competition, and advantages of being the first hospital in the area to deliver images to the referring physicians' desktop computers. They agree that it is reasonable to assume a potential 5% shift in the outstanding market share to Community Hospital if digital image service is available and effectively marketed to the referring physician base. Even with a current outpatient volume of only 75,326, a 5% shift of the outstanding market share to Community Hospital would result in an additional 2,511 exams annually.

To estimate potential revenue, Community Hospital multiplies the potential outpatient exam volume by its charge per exam and then by its average collection rate. The result is the potential revenue opportunity that might be available by effectively marketing a PACS in a competitive market. Table 3.3 illustrates the calculation for potential revenue. Table 3.3 indicates that a 5% shift in market share could potentially lead to over a half a million dollars in additional revenue. The model assumes that the mix of exams will remain the same.



The next question typically asked is, “What additional resources would be required to accommodate the potential additional business?” The answer lies in the equipment and technologist productivity gains that may accrue from a PACS.

## MEETING FUTURE DEMAND NEEDS THROUGH PACS OPPORTUNITIES

Increasing the radiology equipment base is expensive, and there is a national shortage of technologists. As discussed in the technologist’s workflow sections above, hospitals and imaging centers that use technologists to process, label, package, and move and hang film will achieve the greatest increases in productivity and capacity gains as a result of a PACS implementation.

A potent argument for PACS in the strategic plan is to compare capacity and productivity gains that can be expected to accrue from a PACS implementation to the future increases in demand for radiological services due to normal growth and increased market share. Often, the additional capacity recouped by eliminating film handling by technologists will be more than adequate to accommodate the additional market share. This is a potent argument; not only because there is a potential for an increase in market share

**TABLE 3.3**

Estimated Additional Revenue Based on a 5% Shift in Outstanding Market Share

<b>Modality</b>	<b>Added Volume</b>	<b>Estimated Average Change per Exam (\$)</b>	<b>Average Collection Rate</b>	<b>Total Additional Outpatient Net Revenue (\$)</b>
CT	449	679	60	182,922.60
Diagnostic	1,706	167	60	170,921.16
MRI	157	1,200	60	113,232.00
US	199	350	60	41,748.00
				508,823.76

CT indicates computed tomography; MRI, magnetic resonance imaging; US, ultrasound.

and additional revenue generation, but because most, and sometimes all, of the revenue can be realized with little or no additional costs.

In the example in Table 3.4, the marginal or additional demand for new exams by year 5 is shown in line A. This is a combination of the demand from new market share due to improved service from the PACS and a 5% estimated annual growth rate. Line B is the total capacity that is expected to become available when film-handling activities are eliminated and when data

<b>TABLE 3.4</b>					
<b>Ability to Meet Growth and Market Demand Through Reduced Film Handling and Worklist Management</b>					
<b>Demand Projections</b>					
	<b>Plain Film</b>	<b>Ultrasound</b>	<b>CT</b>	<b>MRI</b>	<b>Nuclear Medicine</b>
Estimated annual growth	1,189	354	603	147	82
Additional market share potential (outpatients)	156	49	56	132	23
Total additional exams (A)	1,345	403	659	279	105
<b>Recoverable Capacity</b>					
	<b>Plain Film</b>	<b>Ultrasound</b>	<b>CT</b>	<b>MRI</b>	<b>Nuclear Medicine</b>
Recovered exam capacity from film handling	0	1,349	1,678	261	261
Recovered exam capacity from worklist management	12,509	590	804	109	136
Total annual potential recovered exams (B)	12,509	1,939	2,483	370	397
<b>Ability to Meet Demand</b>					
B/A	930%	481%	377%	132%	380%

CT indicates computed tomography; MRI, magnetic resonance imaging.

entry is reduced through worklist management. The last line, B/A, shows the percentage of the new demand by year 5 that can be met using existing human and physical resources. The result can be translated to cost avoidance and may be included in the PACS ROI.

Many other benefits to the enterprise that can be attributed to PACS may be difficult to quantify. Clinicians will experience improved productivity since they will no longer need to spend time looking for film or visiting the film library. Images and reports will be available in a more timely fashion, providing improved support for clinical decisions. Improved turnaround for radiological reports can reduce inpatient length of stay. Deployment of CR or DR and electronic viewing in the OR can reduce the amount of time required for film processing and therefore the amount of time required for the operation, reducing the risk to the patient and potential costly downtime of the OR.

Some PACS-related costs may also be difficult to quantify. Training and support are extraordinarily important to the success of PACS, and the resources required for effective delivery are consistently underestimated by potential PACS buyers. A marketing campaign informing referring physicians about the availability and benefits of electronic image distribution is crucial to the realization of film cost savings and maximizing the competitive advantage this can provide.

## **PACS STRATEGIC PLANNING: TECHNICAL**

The technical strategic plan includes an assessment of IS readiness to implement a PACs and an assessment of the network infrastructure and PC base with respect to supporting a PACS and electronic image distribution. The section should conclude with a discussion of PACS management issues.

### **TECHNICAL READINESS**

The chief information officer and the network engineering department can describe their existing networks, their PC support, and their general readiness to support a PACS as well as other information technology (IT) initiatives. In some cases, interviews with IT professionals can be a potent opportunity to allay any apprehensions about the impact of a PACS on their existing operations.

The interviews with the IT staff should include an evaluation of the current local area and wide area network infrastructure and its ability to

support the networking traffic that PACS is expected to introduce, with recommendations for any upgrades that may be required. In addition, it should be determined whether the existing modalities are currently networked via a private network or connected to the networking infrastructure that will be used for PACS (i.e., the hospital network), as this migration may need to be performed prior to implementation.

The installed PC base should be evaluated as to its suitability for enterprise electronic distribution. A minimum of Pentium II processor, 500 MHz with 128 MB RAM, with a 17-in monitor is generally recommended. Newer machines may be required by some PACS vendors.

Finally, suitable space for the PACS core components (e.g., a data center) must be available in order to deploy PACS. Estimates of space requirements should be included, and any needs that are not easily accommodated should be identified. Anticipated deficiencies in the reading areas should be described, and anticipated construction costs and workstation furniture should be included.

The goals of the interviews with the RIS manager are to understand the RIS-related workflow and any expected equipment upgrade paths. Workflow with regard to entry of patient demographics, scheduling, order entry, and the systems utilized for each of these steps should be reviewed to determine whether any changes will be required to support PACS workflow and functionality.

The ability of both the PACS vendor and the RIS vendor to provide the necessary integration or interfaces between systems should be evaluated and any costs associated with this task should be determined. If there are multiple (different) HIS-RIS deployed across the enterprise, then this should be documented for the vendor in the request for proposal (RFP). It is important to specify the exact version of the RIS(s) when making this determination. It is helpful to identify another site with a similar combination of RIS-PACS and contact that site as a reference.

## PACS MANAGEMENT

The strategic plan should include a discussion of the human resources required to successfully deploy and support the PACS. Consider the skill sets required and the duties that will be expected of each individual involved in the project and in support for the project. Services that are expected to be outsourced from the PACS or other vendor should be described and eventually included in the RFP. Services that will be needed on an ongoing basis

(e.g., an on-site service representative) will need to be included in the system maintenance agreement negotiated with the PACS vendor.

A downtime contingency plan is crucial to a successful implementation. Although a detailed operational plan does not need to be developed until the system implementation is in progress, it is important to include an estimate of any capital investments that will be required to support the contingency plan (e.g., backup CRs and laser printers).

## FINANCIAL ANALYSIS

The financial analysis should include an ROI analysis. All quantifiable cost savings, cost avoidance, revenue opportunities, and revenue protection estimates should be included in the analysis. The cash flows should be based on the implementation plan. For example, if CR is not expected to replace conventional radiography until year 2, then any film savings from conventional radiography should not be included until year 2.

As discussed in the operational plan, cost savings may include film savings, film library, technologist productivity, and off-site storage, including transportation. Revenue opportunities are based on estimates of increased exam volume due to competitive marketing of improved service related to PACS.

Capital expenses include the estimated costs of the PACS equipment, computed radiography, required networking upgrades, and services (e.g., project management, implementation support, and training) required to support the implementation. Capital expense estimates should also include cost estimates for system upgrades, such as additional storage purchased as required.

Operating expenses include the annual cost of the system maintenance agreement with the PACS vendor and the additional human resources (full-time equivalents) required to support the PACS.

The financial analysis should be projected for at least 5 or 6 years and can be extended further as necessary to show an eventual payback.

A Web server, typically included in most PACS implementations for the electronic distribution of images to clinicians, can also provide images to radiologists, providing remote night-call via home teleradiology. If this is a requirement, an assessment of the availability of cable modem or digital subscriber line service at the radiologists' homes should be conducted to determine the feasibility of this approach. In addition, a preliminary assessment of the infrastructure necessary to support secure access from outside

the firewall to an imaging Web server via, for example, a virtual private network server, should be included.

## **CONCLUSION**

In conclusion, an effective PACS strategic plan serves as a marketing tool to communicate the benefits of PACS to senior administration and all stakeholders and educates key decision makers regarding the scope of the issues necessary to make informed decisions. The plan must make a serious attempt to quantify expected benefits in a way that is convincing to the key decision makers in order to provide a basis for allocating the necessary funding. The plan should also identify all costs associated with converting to a PACS, including training and PACS administration. Last, a formal strategic plan establishes the basis to move forward with a comprehensive RFP document, typically the next step in the PACS acquisition process.

# CREATING THE PACS REQUEST FOR PROPOSAL AND SELECTING A VENDOR

---

---

ALAN L. SCHWEITZER • GORDON SMITH

**T**he selection of a PACS vendor that can meet the needs of the healthcare enterprise can be a complex and painstaking process. A well-written request for proposal (RFP) is a key step in this process. Although it may be tempting to short cut this process and simply request proposals and/or quotations from the vendors in which you are interested, it is important to understand that a well-written RFP should satisfy the following objectives:

1. The RFP should provide information about the site to enable the vendor to provide a solution that best matches the unique requirements of the site. This should minimize the time required to communicate requirements verbally to each vendor individually.
2. The RFP should establish a framework for contractual requirements related to system function, implementation, training, and service and support.
3. The RFP should create a format for responses that facilitate vendor comparisons. This provides a mechanism for “leveling the playing field” among vendors.

All of these objectives should all be kept in mind as guidelines when drafting the RFP. In addition to these objectives, it is useful to consider a few further guidelines.

1. The RFP should be written as a functional specification. Specifications are written as a list of requirements (e.g., “The monitor shall be blue.”) It is helpful to structure the RFP so that each requirement is defined by a single statement in a uniquely numbered paragraph. Additional clarifying language may be used to help in the interpretation of the requirement.
2. The RFP should not overprescribe or engineer the solution. It is important to distinguish between requirements and design. The key is to clearly describe your operational requirements and allow the vendor to describe how his solution meets the requirement. If you have no specific requirement regarding some aspect of the specification but rather have either a preference or a desire to simply know the specification (e.g., in order to compare it with other vendors’ offerings), it is appropriate to request the vendor to define the specification for the system being proposed.
3. Before writing the RFP, it is important to consider what your primary evaluation criteria and process for vendor selection will be (see “Vendor Evaluation and Selection Process” later in this chapter). The evaluation and selection process will be facilitated if you can structure your RFP around these criteria and include content in the RFP that will solicit responses that easily differentiate the vendors from each other.
4. It is helpful to include forms that encourage vendors to summarize and condense their responses so that you can compare vendor responses side by side. This format, however, may constrain the



responses to the extent that you may not get as much detail or explanation as you would like. It is best to provide formats for both summary and detailed responses.

## RFP CONTENT

### INTRODUCTION

The introduction should include general information about the healthcare enterprise (e.g., descriptions of each site, number of beds, medical specialties, and any plans for expansion). General information can also include a statement of the enterprise business strategy and a description of the healthcare market in the local area (e.g., population, competition, etc.). This section should also include information about the Radiology department (e.g., the imaging services offered at each site, total number of procedures), a list or table of imaging modalities at each site, and general information about the professional practice (e.g., number of radiologists in the practice, number of radiologists reading during peak hours, etc.).

This section is provided as information to the vendor and generally does not require a response.

### STRATEGIC GOALS

You will want to include a section outlining the strategic goals you hope to achieve with PACS and how the realization of these goals will contribute to strategic objectives of the enterprise as a whole. This section should communicate the expectations the stake holders have of PACS and a sense of the prioritization of these goals. This section is provided as information to the vendor and generally does not require a response.

### CLINICAL OPERATIONS OVERVIEW AND REQUIREMENTS

The vendor needs to understand the unique aspects of your clinical operations and workflow and any general requirements or expectations that you have of a PACS to support your workflow. This section focuses on those aspects of your clinical operations that you believe will either impact or be

impacted by PACS. You will want to describe both your current workflow and how you envision the workflow in a PACS environment. You should include scenarios to describe both the pre- and post-PACS workflow and request that the vendor describe how his system will either optimize this workflow or be impacted by it.

This section is intended to both provide information to the vendor and to solicit general responses from the vendor that describe features that might not easily be described in a technical specification. In general this will break the rules of defining functional specifications in simple normative statements, but this section can give the vendor an opportunity to describe features of his product offering that you may not anticipate in the technical specification and that may, in fact, provide value in your environment.

### SITE-SPECIFIC OPERATIONS

For multisite operations, characteristics of the workflow that are unique to each site should be described. If each site has a unique HIS and/or RIS or independent master patient index, this should be highlighted. Patient registration, scheduling, and exam order entry should be described for each site.

### IMAGING-MODALITY BASED OPERATIONS

For each imaging modality, the workflow description should include exam scheduling, patient registration, exam order entry, patient identification, image acquisition, quality assurance, introduction of images to PACS, and any unique requirements (by modality) for diagnostic review and reporting. Workflow and processes that are unique to handling STAT exams should be described. Any paper processes that are in place should be described with an eye towards replacing these processes with an electronic analog. Some measure of peak throughput and staffing should be provided as a part of the description. The workflow should be evaluated and described for each site and area (e.g., outpatient vs. inpatient vs. ED).

Imaging modalities described should include the following:

- ▶ Diagnostic X-ray
- ▶ Portable X-ray
- ▶ Computed tomography (CT)

- ▶ Magnetic resonance imaging (MRI)
- ▶ Ultrasound (US)
- ▶ Nuclear medicine
- ▶ Special procedures
- ▶ Mammography (if included for PACS)

## DIAGNOSTIC REVIEW

This section can include information that would detail your vision for workstation deployment. This would include how many radiologists could be reading simultaneously, from which locations, and the division of work within the department. Additional workstations to support physical proximity of radiologists to imaging services they support should also be included. The section should also describe how exams would be reported, key image presentation functions to be used, and how the report transcription and approval functions work.

## CLINICAL REVIEW

Most of the requirement for clinical review is typically addressed by general purpose PC's using a web browser to access images in the PACS. Clinicians in areas such as ED and ICU that rely heavily on imaging services and routinely make treatment decisions without the radiologists' final interpretation may want dual monitor workstations with high-brightness monitors to more closely approximate the diagnostic workstations used by the radiologists. This section should describe your expectations of the needs of the clinicians, specific medical specialties, locations, and expected deployment for clinical review workstations.

You may also wish to describe the physical locations of physicians' offices for your major referrers and how you expect to provide access to any physicians that require access to images remotely from the main facility.

## EXTERNAL SYSTEMS

Describe any other systems that will need to have access to images from the PACS, such as radiotherapy, surgical planning, or another PACS.

## TELERADIOLOGY OPERATIONS

Describe how exams would be acquired and transmitted from remote sites to the PACS or viewed remotely by a radiologist at home or a remote tele-radiology service. Describe how reports would be handled. Focus on needs and requirements, not on technology.

## TECHNICAL REQUIREMENTS

### SYSTEM ARCHITECTURE

The RFP should request the vendor to provide an overview of the system architecture and provide specific information regarding the architecture that would help to differentiate vendors' solutions.

Examples are:

- ▶ Platform (e.g., Unix, Linux, MS Windows)
- ▶ Web-based vs. Hybrid (Client/Server diagnostic workstations + Web distribution)
- ▶ All images online vs. online and nearline hierarchical storage
- ▶ Redundancy features
- ▶ Architecture for multiple sites

### CORE SYSTEM

The core system of the PACS includes all hardware and software necessary to support the acquisition of images, image storage/archive, database management, image management and image retrieval. The RFP should describe these components in a general way, specify the requirements for each, and ask the vendor to describe specifics regarding each of these.

**IMAGE ACQUISITION** The RFP should list all current and planned imaging modalities, including vendor, model number, age, software revision level, and supported DICOM services classes for each. The RFP should require the vendor to assume responsibility for the success-

ful integration of all modalities. For modalities that cannot be made “DICOM ready,” with a software upgrade, the RFP should request the vendor to propose a solution to interface to the modality. The RFP should request that the vendor describe the architecture used for image acquisition and to describe the upgrade path for adding additional modalities, for example, if additional acquisition hardware is required. The vendor should be asked to describe the mechanism by which validation of image data against RIS data occurs for modalities that do not have DICOM modality worklist functionality. The vendor should also describe if the technologist is provided with feedback when validation fails and the means to correct any exams that fail validation.

**ONLINE STORAGE** Online storage, typically RAID storage, is the primary storage component of the PACS and is used to store images that are available for fast retrieval of newly acquired studies. Current storage costs have made it economically feasible for many vendors to configure PACS so that all images are available online, expanding storage as needed to accommodate newly acquired images. The RFP should estimate projected storage requirements over the life of the PACS or provide enough information for the vendor to make this estimate. Online storage capacity is heavily dependent upon the ability of the vendor to store priors in lossy compressed format (while preserving the original uncompressed or lossless compressed image in long-term storage) and the willingness of the site to utilize this technique to reduce storage costs. If this is a desirable strategy, the RFP should state this as a requirement.

Many sites are beginning to consider an enterprise storage strategy for all their storage needs, purchasing storage directly from a storage vendor. This decision is frequently driven by PACS. If this is the direction to be taken, the RFP should specify that the PACS be compatible with the preferred storage vendor. If storage is to be purchased from the PACS vendor, the RFP should specify how much online storage is required initially.

**LONG-TERM ARCHIVE AND DISASTER RECOVERY** Long-Term Archival (LTA) storage is required for legal archive, backup, and disaster recovery. The RFP should require that the vendor specify the total capacity of the storage device used. If the site has a preference for the technology to be used (e.g., DVD, tape, content-addressed storage), the RFP should

specify this. You will want to include a requirement for the vendor to describe their disaster recovery plan including an estimated length of time to restore the system to operation and the length of time required to restore access to prior exams. In addition, the vendor should describe how offline (shelf) storage is managed.

If the online storage will not be expanded to accommodate all prior exams and the LTA is to be used as a nearline storage device, the RFP should require that the system automatically retrieve images from the LTA if unavailable in online storage in response to an ad-hoc query or selection from a worklist. In addition, pre-fetch of relevant priors should be supported and require the vendor to include a description of the algorithm used to pre-fetch prior exams.

The RFP should request that the vendor describe available image compression used in conjunction with the LTA and if compression is a requirement, the RFP should state this.

**DATABASE MANAGER** The database manager in PACS systems is utilized to store the patient and exam data, maintain pointers to the image data to permit efficient retrieval, track exam statuses (e.g., acquired, validated, and dictated), store user account information, and maintain system information (e.g., DICOM parameters for each modality). It effectively serves as the “memory” of the system with the Image Manager serving as the “intelligence” of the system.

PACS vendors typically imbed a commercial off-the-shelf database management product to implement the PACS database (e.g., Oracle, MS/SQL, Sybase) and if it is important to you to know which one, and/or you have a preference, the RFP should state this.

The database manager is a single point of failure in a PACS and if you want optimal reliability, you will want the vendor to specify redundant database servers with automatic failover. The RFP should request the vendor to describe included or optional redundancy features.

A unique feature of a PACS database is that it grows indefinitely as exams are acquired. System performance may be adequate at the time of installation, but as the database grows, if the database manager hardware and software is not specified and configured to support the potential growth in the database size, system performance can degrade over time. The RFP should require that the database manager maintain system performance for at least five years of operation.

The vendor should provide the hardware and software necessary to automatically backup the database to removable media with no human inter-

vention. The RFP should also request the vendor to describe the database restoration procedure.

The vendor's response should describe how data that was entered incorrectly can be corrected, and what tools are available to effect these corrections.

The RFP should require the database manager, in conjunction with all applications that access the PACS database, to be compliant with all regulations associated with HIPAA, including security and auditing.

**IMAGE MANAGER** The Image Manager typically handles functions related to how images are introduced and moved through the system. The RFP should request that the vendor describe these and specify features that you consider to be a requirement. Examples are as follows:

- ▶ Automatic archiving to the near-line archive
- ▶ Automatic purging of the online storage archive
- ▶ Automatic retrieval from nearline archive in response to ad hoc query
- ▶ Pre-fetch of prior exams from nearline archive
- ▶ Validation of data against exams scheduled in the radiology information system (RIS)
- ▶ Autorouting to an external device
- ▶ DICOM query/retrieve
- ▶ DICOM Copy

#### RIS INTERFACE

A robust interface to the RIS is key to supporting the overall radiology workflow. The RIS interface is necessary to support the following PACS functionality:

- ▶ Modality worklist management support for any modality that has DICOM worklist management as a feature.
- ▶ Validation of data sent to PACS from any modality by comparing key data fields in the image header against data fields from the RIS.
- ▶ Display of diagnostic reports on PACS workstations.
- ▶ Automatic pre-fetch from the nearline archive (jukebox) based on scheduled exam information from RIS.

The RFP should specify the site's RIS, including the software revision level. If scheduling is done in a different system, and support for pre-fetch is required, the scheduling system should also be specified. The RFP should require that an interface to the site's RIS be included in the proposal and request a complete description. In addition, the RFP should specify that the functions listed above be supported.

You will want to specify if you want film-based exams and their associated reports to be available in the PACS. Exams completed prior to the PACS implementation may require a historical data upload. You should specify both requirements in the RFP.

## DIAGNOSTIC REVIEW WORKSTATION

The diagnostic review workstation is by the radiologist for primary interpretation and is one of the most important components of the PACS system. Its functionality will significantly impact the radiologists' productivity and it is therefore important to carefully specify the requirements for this component. A suggested organization for specifying these requirements is as follows:

- ▶ General System Requirements
- ▶ Monitors
- ▶ User Interface and Profiles
- ▶ Worklists and Queries
- ▶ Diagnostic Report Display
- ▶ Examination Display and Arrangement
- ▶ Image Display and Paging
- ▶ Grayscale Operations
- ▶ Image Orientation, Zoom, Pan, and Magnifying Glass
- ▶ Region of Interest, Distance and Angle Measurement
- ▶ Image Annotation
- ▶ Image Identification
- ▶ 3D Processing
- ▶ Hard Copy Printing
- ▶ Speech Recognition
- ▶ Scanned Documents



In addition to the explicit response to the requirements, the RFP should invite the vendor to describe other options available either directly from the vendor, or via a third party, for example, advanced 3D processing, Orthopedic templates, Nuclear Medicine, Computer-Aided Diagnosis, etc.

A more detailed discussion regarding diagnostic workstation functionality and a source for deriving requirements can be found in Chapter 17.

### IMAGE DISTRIBUTION VIA WEB SERVER

Most enterprise PACS deployments include the ability to provide images to users outside of the department of radiology. The use of a Web server in conjunction with the hospital Intranet and Internet allows for distribution inside and outside the hospital walls. The Web server can secondarily support radiologists providing off-hours coverage by making images available for review on a home PC. Some PACS vendors now have Web-based PACS implementations where there is no distinction between the diagnostic workstation and the functionality provided to the clinician other than the monitors used and the privileges granted to the user. Many vendors, however, have a client-server application for the diagnostic workstation and a separate Web-enabled application used for enterprise distribution of images to desktop PCs. In general, the functionality of the Web-enabled application will be a subset of that offered on the diagnostic workstation. Many vendors are moving toward a common user interface between the two products distinguished only by the inability of the Web product to mark an exam as having been dictated and the lack of integration with third-party software packages (e.g., advanced 3D).

The RFP should specify the minimum functionality required for the Web-enabled image distribution subsystem and invite the vendor to fully describe the functionality of their Web distribution offering.

### CLINICAL REVIEW WORKSTATION

In clinical areas that are heavy users of radiological services, such as the ED and ICU, it may be useful to deploy dual-monitor viewing stations to provide the ability to view AP and lateral, or current and prior true-size chest images simultaneously. Vendors whose diagnostic workstation products differ

significantly from their Web-distribution products will sometimes offer an “intermediate” clinical review workstation product that more closely resembles their diagnostic workstation for clinicians whom feel they need functionality equivalent to that of the radiologist. The decision regarding which software product to deploy in these areas should be made after selecting the vendor based upon the suitability of the Web-distribution product to each area’s needs. If any of the vendors you are soliciting have products that are intended for this application, you will want to include a specification of the functionality required for these clinical areas. The RFP can invite the vendor to simply respond with how their clinical workstation offering differs from the diagnostic workstation.

### TECHNOLOGIST Q/A WORKSTATION

Depending on your intended workflow, it may be helpful for technologists to confirm the successful transmission of studies to the PACS, to have the ability to “fix” study information that does not correlate with corresponding RIS data, view historical exams on PACS, and to print images from PACS in response to requests from referring physicians. Ad hoc printing is a capability you may also wish to provide to the film library. Some vendors provide these capabilities via a web client which can be accessed from any PC, however many vendors require at a minimum a software license for each workstation or PC that has this software installed. You will want to specify the minimum functionality required by the technologists and film library.

### NETWORK

The RFP should include a description of each site’s networking infrastructure, including both the local area network (LAN) and the wide area network interconnecting the sites that will have PACS deployed or will be utilized for enterprise distribution. The description should include the vendor(s) and models used for the core routers and switches, bandwidth, and services that provide the wide area network connections. The PACS vendor should be asked to respond regarding the suitability of the existing networking infrastructure and to propose any upgrades they believe necessary to achieve acceptable performance. In addition, the RFP should request the solicited vendors to propose any additional storage cache hardware that would be needed to minimize traffic on the wide area network connections.

## SYSTEM THROUGHPUT AND PERFORMANCE

System performance in PACS is of importance primarily to the radiologist using the display workstation. This section should define performance requirements for the PACS that the vendor will commit to. Performance requirements should be defined based on a reference set of images which would define a typical study for CR, CT, and MRI. Ultrasound and Nuclear Medicine studies are typically less demanding so these do not need to be included. Performance benchmarks defined can be as follows:

- ▶ Image load time from selection from the worklist to appearance of the first image
- ▶ Time to display the complete study
- ▶ Time to display the results of a database query (note that this is for display of a list of studies; not the images in a study)
- ▶ Time to display images retrieved from nearline archive (if this is applicable)
- ▶ Time to send a complete study from the modality and display on a workstation

Most vendors will hedge their response based on their inability to control network traffic, so it is appropriate to include language that requires that the defined benchmarks be met presuming that non-imaging networking traffic is negligible.

## IMPLEMENTATION PLAN

It is important that the vendor understands your environment and has an overview of your vision for the rollout of the system. Some departments will want the “big-bang” approach where the entire system is installed in the department and the transition covers every area of the department at once. The disadvantage to this method is the disruption to the department and the demands that are placed upon the deployment staff to ensure the process is successful. In addition, this approach may place demands on the financial resources of an institution that cannot be met. An alternative to the “big bang” is a phased implementation with each phase focusing on a specific objective. This type of conversion has less of an impact on operations and allows the staff to be trained sequentially as each phase is rolled out.

Typically the first phase focuses on project planning, communication, and implementation of the infrastructure, including networking, PACS core components, HIS/RIS interface, modality upgrades and DICOM integration, and EMR interface.

The following phases focus on implementation within the radiology department. This can proceed either by modality or by site, with digital modalities (CT, MRI, US, NucMed) being implemented first, followed by general radiography (plain film x-ray) and mammography. For most sites, implementation of general radiography requires conversion to computed radiography or digital radiography, and implementation of mammography requires conversion to digital mammography, both of which represent a major investment. Postponing implementation of these modalities until later phases can ease both the cash flow and the demand on human resources.

It is typically recommended that electronic distribution of images throughout the enterprise be planned as the final phase of the PACS implementation. This gives radiology the opportunity to fully absorb the technology, refine processes and procedures, and adjust workflow to optimize the use of the new technology before having to address the change management required to convert the whole enterprise to utilization of soft-copy distribution of images.

The RFP should present a high-level implementation plan to the vendor to provide an understanding of the resources you will expect from the vendor for a successful deployment. The vendor should be required to present a proposed implementation plan that meets your expectations and includes the following elements:

- ▶ Vendor support to be provided, including specific personnel.
- ▶ Qualifications of staff assigned to the implementation
- ▶ Equipment to be installed during each phase.
- ▶ Amount of time the vendor expects to spend on-site.
- ▶ Staff support required by the department
- ▶ Proposed implementation timeline
- ▶ Costs associated with each implementation phase.

## TRAINING

A robust training program is an absolute necessity for a successful PACS implementation. This section of the RFP should define your training requirements and give the vendor enough information to realistically esti-

mate the resources that will be required. Training should include the PACS Administrator, radiologists, technologists, clerical staff (including film librarians), and clinicians. The RFP should estimate the number of personnel in each discipline that will need to be trained.

The RFP should solicit the vendor to describe their training methodology for each user category and outline the number of hours of training to be provided. The vendor should be asked to specify if training is to be provided in a classroom or one-on-one setting and which training modules include hands-on training. The vendor should be required to specify if they intend to provide end-user training or follow a train-the-trainer strategy.

Required reference materials, such as manuals, online help, computer-based training, reference cards, etc. should be specified by the vendor in this section. The vendor should be asked to specify if their training includes a competency-based evaluation to validate that trainees have effectively absorbed the material covered.

## SUPPORT AND MAINTENANCE

Vendor support and maintenance throughout the implementation and life of the PACS is a key component of the services that a PACS vendor must provide in order to maximize the benefits that PACS can provide. The RFP should specify your expectations for the following services:

- ▶ Project planning and installation
- ▶ System reliability, uptime and response time
- ▶ Warranty
- ▶ Maintenance and support

## PROJECT PLANNING AND INSTALLATION

This section should solicit the vendor(s) to describe their implementation methodology and explicitly state the site's expectations regarding the vendor(s) responsibilities.

The following are some specific areas to cover in the RFP:

- ▶ Project management services
- ▶ Installation of all vendor-supplied system components

- ▶ Interfacing of imaging modalities
- ▶ Interfacing to the existing laser printers
- ▶ RIS interface validation
- ▶ User training

### SYSTEM RELIABILITY, UP TIME, AND RESPONSE TIME

This section of the RFP is most likely be the area of greatest contention in the entire process. This is where the customer will be reducing the profitability of the sale by increasing the service levels that the vendor is being held against. This is where the customer needs to insist and only bend if the vendor offers another area of savings that is just as advantageous. Be careful here; what may seem as a good deal financially may be at the sacrifice of prudent clinical services. To ensure these requirements force compliance, financial penalties should be assessed for each violation. It is imperative that the practice keep independent records of downtime in order to ensure compliance.

This section of the RFP specifies your expectations regarding system up-time and the vendor's response time to resolve critical problems. For the new initiate, it is important to note that only a few tenths of a percentage of guaranteed uptime can work out to a significant amount of time. For example, if the customer agrees to 99.5% uptime, this means that the system can be off-line for only 4 hours a month whereas when the uptime guarantee is 99.95%, downtime is restricted to 20 minutes per month. In most practices, a 4-hour outage per month would be unacceptable, particularly if it was a single 4 hour outage that occurred during peak hours.

The industry standard is to distinguish between downtime that renders the entire system unusable and downtime affecting a single component such as a workstation or single modality interface. You will want to require that the core system be held to 99.9% uptime at a minimum. The other aspects of the PACS system relating to clinical viewing, web-distribution, and near-line storage can be held to a different service level agreement such as 99% uptime.

The vendor(s) responses will also typically distinguish between planned (e.g. software upgrades) and unplanned downtime, with planned downtime being exempt from the requirements. Distinctions may also be made between downtime experienced during "normal business hours" and "off-hours" where the impact to the operation may be somewhat less severe.

This section should also solicit the vendor(s) to commit to a maximum response time to reported problems. The RFP should distinguish between remote support and onsite support. The majority of problems with PACS can typically be handled remotely, but hardware problems require onsite support, and most vendors will not commit to less than 4 hours to dispatch a service technician for onsite support. It is important to keep in mind the geographic location of the vendor's support team closest to your institution and the RFP should request that the vendor(s) indicate where their onsite support personnel will be dispatched from.

To ensure these requirements force compliance, the RFP should define the financial penalties to be assessed for each violation, typically in the form of an extension of the warranty or service contract or a replacement option for hardware components that repeatedly cause downtime. It is imperative that the site keep independent records of downtime in order to ensure compliance.

## WARRANTY

Standard warranty coverage for PACS is typically one year, with some vendors offering only 90 day coverage on hardware: The RFP should define the site's expectations regarding the warranty. Examples of areas to cover are:

- ▶ When the warranty starts
- ▶ Length of the warranty
- ▶ Coverage (e.g., parts and labor)
- ▶ Service response time during the warranty period
- ▶ Software upgrades to be included during the warranty period
- ▶ Penalty clauses for failure to meet the requirements of the RFP

## MAINTENANCE AND SUPPORT

Maintenance and support should always be negotiated at the time of the initial PACS purchase since this is when the customer has the most leverage. The RFP, as a precursor to the contract should therefore define the site's expectations and attempt to get the vendor to commit to a service pricing for coverage for the expected life of the system. The vendors quotations should include options to purchase on an annual basis or commit to a more

extended coverage period in return for guaranteed pricing. This is the area of the RFP that will define the ongoing relationship between the customer and the vendor. The most important negotiation point that can be passed along in this book is that **NOTHING IS STANDARD**, especially in the PACS industry.

Below are some key points to be covered in the RFP are:

- ▶ Request for quotation of one year and four year contract
- ▶ Specification of ordering and payment terms
- ▶ Penalty clauses for failure to meet the requirements of the RFP
- ▶ Coverage details (e.g., parts, labor, software upgrades)
- ▶ Expected software release schedule (e.g., quarterly, biannually)
- ▶ Operating system(s) security patches and antivirus software updates
- ▶ Support mechanisms (e.g., telephone, Web submissions, e-mail, etc.)
- ▶ Priority levels for incident reporting (e.g., critical, urgent, high) and the associated guaranteed response times
- ▶ Call escalation procedures
- ▶ Maintenance activities, if any, that are the responsibility of the site
- ▶ Costs associated with on-site technical support
- ▶ Minimum qualifications of staff assigned for technical support

## VENDOR INFORMATION AND SELECTION SCHEDULE

This section gives the vendor an overview of the response expectations and details regarding the process for the response and final vendor selection.

- ▶ Confidentiality and nondisclosure
- ▶ Format for the response
- ▶ Selection schedule
- ▶ Remittance of proposals
- ▶ Contact and procedure for submitting questions regarding the RFP
- ▶ General response requirements
  - ▶ Primary vendor contact person
  - ▶ Overall responsibility for delivery, implementation, and maintenance of hardware, software, and services



- ▶ Conformance to federal, HIPAA, state, local, JCAHO, and American College of Radiology requirements.
- ▶ Delivery of works-in-progress
- ▶ Evaluation criteria
- ▶ Evaluation process
- ▶ Disclaimers

## APPENDICES

The appendices provide information to the vendor that may not be conveniently presented in the body of the RFP or that is cross-referenced by multiple sections of the RFP. Recommended examples are:

- ▶ List of modalities and their supported DICOM service classes
- ▶ List of laser printers
- ▶ Exam volume broken down by site and modality
- ▶ Growth projections broken down by site and modality
- ▶ Estimated storage requirements broken down by site and modality

## RESPONSE FORMS

In order to facilitate the comparison of multiple vendors, it is helpful to provide a response form that vendors are required to complete. The response forms should force the vendor to respond in a tabular format with summarized responses that facilitate comparison of critical requirements, features, and vendor capabilities. In preparation for designing these forms, it is helpful to map out the criteria you will use for evaluating the vendors. Once this has been determined, the forms can be designed to assist with the differentiation of vendors around these criteria. Some suggested response forms are as follows.

## GENERAL DESCRIPTION OF THE VENDOR

Include year established, company ownership (e.g., private, public), parent company, number of personnel in R&D, Sales & Marketing, Service & Support, gross revenues and net income from PACS, R&D investment.

## CLIENT BASE

Include number of sites live, implementing, and contract signed for current year and two prior years. Also include number of sites which are more than 70% filmless.

## CLIENT REFERENCE

Include site profile (e.g., institution name, number of beds, exam volume, HIS/RIS) and contact information.

## TECHNICAL REQUIREMENTS, TRAINING, AND SERVICE AND MAINTENANCE

These forms can be structured to mirror the corresponding sections of the RFP. The response forms should not necessarily replace a comprehensive response to these sections of the RFP, which is used to insure compliance to a specification, but should rather be designed to highlight major differences among vendors.

## ITEMIZED PRICING

This form is perhaps the most important of all, as it not only facilitates comparison among vendors, but requires the responses to provide line-item pricing for each major system component. Forcing the vendor to expose the detailed cost structure of the purchase gives the site added leverage in the negotiation.

It is recommended that pricing for software and hardware be broken out separately, and that line-item pricing for each major component described in the RFP (e.g., PACS core, Web server, modality interfaces, HIS/RIS interface, diagnostic workstations, clinical review workstations) be required. The form should also include pricing for systems integration, implementation, training services, and extended maintenance services. This gives you the ability to compare pricing for a software-only purchase, for the grand total including services, and for the total cost of ownership over the life of the PACS. Line-item pricing also gives you the ability to estimate costs if you wish to adjust the purchase to either add or eliminate individual line items or, for example, to increase the number of workstations.

## DISTRIBUTING THE RFP

The first temptation after all the hard work that went into the creation of the RFP is to send it out to every PACS vendor the site can think of. Remember that for every RFP you send out you will receive a response ranging from 100 to more than 500 pages. Reviewing these can be challenging.

Based upon the work that has gone into the RFP, the site should have a very good idea of its requirements. With some homework, you should be able to narrow down the number of possible vendors to a manageable number. You can further refine that list by requesting vendor presentations at your site or by visiting any of the tradeshow featuring radiology vendors. Once a manageable list of vendors (4 to 6) is determined, the RFPs should be sent out with a firm due date giving the vendors 4 to 6 weeks to respond.

Vendors should be given a contact for clarification of any questions they may have regarding the RFP. If inconsistencies or ambiguities are uncovered during this process, it is best to provide clarification to all vendors to insure a level playing field among all vendors, and that individual vendors are not making erroneous assumptions based on their interpretation of the RFP.

## VENDOR EVALUATION AND SELECTION PROCESS

### REVIEW PROPOSALS

Once the proposals are received the process of reviewing each in detail begins. The review should ensure that all the requirements outlined in the RFP have been responded to properly and that the answers to the questions are relevant to the requests made of the vendor. This can be a tedious task, but it is one that is quite important to ensure that improper assumptions are not carried forward into the evaluation process.

Selection of a PACS is typically done by a committee in order to ensure that the views of all critical stakeholders are incorporated in the decision process. Few members of the committee will have the time to review in detail the entire response from all vendors, so it is helpful to summarize the responses in a tabular format, laying out all vendors' responses side-by-side in a comparison table. This can be done by utilizing the responses to the forms included in the RFP. The summaries should then be distributed to the members of the selection committee for their review.

## CLARIFY QUESTIONS

Prior to proceeding with a decision, it is important to clarify any omissions or ambiguities in the vendor responses and to ask the vendor to revise their response appropriately. If discrepancies are discovered between the responses and the RFP, it is a good idea to review how the request was stated and if it needs further clarification. This clarification should then be distributed to all vendors as it would be if the ambiguity was discovered before the responses were due.

## MODEL THE DECISION

Once the responses have been evaluated and summary information compiled, it is suggested that some form of decision model be used to help facilitate the decision process, and to make it as objective as possible. The decision model forces the stakeholders who are involved in the decision to base their decision on an agreed-upon set of criteria and can help prevent a decision based on a single criteria that may seem to be overwhelmingly important, or the bias of a single member of the selection committee.

The recommended process is to establish a set of criteria by which the vendors will be evaluated. A set of “attributes” or subcategories for each major criteria should also be established to facilitate the rating of each criteria. Each major criteria should be weighted in relationship to other criteria based upon the perceived value to the site. For example, if “technology” is an agreed upon major category and it is perceived as the most important, then it can be weighted as a “10.” If “price” is the next most important criteria, but is considered to be somewhat less important than “technology” then it can be weighted as an “8.” Other criteria can then be weighted in comparison to these criteria. Attributes within each category should also be weighted in comparison to one another. The total of all attribute weights within a single criteria should be consistent across all major criteria.

The decision model, including the criteria, attributes, and assigned weights should be developed by consensus of the selection committee. Once the model has been developed, vendors should be rated against each attribute in relationship to one another. It is suggested that a rating system of 1 to 5 be used, with a score of 1 indicating a weak rating for a given attribute and a score of 5 indicating a strong rating. Rating can be done independently by each selection committee member, and then averaged across all committee members’ responses. Each attribute score is then weighted by the agreed-upon model and the weighted scores totaled for each vendor. The highest score “wins.”

Scoring is typically based on the responses to the RFP, but committee members should also rely on their knowledge of vendors' products and reputations gleaned from demonstrations at trade shows, onsite presentations, reference checks, experience of colleagues, prior experience with individual vendors, and so on.

An example of the decision model described is shown Table 4.1 below. It is helpful for the model to be created in a spreadsheet format, so that weights and ratings can easily be changed.

During this process it can be very helpful to employ an independent facilitator who is not invested in the outcome to guide the development of the decision model and facilitate a discussion of the scoring. If the facilitator is involved in the RFP review process as well, and is knowledgeable about the vendors being considered, the process can potentially be streamlined by using the facilitator to provide a "straw horse" scoring model that can then be adjusted by the consensus of the committee.

If, after the model is developed, the ratings established, and the total scores tallied, the results are not consistent with the apparent leanings of the committee, it may be appropriate to re-evaluate the model, adding criteria and adjusting the weights and ratings to be consistent with the apparent consensus of the committee. It is important to keep in mind that the purpose of the model is to facilitate objective discussion and reach consensus and not to single-mindedly drive the process. The reasons should be objectively evaluated and the ratings or weights possibly changed.

## NARROW TO TWO VENDORS

The objective of the decision model is to choose two semifinalists, either of which is acceptable to the selection committee. Limiting the number of semifinalists will help make the planning of site visits and contract negotiation more manageable, but it is important to proceed to negotiation with more than one vendor to insure that the vendors stay competitive, even if there is a clear preference.

## CONDUCT SITE VISITS

This is probably one of the most important steps in the decision process. This is the opportunity for the physicians and staff to see the system in action and ask detailed questions about the advertised functionality versus reality. The most important thing to remember in this step is not to allow the vendor to escort you during the entire visit. Many of the vendors arrange the visits;

TABLE 4.1  
Decision Model

Vendor Category	Weight	Vendor 1		Vendor 2		Vendor 3		Vendor 4	
		Raw Score	Weighted Score	Raw Score	Weighted Score	Raw Score	Weighted Score	Raw Score	Weighted Score
<b>TECHNOLOGY</b>	10.00	22	38.00	23	40.00	27	47.00	19	31.00
Compliance to Specification	0.10	2	0.20	4	0.40	4	0.40	4	0.40
Perceived Workability	0.20	4	0.80	3	0.60	5	1.00	4	0.80
Use of Current Technology	0.30	4	1.20	5	1.50	5	1.50	3	0.90
Flexibility	0.20	4	0.80	4	0.80	5	1.00	2	0.40
Adherence to Standards	0.10	4	0.40	5	0.50	4	0.40	4	0.40
RIS/PACS Integration	0.10	4	0.40	2	0.20	4	0.40	2	0.20
<b>PRICE</b>	10.00	15	39.00	13	31.00	19	48.00	8	21.00
Software	0.20	4	0.80	3	0.60	4	0.80	3	0.60
Hardware	0.10	5	0.50	4	0.40	5	0.50	2	0.20
Services	0.10	2	0.20	3	0.30	5	0.50	1	0.10
PACS Grand Total	0.60	4	2.40	3	1.80	5	3.00	2	1.20
<b>VENDOR VIABILITY</b>	10.00	18	36.00	12	24.00	22	44.00	21	42.00
Vendor Overall Experience	0.20	3	0.60	2	0.40	4	0.80	5	1.00
Vendor/Product Stability	0.20	4	0.80	1	0.20	5	1.00	4	0.80
Site Sophistication	0.20	4	0.80	4	0.80	5	1.00	4	0.80

PACS Implementation Experience	0.20	2	0.40	2	0.40	4	0.80	5	1.00
Perceived Responsiveness	0.20	5	1.00	3	0.60	4	0.80	3	0.60
<b>TRAINING</b>	2.00	6	4.00	13	8.67	13	8.67	13	8.67
Overall Philosophy	0.33	3	1.00	3	1.00	5	1.67	3	1.00
Manuals	0.33	2	0.67	5	1.67	5	1.67	5	1.67
Training Classes	0.33	1	0.33	5	1.67	3	1.00	5	1.67
<b>SUPPORT</b>	4.00	10	13.33	11	14.67	11	14.67	14	18.67
Support Strategy	0.33	3	1.00	3	1.00	3	1.00	4	1.33
Service Responsiveness	0.33	4	1.33	5	1.67	4	1.33	5	1.67
Uptime Guarantee	0.33	3	1.00	3	1.00	4	1.33	5	1.67
<b>CLIENT RELATIONSHIP</b>	2.00	8	5.33	8	5.33	10	6.67	7	4.67
Existing relationship	0.33	3	1.00	1	0.33	2	0.67	2	0.67
Contact during RFP process	0.33	3	1.00	2	0.67	4	1.33	2	0.67
RFP Quality	0.33	2	0.67	5	1.67	4	1.33	3	1.00
<b>Summary</b>			<b>136</b>		<b>124</b>		<b>169</b>		<b>126</b>

however it is very important that they not interfere with the candid responses from the site you are looking for. The site visits should not be limited to just the physicians; support and I.T. staff should also be part of the visit to evaluate the operational and technical issues.

Try to schedule the site visits as close to each other as possible. It is a good idea to perform a review session with the visit team after each visit and to compare your experiences against the responses and results of the RFP measurements. You are making a decision that represents a major investment and will be very difficult to change once the system is deployed. This does not mean that the process should get bogged down into analysis paralysis. Once two vendors have been selected, it is time to move to the next step.

## NEGOTIATE THE CONTRACT

The RFP process can be long, difficult, and sometimes very frustrating. The final decision for two vendors has been made, and it is time to “put the two vendors in a room and let them fight it out.” This is where two vendors are played off against each in order to get the best value. Remember that value includes not only price but system options, hardware, software licensing, service contracts, committed functionality, all of which should be included in the negotiation.

The basics of negotiation skills are to understand that any negotiation breaks down into three principal focus areas (known as the dimensions of negotiation); tactics, deal design and setup. Tactics are based upon people and processes. Deal design is based upon value and substance, and setup is the scope and sequence of the deal.

Barriers to successful tactics are interpersonal issues, poor communication and “hard ball” attitudes. A barrier to deal design is the lack of feasible or desirable arguments. The barrier to setup is that the parties do not support a viable process or valuable agreement.

The approach that can help resolve tactical issues is to act “at the table” to improve interpersonal processes and tactics. Deal design issues are best resolved by redesigning a deal that unlocks value. For a successful setup, a change in focus needs to be made away from the table to create a more favorable scope and sequence for the approach.

How does this information relate to negotiating a PACS deal? Being able to identify an area of weakness in the focus areas described above will help provide a very general way of measuring how you should initiate and change course as you head through the negotiation process.



In addition to the theory presented, the best practical advice available is to come to the table prepared. Make sure when you begin the negotiation session you have a clear set of objectives for that session. Realize that you may not make it past the first objective in that meeting and that you do not have to solve everything before leaving the table. It is important to understand that the person speaking is usually on the defensive and is trying to convince the other party.

A “hard ball” attitude is very likely to fail because a negotiation is a process of coming to a mutually beneficial arrangement. It is not realistic to expect a PACS vendor to enter into a business relationship where they will not make a profit. Keep this in mind when the deal looks like it may be going south. If the deal design is breaking down a re-thinking of the deal may be needed. This can be especially true as you continue to play the competing vendors against each other.

In the long run, the goal is to get the deal you can feel comfortable with. If you are not comfortable, then there is doubt about the value of the deal. It needs to be reevaluated or even renegotiated until the doubt is eliminated.

## CONCLUSION

The process of creating an RFP and selecting a vendor requires attention to detail, patience, and the willingness to look deep within your own institution to realistically determine your requirements and expected benefits. Your understanding your institution’s objectives, expected benefits, and your operation is used to develop a vision of the future. This must be translated into a set of requirements that describes the operations of the department and defines the technical, implementation, training, and support needs in order to best achieve a solid Request for Proposal that best reflects the needs of organization and obligates the vendor to meet those needs.

To put it simply, when creating an RFP the end result is proportional to the effort put into it. The RFP and the decision process that follows set the stage for a favorable contract negotiation and a successful implementation and should be undertaken with care.

## REFERENCE

1. Lax DA, Sebenus JK. 3-D negotiation. *Harvard Bus Rev.* 2003;81(11):65–74

# REENGINEERING WORKFLOW: A FOCUS ON PERSONNEL AND PROCESS

---

---

BRUCE I. REINER • ELIOT L. SIEGEL

**D**iagnostic radiology departments and practices share a dual mission. One is to provide the highest quality of care possible for patients, a goal that involves a well-informed selection of imaging techniques, skilled interpretation, timely and clear communication of findings to clinicians and referring physicians, and the most efficient and effective use of available personnel. The second is to maximize income (or at least to minimize loss).

Although some healthcare providers may balk at the comparison, these goals are not dissimilar from those that govern industrial assembly lines, where efficiency and productivity are demanded along with high levels of quality and consistency and where innovation and economies of scale drive cost decisions.

In the last decade, the increased penetration of managed care, growing competition among diagnostic imaging providers, and an assortment of economic factors have caused unpredictable fluctuations and, in many cases, decreases in reimbursement rates for imaging services. For most imaging practices and departments, this resulted in the need to decrease operating expenses across the board at the same time that the volume of studies to be read increased. The assumption was that these reductions could be made without compromising the quality of services rendered.

The initial reaction to these financial constraints at many sites was to suggest that radiologists and other imaging personnel could work “harder, faster, and longer” while trimming excess fat from expenses—a strategy that ultimately proved self-defeating as these enforced economies quickly reached their limits and took an inexorable toll on staff morale. Moreover, these economies have been put in place during a time of rapid and revolutionary change, with the implementation of filmless imaging, integration of picture archiving and communication systems (PACS) with other radiology and hospital information systems (RIS and HIS, respectively), and the advent of new and increasingly complex and sophisticated image sets. This changing environment has also been affected by a shortage of trained radiologic technologists (RTs) and by increased demands for accountability and monitoring of medical error. It is clear to some observers that the most promising way to realize further reductions in costs and maximize efficiency without compromising patient care is to adapt methodologies from industrial engineering to study, modify, and, if necessary, completely redesign the workflow process in the imaging department.

Workflow engineering is a highly regarded and well-developed field within the industrial and manufacturing communities. Simply put, it involves the analysis and breakdown of individual steps that occur in the performance of a multistep task, such as the assembly of an automobile or the acquisition and interpretation of a radiograph. Although the engineering literature contains thousands of references to workflow analysis, which has been adopted widely in other disciplines, radiology has not embraced the concept. Lessons learned from industry may govern the supply chain in a clinic, the pharmacy ordering and distribution system in a hospital, or the timing and flow of cafeteria service in a medical center, but they are rarely applied in the imaging suite.

Unlike most industrial assembly lines, radiology departments are notoriously inefficient. An inordinate number of personnel steps are typically required from the time a patient is registered until an imaging report is made available to the referring clinicians. These functions are performed by a variety of clerical, technical, and professional personnel in a process that often evolves over a period of years into a confusing and disorganized routine dictated more by tradition than by logic. Even after the introduction of filmless radiology and PACS, many of these inherited routines remain fixed but vestigial parts of departmental workflow. The challenge is to step back and look at the big picture of workflow as it relates to well-established priorities for any imaging center, including quality of service, timeliness of service, quality of product, and economic efficiency.

In a truly industrialized approach, the optimal workflow strategy to attain high levels for these priorities might be to create the equivalent of a radiology assembly line, in which each staff member concentrated on one and only one task. The technologist, for example, would perform only tasks directly related to image acquisition. Additional tasks, such as patient escort, room preparation, exam scheduling, accessing data, and retrieval of historical exams, might be allocated to radiology aides—with individual aides specializing in specific tasks.

This brave new world of radiology is not yet within the bounds of feasibility. This is perhaps fortunate, because the essential element left out of a truly industrialized view of radiology workflow is, in fact, the most important: the patient. In the industrial analogy, the patient becomes the product. In too many imaging settings, the implicit goal has already become to reduce the patient to a conglomerate of pixels as quickly as possible. The process of optimizing workflow must keep a focus on the patient and his or her needs as well as the essential part that satisfaction and well-being play in maintaining efficient and predictable workflow.

In a digital imaging environment, the greatest changes in workflow optimization involve three simple tasks: automation, integration, and simplification. Automation, experienced by every imaging department in many and ongoing forms, is the replacement of manual tasks by computer-generated tasks. An example would be that of modality worklist software, which automates the process of linking patient demographic information between the HIS and RIS and effectively eliminates manual data entry by the technologists. Integration is the linking of disparate computer systems, eliminating the time-consuming steps of accessing and transferring data from one computer system to another. Finally, simplification is the ability to convert complex, time-consuming tasks into more straightforward ones. An example of simplification would be the development of radiologist-specific

hanging protocols that are customized to the needs of the radiologist and eliminate the time-intensive and frustrating tasks of manually retrieving comparison studies in the interpretation process.

In this chapter, we look at the changes in workflow that have accompanied the switch from film to PACS at many institutions and at ways of analyzing imaging workflow to improve operational efficiency and productivity. We address the personnel and processes involved in the steps that precede the interpretation process, including exam scheduling, patient arrival, image acquisition, processing, and transfer/storage. Chapter 6 focuses on radiologist workflow as it pertains to the processes of image display, interpretation, and reporting.

## **ANALYZING WORKFLOW FROM THE PERSONNEL PERSPECTIVE**

In the 1980s, several investigators attempted to estimate the average time required to perform various radiological examinations. The reported procedure times and their constituent components varied widely (by 300% or more for the same exam) because of institutional idiosyncrasies, differences in equipment, and even differing definitions of when an examination actually began and ended. The introduction of PACS provided opportunities to reassess workflow within a number of institutions worldwide that subsequently reported on the effects of the film-to-filmless transition on time, cost, and quality of product.

The most widely used approach to analyze, model, and reengineer workflow in the medical setting has been task analysis, where the performance of duties of each employee is described, timed, and integrated into a workflow diagram. Such studies can involve time-motion investigations, such as one conducted at the Baltimore Veterans Affairs Medical Center (BVAMC), before the introduction of PACS. Management consulting firm Booz Allen Hamilton enumerated more than 50 steps in the process from initiation of a clinician order for a radiology study until the transcribed radiology report was available in the patient's chart for review. As more departments have gone filmless, the number of consultants who provide similar task-specific analyses has increased. Many radiologists and administrators were amazed to see how inefficient their departmental and intra-institutional practices had become over time. However, after a careful redesign of the workflow process and the installation of a well-integrated PACS and RIS, the lessons learned from the reengineering have, in many cases, been lost. It

is useful to look once again at the effects on personnel and task-specific workflow in the transition from film to PACS to pinpoint those areas in which additional economies of time and labor might be realized and incorporated into additional reengineering for high-quality, patient-centered care.

## CLERICAL PERSONNEL

### BEFORE PACS

The time-motion study developed for the BVAMC's imaging department before the introduction of PACS demonstrated that a large percentage of the steps in the workflow process were clerical in nature. Among these were a number of steps in the completion, submission, and handling of an order or requisition for the examination, as well as in the processing of that request in the radiology department and in communication between clerical staff and RTs. Clerical functions also included the many steps involving film handling and movement throughout the medical center. One particularly labor-intensive task was the daily requirement to retrieve more than 500 film jackets for patients who were to be seen in the various outpatient departments and then to attempt to retrieve each of these studies so that they could be returned and refiled in the film library. Film library staff was also responsible for finding old patient film jackets and matching these to new studies, so that old and new examinations could be transported together to reading rooms for interpretation by radiologists. After studies had been interpreted, the clerical staff was again involved in transporting films back to the library or to other areas of the medical center. They were also responsible for delivery of report dictation audiotapes from the reading room to the report transcription area.

### AFTER PACS

The transition to the use of an HIS-RIS and PACS at the BVAMC had a profound effect on workflow in the clerical areas of the department. The use of the computer systems resulted in a dramatic reduction in the number of steps required and the amount of time required to perform the remaining steps. The use of an integrated HIS-RIS reduces the process of submission of a radiology request to a much faster and simplified process in which the clinician identifies the patient to be studied and the exam to be performed and enters the reason for the examination on a computer workstation using

a graphical user interface. Other patient information, such as age, sex, location within the hospital, requesting clinician, and primary care provider, are automatically sent by the HIS to the radiology department and the PACS. This information is also available to RTs, who interact with the system to schedule the patient examination and (when appropriate) edit the request.

When ordering information is sent to the PACS, it creates an electronic image “folder” entry in its database and initiates a request to the long-term image storage device to retrieve previous examinations for comparison in cases in which they are not already available in short-term storage. This pre-fetching of old studies from the long-term archive is thus initiated before the new study is performed, well in advance of interpretation by the radiologist or review by the clinicians. Depending on institutional requirements or preferences, this pre-fetching of comparison studies can be initiated by other events, such as patient admission, patient transfer, or a scheduled appointment for an outpatient visit. Any patient images can be routed from long- to short-term storage or can be sent to a specific workstation or group of workstations according to predefined rules. This may be required for applications such as at-home teleradiology, where bandwidth is limited.

The elimination of film and integration of the PACS with the HIS-RIS have obviated the need for a paper request for the ordering clinician. At the same time, the labor- and time-intensive task of physically moving films around the hospital has been eliminated. At our institution, only mammography uses film (but will doubtless go filmless here and in most institutions in the near future). Although films still come from other institutions, the dramatic reduction in the number of films has reduced the need for film file room personnel.

If one of the goals of workflow analysis is to determine the most efficient use of personnel in the face of changing technologies, then one of the inevitable results is that some personnel roles will change and others will be eliminated. At the BVAMC, all but one of the film room librarians were eliminated and the number of other clerical assistants reduced. It is worthwhile to note that new assignments were found for each of these employees within the hospital. Every department looking seriously at reengineering workflow should plan for the possibility of such reassignments. If innovation becomes associated with employee termination and new technology with career obsolescence, then change will be met with increasing resistance by staff throughout the medical enterprise.

At the BVAMC, as at other institutions, the switch to PACS brought unexpected benefits to transcriptionists using a digital dictation system in combination with the HIS. Transcriptionists now work at home and connect

to the digital dictation system using dedicated phone lines. Reports are typed directly into the HIS-RIS radiology reporting module and are then available immediately to clinicians, who may review the preliminary report, and to radiologists, who verify and finalize their reports. This process has resulted in substantial reductions in report turnaround times, from 24 to 48 hours to approximately 2 hours.

The elimination of the majority of steps previously required for clerical operations has resulted in the reduction in our clerical staff by 56%, with significant cost savings. Clerical personnel savings alone, in fact, are greater than those achieved by the elimination of film in the department. This degree of savings for clerical staff is similar to that reported by Saarinen et al., who performed time-motion analyses with activity sampling and interviews to estimate savings in film handling times associated with PACS. The estimated time savings for film library (clerical) staff was 55%, for RTs 10%, and for radiologists 10%.

The substitution of a digital storage system for film has also resulted in eradication of one of the most vexing and time-consuming problems in any imaging department. This is the phenomenon of “lost” or missing films. Our rate of missing (more specifically, rate of undictated studies) dropped from an unacceptable 8% to less than 1% within a year after the implementation of the PACS and has subsequently dropped to approximately 0.3% (Figure 5.1).

This 0.3% rate represents a drop of more than 25-fold but indicates the continued presence of unreported cases that “slip through the cracks.” At our institution, a management report using the RIS identifies these unreported studies, including occasional identification of an inadvertently unread case, network transmission failures, and other miscellaneous causes—all potent reminders that no amount of automation will ever entirely eliminate error.

## RADIOLOGIC TECHNOLOGISTS

### BEFORE PACS

Before the transition to the use of PACS and the HIS-RIS at the BVAMC, radiology technologists (RTs) had a number of responsibilities that overlapped with the clerical and film library staff. Technologists routinely performed a large number of manual processes that added to the number of workflow steps in this section of the department. Moreover, as the most significant point of workflow continuity (from welcoming the patient to acquir-



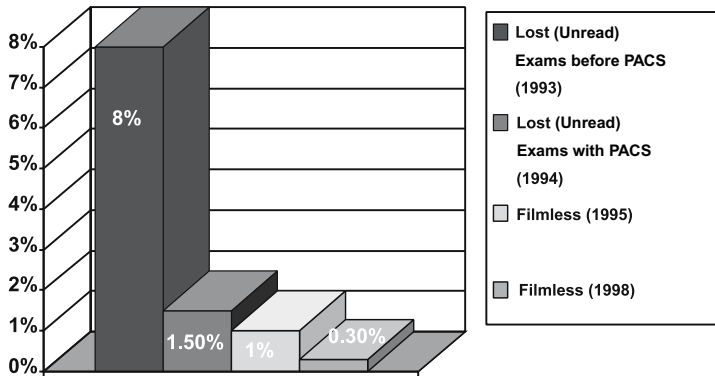


FIGURE 5.1

“Lost” films. Examinations not interpreted by radiologists.

ing the image to hanging films for interpretation), each RT could perform thousands of individual steps in a given workday. Technologists were also responsible for reentry of patient information from physician order forms into computers associated with the various imaging modalities. The result was a relatively high rate (~15%) of errors, including spelling of names and patient identification numbers, which could result in additional time delays in locating and correctly identifying studies.

### AFTER PACS

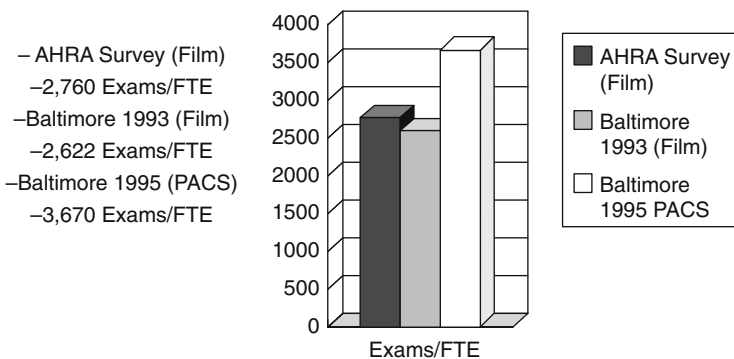
At our institution and elsewhere, the elimination of film and the transition to the PACS and HIS-RIS resulted in extraordinary improvements in the workflow of the technologists. Among the most significant of these was a 40% increase in technologist productivity for general radiography, which accounts for 65% of the total number of studies performed in the imaging department (Figure 5.2).

To assess the ways in which the transition to PACS had effected this change, we performed a study in which we reverted to conventional, film-based operation for a period of 1 week to perform detailed time-motion studies of technologists performing computed tomography (CT). We compared data gathered over the course of the week with similar time-motion studies performed with PACS in the filmless environment. The study documented that filmless operation resulted in the elimination of a large number of steps previously used in the CT suite (Figure 5.3). Eliminated steps

included those related to the creation of multiple versions of images in different window and level settings (a task that can be performed to personal preferences by the radiologist at the workstation) and those related to the handling and distribution of films. The elimination of these steps resulted in a 45% reduction in the amount of time required for a CT technologist to perform an examination (Figures 5.4 and 5.5).

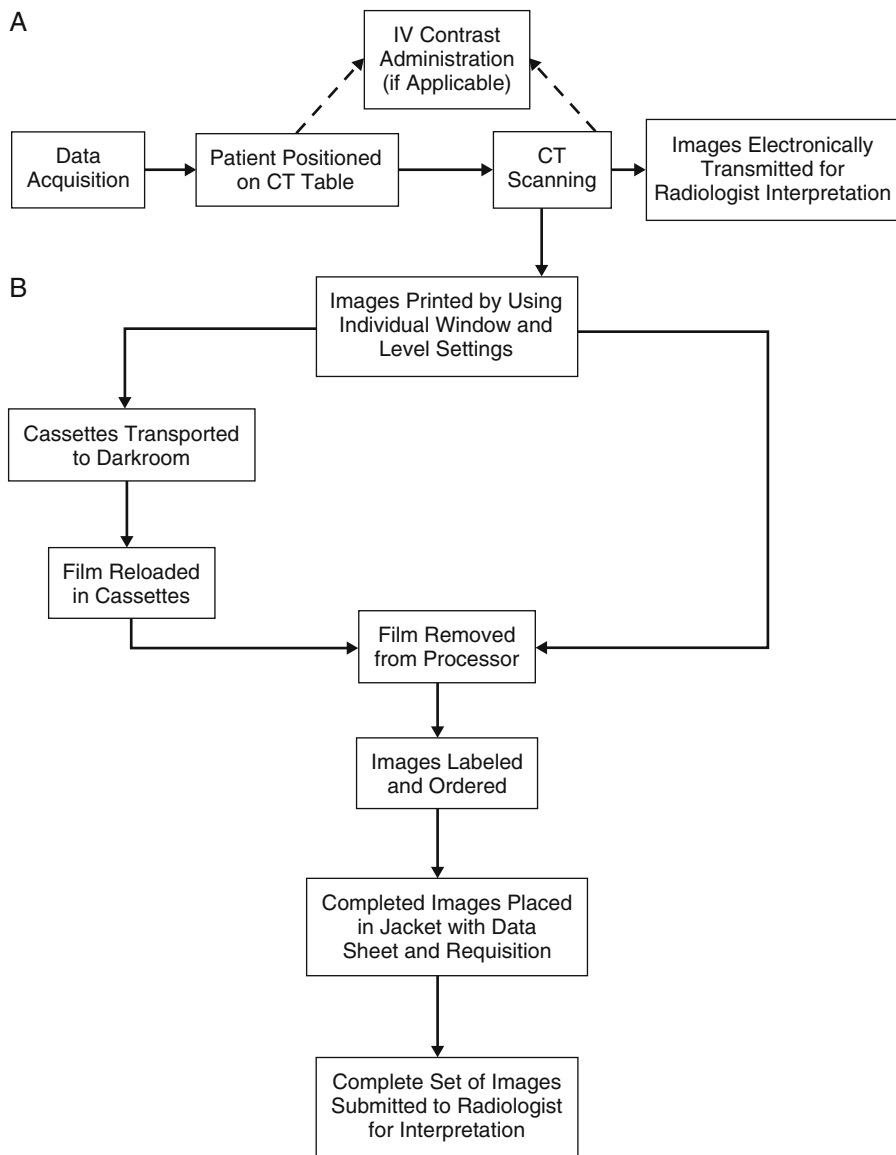
In addition, a dramatic reduction in examination retake rates was noted for general radiographic studies performed in the department. This reduction has been the result of the very wide dynamic range facilitated by computed radiography (CR) and the ability to modify the window and level (contrast and brightness) settings at the computer workstation. This drop from a 5% to a 0.8% repeat rate represents an 84% drop in the retake rate—with significant benefits realized in RT workflow.

One of the benefits of integration of an imaging modality such as a CT scanner, an HIS-RIS, and a PACS is the ability to reduce workflow steps and improve accuracy using the modality worklist feature. Using this feature, the imaging modality can communicate with a PACS or HIS-RIS to obtain the list of examinations to be performed and generate a worklist that can be displayed on the technologist operator console. The technologist can then easily select a specific examination or combination of studies, speeding entry of patient information and increasing the accuracy of data. Incorporation of this feature reduced modality transmission error rates below 1.5% in comparison with the approximately 15% rate of patient entry errors by technologists performing manual entry (Figure 5.6).

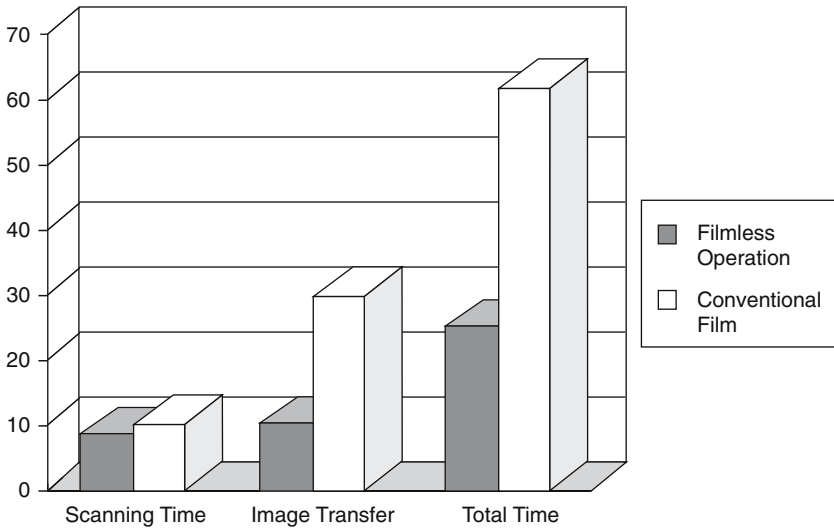


**FIGURE 5.2**

Technologist productivity increased by 40% after the transition to the use of computed radiography and PACS more than a decade ago at our institution.

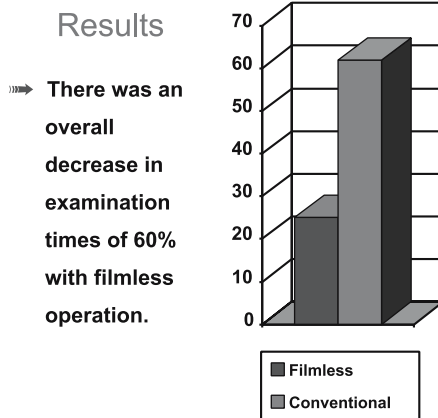
**FIGURE 5.3**

CT technologist workflow diagram comparing (A) filmless and (B) film operation.



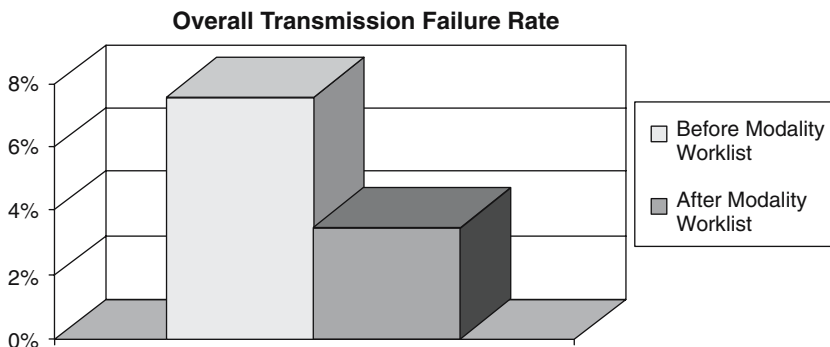
**FIGURE 5.4**

Comparison of time required to scan and transfer images and total time for film-based and filmless operation.



**FIGURE 5.5**

Reduction in CT technologist scanner time with filmless operation.



**FIGURE 5.6**

CT transmission failure rate before and after modality worklist.

## WHEN TECHNOLOGY LIMITS PERSONNEL

The potential personnel hours saved by PACS and the incremental benefits in fewer errors and improved patient care have been obvious from the beginning. However, in many cases, actual workflow benefits have been limited by difficulties in integration and communication between the HIS-RIS, the PACS, and the imaging modality, with a large body of literature offering firsthand experiences and advice. The two most common standards used in this communication are Digital Imaging and Communications in Medicine (DICOM) and Health Level Seven (HL7). Unfortunately, the majority of HIS-RIS's do not adequately support DICOM, and specific implementations and interpretations of both DICOM and HL7 vary widely. The Radiological Society of North America and the Health Information Management Systems Society sponsored a "phased series of public demonstrations of increasing connectivity and systems integration," which has brought together imaging vendors with HIS-RIS vendors. This very important effort, known as the Integrating the Healthcare Enterprise (IHE) initiative, has begun to facilitate more efficient and predictable communication among imaging modalities, the HIS-RIS, and PACS, and potentially among medical facilities as well. Additionally, several vendors now offer consolidated RIS-PACS or RIS-HIS-PACS solutions, or both. The lesson to be learned from these lingering communication bottlenecks is that even when personnel workflow is reengineered to keep up with evolving technologies, there may be times when some aspects of those technologies lag behind.

## INTRODUCING NEW PERSONNEL

One effect of the introduction of PACS into the radiology workplace has been the identification of new roles for personnel. The most prominent has been that of PACS administrator, which has risen from a hypothetical identity to an emerging profession in less than 5 years. One of the main functions of this technical/administrative position is to identify bottlenecks in workflow, both human and automated, and evaluate methods to streamline cost and maintain the quality of care in an atmosphere of continuously evolving technology.

Others have posited the need for a new kind of technologist, one dedicated to the PACS in the same way that RTs are (ideally) dedicated to image acquisition. Still others have noted that with the ever-decreasing clerical staff in radiology departments, a number of tasks (such as patient escort and data entry) could be performed effectively by full-time radiology aides. Many radiology departments are experimenting with different combinations of employees in different configurations. It seems clear that as technology changes, so will the duties and identities of radiology personnel. The challenge is to build this flexibility and innovation into workflow analyses to accommodate such change.

## ANALYZING WORKFLOW FROM THE QUEUING PERSPECTIVE

Queuing theory is a branch of mathematical reasoning and analysis that seeks to theoretically describe production processes along with the statistical/probabilistic techniques to account for varying dynamic patterns within the stages of those processes. Mathematicians have long been fascinated by the seemingly random but ultimately predictable flow dynamics of lines of cars in traffic or humans waiting to see a movie. The field received a boost when high-speed computing created “queues” of incoming and outgoing data that required analytic predictive metrics. Today, the results of queuing analysis are incorporated into a wide range of applications, from waiting-line configurations outside roller coasters to sorting and prioritizing of high-speed genomic computing tasks.

A number of authors have incorporated queuing analysis—based methods into their investigations of patient scheduling, hospital or institutional systems analysis, and medical workplace efficiency. Queuing theory, which breaks up any process into constituent time-span components, is

particularly suitable for looking at radiology workflow, where the patient (and, later, data acquired from the patient) can be tracked through a series of time-specific fields from arrival to delivery of images to the clinician. Although individual work assignments are taken into perspective in queuing analysis, the focus is actually on the overall span of time and the true flow of activity.

Queuing analysis poses specific questions about the process of workflow. In the radiology setting, these questions are broken into individual sub-components of arrival, service, and queue characteristics.

## ARRIVAL

The key questions to be asked are: How do customers (patients) arrive for service (individually or in batch mode)? How are arrivals distributed over time (episodic or continuous)? And how do arrivals differ according to hour of the day, day of the week, or even season? In a radiology department, some arrivals are predictable and others are not. When analyzing the queue for general radiography, certain variations in arrival can be predicted reliably. For example, in large hospitals, orthopedic clinics tend to be high-volume users. It is possible to predict from experience with reasonable certainty that during the hours of orthopedic clinic operation a batch influx of patients will arrive for radiographic exams. Some of these patients will have been previously scheduled and can be anticipated, whereas others are unscheduled. Yet all of these arrivals can be predicted as part of planned departmental workflow, because the day of the week and time of the day for orthopedic clinic operation are known in advance.

Emergency department (ED) utilization of radiology services, on the other hand, is not as predictable. Referrals can be individual or in groups and can be episodic or continuous. However, a review of weekly or monthly statistics on these referrals may reveal patterns. Certain ED physicians regularly on duty may order radiology studies with more frequency than others. Certain times in the ED (trauma-heavy Friday and Saturday evening shifts, for example) may routinely produce more referrals. All these data can be compiled to provide a degree of predictability.

The simplest arrival process is one of predictable and regular arrivals, which translates into a constant time interval between successive arrivals. Only in rare instances in medicine (outpatient psychiatric therapy, for example) is the arrival distribution this tidy. Queuing analysis provides mathematical models, such as the Poisson distribution for arrivals, that can take

seemingly random data—the rate of real-life arrivals in a radiology department—and transform these into a single average arrival rate for specific days and times.

## SERVICE

The key questions to be asked are: How long does the service take for completion of the desired task (service time distribution)? What are the required resources needed to perform the service? Is the service performed in parallel (one queue for all servers) or in series (each server with his or her own individual queue)? And is preemption allowed (can a server reprioritize the queue)?

To analyze radiology workflow, we must assess how long a given exam will take, realizing that this time will differ depending on a dauntingly large number of factors, such as patient characteristics (inpatient on stretcher vs. ambulatory outpatient, for example), type of exam (5-view vs. 2-view lumbar spine radiographic exam, for example), individual technologist skills and habits, and technology (wall-mounted bucky unit vs. table unit, for example). Required resources can vary according to these and other variables. For example, an immobile stretcher patient will require multiple staff members to assist with transfer to and from the table, whereas the ambulatory outpatient can be accommodated quickly and efficiently by a single technologist. Preemption commonly occurs based on the clinical urgency of the individual patient. ED patients, for example, often take precedence over electively scheduled outpatients. Such preemption has significant effects on the queue status when ED utilization rates are high.

Although similar in some ways to time-motion studies, the queuing perspective on service takes into account multiple variables and again facilitates useful predictions about resources that will be needed at specific times.

## QUEUE CHARACTERISTICS

The key questions to be asked are: How does one choose the order of service (random, by order of arrival, priority according to clinical need, or the “squeaky wheel” phenomenon)? Do some customers (patients or referring clinicians) leave the queue if the wait becomes too long? Do customers (both patients and referring clinicians) represent a captive audience that is forced to remain in the queue regardless of service limitations?



Queue characteristics vary in institutional and imaging settings. For example, an outpatient imaging center's business model is largely predicated on rapid service. To remain competitive, queues must remain small with rapid service delivery. If they fail, they risk alienating patients, their families, and referring physicians. A VA hospital, on the other hand, deals with a capitated patient population that is in effect a captive audience. Longer queues are tolerated in such an environment because the resources are limited and the demand will remain high, even with longer queues.

## QUEUING ANALYSIS

To use queuing analysis to provide insight into management decision making in the allocation of personnel and technology, five key questions can be asked:

1. What is the average patient waiting time?
2. What is the average time to exam completion?
3. What is the probability of a patient having to wait longer than a specified time interval before service is initiated?
4. What is the average length of the queue? How does it vary over time?
5. What is the expected utilization rate of the individual server?

One example of the results of asking these questions can be found in a recent study comparing CR and direct radiography (DR) workflow costs. Despite the fact that DR was found to be significantly faster than CR in average examination time, the cost differential favored CR in most imaging settings. The differentiating factor was not employee time or other costs but the rate of flow of patients. The more expensive technology could be cost justified only when utilization rates per radiographic room exceeded 80%. Such rates occur only in busy outpatient centers with continuous, high-volume, scheduled patient arrival rates. For most hospital practices, the episodic and intermittent nature of patient arrival rates produces lower utilization rates and shorter queues in which the less expensive technology would be more cost-efficient. Only by analysis of waiting time and of the flow of the patient from waiting time to completion of study can such information be integrated with other workflow analyses to yield a complete picture of departmental productivity and efficiency.

## THE ROLE OF QUALITY ASSURANCE IN WORKFLOW AND PRODUCTIVITY

During the course of our professional training, all of us are taught the importance of quality. No other profession should be more quality oriented than medicine, where a lapse can affect every participant and every subsequent step in the medical imaging chain, from the scheduling of an examination to clinical management decisions made by the referring clinician based on imaging findings. Missteps can result in adverse outcomes as minor as longer waiting times for imaging and as major as misdiagnoses. Rigorous quality assurance (QA) measures, originally borrowed from industrial models, are now widely acknowledged as important but are not uniformly implemented throughout the medical enterprise. Quality assurance is often cited as the best way to ensure optimal patient care and decrease possible medicolegal consequences. An often overlooked benefit of QA, however, is its potentially positive effect on productivity and workflow.

When evaluating productivity and workflow within the medical imaging department, objective measures are typically quantity oriented rather than quality oriented. Productivity and workflow can be measured in a number of different ways, depending on both perspective and the intended uses for the information. For RTs, productivity is most often measured in exams per full-time equivalent, whereas for radiologists it is measured in relative value units. For the administrator, who takes a more global perspective on productivity and workflow, the foci of objective measures include time and effort spent in scheduling backlog and exam/report turnaround time. An emphasis on such quantitative measures (and the need for speed and efficiency) often has the unintended effect of worsening rather than enhancing quality. In fact, many supervisors and decision makers within the imaging department and healthcare enterprise view QA programs as cost inefficient.

The only imaging modality in which QA is consistently performed is mammography, because it is federally mandated through the Mammography Quality Standards Act. For most other modalities, QA is highly variable. Short of additional federal or state legislation, additional QA programs require some sort of economic justification for widespread adoption. Examples of such justification include reduced medicolegal risk (by lowering insurance premiums and decreasing the likelihood of large payouts), increased revenue (through improved customer satisfaction and increased referrals), or decreased expenses (by improved operational efficiency and productivity). Although the last justification is not commonly thought of as a direct result

of proactive QA programs, it is indeed a distinct possibility that can be validated though clinical research.

In an unpublished study evaluating digital radiography workflow at the BVAMC, 81% of total time differences between computed and direct radiography (CR and DR), was found to be the result of differences in post-acquisition processing, which includes image QA. This pilot study led in turn to a larger multi-institutional study, which took into account differences in digital technology, systems integration, and functionality. When controlling for these additional variables, 60% to 100% of the total time differences between CR and DR was attributed to postacquisition processing. At the same time, 33% to 57% of the total exam time for CR (which was the less efficient technology) was accounted for by postacquisition processing (Figure 5.7). This suggests that QA, as currently practiced, is an inefficient and time-intensive process. One way to reengineer workflow within the filmless imaging department would be to focus on QA, with a net result in considerable time and cost savings (and, of course, the important benefits of improved quality). In addition, these time savings would be felt most significantly at the all-important level of the understaffed technologist. To accomplish this, it is necessary to rethink several biases long associated with QA in film-based operation and think “outside the light box.” Three options for QA in the filmless environment are:

1. Eliminate QA entirely
2. Outsource QA responsibilities to a dedicated QA specialist
3. Automate QA with the development of automated QA software




FIGURE 5.7

Quality assurance and post-processing as a percentage of total examination time.

While the first option may seem heretical, in reality it is highly practical. As previously noted, the retake rate at BVAMC went down to only 0.8% when transitioning from film to filmless operation. In such a scenario, only 1 out of every 125 patients would effectively require a callback because of technical deficiency. Elimination of QA would free up 33% to 57% of technologist time, thereby improving technologist productivity, departmental workflow, and operational efficiency.

The second option requires the training of a dedicated QA specialist, which has the mixed effect of enhancing technologist productivity at the expense of creating a new full-time employee. The productivity gains of multiple technologists would more than cost justify this position. At the same time, changing QA workflow from single to batch mode would significantly improve departmental workflow while improving the overall consistency of the QA process (and potentially improving overall image quality).

The third option combines the best of all worlds and offers the future opportunity to further automate an existing manual process. Automation of QA has the potential to simultaneously improve image quality, education, and training (via a referenceable QA database), reduce payroll expenses, and improve productivity and workflow.

In the end, optimizing QA and workflow can be a synergistic process. To accomplish this, however, one needs to purge the biases of film-based operation and understand the capabilities of digital imaging. By removing QA as a back-end technologist function, significant productivity gains can be realized while improving image consistency and overall quality.

## CONCLUSION

Many, if not most, imaging departments have not significantly altered fundamental patterns of workflow since their film-based days. Despite significant changes in technology and personnel, these patterns have remained fixed. Technology, however, is more than the focus around which workflow changes must be made. It provides the means with which to improve workflow and optimize the contributions of the humans who, in the end, are still responsible for the work and the end product. If workflow (along with resulting patient care, personnel satisfaction, and overall profitability) is to be truly optimized, expectations must be changed. Individual tasks should be automated whenever possible and specialized when not. To accomplish this, each “assembly line worker” must have clearly defined tasks that he or she is trained to perform with a greater degree of efficiency than could otherwise be realized in the current environment of interchangeable parts and multitasking.

We must never forget, of course, that what distinguishes us from a truly industrial endeavor is that our end product is not a gadget or inanimate object but a dynamic entity, derived from a living, breathing person. Efficient workflow must take the human source into consideration and remember that the perception of service delivery is also predicated on the quality and timeliness of service delivery.

Nor should we forget that the lessons of QA borrowed from industry take on a special urgency and importance in medical imaging. No matter how efficient a process, it has failed if the end product is of inferior quality. We must continuously analyze and reanalyze the overall process, looking to refine any glitches that may have an adverse effect on quality. Current standards of practice provide valuable guidelines, but standards such as DICOM and HL7 can take us only so far. True integration requires a consensus such as that provided by the IHE effort about how we will use these standards to solve practical challenges, such as smooth workflow, communication of information, uniformity of image presentation, unidentified emergency room patients, and security.

## ► REFERENCES

- Andriole KP, Luth DM, Gould RG. Workflow assessment of digital versus computed radiography and screen-film in the outpatient environment. *J Digit Imaging*. 2002;15(suppl 1):124–126.
- Bartholmai BJ, Erickson BJ, Hartman TE, et al. The electronic imaging technology specialist: the role of a new radiology subspecialty for the 21st century. *Semin Roentgenol*. 2003;38:256–269.
- Beasley JE. Queueing theory. Operations research notes from Imperial College. Available at: <http://people.brunel.ac.uk/~mastjjb/jeb/or/queue.html>. Accessed on: August 3, 2005.
- Boochever SS. HIS/RIS/PACS integration: getting to the gold standard. *Radiol Manage*. 2004;26:16–24.
- Cabrera A. Defining the role of a PACS technologist. *J Digit Imaging*. 2002;15(suppl 1):120–123.
- Gocke P, Debatin JF, Durselen LF. Process management and controlling in diagnostic radiology in the hospital [in German]. *Radiologe*. 2002;42:332–343.
- Gross-Fengels W, Miedeck C, Siemens P, et al. PACS: from project to reality. Report of experiences on full digitalization of the radiology department of a major hospital [in German]. *Radiologe*. 2002;42:119–124.

- Guthrie BL, Price C, Zaleski J, Backensto E. The Digital Imaging and Communications in Medicine (DICOM) archive is a dynamic component of a clinician image-related workflow solution. *J Digit Imaging*. 2001;14(2 suppl 1):190–193.
- Hayt DB, Alexander S, Drakakis J, Berdebes N. Filmless in 60 days: the impact of picture archiving and communications systems within a large urban hospital. *J Digit Imaging*. 2001;14:62–71.
- Health Level Seven, Inc. HL7. Available at: <http://www.hl7.org>. Accessed August 15, 2004.
- Honeyman-Buck J. PACS adoption. *J Digit Imaging*. 2002;15(suppl 1):184–188.
- Janower ML. 1988. Productivity standards for technologists: how to use them. *Radiology*. 1988;166:276–277.
- Kolb G. Going filmless: lessons from a Swedish breast center. *Decis Imaging Econ*. [serial online]. August 2004. Available at: <http://www.imagingeconomics.com/library/200408-08.asp>. Accessed August 15, 2004.
- Kuzmak PM, Dayhoff RE. The use of digital imaging and communications in medicine (DICOM) in the integration of imaging into the electronic patient record at the Department of Veterans Affairs. *J Digit Imaging*. 2000;13(2 suppl 1):133–137.
- Kuzmak PM, Dayhoff RE. Experience with DICOM for the clinical specialties in the healthcare enterprise. *Proc SPIE Med Imaging*. 2003;5033:18–29.
- Langen HL, Bielmeier J, Wittenberg G, Selbach R, Feustel H. Workflow improvement and efficiency gain with near total digitalization of a radiology department [in German]. *Rofo Fortschr Geb Rontgenstr Neuen Bildgeb Verfahr*. 2003;175:1309–1316.
- Locko RC, Blume H, Goble JC. Enterprise-wide worklist management. *J Digit Imaging*. 2002;15(suppl 1):175–179.
- MacEwan DW. Radiology workload system for diagnostic radiology: productivity enhancement studies. *Radiology Manage*. 1982;5:24–38.
- May GA, Deer DD, Dackiewicz D. Impact of digital radiography on clinical workflow. *J Digit Imaging*. 2000;13(2 suppl 1):76–78.
- Mazurowski J. Informatics management and the PACS administrator. *Radiol Manage*. 2004;26:41–42.
- Mulvaney J. The case for RIS/PACS integration. *Radiol Manage*. 2002;24:24–29.
- National Electrical Manufacturers Association. DICOM: Digital Imaging and Communications in Medicine. Available at: <http://medical.nema.org>. Accessed August 15, 2004.
- Nissen-Meyer S, Holzknrecht N, Wieser B, et al. Improving productivity by implementing RIS and PACS throughout the clinic: a case study [in German]. *Radiologe*. 2002;42:351–360.

- Ouvry A. PACS implementation dramatically impacts people and radiology work processes. *Proc SPIE Med Imaging*. 1997;3035:602–603.
- Radiological Society of North America. Integrating the Healthcare Initiative. Available at: <http://www.rsna.org/IHE>. Accessed August 15, 2004.
- Redfern RO, Horii SC, Feingold E, Kundel HL. Radiology workflow and patient volume: effect of picture archiving and communication systems on technologists and radiologists. *J Digit Imaging*. 2000;13(2 suppl 1):97–100.
- Redfern RO, Langlotz CP, Abbuhl SB, Polansky M, Horii SC, Kundel HL. The effect of PACS on the time required for technologists to produce radiographic images in the emergency department radiology suite. *J Digit Imaging*. 2002;15:153–160.
- Reiner B, Siegel E, Carrino JA. Workflow optimization: current trends and future directions. *J Digit Imaging*. 2002;15:141–152.
- Reiner B, Siegel E, Carrino JA, McElveny C. SCAR Radiologic Technologist Survey: analysis of technologist workforce and staffing. *J Digit Imaging*. 2002;15:121–131.
- Reiner B, Siegel E, Scanlon M. Changes in technologist productivity with implementation of an enterprisewide PACS. *J Digit Imaging*. 2002;15:22–26.
- Reiner BI, Salkever D, Siegel EL, et al. Multi-institutional analysis of computed and direct radiography. Part II. Economic analysis. *Radiology*. 2005;236:420–426.
- Reiner BI, Siegel EL, Carrino JA, Goldburgh MM. SCAR Radiologic Technologist Survey: analysis of the impact of digital technologies on productivity. *J Digit Imaging*. 2002;15:132–140.
- Reiner BI, Siegel EL, Hooper FJ, Glasser D. Effect of film-based versus filmless operation on the productivity of CT technologists. *Radiology*. 1998;207:481–485.
- Roos JE, Desbiolles LM, Willmann JK, Weishaupt D, Marincek B, Hilfiker PR. Multidetector-row helical CT: analysis of time management and workflow. *Eur Radiol*. 2002;12:680–685.
- Saarinen AO, Wilson MC, Iverson SC, et al. Potential time savings to radiology department personnel in a PACS-based environment. *Proc SPIE*. 1990;1234:261–269.
- Sacco P, Mazzei M, Pozzebon E, Stefani P. PACS implementation in a university hospital in Tuscany. *J Digit Imaging*. 2002;15(suppl 1):250–251.
- Schmidt J, Meetz K, Wendler T. Workflow management systems—a powerful means to integrate radiologic processes and application systems. *J Digit Imaging*. 1999;12(2 suppl 1):214–215.
- Schoppe WD, Hessel SJ, Adams DF. Time requirements in performing body CT studies. *J Comp Assist Tomogr*. 1981;5:513–515.
- Sferrella SM, Story CP. The impact of tech aides in radiology. *Radiol Manage*. 2004;26:22–30.

- Siegel EL, Diaconis JN, Pomerantz S, et al. Making filmless radiology work. *J Digit Imaging*. 1995;8:151–155.
- Siegel EL, Reiner BI. Challenges associated with the incorporation of digital radiography into a picture archival and communication system. *J Digit Imaging* 1999;12(2 suppl 1):6–8.
- Siegel EL, Reiner BI, Severance S, et al. Comparison of technologist productivity for chest radiography when using CR and DR: importance of workflow optimization. Paper presented at: Annual Meeting of the Society for Computer Applications in Radiology; May 3, 2002; Cleveland, OH.
- Trisolini MG, Baswell SB, Johnson SK, et al. Radiology work-load measurements reflecting variables specific to hospital, patient, and examination: results of a collaborative study. *Radiology*. 1988;166:247–253.
- Wendt G, Peppler W, Edwards W, Parisot C, Lindop C. Clinical implementation of the IHE presentation of grouped procedures integration profile in a multi-vendor environment—workflow modification and barriers to implementation. *J Digit Imaging*. 2002;15(suppl 1):64–66.
- Wong STC, Tjandra D, Wang H, Shen W. Workflow-enabled distributed component-based information architecture for digital medical imaging enterprises. *IEEE Trans Inf Technol Biomed*. 2003;7:171–183.



# REENGINEERING WORKFLOW: THE RADIOLOGIST'S PERSPECTIVE

---

---

ELIOT L. SIEGEL • BRUCE I. REINER • NANCY KNIGHT

*Busy day in the department. Morning: read several kidney studies and consulted on a fx. Home for lunch: stopped at drug store. Cut grass. Afternoon: read 3 studies. Organized cases. Home at 4:00. Took Lucille to see Forbidden Planet.*

May 1956 entry in the diary of Col. William LeRoy Thompson, MD  
First Chair and Registrar of the Department of Radiologic Pathology  
at the Armed Forces Institute of Pathology, Walter Reed Army  
Medical Center

**A**sk almost any observer what has changed most about the practice of radiology in the last half century and the immediate answer will be the technology. New modalities, interventional techniques, and the digitalization of almost every aspect of image acquisition, retrieval, processing, reporting, and

archiving have profoundly altered the look of the imaging department. Ask a radiologist the same question, however, and he or she will respond that the most fundamental and challenging change in actual practice has been in the pace of work demanded of the individual who interprets the images. Gone are the days of a few studies in the morning, home for lunch, and a few more in the afternoon, as Colonel Thompson outlined in his diary at Walter Reed in 1956. Almost gone are the days of single- or two-view studies, interpreted on film and returned with dictated and manually transcribed reports to the referring physician.

Instead, an extraordinarily rapid and exponential growth in the number of images that constitute a single study, as well as the performance of numerous studies per patient, has multiplied the daily total of images presented to many radiologists by factors of not tens or hundreds but thousands. Tenacity and creativity have joined sensitivity and specificity as necessary metrics of radiologic success, as once-daunting 8-channel computed tomography (CT) datasets are replaced by those generated from 64-channel studies, and modalities from magnetic resonance (MR) to ultrasound (US) to fusion techniques in nuclear imaging yield increasingly large and complex groups of images. As Horii and others have pointed out, the size of these datasets has actually accomplished what logic and documented successes sometimes failed to do. These images form the “first group of examinations that cannot practically be printed to film for interpretation.” A 1000-image CT examination—now commonplace in many institutions—would require a minimum of 67 film sheets to print (using a  $4 \times 4$  matrix on film) and would take up 17 panels of a 4-light-box mechanical film changer for each window/level combination. The image explosion has made the transition to filmless imaging mandatory for many who once considered it an interesting future option.

The result over the past 5 years has been an increasing focus on workflow issues relating to the essence of radiologic practice: the process that occurs at the interface between the interpreter and the image. The literature documenting radiologist workflow issues is growing. Of course, the tasks of the radiologist at the workstations are part of a larger workflow and depend on a number of factors, including effective picture archiving and communication systems (PACS) integration with the radiology and hospital information systems (RIS and HIS, respectively), worklist management, workstation design, and innovations in the interpretation process, reporting, and interactions with clinicians and the larger medical enterprise. Less studied but equally important are the effects of room design and ergonomics on radiologist workflow and productivity as well as the need for reliable metrics and tools by which such productivity can be assessed and compared.

Because PACS serve increasingly as the nexus and conduit for the work of the radiologist, a number of other chapters in this book cover in greater detail the technical elements that make up routine workflow. Our goal in this chapter is to provide background, overview, and resources on current challenges and benefits associated with various elements of radiologist workflow.

## BACKGROUND

For a number of reasons, including the fact that radiologists are the most expensive members of the imaging department's staff, interpretation workflow has been studied for many years to determine which factors influence productivity and accuracy. These studies took on additional importance when quantifiable results were needed to bolster the transition from film-based to filmless imaging in many institutions. Studies performed at the Baltimore Veterans Affairs Medical Center (BVAMC) documented an increase in radiologist productivity by more than 50% over the course of such a transition. This marked improvement occurred despite the fact that radiologist reading times decreased only slightly, by approximately 8% to 15%. The improvement was believed to be the result of a combination of complex factors, including more effective sharing of the workload by the radiologists, fewer interruptions, immediate availability of old examinations and reports for comparison, and the elimination of the film library and the inefficiencies and time delays associated with it (see Chapter 5 for information on departmental workflow effects that contributed to overall increases in productivity). Since the time of our original report, many other groups of radiologists have reported similar increases in productivity during the transition to filmless operation with an enterprise-wide PACS. Although this transition has been made in a variety of settings, from academic to private practice and in countries with varying reimbursement, personnel, and patient characteristics, some elements of success—as well as continuing challenges—remain constant.

## PREPARING FOR INTERPRETATION: WORKLISTS, ARCHIVE ACCESS, AND INITIAL DISPLAY

Among the most important determinants of radiologist and clinician performance is the time required for a PACS workstation to retrieve one or more imaging studies, to display them for interpretation, and to present a

suitable choice of workstation tools that allow the radiologist to glean as much clinically significant information as possible about the images. Studies that we have performed at the BVAMC and the University of Maryland have suggested that the first image from an 8-megabyte computed radiography study should take less than 3 seconds to display. Cross-sectional images, such as  $512 \times 512$ -pixel CT images, should display at a rate of at least 5 per second or faster. Our data suggest that image display performance that is significantly slower may result in a significant decrease in radiologist reading speed and produce increased levels of fatigue.

Manufacturers of PACS originally took one of two general approaches to the delivery of images to a radiologist's workstation. In the first, images were stored on a single, large, short-term storage unit, typically a redundant array of inexpensive disks (RAID). With this approach, workstations communicated directly with this centrally located storage device over a very fast network for image retrieval. The major advantage of this approach was the flexibility to rapidly retrieve any images at any location for both radiologists and clinicians. The major disadvantage was that it required a fast network in which single points of failure posed threats to continuous operation. Such systems were vulnerable to a major loss of function in the event of failure of the short-term storage device.

The alternative PACS architecture for delivery of images utilized a model that more closely emulated (and had many of the disadvantages of) a film-based environment in which films were sent to or placed on a film alternator. With a PACS, this was accomplished electronically by routing appropriate images to one or more workstations that were most likely to be used to review those image studies. Images to be read and comparison studies were stored locally on the hard drives (local storage) of the workstations themselves. Images could be intelligently routed to any number of workstations most likely to retrieve these studies. For example, all CT examinations could be routed to one or more workstations dedicated to interpretation of CT examinations or to the workstations or radiologists likely to read those studies on any particular day. Relevant comparison CT, general radiographs, or other studies deemed to be likely to be needed for comparison could also be routed to those workstations automatically, using predefined rules determined by the radiologists. The major advantage of such a system was independence from failure of any one component of the PACS or even the network itself. The disadvantage was the difficulty in selecting rules that would anticipate often unpredictable and spontaneous requirements of radiologists for comparison studies and older imaging studies that could be requested by clinicians for review. The other disadvantage of this architecture was the tendency for workstations to require a greater amount of

local storage as well as the inefficiencies of storing images in multiple locations.

Fortunately, most PACS now combine the two approaches to minimize disadvantages and optimize benefits. (The currently available choices of PACS architectures are discussed in detail in Chapter 13.) For example, PACS that use a central hard drive (or RAID) may also employ mirrored or backup systems to further decrease the likelihood of a general system failure. Those systems that use local workstation storage now create more central “nodes” or short-term image servers that store images for a cluster of workstations. This can result in more efficient storage and retrieval and decrease the need for very specific algorithms for routing images to a particular workstation. It is likely that this trend will continue in the future and will significantly blur the differences between the two approaches to short-term image storage and distribution.

The combination of the use of modality worklists, display default protocols, and a fast image retrieval and display system now provides the ability to customize the radiology workflow process. Given these workflow and performance features, radiologists should be substantially more productive with less fatigue than in the conventional, film-based environment.

## WORKLISTS AND WORKFLOW

Other chapters in this book address questions of workstation accessibility, security, and sign-on. Identifying the appropriate images to be read and connecting these images with the appropriate patient data and relevant priors, however, have always constituted the first major step in the radiologist workflow process. This step has changed profoundly within the last decade.


To maximize the efficiency of the workflow of diagnostic radiologists, early PACS adopters discovered that it was important to achieve paperless as well as filmless operation. Even today, the efficiency of many PACS implementations is hampered by the failure to eliminate paper from the radiologists' (and others') workflow process. In paper-based departments, radiologists began their work when presented with a stack of forms that contained information about the patient, examination, and reason for the study, with limited additional information, such as the name of the ordering physician or service, patient location, and so on. The radiologist then would use the information on these forms to enter patient or study identification manually or use a barcode reader or other mechanical tool to identify the patient and study to the workstation.

In the filmless department, these steps (and the attendant high rate of introduced error) have been eliminated by the implementation of worklists that define the type of unread studies to be presented for interpretation and their order of presentation. Worklists are shared, enterprise-wide rosters of unread studies, which, through the PACS, are integrated with a wealth of additional information. The worklist acts as a database filter that allows a radiologist to view images defined by anatomic regions (e.g., chest, neuro-radiology, or musculoskeletal) or modality (e.g., ultrasonography, nuclear medicine, angiography, or special procedures) or any combination of these. The advantages of worklists for radiologists include the ability to sign in at any location and have full access to all unread examinations in their area or areas of expertise at any time, the ability for multiple radiologists to share responsibility for reading similar types of studies, and the performance improvements associated with the elimination of manually keyed or bar-coded patient or study information.

The physical time savings for all radiology personnel and especially for the radiologist are evident. However, smooth implementation of shared worklists across the enterprise were hindered for a number of years by problems with interfaces between and among the RIS, HIS, PACS, and other hospital or practice information systems. The development and acceptance of the Digital Imaging and Communications in Medicine (DICOM) worklist standards, which offer guidelines for electronic communication between the imaging department and other parts of the hospital enterprise, provided the missing link that many institutions needed to begin to successfully implement worklists at the radiology workstation. The Integrating Healthcare Enterprise (IHE) initiative of the Radiological Society of North America and the Healthcare Information and Management Systems Society has addressed many of the remaining challenges associated with the lack of plug-and-play compatibility that resulted from the substantial flexibility (or looseness) of the Health Level Seven (HL7) and DICOM standards and those of other information technology systems.

Implementation of DICOM modality worklist management software has been reported to reduce input errors from 6.4% to 0.1%. In another study, a pre-PACS CT transmission failure rate of 7.6% (largely the result of human error in data entry) was reduced to 3.5% after the addition of DICOM worklists, with a much smaller portion of that percentage accounted for by human error (Figure 6.1).

Worklists can also be generated using an algorithm designed to prevent studies from being read and reported more than once (overreading), thus increasing overall radiologist productivity and reducing the possibility of conflicting recommendations. They can also be customized so that images can be used in resident or teaching review.

**FIGURE 6.1**

The CT transmission error rate was cut dramatically after the introduction of modality worklists.

## PRE-FETCH

One of the workflow rate-limiting steps in the process of image retrieval and display is related to the fact that retrieval of images from long-term storage from an optical, magneto-optical, or tape archive is quite slow. In fact, retrieval times can be 10 to 100 times slower from long-term than from short-term storage, depending on the PACS architecture and equipment. Archiving and retrieval are addressed in other chapters of this volume.

Despite the more widespread use of RAID for both short- and long-term storage associated with substantial cost reductions in “spinning disk” archives, strategies to minimize the likelihood of a delay in image retrieval remain important elements in well-planned workflow. Such strategies use a set of algorithms (rules) that attempt to maximize the likelihood that the required images are available in short-term storage. The goal is to have the optimal number of relevant priors available without initiating unnecessary transfers from long-term storage. This process is an excellent example of the advantage of a PACS-HIS-RIS that forms an “intelligent” system, and numerous algorithms have been suggested and investigated.

One of the most straightforward examples of image pre-fetch is storage of new and historic examinations locally at a workstation. With this PACS architecture, images and predefined prior studies (for example, the last two studies that match both modality and anatomic location) are routed to a particular workstation or workstations. This pre-fetch strategy can also be used in a system in which workstations share a single RAID server. In this type of system, predefined relevant priors are retrieved from long-term storage automatically when a new imaging study is performed.

Other pre-fetch strategies can substantially increase the possibility that images that are likely to be needed by the radiologist or clinician are available on a local workstation or server. For example, image pre-fetches can be triggered based on a scheduled or new admission to the hospital, a scheduled outpatient appointment, or a transfer of a patient from one location to another within the hospital. Our analysis of the RIS database at the BVAMC indicated that a relatively small number of studies can be pulled to achieve a high likelihood that the required studies will be available on local storage. We found that if a PACS retrieved the most recent 30% of a patient's previous examinations into a short-term storage area before an outpatient appointment, there was a 91% probability that the required images would be available on the server rather than in the long-term archive. Such a pre-fetch strategy is the digital and much less labor-intensive equivalent of pulling film jackets in advance of outpatient visits and can be very effective in optimizing radiologist and clinician workflow in the review of imaging studies.

## HANGING PROTOCOLS/DISPLAY

One of the more complex processes in reading conventional film was the arrangement of images from current and previous studies on a view box or film alternator. In a film-based department, radiologists typically functioned in one of two ways. In the first, the radiologist was responsible for taking a new study and finding comparison studies from a film jacket. Often, the film librarian placed these outside the film jacket. The radiologist then took these examinations and arranged them on a series of view boxes. The radiologist then found any relevant old reports, interpreted the study, took the films back down, and (with luck) placed them back into the correct film jacket. In the second mode, the study to be read was placed on a film alternator with any relevant films and reports. The competent film librarian learned how each radiologist preferred to have his or her films arranged for specific studies. Having the fileroom personnel arrange the films resulted in improved workflow for the radiologists but required additional time and fileroom staff. The hanging and removal of the studies also created delays in radiologist workflow, because the radiologist often had to wait for studies to be hung, then taken down, and new studies to be put up. The general in-house criteria that determined how film would be placed on film alternators or view boxes were referred to as "hanging protocols" or "display protocols."

A PACS softcopy workstation automates many of these manual workflow steps, eliminating delays in the display of imaging studies. Hanging pro-



ocols can be executed much more rapidly and reliably and can be customized to the specific demands of individual radiologists at a single workstation or any workstation across the enterprise. PACS hanging protocols can be relatively simple (new studies on the right, older ones on the left) or quite complex. The system can define, for all users or for specific radiologists or clinicians, specific rules that determine which previous studies, if any, are retrieved for comparison and precisely how current and prior studies are displayed for interpretation. Images can be displayed, for example, in frame mode, which closely emulates film (nine images on one monitor), or in stack mode (a series of images displayed sequentially, much like viewing animation by flipping through a stack of cards). The need for these hanging protocols is even greater with more complex examinations, such as a large thin-section CT dataset in which multiplanar, three-dimensional (3-D), or maximum-intensity projection (MIP) views with variable slice thickness may be selected for current and comparison study.

The PACS at the BVAMC uses a series of algorithms for display on a multimonitor workstation. These are known as default display protocols (DDPs). The use of the DDP, which can be toggled off or on, was found to result in an increase in radiologist productivity of between 10% and 20%, depending on the imaging modality (Figure 6.2). In addition, radiologists reported less fatigue subjectively with the use of the DDP in comparison with electronic but manual selection of prior studies to be retrieved and manual (electronic) or nonintelligent placement of the images on the workstation. Reading times were also decreased somewhat by the reduced amount of time required to review previous reports. Using the PACS, previous




FIGURE 6.2

Radiologist reading times for general radiography decreased by 10% using the default display protocol.

reports are organized in chronological or another organized format to make review of priors rapid and their relationships easily understood.

## THE INTERPRETATION PROCESS

Radiologist image interpretation speed has been only one of a number of factors that have resulted in increased productivity, but the workflow improvement that has been associated with PACS in this area has been significant. At the BVAMC, we found that radiologist reading times decreased by 19% in the interpretation of portable chest radiographs from the intensive care unit (unpublished data). Another study performed at our facility demonstrated that radiologists were 8% faster in the interpretation of musculoskeletal radiographs using computer workstations and computed radiography in comparison with interpretation using conventional film. Similarly, radiologists were found to require 15% less time to interpret CT examinations using a computer workstation than using film. This was, for the most part, associated with the decreased amount of time required to display images, particularly in multiple window/level combinations. The advantage of softcopy interpretation over film for CT studies was even greater for examinations in which there were previously performed CT examinations for comparison. This increased speed of CT interpretation was not associated with any decrease in the accuracy of interpretation; in fact, accuracy increased to a statistically significant degree. Others have documented similar decreases in total interpretation time.

## WORKSTATION TOOLS

The retrieval of images to the workstation is only the first of several workflow steps in the interpretation of an imaging study. To extract as much clinically useful information as possible from the images, a number of steps may be helpful:

1. Images must be optimized with regard to window/level (brightness/contrast) settings. There is no optimal window/level for most images. Consequently, continuous dynamic adjustment of window/level settings is often necessary for a conventional radiograph (such as a foot examination). Alternatively, certain presets may be used (as would be typical for a thoracic CT).

2. The method of image display and navigation must be chosen. The simplest of these is static softcopy interpretation, with images displayed on workstations much as they would have been on view boxes. Frame mode is an example of this type of static mode in which images are displayed in a matrix similar to that typically printed to film. Stack mode displays images sequentially in a single window in a movie- or cine-like format. Linked stack mode, a further enhancement that synchronizes multiple stacked images within a single examination or across a current and one or more prior examinations, is increasingly available and utilized, although its value has not been adequately documented in the literature. Most recently, volumetric navigation of isotropic CT datasets permits review of images in any desired two-dimensional (2-D) or 3-D perspective, resulting in a separation of the manner in which images were acquired (the axial plane for CT, for example) from the way in which they are reviewed.

3. Images and portions of images can be zoomed or magnified.

4. Images can be viewed using MIP, which has been documented to be useful in the evaluation of blood vessels and lung nodules.

5. Thin-section images can be combined arithmetically to create a user-selected slice thickness that is a multiple of that reconstructed by the acquisition device.

6. Images can be arranged in a logical format to make it as easy as possible to compare various sequences (e.g., enhanced vs. unenhanced, or T1 vs. T2 vs. contrast-enhanced MR images) and to compare a current study with comparable images from a previous study.

7. Images can also be enhanced with tools, such as edge enhancement, smoothing or interpolation algorithms that “smooth” the image to give it a less boxy or “pixely” appearance, or those that enhance the ability to display a wide range of contrast on an 8-bit monitor or film.

8. Images can also be processed using more sophisticated techniques to achieve spatial frequency and image contrast optimization.

9. Additional tools can be implemented to aid in decision support, including computer-aided detection tools that have been successfully applied in mammography and in the detection of lung nodules on CT or conventional chest radiography.

Each of these steps depends to some degree on the preferences of the radiologist and on the demands of the specific study and modality. For many radiologists, both experienced and in training, one of the greatest current challenges is in identifying these preferences as the range of choices expands

and technology evolves. The “intelligent” workstation and the PACS that supports it not only must be ready to customize different combinations of workstation tools to suit each user but must be configured to seamlessly integrate new software that enhances the interpretation process.

Current PACS workstations vary tremendously in the success of their graphical user interface and in the number of steps required to utilize these and other tools. Most workstations do a relatively poor job of optimizing radiologist workflow. The best of these workstations have a relatively simple (elegant) graphical user interface and require a minimum number of key-strokes and steps to retrieve, optimize, compare, and remove a study, and then proceed to the next imaging study. As the PACS industry continues to develop and mature, vendors are spending an increasing amount of time obtaining feedback and performing studies of radiologist workflow in the interpretation of imaging studies, which has resulted and will continue to result in improvements in the radiologist-machine interface.

We have found that the use of workstation tools by radiologists changes with increasing experience with the system. Radiologists have a tendency to use tools such as image zoom and magnification less frequently as they gain additional experience with the workstation. However, we have found that even experienced radiologists utilize the window/level adjustments in the majority of cases.

## WORKFLOW TOOLS TO COPE WITH IMAGE OVERLOAD

The use of volumetric navigation has been accelerated by the rapid transition to the use of multidetector CT scanners. Radiologists around the world are finding other interpretation routines inadequate for the large numbers of images generated from these systems. A routine CT of the thorax using a multidetector system can generate 30 sheets of film each for lung, mediastinum, and liver settings. Even stack mode is inadequate for review of the 300 to 500 images acquired for a routine CT of the chest or the abdomen and pelvis and is even more so for the 1500 to 2000 images acquired for a CT angiography “runoff” study.

Several strategies are being investigated for dealing with this image overload. The most common is to acquire images using a multidetector scanner using thin-collimation and then reconstruct the images that are sent to the PACS using much thicker (e.g., 5- or 8-mm) sections, resulting in a 3- to 10-fold reduction in the number of images sent to the PACS. Additional reconstructions or renderings can then be performed by the technologists using a dedicated CT workstation. Unfortunately, this approach is

unsatisfactory. It requires a large amount of additional technologist time, especially for angiographic rendering, analogous to the extra time required for technologists to produce films in multiple window/level settings. Because of the complexity and time required, technologists only perform this rendering in a small percentage of cases. In addition, the reconstructed images unnecessarily take up a good deal of archival, network, and workstation memory space.

Radiologists should have flexibility from case to case to determine whether the images should be reviewed in the sagittal, coronal, or oblique planes or using a 3-D perspective. Volumetric navigation frees the radiologist from the limitations of fixed-slice axial images. An image of the spine, for example, can be rapidly and interactively rendered and reviewed as a sagittal or coronal dataset at any desired slice thickness. The viewing perspective can be determined by the area being examined and the clinical history. The pulmonary arteries, for example, can be reviewed using relatively thick-slice coronal or oblique perspectives, with or without the use of MIP rendering. The colon, in our experience, is best depicted in the coronal plane, whereas the liver and spleen may be examples of organs best reviewed in the axial plane but may be improved with the use of MIPs. The vasculature of the thorax, abdomen, pelvis, and other areas may be optimally displayed according to their orientation within the body and are also probably best rendered as MIP images.

Although volumetric navigation has tremendous potential, it poses some unique and daunting challenges as well, especially the concern that we might be trading image volume overload for clinical image content overload. Our abdominal and thoracic subspecialists have asked whether they are now responsible for detailed reports of the musculoskeletal system and spine and of the individual vessels now visualized on a routine body CT study. Should they specifically and routinely comment, for example, on the renal arteries, aortic and iliac arteries, or superior and inferior mesenteric arteries? What are the implications of this on the time required to dictate a study?

Perhaps the biggest barrier to the transition to the use of volumetric navigation has been the lack of integration of this capability in the current generation of PACS workstations. It is not practical for a radiologist interpreting a study using a PACS workstation to walk over to a dedicated 3-D/multiplanar workstation for each case. Another challenge is the fact that image navigation is typically not a linear, sequential process like review of a set of axial images but may be performed in a more haphazard fashion, with a radiologist reviewing a portion of a dataset in one plane and other portions using other views, which could result in portions of a dataset not being reviewed at all.

The transformation of the radiology interpretation process will continue at a rapid pace. Although image navigation and enhancement will continue to improve (including better support for multimodality fusion such as positron emission tomography [PET]/CT), the next major phases will focus on decision support tools such as computer-assisted diagnosis (CAD) and intelligent applications of informatics. Computer-assisted imaging may take many forms, including an overlay in which microcalcifications are circled on a mammogram and lung nodules appear in a color that indicates their probability of malignancy. CAD programs will come into routine use in the next few years, especially in the detection of lung nodules and breast cancers. Clinical information from the electronic medical record, results of previous examinations, and clinical and imaging expert systems will be utilized to optimize image navigation, computer imaging, and CAD programs and to suggest diagnostic possibilities. Comparison with large computerized reference image datasets may also be utilized routinely by radiologists to facilitate more rapid and accurate diagnoses. These future additions to the armamentarium of the radiologist will also create additional challenges that will undoubtedly require the creativity and expertise of the medical imaging community.

## CHANGING THE REPORTING PARADIGM

### SPEECH RECOGNITION

The interactive workstation has facilitated one of the most obvious changes in radiologist workflow over the last decade: the way in which radiology reports are generated, reviewed, and relayed to the referring physician and to the medical record. Machine dictation of radiology reports and subsequent manual transcription were among the first “automated” elements in the radiology workflow. At the same time, manual transcription accounted for the main time lag between image acquisition and delivery of interpretation results to referring clinicians. As early as the 1980s, radiologists were considering the possibility that speech recognition techniques might be incorporated into the reporting process.

In the intervening 2 decades, speech recognition (sometimes misidentified as voice recognition) has become an integral part of the radiology workflow, with many institutions eliminating medical transcription positions entirely. Speech recognition has been investigated more closely than almost any area of digitalization for two reasons: it affects every imaging specialist,

across modalities and subspecializations, and it has often been a “hard sell” to the radiologists who (in theory) would benefit most from its implementation. The literature on speech recognition is voluminous, with many reports of benefits in decreased needs for auxiliary staff, money savings, and time savings in turning around reports for delivery to physicians who ordered studies.

The process of acceptance of speech recognition technologies by radiologists has not been smooth. In part, this was because it was unfamiliar. More important, many radiologists perceived speech recognition reporting to be more difficult, more time consuming, and part of a slippery slope that seemed to be taking clerical tasks out of the hands of paid assistants and turning them into routine parts of the professional interpretation and reporting process. (They were, of course, correct in the last of these assumptions.) Efforts by vendors to simplify enrollment (the process by which an individual imprints his or her speech patterns on the system), increased use of report templates, well-executed training, and the introduction of innovative time-saving features have served to win over many who originally opposed the introduction of speech recognition. Advantages in report turnaround times, the ability to correct and redact reports at the time of dictation or subsequent time of choice, access to the report through the PACS from anywhere in the medical enterprise, and a tendency toward the production of shorter, more organized reports have bolstered support for speech recognition among radiologists. For those who use speech recognition in combination with structured reporting technologies, real-time savings and workflow advantages are being realized.

## STRUCTURED REPORTING

Structured reporting is not new. Since the first “roentgenology” reports in the late 19th century, imaging specialists have sought to simplify their work by eliminating unnecessary duplication from report to report. In the earliest days, there were fill-in-the-blank reporting forms. Later, radiologists read from templates as part of dictated reports, filling in the specific information for each patient. Today, templates and macros integrated into the electronic radiology reporting process allow each radiologist to customize routine reporting and enter large sections of reports using only a few keystrokes. Moreover, templates can be used in “batch mode” for high-volume study reporting.

Structured reporting is now being combined with other workstation interpretation tools to yield what some have called the radiology report of the future: a multimedia package that includes not only the traditional report but embedded annotated images, lists of and links to additional informational and visual resources, and cues and guidance provided by computer-aided decision support strategies. The possibilities for entirely transforming both the radiology reporting process as well as elevating the level of content in reports are truly exciting. The challenge will lie in exploring these possibilities in ways that enhance rather than add to the work of the radiologist.

## THE EFFECT OF CHANGED RADIOLOGIST WORKFLOW ON ENTERPRISE INTERACTIONS

### REFERRING PHYSICIANS AND CLINICIANS

The transition to filmless operation at the BVAMC was associated with an 82% reduction in in-person consultation rates in general radiography and a 44% reduction for the cross-sectional imaging section, despite an increase in the volume of studies (Figure 6.3). This decrease in the general radiography consultation rate from 13% (pre-PACS) to 2.4% was greater than we had anticipated. The ability of clinicians to remain in patient care areas and not make time-consuming trips to the radiology department has significantly changed their workflow as well. The direct consultation process has been




FIGURE 6.3

The advent of PACS and image distribution throughout the enterprise brought an 82% reduction in clinical consultations on general radiology cases.



transformed into an electronic process that relies to a greater extent on digital annotation of images viewed at workstations, access to digital dictations over the telephone, and increased use of e-mail and physician alerts available in the EMR.

The increased frequency of image review by physicians at our institution was associated with an unexpected increase in radiology utilization rates and, perhaps ironically, in a consequent increase in overall radiology workload. The number of studies ordered per patient admission at our institution increased by 43% after the implementation of digital distribution of images and reports, compared with a 0% increase for the rest of the VA hospitals throughout the country during the same period. Outpatient utilization at our institution increased by 21% during the same period, while national VA outpatient utilization decreased.

Clinicians adjusted with great alacrity to decreased report turnaround times after radiologists began reading studies within a few minutes of acquisition and came to expect these accelerated turnarounds as a part of radiology department services. Ninety-eight percent of clinicians surveyed at the BVAMC indicated that the use of the PACS contributed to more effective use of their time. This was largely due to the improved access to current and previous imaging studies and convenient availability of access to these images in patient care areas. Clinicians indicated that they accessed the PACS 3 to 5 times per day, with 22% accessing the system more often. The average estimate of the amount of time saved because of the PACS, according to clinician surveys, was approximately 50 minutes, suggesting that the system substantially enhanced their workflow. In some institutions, clinicians can now access the digital dictation system by phone for radiology reports as well as read near real-time reports distributed after speech recognition sign-off on the PACS.

A not inconsiderable amount of concern has been expressed (for the most part in editorials and not in data-driven studies of the problem) about the result of the changed radiologist workflow and output on the quality of interaction with clinicians. It is true that the number of face-to-face interactions over specific cases has been greatly diminished in the all-digital radiology department. However, other technological innovations may be serving to compensate for what some perceive as the radiologist's diminished presence in the diagnostic and treatment equation. For example, radiologists are much more likely to use the telephone to discuss urgent findings, such as a pneumothorax in the intensive care unit (ICU), than before the transition to filmless operation. This is because the decreased time between study acquisition and review makes it much more likely that the radiologist will be the first to review the images. Thus, many clinicians are receiving real-time

reports from their colleagues in radiology about their most urgent and/or problematic cases.

Moreover, in a reversal of the old consultation pattern that had the clinician coming to the radiologist, the radiology department is now, in effect, distributing itself throughout the enterprise for immediate access to a wide range of practitioners. PACS workstations with integrated access to the RIS in surgical suites, ICUs, emergency departments, and satellite clinics have eliminated an entire segment of routing workflow and put the results of the radiologist's work at the immediate service of the clinician where he or she needs it most. DeSimone et al. evaluated the impact of PACS on clinical practice in the ICU setting and showed statistically significant reductions in time to perform clinical actions after the diagnostic examination. Using PACS, significant alterations were demonstrated in the processes of obtaining radiologic information, viewing exams, and consulting between ICU physicians and radiologists. The results of the study suggest that a PACS has a major effect on both patient management and radiology department workflow in the ICU setting.

One example of this improved workflow has been in the communication of abnormal findings by emergency room (ER) or emergency department physicians. One study reported that even in a non-PACS environment, the introduction of a single workstation in the ER reduced time to delivery of the radiology report from an average of 40 to 16 minutes by eliminating printing and transport. One of the challenges in a paperless and filmless environment has been the communication of preliminary impressions of ER physicians to radiologists. This has been addressed at our institution by giving the ER clinicians the ability to view a field on their workstation that allows them to determine whether a study has been interpreted by a radiologist. For those studies that have not yet been dictated by a radiologist (a minority during the working day) by the time the ER physician reviews the imaging study, the ER physician can type a preliminary impression directly into the PACS electronic display in the section that lists the reason for the study. The radiologist is then able to alert the ER physicians when there is a discrepancy between their preliminary impressions and the radiologist's interpretation of the images.

These advantages, when combined with the incremental positive effects of collegial e-mails, voice mails, augmented teaching sessions and conferences, and the generally increased utilization of a growing number of different imaging services, suggest that the benefits of digital imaging will enhance rather than detract from the radiologist's standing in the larger medical enterprise.

## WORKFLOW AND MULTI-INSTITUTIONAL PACS

The market forces that have created the strong impetus to increase efficiency and productivity have also fostered the formation of imaging networks in which large radiology groups provide imaging services for multiple facilities. This can result in substantial savings by taking advantage of centralized administration, scheduling, and staffing and economies of scale in supplies, furniture, and even imaging apparatus. However, imaging networks pose challenges that include issues related to distance, different equipment and information systems, communications, and personnel with different “cultures” and a variety of approaches to the departmental operation. Optimizing RIS and PACS workflow management can be difficult across such a disparate landscape but carries with it a number of long- and short-term benefits. Radiology coverage and subspecialty expertise can be shared across multiple hospitals. This has been particularly helpful in situations in which one or more radiologists provide overnight coverage for multiple facilities. The ability of a single radiologist to provide network coverage across multiple facilities has been a major impetus for many radiology groups to install teleradiology systems or PACS. Such efforts have their own difficulties in tracking multiple patient identification systems, interfacing adequately with several vendors, and dealing with communications problems between hospitals with different systems. Despite these challenges, the potential workflow advantages of a multifacility shared or complete virtual radiology department are tremendous, both to the radiology department and to the clinician. The potential to share the radiology caseload in a more effective manner made possible by PACS in a single institution is even greater in a wide area networked virtual department, particularly with regard to subspecialty expertise. The ability to access images obtained at other institutions within the network can eliminate many of the delays associated with film transportation as well as decrease the number of unnecessarily repeated examinations.

In the VA Maryland HealthCare System, the transition to a wide area network virtual radiology department has resulted in savings of approximately \$800,000 to \$1,000,000 per year, largely in personnel costs. The network is set up in a hub-and-spoke configuration in which images are sent to Baltimore for storage on the VA Baltimore commercial PACS and are then made available throughout the healthcare network. This hub-and-spoke configuration is also used for the HIS and RIS, resulting in the need for a central computer system in Baltimore and reliable high-speed networks connecting the facilities. This “central” architecture for the PACS

has been successful in our environment, with four facilities connected to the BVAMC.

Perhaps the biggest challenge with regard to integration of multiple healthcare facilities is the need to have a common, agreed-on method for exchange of patient images and other patient information. Recent federal government mandates in the United States have hastened efforts to codify standards that would make processes such as patient identification, modality worklists, and management of image interpretation automated in a multi-vendor, multi-HIS-RIS hospital system or between two or more healthcare systems. At the same time, such improved communication is under new restrictions designed to protect patient privacy. The IHE initiative is at the forefront of efforts to bring modality imaging vendors to the table with both RIS and HIS vendors to formulate solutions to these challenges. True integration across one or multiple healthcare enterprises is a much more practical and easily achievable goal than it seemed only a few years ago.

## **THE READING ROOM ENVIRONMENT AND WORKFLOW**

One of the more hotly debated issues with regard to optimization of radiologist reading performance and workflow has been the question of the optimal number of monitors that are required when using a PACS workstation for various modalities such as computed radiography, CT, MR, sonography, and so on. This is particularly important given the substantial expense of these monitors and the high percentage of the total workstation cost associated with the number of monitors. At the BVAMC, we performed a prospective study of the impact on radiologist performance and levels of fatigue as a function of the number of monitors. We found an approximately 25% increase in radiologist reading speed for a 4-monitor in comparison with a 1-monitor workstation in the interpretation of portable chest radiographs performed using computed radiography when we took into account the number of prior studies reviewed. Interestingly, there was a decrease in the number of historical studies reviewed as the number of monitors decreased. There was very little difference in the amount of time required to read the studies when comparing a 2- with a 4-monitor workstation, and the largest increase in performance was seen between a 1- and a 2-monitor workstation. Although we have not yet performed this study, our expectation would be that the use of stack mode for CT and MR studies would substantially decrease the added value of 4- or even 2-monitor workstations for the interpretation of these studies. Anecdotally, this seems to be particularly

true when the workstation permits images from multiple examinations to be linked according to anatomic section, which facilitates easy comparison of current and previous cross-sectional images.

The number of monitors is only one factor in designing an ideal radiology reading room, a goal that has only recently been appreciated as a potential contributor to improvements in radiology workflow. Despite transitioning to filmless or almost-filmless imaging, most institutions maintained without question the traditional configuration of the film reading room. And, given the extensive attention that has been paid to PACS monitors and workstations, surprisingly little attention has been paid to the radiologist (and clinician) reading room environment. Our research at the BVAMC has indicated that a number of factors, including monitor and ambient room lighting levels, among others, play a critical role in radiologist productivity and fatigue. We found that radiologist performance decreases significantly and that fatigue increases as monitor brightness drops or as ambient room lighting increases.

Several research laboratories have now documented the importance of additional factors, such as workstation chair design, the availability of individual lighting and temperature controls, and room acoustics, on radiologist performance and fatigue. Architects, who have responsibility for designing workplace environments, have also recognized the vital role of workstation ergonomics.

## **AUDITING THE RADIOLOGIST WORKFLOW PROCESS**

Every imaging department, regardless of the quality of staff, integrity of findings, or excellence of technical facilities, shares a usually unspoken secret: their way of doing radiologic work evolved over time and without a comprehensive workflow design strategy. Even the most carefully planned of all-digital departments still contain elements of workflow that are purely vestigial. With the rapid turnover of technology and personnel, potential workflow improvements are available to every department and could not only realize time and money benefits but enhance patient outcomes and quality-of-work issues for both radiologists and technologists. The problem, of course, is in identifying areas of potential improvement before the next round of change brings in new variables.

The digitalization of the entire acquisition, processing, reading, reporting, and archiving process in radiology presents extraordinary opportunities for quantifiable study. Some variables are easily extracted from the informa-

tion technology record (speech recognition dictation and correction times, for example), whereas others require more definition (such as reports on how many times a radiologist looks at a specific type of image and/or refers to others).

Radiologist workflow cannot be studied without adequate tools to measure performance (including the use of workstation tools of all types), accuracy, comfort, fatigue, and satisfaction. We have recently worked with industry to extend error logging tools built into workstation software into more comprehensive workstation “audits” that generate very large amounts of data (approximately 50 pages per minute) and have used these to produce detailed analyses of radiologist interpretation workflow. These audit databases constitute as yet untapped gold mines of fascinating data and should lead to new insights into the ways in which radiologists utilize conventional and multiplanar 3-D workstations in rendering routine diagnoses. These types of investigations should provide the data that will drive future workstation enhancement and influence the next generation of intelligent (or less dumb) workstations that analyze each radiologist’s interpretation habits and then adapt responses to help the radiologist achieve enhanced efficiency and accuracy.

## CONCLUSION

The most frightening and most promising word that currently defines radiologist workflow is “more”: more images generated from more modalities to be read at a faster pace, more tools that support image processing and reporting, more ways to shape interactions with others in the medical enterprise, and more possibilities to enhance diagnostic capabilities and the well-being of patients. The problem, of course, is that there are no more hours in the day.

The importance of an understanding of workflow for RIS and PACS vendors has resulted in substantial improvements in the development of intelligent software and use of integration with other information systems. This trend will undoubtedly continue. Universal adoption of integration protocols such as the IHE and standards such as DICOM and HL7 will continue the trend toward the elimination of paper and will result in further reductions in the number of steps in the flow of information to and from the imaging department. Computer-assisted detection will provide both a pre-screen and a double-read for radiologists in the interpretation of a much wider array of imaging studies. Workstation innovations will continue to improve on the ways in which radiologists can access and compare current

and previous examinations and will permit a greater degree of interactivity with the images themselves. New strategies will doubtless be offered for dealing with increasing datasets.

The challenge, of course, is to be able to incorporate such new strategies into existing PACS configurations and to evaluate their utility in both workflow benefits and diagnostic outcomes. PACS are well beyond the “early adopters” phase and have become an integral part of all aspects of radiology workflow. In the process, the pace and nature of the radiologist’s work has changed profoundly, and we are only now beginning to investigate the effects of these changes and the ways in which future change may continue to alter training, practice, and workflow.

## REFERENCES

- Andriole KP, Avrin DE, Yin L, Gould RG, Luth DM, Arenson RL. Relevant priors prefetching algorithm performance for a picture archiving and communication system. *J Digit Imaging*. 2000;13(2 suppl 1):73–75.
- Andriole KP, Morin RL, Arenson RL, et al. Transforming the radiological interpretation process: the SCAR TRIP initiative. *Proc SPIE Med Imaging*. 2004;5371:1–7.
- Bennett WF, Vaswani KK, Mendiola JA, Spigos DG. PACS monitors: an evolution of radiologists’ viewing techniques. *J Digit Imaging*. 2002;15(suppl 1):171–174.
- Boochever SS. HIS/RIS/PACS integration: getting to the gold standard. *Radiol Manage*. 2004;26:16–24.
- Cavanaugh BJ, Garland HT, Hayes BL. Updating legacy systems for the Integrating the Healthcare Enterprise (IHE) initiative. *J Digit Imaging*. 2000;13(2 suppl 1):180–182.
- Dackiewicz D, Bergsneider C, Piraino D. Impact of digital radiography on clinical workflow and patient satisfaction. *J Digit Imaging*. 2000;13(2 suppl 1):200–201.
- DeSimone DN, Kundel HL, Arenson RL, et al. Effect of a digital imaging network on physician behavior in an intensive care unit. *Radiology*. 1988;149:4–44.
- Donnelly JT, Anderson Q. Automating clinically relevant prefetch processes in picture archiving and communication systems solutions. *J Digit Imaging*. 2001;14(2 suppl 1):135–139.
- Gale B, Safriel Y, Lukban A, Kalowitz J, Fleischer J, Gordon D. Radiology report production times: voice recognition vs. transcription. *Radiol Manage*. 2001;23:18–22.
- Gale ME, Gale DR. Modality worklist. An essential component in a PACS environment. *J Digit Imaging*. 2000;13:102–108.

- Gay SB, Sobel AH, Young LQ, Dwyer SJ III. Processes involved in reading imaging studies: workflow analysis and implications for workstation development. *J Digit Imaging*. 2002;15:171–177.
- Ghosh S, Andriole KP, Avrin DE. Viewport: an object-oriented approach to integrate workstation software for tile and stack mode display. *J Digit Imaging*. 1997;10(3 suppl 1):177–179.
- Graf B, Mendenhall R, Suero M, Wiggins R. Study status consistency and duplicate-read protection in a distributed architecture. *J Digit Imaging*. 2001;14(2 suppl 1):60–62.
- Gross-Fengels W, Miedeck C, Siemens P, et al. PACS from project to reality. Report of experiences on full digitalization of the radiology department of a major hospital. *Radiologe*. 2002;42:119–124.
- Hayt DB, Alexander S. The pros and cons of implementing PACS and speech recognition systems. *J Digit Imaging*. 2001;14:149–157.
- Hirschorn D, Eber C, Samuels P, Gujrathi S, Baker SR. Filmless in New Jersey: the New Jersey Medical School PACS Project. *J Digit Imaging*. 2002;15(suppl 1):7–12.
- Hirschorn DS, Hinrichs CR, Gordon M, Shah K, Visvikis G. Impact of a diagnostic workstation on workflow in the emergency department at a level I trauma center. *J Digit Imaging*. 2001;14(2 suppl 1):199–201.
- Horii S, in Reiner BI, Siegel EL. The cutting edge: strategies to enhance radiologist workflow in a filmless/paperless imaging department. *J Digit Imaging*. 2002;15:178–190.
- Jager D, Kotter E, Langer M. RIS modality interfaces. From proprietary solutions to DICOM worklist management [in German]. *Radiologe*. 1999;39:286–291.
- Junck KL, Berland LL, Bernreuter WK, et al. PACS and CR implementation in a level I trauma center emergency department. *J Digit Imaging*. 1998;11(3 suppl 1):159–162.
- Kim Y, Fahy JB, DeSoto LA, et al. Development of a PC-based radiological imaging workstation. *J Digit Imaging*. 2003;16(1):104–113.
- Krupinski EA, Radvany M, Levy A, et al. Enhanced visualization processing: effect on workflow. *Acad Radiol*. 2001;8:1127–1133.
- Kuzmak PM, Dayhoff RE. Experience with DICOM for the clinical specialties in the healthcare enterprise. *Proc SPIE Med Image*. 2003;5033:18–29.
- Langer HL, Bielmeier J, Wittenberg G, Selbach R, Feustel H. Workflow improvement and efficiency gain with near total digitization of a radiology department [in German]. *Rofu Fortschr Geb Rontgenstr Neuen Bildgeb Verfabr*. 2003;175:1309–1316.
- Langer S. Impact of tightly coupled PACS/speech recognition on report turnaround time in the radiology department. *J Digit Imaging*. 2002;15(suppl):234–236.



- Langer S. OpenRIMS: an open architecture radiology informatics management system. *J Digit Imaging*. 2002;15:91-97.
- Langer S. Radiology speech recognition: workflow, integration, and productivity issues. *Curr Probl Diagn Radiol*. 2002;31:95-104.
- Locko RC, Blume H, Goble JC. Enterprise-wide worklist management. *J Digit Imaging*. 2002;15(suppl 1):175-179.
- Mack S, Holstein J, Kleber K, Gronemeyer DH. New aspects of image distribution and workflow in radiology. *J Digit Imaging*. 2000;13(2 suppl 1):17-21.
- May GA, Deer DD, Dackiewicz D. Impact of digital radiography on clinical workflow. *J Digit Imaging*. 2000;13(2 suppl 1):76-78.
- Moise A, Atkins S. Interaction techniques for radiology workstations: impact on users' productivity. *Proc SPIE Med Imaging*. 2004;5371:16-22.
- Moore SM. Using the IHE scheduled workflow integration profile to drive modality efficiency. *RadioGraphics*. 2003;23:523-529.
- Morioka CA, Valentino DJ, Duckwiler G, et al. Disease-specific intelligent pre-fetch and hanging protocol for diagnostic neuroradiology workstations. *Proc AMIA Symp*. 2001;468-472.
- Mulvaney J. The case for RIS/PACS integration. *Radiol Manage*. 2002;24:24-29.
- Nagy P, Siegel E, Hanson T, Kreiner L, Johnson K, Reiner B. PACS reading room design. *Semin Roentgenol*. 2003;38:244-255.
- Offenmueller W. Managed workflow between HIS, RIS, modalities, and PACS: now with DICOM or beyond? *Proc SPIE Med Imaging*. 1996;2711:144-155.
- Okura Y, Matsumura Y, Harauchi H, et al. An inductive method for automatic generation of referring physician prefetch rules for PACS. *J Digit Imaging*. 2002;15:226-231.
- Oosterwijk H. DICOM vs. HL7 for modality interfacing. *J Digit Imaging*. 1998;11 (3 suppl 1):39-41.
- Page D. Voice recognition saves time, money, but rads resist it. *Diagn Imaging*. 2000;22:6-7.
- Redfern RO, Horii SC, Feingold E, Kundl HL. Radiology workflow and patient volume: effect of picture archiving and communication systems on technologists and radiologists. *J Digit Imaging*. 2000;13(2 suppl 1):97-100.
- Reiner BI, Siegel EL. Effect of filmless imaging on the utilization of radiologist services. *Radiology*. 2000;215:163-167.
- Reiner BI, Siegel EL. The cutting edge: strategies to enhance radiologist workflow in a filmless/paperless imaging department. *J Digit Imaging*. 2002;15:178-190.
- Reiner BI, Siegel EL. Evolution of the digital revolution: a radiologist perspective. *J Digit Imaging*. 2003;16:324-330.
- Reiner BI, Siegel EL, Carrino JA. Workflow optimization: current trends and future directions. *J Digit Imaging*. 2002;15:141-152.

- Reiner BI, Siegel EL, Hooper FJ. Accuracy of interpretation of CT scans: comparing PACS monitor displays and hard-copy images. *AJR Am J Roentgenol.* 2002;179:1407–1410.
- Reiner BI, Siegel EL, Protopapas Z. The impact of filmless radiology on the frequency of clinician consultation with radiologists. *Am J Radiol.* 1999;173:169–172.
- Reiner BI, Siegel EL, Shastri K. The future of radiology reporting. In: Reiner BI, Siegel EL, Weiss DL, eds. *Electronic Reporting in the Digital Medical Enterprise.* Great Falls, VA: Society for Computer Applications in Radiology; 2003:83–104.
- Reiner BI, Siegel EL, Weiss DL, eds. Introduction: the evolution of the radiology report and the development of speech recognition. In: *Electronic Reporting in the Digital Medical Enterprise.* Great Falls, VA: Society for Computer Applications in Radiology; 2003:1–7.
- Reiner B, Siegel E, Kuzmak P, Severance S. Transmission failure rate for computed tomography examinations in a filmless imaging department. *J Digit Imaging.* 2000;13(2 suppl 1):79–82.
- Robbins AH, Horowitz DM, Srinivasan MK, et al. Speech-controlled generation of radiology reports. *Radiology.* 1987;162:569–573.
- Roos JE, Desbiolles LM, Willmann JK, et al. Multidetector-row helical CT: analysis of time management and workflow. *Eur Radiol.* 2002;12:680–685.
- Rosenthal DR, Chew FS, Dupuy DE, et al. Computer-based speech recognition as a replacement for medical transcription. *AJR Am J Roentgenol.* 1998;170:23–25.
- Rostenberg B. *The Architecture of Imaging.* Chicago, IL: American Hospital Publishing Inc; 2002:173–181.
- Rostenberg B. Reading rooms should be designed to accommodate future changes. *Diagn Imaging.* 2002;24:11–15.
- Siegel EL, Reiner BI. Recommendations for image prefetch for film digitization strategy based on an analysis of an historic radiology image database. *J Digit Imaging.* 1998;11:94–99.
- Siegel E, Reiner B. Filmless radiology at the Baltimore VA Medical Center: a 9 year retrospective. *Comput Med Imaging Graph.* 2003;27:101–109.
- Siegel EL, Reiner BI. *Digital Eye for the Analog Guy* [film]. Waukesha, WI: GE Medical; 2004.
- Siegel EL, Diaconis JN, Pomerantz S, et al. Making filmless radiology work. *J Digit Imaging.* 1995;8:151–155.
- Sistrom CL, Honeyman-Buck JC. Template and macro management for a radiology speech recognition system. In: Reiner BI, Siegel EL, Weiss DL, eds. *Electronic Reporting in the Digital Medical Enterprise.* Great Falls, VA: Society for Computer Applications in Radiology; 2003:73–82.
- Thomas MA, Rowberg AH, Langer SG, et al. Interactive image enhancement of CR and DR images. *J Digit Imaging.* 2004;17(3):189–195.

- Walter LeRoy Thompson Collection. Archives of the American College of Radiology. Reston, VA. Accessed April 1996.
- Weiss DL. A users guide to speech recognition technology. *Decis Imaging Econ.* 2001;14:32-34.
- Weiss DL. Speech recognition technology improves workflow efficiency. *Diagn Imaging.* 2002;24:48-51.
- Weiss DL. Speech recognition: evaluation, planning, installation, and use: purchasing a speech recognition system. In: Reiner BI, Siegel EL, Weiss DL, eds. *Electronic Reporting in the Digital Medical Enterprise.* Great Falls, VA: Society for Computer Applications in Radiology; 2003:27-42.
- Yoshimura H, Inoue Y, Tanaka H, et al. Operating data and unsolved problems of the DICOM modality worklist: an indispensable tool in an electronic archiving environment. *Radiat Med.* 2003;21:68-73.

# FINANCIAL MODELING

---

---

SYRENE R. REILLY • DAVID AVRIN

**T**his chapter provides an overview of the financial concepts and tools that are useful in the financial evaluation of picture archiving and communication systems (PACS). The first section discusses various analysis methods and makes a case for using net present value (NPV) methodology. The second section looks at the major cost elements that should be considered and quantified. The third section explores the cost-saving opportunities and nonfinancial benefits of implementing PACS. These three sections should help you on your path to justifying PACS financially.

## ANALYSIS METHODS

There are numerous ways to evaluate a capital investment such as PACS. It is worthwhile to understand all of them and to determine which methods are most widely used and respected at your organization, especially by those with decision rights. It is often helpful to use several methods, as each provides a different lens through which you can analyze your investment opportunity. Different methods appeal to different constituencies. The

nonfinancial benefits need to be considered as clearly as do the financial benefits if you are to fully evaluate any investment opportunity, especially in a healthcare environment. To add credibility and ensure quality, this financial analysis is best done by an impartial person who has business analysis skills and credentials.

It is important to define the objectives of the financial analyses at the outset. Objectives can be any or all of the following:

- ▶ Determining whether investing in PACS makes sense financially
- ▶ Obtaining organizational approval
- ▶ Negotiating discounts with PACS vendors
- ▶ Analyzing different scenarios and performing sensitivity analyses
- ▶ Developing budget estimates
- ▶ Tracking results

It is possible to incorporate all of the above features in one model. The best financial models are those that clearly lay out assumptions and sensitivities to those assumptions and assign cost-savings responsibilities to parties who control the costs, for example, use of film and the film library. Cost savings produced by eliminating conventional film systems are discussed in detail later in this chapter.

## CASH

Most investment analysis methods are based on cash flow. A major difference between accounting income and cash flow is the treatment of capital assets. For accounting income, the cost of a capital asset is allocated via depreciation expense to the periods that benefit from the asset. For cash flow, each year reflects cash spent on the capital asset. To evaluate a capital project, you will want to weigh the capital cash outlays associated with the project against the benefits in terms of cash returned to the enterprise.

Example: A company purchases a \$10 million asset that produces \$2 million of annual income (cash) and has an expected life of 10 years. Accounting income spreads the cost of an asset over the asset's useful life and matches the cost of the asset to the income it produces. This is the theory behind depreciation. If the asset continues to produce \$2 million in revenue in the 11th year (as shown in Tables 7.1 and 7.2), there is no depreciation expense because the asset has been fully depreciated over the prior 10 years.

**TABLE 7.1**  
Cash-Flow Method (in million \$)

<b>Year</b>	<b>Cash Outlay</b>	<b>Cash Inflow</b>	<b>Annual Net Cash Flow</b>	<b>Cumulative Cash Flow</b>
0	10		(10)	(10)
1		2	2	(8)
2		2	2	(6)
3		2	2	(4)
4		2	2	(2)
5		2	2	0
6		2	2	2
7		2	2	4
8		2	2	6
9		2	2	8
10		2	2	10
11		2	2	12
<b>Total</b>	<b>10</b>	<b>22</b>	<b>12</b>	

### SUNK COSTS

The purpose of all these techniques is to evaluate a possible capital investment. A sunk cost is a cost that has already been incurred and cannot be changed. Sunk costs are irrelevant to the decision of whether to make an investment. Thus, the cost justification effort is less burdensome for those who have already made past investments in digital equipment, information systems, and hardware.

### IRRELEVANT COSTS

Costs that would be incurred regardless of the implementation of PACS should be ignored. This is particularly appropriate for organizations that already plan to implement computed radiography. Such costs are not relevant in the financial justification of PACS. Similarly, the decision to invest

**TABLE 7.2**  
Accounting Income Method (in million \$)

Year	Depreciation Expense	Revenue	Annual Income	Cumulative Income
0				
1	1	2	1	1
2	1	2	1	2
3	1	2	1	3
4	1	2	1	4
5	1	2	1	5
6	1	2	1	6
7	1	2	1	7
8	1	2	1	8
9	1	2	1	9
10	1	2	1	10
11	—	2	2	12
Total	10	22	12	

in voice recognition technology is separate from the PACS decision and should be analyzed separately.

### PAYBACK PERIOD

Payback period represents the number of years it takes for an organization to recover its initial investment via the cash flows generated from the investment, without adding the cost of capital (interest). This is also the point at which the project breaks even on a cumulative cash-flow basis. Some organizations establish required payback periods in addition to other financial hurdles (described later). This method offers ease of use and simplicity of application, but it does not help determine the true value of the investment over its lifetime or its value relative to other investment opportunities (see Table 7.3). The example in Table 7.1 illustrates a payback period of 5 years, the point at which the cumulative cash flow is \$0.

## NET PRESENT VALUE

The net present value (NPV) method assesses the worth of a project by bringing all cash inflows and outflows associated with the project into one value in today's dollars. With a 10% interest rate, an investor with \$1.00 today can generate a future value of \$1.10 in 1 year. Alternatively, this investor would value a riskless payment of \$1.10 in 1 year at \$1.00 today, in "today's dollars." In this example, \$1.00 is the present value, \$1.10 is the future value, and the discount rate is 10%. Net present value is the current value of the cash inflows less the current value of the cash outflows. For example, suppose this investor were offered an alternative project in which he would get \$1.10 at the end of the year if he invested \$.98 today. Since the \$1.10 in the future is worth \$1.00 to him today, and the cost of the investment is only \$.98, he gains \$.02 by accepting this project versus his first alternative. This \$.02 return is the NPV of the project. An investment is worth making if it has a positive NPV; an investment is not worth making if it has a negative NPV. This is the most widely accepted and respected analysis method.

The underlying concept of NPV can best be understood by the following example: Assume that someone promises to pay you \$100 1 year from now. What would you be willing to loan that person today? If the loan is to someone you know and trust to pay you back, you would be willing to give that person market rate, or approximately \$93 at 8% (= \$100.44) for 1 year. On the other hand, if you do not know the person, the risk is substantial,

---

---

**TABLE 7.3**

Payback Analysis: Advantages and Disadvantages

*Advantages:*

1. Easy to do
2. Quick financial reality check
3. Helps identify capital costs
4. Helps identify sources and magnitudes of savings

*Disadvantages:*

1. Does not take account of the cost of capital (current market interest rules)
  2. Does not account for risk of project
  3. Does not quantify the investment value of the project
- 
-



and you may be willing to give that person only \$80, or even \$50, based on the higher risk of not being paid back.

Furthermore, if the term of the promise were longer, say 5 years instead of 1, like a savings account in reverse, the interest would have to be compounded over the term of the investment, usually on an annual or monthly basis. The formula for this process, NPV, is similar to the familiar interest compounding formula, but with the compounding portion in the denominator, as shown:

$$\text{NPV} = \frac{P}{(1+d)^n}$$

where:

P = future value (being discounted)

d = discount rate per period

n = the number of periods

---

---

**TABLE 7.4**

Risks of PACS Implementation

*Technology:*

Integration/interoperability

1. Modalities—DICOM compliance
2. RIS-HIS

Software: stability/(robustness)

Scaling

Network infrastructure

1. Institution
2. Community

Disaster Protection

*Organization:*

Acceptance

User interface

Radiologists

Clinicians

Realization of film and personnel savings

Lack of in-house expertise

---

---

The discount rate has two components: (1) the underlying (riskless) market rate or cost of capital for the term (including inflation), and (2) an estimate of the risk premium, or interest rate related to the risk of the project. Risk of PACS project implementation is a complex topic that we do not discuss in detail, other than to consider that the risk factor should encompass all the assumptions of the project: costs of implementation, timeliness of implementation, and realized cost savings. Some of the risks that should be considered are listed in Table 7.4.

There are many ways for a PACS project to get off track. Major obstacles or risks with major or even disastrous consequences are often referred to as “showstoppers” by information technology (IT) professionals. Note that because the discount rate includes an inflation factor, the cash flows it is applied to should also include an inflation factor so the analysis compares “apples to apples.”

## RISKS OF PACS IMPLEMENTATION

When an organization has more projects than capital, the discount rate should be set at the risk-adjusted return that the funds could generate on a competing project, as a hurdle rate (e.g., build an operating room suite vs. implement PACS). A certificate of deposit bearing 7% offered by a bank insured by the FDIC has a risk-adjusted return of 7% because there is no risk. A PACS implementation expected to generate 20% returns if the implementation is flawless may have a risk-adjusted return of 12% to 15% to reflect the risk that savings might not materialize or additional revenue might not be generated. In this case, enterprises generally set the discount rate or hurdle rate at the corporate level. That rate is usually 15% to 20%, depending on the risk profile of the enterprise. In health care, IT projects are often assigned higher risk rates because they have a reputation of not being able to produce the desired return; PACS falls into this category. For certain IT projects in health care, there are often other enterprise-wide strategic reasons to proceed, even if the expected returns do not overcome the hurdle rate.

If the investment and/or savings occur at different times (years) and/or in differing amounts, the NPV calculation is the sum of each value for the specific length of time from the time of investment into the future:

$$\text{NPV} = \sum_{i=1}^n P_i / (1 + \text{discount rate per period})^i$$

i = 1 to n, where n = number of periods

The simplest way to calculate the NPV is to discount the annual net cash flow, or the sum of capital outlays (termed “investment”), and cost savings (termed “incremental cash flow”), as demonstrated in Tables 7.5 and 7.6. At a discount rate of 15%, cash flows beyond 10 years have a marginal impact, as evidenced in the table examples, in which the \$2 million in cash flow in year 10 has a present value of \$490,000, or 25% of its future value.

It is useful to project out as many years as it takes to reach steady-state cash flows or to the point at which no further benefits are expected from the investment. The capital outlay occurs in the first year(s). Operating costs will ramp up as the system reaches completion and should be adjusted each year for inflation. Savings ramp up as the enterprise discontinues its use of film. So, if the organization plans to implement PACS over 2 years and to take 5 years from implementation before achieving its film elimination targets, the analysis should be carried over for 7 years (2 years to implement plus 5 years to achieve full cost savings). The PACS life expectancy would serve as a time-cap on this exercise.

**TABLE 7.5**  
Net Present Value Example with Initial Investment of \$10 (in millions)

Year	Investment	Incremental Cash Flow	Net Cash Flow	Discounted Cash Flow
0	\$10.00		(\$10.00)	(\$10.00)
1		\$2.00	\$2.00	\$1.74
2		\$2.00	\$2.00	\$1.51
3		\$2.00	\$2.00	\$1.32
4		\$2.00	\$2.00	\$1.14
5		\$2.00	\$2.00	\$0.99
6		\$2.00	\$2.00	\$0.86
7		\$2.00	\$2.00	\$0.75
8		\$2.00	\$2.00	\$0.65
9		\$2.00	\$2.00	\$0.57
10		\$2.00	\$2.00	\$0.49
			NPV	\$0.02

TABLE 7.6

Net Present Value Example with Initial Staggered Investments of \$6.00, \$2.00, and \$2.00 (in millions)

Year	Investment	Incremental Cash Flow	Net Cash Flow	Discounted Cash Flow
0	\$6.00		(\$6.00)	(\$6.00)
1	\$2.00	\$2.00	\$0.00	\$0.00
2	\$2.00	\$2.00	\$0.00	\$0.00
3		\$2.00	\$2.00	\$1.32
4		\$2.00	\$2.00	\$1.14
5		\$2.00	\$2.00	\$0.99
6		\$2.00	\$2.00	\$0.86
7		\$2.00	\$2.00	\$0.75
8		\$2.00	\$2.00	\$0.65
9		\$2.00	\$2.00	\$0.57
10		\$2.00	\$2.00	\$0.49
			NPV	\$0.77

In reality, an organization's capital projects with positive NPVs may exceed the capital available. As a result, projects with the highest return win in the battle for capital. Sometimes political and other nonfinancial considerations increase or decrease the financial value of a project. Those who prepare on all fronts increase the likelihood that the capital project will be approved.

## INTERNAL RATE OF RETURN

The internal rate of return (IRR) is the discount rate at which the NPV of a project is 0. Instead of solving for a project's worth in dollars after applying a predetermined hurdle rate, the formula is solved for the discount rate itself, specifically the rate at which the NPV equals 0. This method

offers one of the most common ways enterprises evaluate portfolios of opportunities, particularly if the decision is made on a financial basis only. This approach is somewhat shortsighted, since some of the costs (savings) are difficult to measure, particularly those that accrue outside the radiology department, and there are enterprise-wide strategic reasons to invest in PACS.

Tables 7.7 and 7.8 demonstrate the IRR method for the preceding example of a phased investment in PACS in years 0 through 2. The IRR calculation in a spreadsheet function, such as the one in Excel, solves for the unknown rate of return by using iterative or repeated calculations of the NPV formula. One actually has to make a guess or initial estimate of the rate, but usually any starting point between 0 and 10% will work. The NPV is calculated and driven to 0 by repeated adjustments to the rate, until the NPV is close to 0. This then yields the calculated IRR for the project or the rate at which future discounted cash savings balance the initial and future discounted investments in the project.

---

---

**TABLE 7.7**

Internal Rate of Return at 7 Years (in million \$)

<b>Year</b>	<b>Investment</b>	<b>Savings</b>	<b>Net Cash Flow</b>	<b>Discounted Cash Flow</b>
0	\$6.00	\$0.00	(\$6.00)	(\$6.00)
1	\$2.00	\$2.00	\$0.00	\$0.00
2	\$2.00	\$2.00	\$0.00	\$0.00
3		\$2.00	\$2.00	\$1.46
4		\$2.00	\$2.00	\$1.32
5		\$2.00	\$2.00	\$1.19
6		\$2.00	\$2.00	\$1.07
7		\$2.00	\$2.00	\$0.96
			NPV	(\$1.32)
			IRR: 11.0%	

---

---

**TABLE 7.8**  
Internal Rate of Return at 10 Years (in million \$)

Year	Investment	Savings	Net Cash Flow	Discounted Cash Flow
0	\$6.00	\$0.00	(\$6.00)	(\$6.00)
1	\$2.00	\$2.00	\$0.00	\$0.00
2	\$2.00	\$2.00	\$0.00	\$0.00
3		\$2.00	\$2.00	\$1.23
4		\$2.00	\$2.00	\$1.05
5		\$2.00	\$2.00	\$0.89
6		\$2.00	\$2.00	\$0.76
7		\$2.00	\$2.00	\$0.65
8		\$2.00	\$2.00	\$0.55
9		\$2.00	\$2.00	\$0.47
10		\$2.00	\$2.00	\$0.40
			NPV	(\$0.00)
			IRR: 17.5%	

Same data as Table 7.7, but with return extended out 10 years.

## BREAKEVEN ANALYSIS AND FIXED AND VARIABLE COSTS

It is useful to compare the fixed and variable costs of the organization's film-based system to those of PACS to determine the volume level at which PACS produces lower total costs than do conventional methods. Fixed costs are costs that do not change as volume changes. Variable costs vary directly with volume and are 0 if nothing is produced. Because the objective is to solve for the volume, it is best to do this as a 1-year snapshot. To arrive at an annual cost, spread the capital costs over the useful life of the asset. Most of the capital costs are fixed, although one could argue that the cost of storage varies with volume. The personnel required to manage the PACS is also somewhat

fixed. Variable costs are minimal. A conventional system's fixed costs are lower, since there is less capital equipment. The conventional system's variable costs consist mainly of film (and other disposables) and film library support activities (personnel). Although these are the major ingredients, you could try to capture numerous other costs that are more difficult to quantify. (We discuss those more fully later in this chapter.) These economic relationships are depicted in Figure 7.1.

In Figure 7.1, the dashed line (traditional fee-for-service income) no longer exists as such but is replaced by an underlying linear demand line for imaging services, to which a value can be assigned or ascribed. For example, in a managed care or capitated healthcare enterprise, a demand for imaging services is some function of the number of insured lives (linear), demographics (nonlinear), utilization profile of the referring physicians (complex), and possibly other factors. Some generalizations can be made, however. If the horizontal axis is labeled "Insured Lives," then the slope of the demand line is proportional to the diagnostic imaging utilization profile and determines the volume of examinations. An institution still has to provide this volume of services. However, the important differential is not between the demand line and the cost line (digital or conventional), but between the conventional and the digital, where the crossover occurs at some volume level.

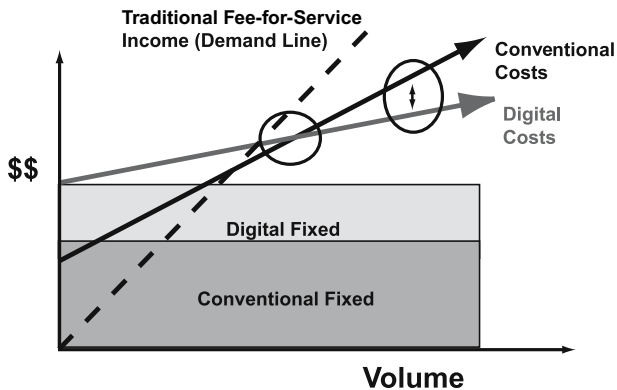


FIGURE 7.1

Breakeven analysis.

That is because the incremental or variable costs of a digital study are lower, particularly for the digital modalities.

## **COSTS**

### **FACTORS DETERMINING COSTS**

There is no boilerplate solution for how to determine the costs of implementing PACS. The costs depend on the sophistication of the enterprise's existing information system network and its imaging equipment inventory and needs. These capital costs, together with ongoing costs for operating the system, determine what levels of savings are required to justify PACS. Investing in PACS represents a trade-off: decreased operating costs (film and film personnel) versus increased capital costs together with PACS maintenance and personnel costs. To produce a credible financial analysis, it is best to err on the side of overstating costs and understating savings to the extent that the results allow.

Determining the cost to acquire, move, and store images is critical. An equipment inventory assessment must be done by a technician who understands how each radiology practice operates, what equipment exists, and what PACS equipment is needed. As the cost of software and maintenance is often in question and the discounts are flexible, the price to pay for the system could be calculated by using the number that generates a positive NPV. The required discount could be calculated by comparing this number to the list price offering. Equipment vendors can easily supply list prices and customary discounts. This discount, which can be substantial, is influenced by negotiation, size of purchase, and reputation value of the enterprise to the vendor. In addition, list prices are in a deflationary mode as technological advances and competition drive down prices.

### **CAPITAL EXPENDITURES**

The initial capital outlay consists of the categories of expenditure listed in Table 7.9, the technical nature of which are discussed more fully in the following chapters. Archive capital costs will continue to decrease, and in spite of early skepticism, creative methods for management of hierarchical storage promise to decrease storage costs even further in the near future.



---

---

**TABLE 7.9**  
Categories of Expenditures

*Imaging equipment:* Captures image in digital form

*Workflow managers/servers:* Store, retrieve, and distribute images

*Archive:* Longer-term storage of images

*Display stations:* Display images to radiologists and clinicians throughout the enterprise

*Facility upgrades:* temperature-, humidity-, and security-controlled environment for equipment; furniture and lighting changes for reading areas

*Clinical distribution and viewing*

---

---

## **COST REDUCTION AND REVENUE ENHANCEMENT**

Once the capital and operating costs are defined, determine the cost savings and revenue enhancements that will result from implementing PACS. There is a credibility continuum, with hard costs such as film and associated costs being the most credible, and soft ones such as improved patient outcomes being the least credible. An analysis that financially justifies PACS without including savings, which are more difficult to quantify and demonstrate, will be better received than one that shows an impressive financial impact but is built on extensive, unproven assumptions. In other words, proceed along the credibility continuum only until the cost is justified. Doing so also eases the follow-up analysis that may or may not be required to demonstrate postimplementation outcomes.

Partners HealthCare System, Inc. (Partners), in Boston provides an example of how one organization proceeded along the credibility continuum, using the financial techniques outlined earlier in this chapter, until PACS was justified. At Partners, founded in 1993, by the Massachusetts General Hospital and Brigham and Women's Hospital, PACS was financially justified based on savings from decreased film and film library costs alone, and no further analysis of cost savings or revenue opportunities was necessary. The Partners system is armed with a world-class information system infrastructure, consisting of the largest integrated Intel/Microsoft platform in the

world connected to more than 30,000 desktop computers for almost the same number of employees. The two hospitals were also well on their way to converting to computed radiography when this analysis was conducted.

Partners arrived at an implementation cost of approximately \$12.6 million, along with operating costs of approximately \$1.5 million per year, together composing the cost to be justified. The analysis was based on an annual volume of 775,000 radiology examinations per year. This represented 2.7 million films, at a film cost of \$3.5 million and a film library cost of \$1.7 million. An 8-year analysis was performed to cover 3 years of investment and implementation, 3 years to break even, and 2 years to reach steady-state savings. All these factors resulted in an NPV of 0 dollars, or breakeven (using a discount rate of 10%, required by the Partners treasury department). Likely but difficult-to-quantify cost savings and revenue enhancement opportunities would clearly produce a positive financial return, not to mention all the nonfinancial benefits such as improved clinical outcomes.

On a per-unit (per-exam) basis, Partners estimated that it would save \$8 per exam for film and film library expenses on an annualized basis, for an additional PACS operating cost of \$2 per exam, resulting in a net savings of \$6 per exam. This, however, required a one-time capital investment in PACS infrastructure of \$16 per annualized exam but only \$3.20 per exam, assuming a useful life of PACS investment of 5 years.

Mayo authors divided personnel costs associated with film into those occurring inside the radiology department and those occurring outside. These costs are incurred by nursing and clerical staff when engaged in both the “film search game” and traditional methods of requesting and managing exams needed in the clinic or operating rooms. Mayo arrived at \$15.82 per exam, as shown in Table 7.10. The Mayo authors also made the comment that “[o]ur estimated cost of film per exam per year is most likely an underestimation of real costs when compared to other institutions.”

**TABLE 7.10**  
Mayo Study: Estimated Film Cost per Exam

Film	\$6.25
Supplies	\$1.46
Personnel	\$5.91 (direct) \$2.20 (indirect)
Total	\$15.82 per exam

## FILM COST

To capture film costs it is necessary to develop assumptions about the number of annual exams, films per exam, and cost per film over the life of the capital investment in PACS. Annual savings is the product of annual exams multiplied by number of films per exam multiplied by the expected film cost per sheet. For example, an enterprise that generally conducts 10,000 annual MRI exams using 8 films per MRI at a cost of \$1.50 per film would save \$120,000 if it eliminated 100% of its film use. It is easiest to combine all associated film costs, including chemicals, processing, folder jackets, and so forth, with the film commodity cost for simplicity.

For most enterprises, it is necessary to ramp up film elimination from 0% to 90% or so over some number of years. It is difficult to eliminate film entirely (thus the 90%) because of the need to produce films for clinicians outside of the enterprise, for legal proceedings, and so on. Nevertheless, it is necessary to reach a fairly aggressive target quickly in order to justify PACS currently. The rapidity with which film use is eliminated is the key factor in cost savings. A long implementation perpetuates dual systems and processes, delays savings, and destroys value. A commitment must be made by the clinical enterprise that film use will be eliminated as soon as PACS is implemented. To make these assumptions real, keep in mind that a replacement for image distribution must be in place as well as a PACS prior to successful film elimination.

## FILM LIBRARY COST

The film library cost consists mainly of personnel managing the contents of the film library. As with film costs, the savings here would ramp up and shadow film reduction. The analysis could also include costs saved by reducing the space required for film storage. For many institutions, the space saved depends on legal requirements for film storage, which may take several years to develop, as the law generally follows practice, and these, from a legal perspective, are uncharted waters. For the analysis to capture space savings, the organization must have an alternate need for the space, and by gaining the film storage space, be able to avoid leasing additional space. It may be easier to treat such space savings as a wash when anticipating the increased space required for the PACS equipment and its staff, but this needs to be determined on a case-by-case basis.

The analysis can also phase in a reduction in warehouse costs for film storage that shadows the implementation phases. This reduction would also have to respect the film storage time required by law.

## LOST EXAMS

The financial impact study could also include the elimination of incremental costs and lost revenue associated with misplaced films. Savings may materialize from a decrease in staffing required of practitioners and administrative personnel to serve existing volumes, or as increased throughput (revenue less incremental costs). This impact is difficult to quantify, and the inclusion of these costs depends on whether the institution tracks this information.

It is also difficult to quantify the cost to the enterprise of not producing a film for a legal proceeding, or the cost of the department's and institution's reputations in not being able to produce a film for a patient or physician who needs it. The nonquantifiable cost to the patient might be staggering if a previous study is needed for comparison with a present study before a clinician can make an informed diagnosis. These situations can be enumerated in the analysis as nonquantifiable benefits.

## REDOS

PACS virtually eliminates the need for redos for two reasons. First, computed radiography imaging modality has a very wide range of latitude for exposure error, compared to film. Second, the rate of lost exams in a well-engineered PACS is very low compared to the estimated 10% to 15% temporary or permanent loss rate in a conventional film library. To calculate this savings, estimate the cost of redos to the organization in terms of time and materials. The savings in time depends on whether the clinician would be serving other patients instead of repeating the process; the savings in cost depends on whether a reduction in staffing would result. Such savings apply not only to radiologists but also to clinicians who are detained by redos. Savings on materials are calculated by the number of redone exams multiplied by films per exam multiplied by the cost per film.

Again, it is difficult to quantify the cost to patients associated with the delay caused by a redo. At the risk of being overly dramatic, we can say that a savings in time can make the difference between life and death for patients whose critical condition may depend on a rapid diagnosis.

## SAVED TIME FOR PRACTITIONERS AND ADMINISTRATIVE STAFF

PACS makes image distribution faster, easier, and more reliable. This feature translates into a cost reduction if staff are eliminated or into an additional financial contribution (incremental revenue less incremental costs) if an unmet demand for services (additional volume) exists. This impact will not be felt until the PACS is fully implemented and all radiologists and clinicians are proficient in its use. This time saved is difficult to measure without comparing the task time today versus the task time in a carefully projected environment, but few data are available for such comparison.

Faster turnaround time will likely translate into shorter stays which, in turn, can reduce costs of care or produce additional income if additional patients can be served. The potential impact on length-of-stay and increased admissions would be difficult to substantiate. It is difficult to know or measure how PACS contributes to shortening length of stay because there are so many complex factors that contribute to length of stay; it is difficult to segregate PACS as a single component.

## SITE OPPORTUNITIES

Just as the electronic revolution makes it possible for millions of people to spend more time working from home, electronic imaging makes the locus of work far less important for radiology services. PACS enables diagnostic images to be available anytime, anywhere they are needed, with little or no human intervention. This eliminates the necessity and cost of having radiologist coverage in multiple sites within an entity and in many entities within a system. The mobility of images created by PACS facilitates peer or expert review of images inter- and intra-network. This mobility reduces the potential number of radiologists required to serve a given population and also the time in which those services can be provided. Enterprises that take advantage of these site opportunities will be able to serve existing patients with fewer resources (reduce costs) or serve additional patients with existing resources (increase revenue).

## MULTISITE IMAGE READING

The peaks and valleys of demand can be better managed by diverting image reading to alternate sites. Diversion allows for more effective use of

resources, faster turnaround times, and improved patient outcomes. It also lets any appropriate radiologist read images for a clinician anywhere within the defined network, thereby allowing patients to receive care in their own locales and in some instances, allowing clinicians to receive radiology services with only a technician, rather than a radiologist, on-site.

## IMPROVED PATIENT OUTCOMES

Perhaps the most difficult benefit to quantify is improved patient outcomes, yet such benefits represent perhaps the most compelling argument for PACS implementation. Improved outcomes are the product of many factors: image clarity, fewer lost exams and redos, multi-availability of digital images, and, most important, turnaround time. The latter is especially true where distance is involved. No simple quantitative value can be placed on improved detection of disease or image availability, nor is there a simple way to assess the value of a secure and fast repository of images. These factors will have a huge impact on the way medicine is practiced and the quality of care patients receive.

## COMPETITIVE IMPACT

A financial analysis could attempt to capture whether implementing PACS would have an impact on the organization's overall revenue and admissions. Many enterprises, especially integrated delivery networks, will see the decision to implement PACS simply as a necessary step in maintaining their market position. If PACS is financially justified and greatly improves patient outcomes, the entity or system that adopts it first will have a competitive advantage.

## REFERENCES

- Avrin DE. PACS economic issues and justifications. American Association of Physicists in Medicine 1999 Summer School: Practical Digital Imaging and PACS. Sonoma State University; Rohnert Park, CA. Madison, WI: Medical Physics Publishing.
- Avrin DE, Andriole KP, Yin L, et al. A hierarchical storage management (HSM) scheme for cost-effective on-line archival using lossy compression. *J Digit Imaging*. 2001;14(1):18-23.

- King BF, Ward S, Bruesewitz RT, et al. Cost of film: purchasing, processing, packaging, storing and disposal over the lifetime of a film examination in a large radiology department. *SCAR Proceedings*. 1996:152–156.
- Pratt HM, Langlotz CP, Feingold ER, et al. Incremental cost of department-wide implementation of a PACS and CR. *Radiology*. 1998;206:245–252.
- Spiro HT. *Finance for the Nonfinancial Manager*. 2nd ed. New York: John Wiley and Sons; 1982.

# LEGAL ISSUES AND FORMAL POLICIES

---

---

GORDON SMITH • DAVID S. HIRSCHORN

**T**he past decade has brought explosive growth in PACS technology, making large-scale teleradiology an integral part of many radiology practices. This reality has left legislatures, the courts, and a wide variety of organizations that formulate healthcare policy scrambling to keep pace with an ever-changing practice environment. The result is a patchwork of laws, court decisions, and formal policies formulated by the American College of Radiology (ACR) and others that address a wide variety of issues relating to teleradiology. Some issues, such as the medical licensure and institutional credentials necessary to practice teleradiology in a given jurisdiction, are fairly well defined. However, the majority of legal and policy issues that confront this increasingly important aspect of radiology practice are unsettled or not even addressed, leaving a broad range of unanswered questions. In addition to the challenges of addressing litigation and standards for the clinical aspects of radiology, the Health Insurance Portability and Accountability Act (HIPAA) has brought upon radiology a level of compliance complexity never seen before. The regulations penetrate every area of the department, and the policies further define and expand the obligations within



HealthCare Financing Administration (HCFA) regulations that were based on the Privacy Act of 1974.

## THE CURRENT SITUATION

We are currently in the midst of an era of rapid technical innovation and change that is unequaled in recorded history. Computing power that required whole rooms a few short decades ago can now be conveniently carried. The cost of this computing power has fallen dramatically, making possible the widespread use of powerful computing platforms. The accompanying dramatic innovation in communications technology has allowed the inexpensive transfer of large volumes of data over long distances.

These innovations have combined to make teleradiology and picture archiving and communication systems (PACS) technology available in many medical settings. Digital images are acquired, transmitted, displayed, and stored in a wide variety of settings. These range from purely local exercises, such as the interpretation of computed tomography (CT) images at a scanner's dedicated workstation, to transmission of images over hundreds or thousands of miles for official interpretation and storage. The activities currently possible through available technology are in many ways limited only by the creativity of those who use it.

Like most activities with the potential to impact the health of the population, teleradiology and PACS technology are subject to controls established by the law and policies developed by various organizations. Laws enacted by Congress and by the various state legislatures are implemented by the executive branch through administrative agencies. These agencies draft regulations that define the day-to-day operation of the legislation, providing detail that is often absent from the law itself. The Food and Drug Administration (FDA) and the HCFA are well-known examples of federal agencies; a variety of state administrative agencies perform similar functions for state laws.

The content, meaning, or appropriateness of laws and regulations is often subject to dispute. Parties may contend that administrative agencies misinterpreted the law when drafting regulations, or perhaps overstepped the discretion that the law allowed them. In extreme cases, there may be questions as to whether Congress or the legislature possessed the authority to pass the law itself. In any such dispute, it is the courts at the federal, state, and local level that serve as the final arbiter of the law. In this role, they shape the final enforcement of any legislation.

Another source of control is policies, guidelines, and standards developed by private organizations with interests in a field. In teleradiology and PACS, the ACR has played a key role in developing standards for both the equipment employed and the role of radiologists and other personnel in applying the technology. While these standards lack the force of law, they serve an important function in defining teleradiology and PACS for those both in and outside the field of radiology. In this context, similar standards have been used by courts in examining disputes involving medical practice.

The various sources of law and policy do not ordinarily prospectively address issues. Typically, legislatures, administrative agencies, and professional organizations develop law and policy after problems have developed that demand resolution. This means that a conflict or problem must first occur and be identified before any action is taken.

Even when the need for a new law or policy is recognized, developing that law or policy is not a quick or easy process. Congress and the state legislatures may take years to draft and enact legislation, and administrative agencies years to define the new law with regulations. Courts may be even slower to resolve new legal problems, as a number of decisions on similar disputes are typically needed to form a body of law. Even professional organizations with vested interests in areas such as teleradiology and PACS, such as the ACR, generally have in place a complex mechanism to develop standards or guidelines, a process that may take years after the need for action is identified.

The result—in rapidly changing, technologically driven fields such as teleradiology and PACS—is a definite disparity between the capabilities of the technology and the institution of laws and policies to govern its use. Today, only a fraction of pertinent teleradiology and PACS issues have been addressed. Although almost every state has explicit or implied licensure requirements for radiologists interpreting teleradiology images from inside its borders, there is a dearth of court decisions addressing the various legal issues that are sure to affect its everyday practice. Furthermore, new laws are passed and new policies established on an ongoing basis, with the pace of these new controls bound to increase as the technology matures and its use becomes even more widespread.

This chapter outlines current law and policy as they pertain to teleradiology and PACS. It also outlines issues with the potential to affect the fields in the near future. It is not a substitute for qualified legal advice, and radiologists engaging in these activities are urged to consult qualified legal counsel before employing these technologies.

## STANDARDS AND POLICIES OF PROFESSIONAL ORGANIZATIONS

The provision of medical services has been a long-standing focus of professional societies, and the activities made possible by the development of teleradiology and PACS technology are no exception. Given the technology's pronounced impact on the practice of radiology, the ACR has taken a leading role in defining what constitutes professionally acceptable teleradiology and PACS services, developing a variety of standards and other policies. The American Medical Association (AMA) has also examined the practice of teleradiology and telemedicine. The standards and policies developed by such organizations do not have the force of law, but they do represent a detailed consensus of expert opinion in the field. As such, they may serve as important indicators regarding what constitutes the professional standard of medical practice in teleradiology and PACS.

### AMERICAN COLLEGE OF RADIOLOGY

The ACR is a leading professional society in radiology, with a membership composed of radiologists, radiation oncologists, and medical physicists. As part of an effort to advance the science of radiology and improve the quality of radiology services, the college has developed a formal mechanism for establishing and revising standards for the various subspecialty areas that make up the profession. Each standard represents a consensus policy statement by the college. Effective January 1, 1999, the ACR established new standards for teleradiology and digital image data management.

### ACR STANDARD FOR TELERADIOLOGY

The ACR standard covers a wide variety of issues related to teleradiology. It stresses that teleradiology must be of sufficient quality to perform the indicated task. When a system is used to perform an official interpretation, there should not be a "clinically significant loss of spatial or contrast resolution from image acquisition through transmission to final image display." From this overriding principle, the document describes in detail the personnel and equipment considered necessary to conduct teleradiology.

Initially, the standard outlines the qualifications of personnel obtaining images at the transmitting site. These individuals must be qualified to

perform the specific examination being performed. In all cases, a licensed and/or registered radiologic technologist, nuclear medicine technologist, or sonography technologist is needed. In addition to appropriate technologists, a qualified medical physicist and an "image management specialist" are desirable to have on-site or as consultants. The document defines an image management specialist as an individual who is "qualified by virtue of education and experience" to provide service to the teleradiology system.

The physician performing the official interpretation of transmitted images must have a basic understanding of the strengths and weaknesses of teleradiology, as well as be qualified to interpret the particular diagnostic modality at issue. With regard to what constitutes adequate qualification, the standard refers to other ACR standards for rendering interpretations on the various imaging modalities. Notably, the teleradiology standard states that this physician should maintain licensure appropriate to the delivery of teleradiology services at both the transmitting and receiving sites. This effectively requires a physician interpreting teleradiology to maintain appropriate licensure in multiple states, if teleradiology is conducted across state lines and the state(s) involved require such licensure. The standard maintains a similar position on staff privileges: If images are transmitted from a hospital, the interpreting physician should be credentialed and obtain appropriate privileges at that institution.

Similar to legal requirements faced by physicians interpreting locally produced images, the ACR teleradiology standard holds the physician providing the official interpretation of teleradiology images responsible for the quality of the images being reviewed. Simply put, this position makes it difficult for physicians providing official teleradiology interpretation to escape potential liability for poor-quality images. Physicians providing official interpretations are also cautioned to consult with their professional liability carrier to ensure coverage in both sending and receiving sites. A large portion of the teleradiology standard addresses technical and legal issues associated with the equipment used and the images displayed and stored by that equipment. All new equipment acquisition should comply with the Digital Imaging and Communications in Medicine (DICOM) standard, developed by the ACR and the National Electrical Manufacturers Association (NEMA). Two matrix categories are established for rendering official image interpretation. A small matrix (512 × 512 resolution with a minimum of 8-bit depth) is deemed sufficient for computed tomography (CT), magnetic resonance imaging (MRI), ultrasound (US), nuclear medicine (NM), digital fluoroscopy, and digital angiography. Computed radiography and digitized radiographs are considered large-matrix studies (a minimum of 2.5 line pairs per millimeter [lp/mm] spatial resolution at a minimum of 10-bit depth).

Image data for teleradiology systems may be obtained by both direct image capture for purely digital images or by secondary image capture for film images that are digitized. Direct image capture is the “most desirable” method of acquisition for primary diagnosis. Regardless of acquisition method, images must have annotation capabilities that allow data such as patient name, identification number, name of transmitting facility, type of examination, anatomic orientation, and method of compression displayed on the image. The standard allows the use of both reversible and irreversible compression, assuming that a qualified supervising physician determines that there is no reduction in “clinically diagnostic image quality.” These compression methods should be reviewed periodically by the supervising physician to “ensure appropriate image quality.” Data transmission is required to have adequate error-checking capability, and there must be no loss of clinically significant data during this transmission.

Display characteristics for the monitors used in officially interpreting teleradiology images are described. These should have a luminance of at least 50 ft-lamberts and be located in areas with suitable room lighting. Image manipulation features should include window and level adjustments, pan and zoom, the capability to rotate or flip images, and the ability to calculate and display accurate linear measurements and pixel values (as appropriate for the modality being interpreted). The images should be accurately associated with the correct patient study and demographic information, and any compression or similar processing should be noted. Requirements for displays not being used for official interpretation are noted to be less stringent, though the exact characteristics are not delineated.

Archiving and retrieving image data receive significant attention in the standard. Prior examinations should be retrievable from the archive in a time frame appropriate to the clinical needs of the facility and medical staff. Any system should provide storage capable of complying with all facility, state, and federal regulations regarding medical record retention. Images stored at either the transmitting or receiving site should meet the specific jurisdictional requirements of the transmitting site. Images interpreted off-site need not be stored at the receiving facility. However, if such data are maintained at the receiving facility, the data retention period must meet the jurisdictional requirements of the receiving jurisdiction as well. All policies relating to the storage of image data should be written and equivalent to policies and procedures that exist for hardcopy medical images.

A teleradiology system should have protections to ensure the security of archived data. Specifically, the confidentiality of patient data must be addressed, as well as measures to safeguard the data from intentional or unintentional corruption. These protections should apply to both the network and the software it employs.

Finally, the standard addresses practical, day-to-day issues of teleradiology. Written policies and procedures to ensure a continuity of care consistent with that for hardcopy images are suggested. Mentioned are internal redundant systems, backup telecommunications links, and a disaster plan. At least monthly image quality control using a test image is described. Spatial resolution at such testing should be consistent with the specific matrix being employed, that is, small or large.

Currently, there is little indication as to how this revised teleradiology standard may be applied in practice. Given the ACR's reputation and the need for minimum standards in clinical teleradiology practice, many of the details of the standard will probably be adopted by radiologists practicing teleradiology. However, given the rapid advancement of technology, it is virtually certain that some of the standard's technical details will shortly be obsolete. The portions of the document calling for appropriate licensure in both the sending and receiving jurisdictions are likely to be considerably more enduring, as are the provisions applying to the archiving and retrieval of teleradiology data.

#### ACR STANDARD FOR DIGITAL IMAGE DATA MANAGEMENT

The ACR maintains a separate standard for digital image data management. Its provisions are applicable to any system of image data management, from single-modality or single-use system to a complete PACS system, as would be used for teleradiology. As a result, there is considerable overlap with the ACR Standard for Teleradiology, which focuses on PACS. Like the teleradiology standard, the digital image data management standard states that the examination that serves as the data source is subject to the specific ACR standard for that modality.

The goals of digital image data management as outlined in the standard include, but are not limited to: (1) initial acquisition or generation of accurately labeled and identified image data; (2) transmission of data to an appropriate storage medium from which they can be retrieved; (3) retrieval of data from available prior imaging studies for comparison; (4) transmission of data to remote sites for consultation, review, or formal interpretation; (5) appropriate compression of image data to facilitate transmission or storage, without loss of clinically significant information; (6) archiving of data to maintain accurate patient medical records in a form that is retrievable in a timely fashion, meets applicable facility, state, and federal regulations, and maintains patient confidentiality; and (7) administration with appropriate database management procedures.

Most of the document itself is devoted to describing in detail how these goals are to be accomplished. Qualifications and responsibilities for personnel, including physicians, electronic/computer assistant, medical image physicist, and image management specialists are outlined, largely paralleling descriptions in the teleradiology standard. Similarly, compliance with the DICOM standard is “strongly recommended,” and image categories for official interpretation are split into those for small and large matrices. The definitions for these matrices and the type of imaging modalities in each type of matrix are identical to the teleradiology standard, as are the descriptions of image acquisition and annotation capabilities. Transmission standards likewise mirror those detailed in the teleradiology standard.

Archiving and retrieval sections of the digital image data management standard also reiterate those found in the teleradiology standard. Storage capacity must be capable of complying with all facility, state, and federal regulations regarding medical record retention, with images stored at either the transmitting or retrieval site complying with the requirements of the transmitting jurisdiction. Storage is not necessary at the receiving site, but if such storage is undertaken, the retention period of that jurisdiction must be met as well. Security to protect the confidentiality of patient identification and imaging data should be present. All policies relating to the achieving and storage of digital image data should be equivalent to those in existence for hardcopy records and should be in writing. For clinical use, any system must allow timely retrieval of archived images, as well as mechanisms to ensure continuity of care.

## AMERICAN MEDICAL ASSOCIATION POLICIES ON TELEMEDICINE AND TELERADIOLOGY

The AMA is the largest medical professional society in the United States, encompassing the spectrum of medical specialties and issues. The growing importance of telemedicine, which includes teleradiology and PACS, has captured the association’s attention at its highest levels. This has led to the issuance of several reports and implementation of certain policies.

In 1996, the AMA published “The Promotion of Quality Telemedicine,” which was jointly issued by the Council on Medical Education and the Council on Medical Service. In this document, the AMA supports the ACR position that physicians providing “authenticated interpretation of images transmitted by teleradiology” should maintain licensure “appropriate to the delivery of radiologic service” at both the transmitting and receiv-

ing sites. As noted previously, this position generally requires that a radiologist interpreting telemedicine studies maintain full licensure in both the transmitting and the receiving jurisdictions. However, if the service provided is “curbside consultation,” a phrase used to describe an informal second opinion where there is no expectation of compensation, the AMA policy recognizes that a full and unrestricted license is not needed.

AMA policy, however, does not recognize the ACR Teleradiology Standard and related standards as such. Under AMA policy for “practice parameters,” as recognized in the AMA Policy for the Promotion of Telemedicine, such parameters serve as “educational tools” and “strategies for patient management that are designed to assist physicians in clinical decision making.” This is distinct from the legal concept of a “standard of care,” the level of medical care established necessary to defeat allegations of negligence in a malpractice action. Generally, this standard is established by physicians, testifying as experts as to the level of care required. Furthermore, a related policy states that “practice parameters developed by a particular medical specialty or specialties should not preclude the performance of the procedures or treatments addressed in that practice parameter by physicians who are not formally credentialed in that specialty or specialties.” Thus, under existing AMA policy, ACR standards on teleradiology and digital image data management serve only an educational purpose and are not acknowledged to establish an actual standard of care.

The AMA has also tracked developments in telemedicine and teleradiology. In 1996, the House of Delegates, the AMA’s governing body, adopted a resolution directing the association to monitor activities of hospitals, specialty societies, and regulatory agencies that affect telemedicine and submit a report. The result of this resolution was the Status Report of Telemedicine, issued at the 1997 interim meeting, a substantial portion of which outlined ACR actions in the area. ACR initiatives such as the DICOM standard, developed in conjunction with the NEMA, were acknowledged. The document also noted that the FDA Center for Devices and Radiological Health had encouraged such collaboration between the clinical community, as represented by the ACR, and manufacturers of diagnostic imaging equipment.

Given the growing importance of telemedicine in general and teleradiology in particular, there is little doubt that the AMA will continue to track developments and generate policy in the area. For the present, it is unlikely that the association will change its stance requiring full and unrestricted licensure in both transmitting and receiving jurisdictions in the setting of teleradiology, or acknowledge that ACR standards represent the professional standard of care.



## GOVERNMENT REGULATIONS

Both the federal and state governments are involved in the regulation of teleradiology and PACS. This regulatory authority stems from legislation that controls medical devices, healthcare benefits, and the practice of medicine, with the regulations themselves drafted by a variety of administrative agencies. Generally, regulation at the federal level is directed at medical devices and the provision of healthcare benefits. At the state level, the dominant activity is regulation of medical practice.

### FEDERAL GOVERNMENT

#### FOOD AND DRUG ADMINISTRATION

The FDA has its regulatory authority for medical devices grounded in the Food, Drug and Cosmetics Act, as amended by Medical Device Amendments of 1976 and other amendments, which requires that products be safe and effective for their marketed indication(s). The definition of a “medical device” under the act is extremely broad—broad enough to include devices employed for teleradiology and PACS. Devices regulated by the agency are broken down into several distinct groups. Initially, all devices are arbitrarily separated into those legally marketed prior to implementation of the Medical Device Amendments on May 28, 1976, and those marketed after that date. These are known as “pre-amendment” and “post-amendment” devices, respectively.

Pre-amendment devices are further divided into 3 classes, based on potential patient risk. Devices with the least risk are placed in class I, which is subject only to “general controls.” Class I products are not individually regulated. Rather, their safety and effectiveness are assured by general controls, which include manufacturing and labeling controls. General controls are considered important for all medical devices. Accordingly, they also apply to class II and III products.

Class II is the intermediate regulatory category for devices with higher risk to patients than class I but not requiring the highest degree of regulation. Products in this class are subject to “special controls,” specific regulations designed to assure their safety and effectiveness. As with class I, these devices are not individually regulated, with each generic product type subject to applicable special controls.

Class III is the most stringent regulatory category. It is reserved for products with either a potentially unreasonable risk of patient injury or with

insufficient data to establish actual patient risk. Devices in this class are technically subject to a premarket approval process, requiring demonstration of safety and effectiveness prior to marketing. However, pre-1976 class III products are “grandfathered” and may be legally marketed until such time as the FDA requests such data and the manufacturer either fails to provide them or the data fail to show safety and effectiveness.

Post-amendment devices are generally subject to a premarket notification process, which generally applies to higher-risk class II and all class III products. This requires that a manufacturer provide the FDA notice of its intention to market a product. If the agency determines that the product is “substantially equivalent” to a pre-amendment device (or a post-amendment device that has been reclassified to class I or II), that device may be legally marketed subject to the regulations currently applicable to its “predicate” device. Should there be no pre-1976 equivalent, the device is automatically placed in class III, subject to the premarket approval process. Lower-risk products may be reclassified to class I or II, although this generally requires evidence that the device’s risk is appropriate to the new classification.

Teleradiology and PACS were not in existence in the pre-1976 world of medical devices. Though these post-amendment devices could have been automatically placed in class III, the FDA treated teleradiology and PACS equipment as accessories to the imaging devices that they serviced, avoiding the premarket approval process. However, this made marketing approval for the devices somewhat complicated, as the products were not themselves classified.

The FDA moved to end this system in 1996, issuing the policy statement “Telemedicine Related Activities.” While reinforcing the agency’s authority to regulate teleradiology and PACS devices, the statement proposed formally classifying the products. Image storage devices and medical image devices were to be placed in class I and exempted from the premarket notification requirement unless irreversible compression was used. Medical image digitizers, medical image hardcopy devices, and PACS systems were to be class II products. General purpose products used in a medical setting were not to be regulated, unless labeled for a medical use. The latter category could include such items as word-processing software employed in a PACS system.

The agency issued its final rule effecting these changes on April 29, 1998. As proposed in *Teleradiology and Related Activities*, these regulations placed medical image storage devices in class I, exempt from the premarket notification requirement unless irreversible compression is used. Medical image digitizers, medical image hardcopy devices, and PACS were made class II devices. A number of “voluntary standards” are to serve as special controls

for these devices: (1) DICOM; (2) Joint Photographic Experts Group (JPEG), which specifies methods for reversible and irreversible compression of digital medical images; and (3) the Society of Motion Picture and Television Engineers test pattern, used to test monitors and printers for acceptance and quality control purposes.

## HEALTHCARE FINANCING ADMINISTRATION

The HCFA oversees the federal Medicare program, disbursing vast sums of money to healthcare providers and institutions nationwide. Given the scope of Medicare, HCFA regulations applicable to Medicare fund recipients have a broad impact on the provision of U.S. health care. HCFA itself is governed by the Privacy Act of 1974, a federal statute that protects the confidentiality of individually identifiable data. In practice, the act requires that HCFA keep the records of its Medicare patients confidential. HCFA is also subject to certain provisions of HIPAA, in which Congress mandated certain security and electronic signature requirements.

Recently, HCFA has become concerned that certain electronic data transmissions have the potential to violate patient confidentiality and hence the Privacy Act of 1974. Its response was the HCFA Internet Security Policy, issued in November 1998. This document applies to what HCFA describes as "HCFA Privacy Act-protected and/or sensitive HCFA information," which includes: (1) all individually identifiable data held in systems of records; (2) payment information that is used to authorize or make cash payments to individuals or organizations; (3) proprietary information that has value in and of itself and that must be protected from unauthorized disclosure; and (4) computerized correspondence and documents that are considered highly sensitive and/or critical to an organization and that must be protected from unauthorized alteration and/or premature disclosure.

The HCFA Internet Security Policy allows covered data to be transmitted via the Internet, as long as "an acceptable method of encryption" is utilized to provide confidentiality and integrity of the data. Furthermore, authentication or identification procedures must be employed to assure that both the sender and the recipient of the data are known to each other and are authorized to receive and decrypt such information. The policy covers all systems or processes that use the Internet or interface with the Internet to transmit sensitive data. However, it does not apply to local data-at-rest or local host or network protections, although it is explicit that such local data must still be protected by "all necessary measures."

The HCFA Internet Security Policy describes in considerable detail the technical specifications of acceptable practices. Minimally acceptable encryp-

tion methods as of November 1998 include algorithms such as Triple 56-bit DES (defined as 112-bit equivalent) for symmetric encryption, 1024-bit algorithms for asymmetric systems, and 160 bits for elliptical curve systems. The agency explicitly reserves the right to increase these minimum levels when “deemed necessary” by advances in techniques and capabilities associated with the processes used by attackers to break encryption.

Acceptable authentication approaches, accomplished over the Internet via an “in-band” process, include: (1) formal certificate authority-based use of digital certificates; (2) locally managed digital certificates, provided that all parties to the communication are covered by the certificates; (3) self-authentication, as in internal control of symmetric “private keys”; and (4) tokens or “smart cards.” Acceptable identification approaches, undertaken outside the Internet via an “out-of-band” process, include: (1) telephonic identification of users and/or password exchange; (2) exchange of passwords and identities by U.S. certified mail; (3) exchange of passwords and identities by bonded messenger; (4) direct personal contact exchange of passwords and identities; and (5) tokens or smart cards.

Entities subject to the HCFA Internet Security Policy must modify their security plan to detail the methodologies and protective measures used if they employ the Internet for transmission of covered data and to adequately test these implemented measures. HCFA reserves the right to audit these organizations and their security policies. Finally, any organization wishing to transmit covered data via the Internet must inform HCFA of its intent to do so.

HCFA is in the midst of promulgating formal regulations addressing security of electronic individual healthcare information, as well as health plan use of electronic signatures.

## DEPARTMENT OF HEALTH AND HUMAN SERVICES

**HEALTH INSURANCE PORTABILITY AND ACCOUNTABILITY ACT** The passage of HIPAA has driven many to demand improvements in the management of the information systems within the healthcare system. When the Privacy Act of 1974 was passed, many saw the legislation as ineffective because it allowed the disclosure of the information without the subject’s approval when the use of the information was routine. The rules that governed the definition of routine use were expanded, and the loopholes in the law continued to increase. Second, the burden of enforcement was placed entirely on the individual via the filing of a civil suit. However, the framework established by the Privacy Act of 1974 and HCFA have led to the development of a framework for HIPAA.

The Health Insurance Portability and Accountability Act is broken down into 3 principal rules:

- ▶ Transaction Rule: The facilitation of the exchange of information between providers and payers
- ▶ Privacy Rule: The empowerment of patients for access and control of their medical information
- ▶ Security Rule: The safeguards for the information exchanged in the transaction and privacy rules

Since computerized information systems drive almost every department within the hospital, the Security Rule will rely heavily on information technology to provide the required support to meet the other rules enforcement. Experts estimate that probably 10% or fewer of private healthcare organizations have adequate security; in other words, 90% or more have inadequate security. The implementation of security within a practice's information system is a complex process ranging from the establishment of dependable secure workflows of most departmental operations to the implementation of many new technical or operational changes to the existing information technology. The details regarding the implementation of policies and procedures to ensure HIPAA compliance are well beyond the scope of this book.

## STATE GOVERNMENT

### LICENSURE

At its most basic level, teleradiology is the practice of medicine. The right of the individual states to license such practice has been settled law in the United States since the turn of the century, when the U.S. Supreme Court upheld a West Virginia statute requiring that physicians practicing in that state obtain a license based on criteria established by the state (*Dent v. West Virginia*). Today, states enforce their licensure prerogative through medical practice statutes, which typically define what constitutes the “practice of medicine” and therefore who is subject to medical licensure. The definition of the practice of medicine is usually broad, as with North Carolina's statute:

any person shall be regarded as practicing medicine or surgery . . . who shall diagnose or attempt to diagnose, treat or attempt to treat, operate or attempt to operate

on, or prescribe for or administer to, or profess to treat any human ailment, physical or mental, or any physical injury to or deformity of another person

Although teleradiology is not specifically mentioned in this and other statutes, there is little doubt that the broad definition of medical practice encompasses the in-state teleradiology practitioner. The impact on out-of-state physicians who consult about patients located within the jurisdiction is less clear. To eliminate this confusion, many states have amended their medical practice statutes to clarify their applicability to out-of-state teleradiology practitioners (Goldberg and Gordon 1998). In states where statutes have not been altered, the impact on out-of-state practitioners remains uncertain.

Many states have various exceptions to their licensure requirement. For example, out-of-state physicians rendering emergency treatment are often exempt. "Occasional" consultants may be exempt, but the definition of what level of activity qualifies differs between states. Several states have "border states exceptions," which exempt licensed physicians in immediately neighboring states from the state's licensure requirement. Given the nature of teleradiology practice, with its typically nonemergent, recurrent nature and broad reach, it is likely that the applicability of all of these exemptions will be limited.

With current medical practice statutes and their exemptions, licensure requirements for out-of-state teleradiology practitioners fall into 1 of 3 general categories: (1) full licensure is either expressly required by statute or presumed because teleradiology and/or telemedicine is not specifically mentioned in the applicable medical practice act and no exemption applies; (2) a "special purpose" license for out-of-state teleradiology practitioners is available; and (3) full licensure is not required, though something short of full licensure may be necessary. The last 2 categories are infrequently encountered.

Given the potential consequences of violating medical practice statutes, it is advisable to exercise caution in all questionable practice situations. Loss of licensure in a practitioner's home state, exclusion from federal Medicare and Medicaid programs, and/or loss of malpractice insurance may all be indirect consequences of practicing without an appropriate license (California Business and Professions Code 1998a; 42 USCA. 1998; NORCAL Mutual Insurance Co. 1997). Interestingly, violation of the medical practice statute itself is typically only a misdemeanor (California Business and Professions Code 1998b).

Licensure requirements, current as of April 1999, for the 50 states appear in Table 8.1. Also included are pertinent, specific state requirements.

**TABLE 8.1**  
Licensure Requirements (1999)

<b>State</b>	<b>Code</b>	<b>Specific Requirements</b>
Alabama	3	Grants a 3-year special-purpose license to nonresident telemedicine practitioners. Excludes informal or uncompensated consultations. Subjects licensee to Alabama medical board jurisdiction and requires licensee's home state to issue reciprocal telemedicine licenses to Alabama physicians.
Alaska	1	
Arizona	2	"Single or infrequent" consultations are exempted.
Arkansas	2	Episodic consultations with Arkansas physicians, provision of services unavailable in Arkansas, or physical travel to the state to provide care are exempted.
California	3	No license required so long as the telemedicine consultant does not have ultimate authority over the patient; requires specific informed consent from the patient to use telemedicine consultation; exempts telephone conversations and e-mail messages between patient and practitioner.
Colorado	2	"Occasional" consultations exempted.
Connecticut	2	"Occasional" consultations exempted.
Delaware	1	
District of Columbia	1	
Florida	2	Full licensure for physicians providing official authenticated interpretations through an ongoing regular arrangement.
Georgia	2	
Hawaii	3	Telepractitioners exempted from licensure if local physician maintains primary control over the patient's care.
Idaho	2	
Illinois	2	Out-of-state physicians practicing telemedicine subject themselves to the jurisdiction of Illinois courts.
Indiana	2	Full licensure for telemedicine on a regular routine or nonepisodic basis.
Iowa	1	
Kansas	2	Exemption for occasional consultation; border states exemption.

TABLE 8.1

*Continued*

State	Code	Specific Requirements
Kentucky	1	
Louisiana	1	No consultation exception.
Maine	1	No consultation exception.
Maryland	1	
Massachusetts	1	Opinion of medical board attorney that full licensure needed.
Michigan	1	
Minnesota	1	
Mississippi	2	Exemption if local physician requests nonresident physician's services. The resident physician must have a prior relationship with the patient being treated via telemedicine.
Missouri	2	Exemption when consulting with local physician.
Montana	3	A bill pending in the legislature would require a telemedicine certificate issued by the medical board; passed House, pending in Senate as of 2/22/99.
Nebraska	2	
Nevada	2	
New Hampshire	1	Bill pending in legislature to explicitly require full licensure for physicians who provide teleradiology services on a regular contractual or frequent basis.
New Jersey	1	
New Mexico	1	
New York	1	Border states exception.
North Carolina	2	Exemption for infrequent consultations. Residents may bring malpractice claims against telemedicine practitioners in North Carolina courts.
North Dakota	1	Bill pending in legislature to require full licensure.
Ohio	1	
Oklahoma	2	Brief consultation exception; telemedicine practitioners submit to the jurisdiction of Oklahoma courts.



**TABLE 8.1**  
Licensure Requirements (1999) *(Continued)*

<b>State</b>	<b>Code</b>	<b>Specific Requirements</b>
Oregon	1	Bill pending in legislature to require a special telemedicine license that is not a limited license but still does not allow the out-of-state physician to practice in the state, except across state lines.
Pennsylvania	1	
Rhode Island	1	
South Carolina	1	
South Dakota	2	Consultation exception limited to maximum 24-hour period in any 1 year.
Tennessee	2	On 5/15/96 the medical board was authorized by the legislature to issue special telemedicine licenses; as of 2/25/99 there is a bill pending in the legislature that would make transmission of patient medical information via telemedicine technology to a person in another state who is not licensed in Tennessee grounds for license suspension or revocation.
Texas	3	The state board of medical examiners is authorized to issue special-purpose licenses for telemedicine; otherwise, full licensure required.
Utah	2	Consultation exception repealed.
Vermont	1	Bill pending to authorize special-purpose license.
Virginia	1	
Washington	1	Bill pending that would require telemedicine practitioner to be sponsored by a local physician.
West Virginia	2	Consultation exception provides that consultant cannot consult for more than 3 months in his lifetime.
Wisconsin	1	
Wyoming	1	

Key: 1: States that have not specifically addressed the telemedicine licensure issue, so that full licensure is presumed.

2: States that specifically include telemedicine in their definition of medical practice and expressly require full licensure.

3: States requiring something other than full licensure, such as a special-purpose license or no license in the state.

Given the myriad of state licensure requirements, some have advocated a more uniform system of licensure for telemedicine/teleradiology. In 1996, the Federation of State Medical Boards suggested that the states adopt limited telemedicine licenses. However, leading national medical organizations, such as the ACR and the AMA, have adopted policies advocating full licensure in each state where a physician practices teleradiology. The states themselves heavily favor full licensure for physicians treating patients within their borders and appear extremely reluctant to surrender any authority to regulate such medical care. In this current climate, it is unlikely that any type of national licensure for teleradiology practice will emerge in the foreseeable future.

### OTHER STATE ISSUES

In addition to licensure, many states have enacted legislation that affects teleradiology. Generally, these laws and regulations address teleradiology/telemedicine initiatives within the state, or attempt to coordinate such activities to achieve a public health goal. For example, some states are actively promoting telemedicine to provide care to their rural populations. A complete description of these nonlicensure activities is beyond the scope of this discussion.

## RELATED LEGAL CONSIDERATIONS

The practice of teleradiology and PACS storage of image data raise a number of legal concerns, mostly related to state law doctrines. These include medical malpractice and record-keeping issues. To date, there are no cases known to the authors or other commentators directly addressing teleradiology and PACS (Caryl 1998). Accordingly, most analysis in this area is by analogy to conceptually similar fact situations.

### MEDICAL MALPRACTICE IN TELERADIOLOGY

#### ESTABLISHING A CLAIM

Teleradiology is medical practice and, as such, exposes a physician to liability under state tort law, commonly known as medical malpractice. Successful malpractice actions require 4 elements: (1) a duty to the patient; (2) a

negligent breach or violation of that duty; (3) patient injury as a result of that negligence; and (4) actual damages from the injury. Assuming that a patient has suffered injury that has resulted in damages, as is the case in most malpractice actions, the question becomes whether the teleradiology practitioner owes a duty to the patient whose images he or she interprets and what constitutes negligence in that interpretation.

There is no definitive case law addressing the existence of duty owed to a patient by a teleradiology practitioner. However, most commentators believe that a doctor-patient relationship exists between a radiologist interpreting teleradiology images and the patient whose images he or she reviews, a relationship that establishes a duty to that patient (Caryl 1998; Cuzmanes and Orlando 1997). A court decision supporting this proposition is *Hand v. Tavera* (1993), in which a physician under a managed care contract who refused to hospitalize a patient was held to have formed a doctor-patient relationship, despite the fact that he had never met or spoken with that patient. The court reasoned that the relationship was established as the patient had paid for the physician's services. Another decision is *McKinney v. Schlatter* (1997), which found that a telephone consultation is sufficient to establish a doctor-patient relationship, when a physician relied on a cardiologist's advice that a clinical problem was not cardiac in nature. Given that a teleradiology practitioner is paid for his or her interpretation, and that interpretation is ordinarily relied on to guide clinical decision making, these cases indicate that typical teleradiology consultations will be sufficient to establish a duty to the patient.

It is less clear that a doctor-patient relationship is established when the consultation is informal, no compensation is received, and no official interpretation is rendered. Specifically, if the teleradiology practitioner is engaged in a "curbside consult," there is the possibility that no relationship will be found (Berger and Cepelewicz 1996). However, if the radiologist receives or expects any compensation from the consult, it is doubtful that any "curbside consult" exception would apply.

A second key requirement of a successful malpractice action is negligent breach of a physician's duty to the patient. Negligence exists when a physician has violated the medical standard of care, a legal concept whose exact definition varies among jurisdictions. Generally, this standard is established by physicians, testifying as experts, as to what constitutes acceptable medical practice in the fact situation before the court. Although these standards were originally based on practice patterns in the local community where injury occurred, there has been a growing trend in medical malpractice to a national standard of care, applicable across jurisdictions. Teleradiology, with its wide geographic sweep and cross-jurisdictional nature, will

almost certainly involve a national standard of care. The exact form this standard takes will depend on case law developed as malpractice cases involving teleradiology inevitably come before the courts.

### CHOICE OF LAW

Medical malpractice is a legal action based in state law—law that may differ greatly among jurisdictions. These differences become problematic when the teleradiology practitioner interprets images of a patient who resides in and was imaged in another state. Here, the question becomes which law, that of the transmitting state or that of the receiving state, to apply.

Although teleradiology and PACS are new technologies, the choice of which state law to apply when a plaintiff and defendant are residents of different jurisdictions is not new for the courts. Under well-established law, a state may exercise jurisdiction on an out-of-state individual or corporation provided that there are “minimum contacts” between the state and the individual or corporation (*International Shoe v. Washington* 1945). Three criteria must be met: (1) the defendant must have purposefully availed him- or herself of acting in the state; (2) the cause of action must have arisen in the state; and (3) the defendant’s acts must have a substantial enough connection to make exercise of jurisdiction reasonable (*Compuserve, Inc. v. Patterson* 1996). In the setting of commercial activity, it is widely acknowledged that committing an act of negligence in a state or doing business in that jurisdiction satisfies these requirements. Commentators examining teleradiology believe that this doctrine will be used to subject practitioners to the laws of the transmitting jurisdiction, although in the absence of applicable court decisions, the question remains unresolved (Caryl 1998). Some states have acted to remove this uncertainty by enacting legislation that specifically subjects out-of-state telemedicine practitioners to the state’s jurisdiction.

The practical implications of a teleradiology practitioner being subject to the laws of the transmitting jurisdiction may be profound. A radiologist could find him- or herself facing a local judge or jury potentially hostile to an out-of-state defendant. Perhaps even more important, applicability of another state’s jurisdiction may destroy protections a physician enjoys in his or her home state, such as award limits on the amount of allowable damages.

### INSURANCE ISSUES

Interstate teleradiology practice raises professional liability insurance coverage issues related to the interpretation of images generated outside the prac-

tioner's home state. Coverage of out-of-state teleradiology activities should not be presumed. Not all insurance carriers are licensed in every state, and underwriting criteria among jurisdictions may vary. Accordingly, many policies specifically exclude coverage for out-of-state incidents, unless a rider has been added to specifically provide such coverage. This means that the unwary teleradiology practitioner subject to an out-of-state malpractice action may find his professional liability carrier reserving coverage rights or completely denying coverage.

## RECORD KEEPING

Data generated from teleradiology and PACS activities are medical records. As such, there are a myriad of considerations regarding data storage, including where the data must be maintained, their form, and the period of retention. Confidentiality of data is another consideration. Laws, regulations, and institutions' policies for film and paper records may serve as a guide, though the vary nature of electronic data will necessarily demand special considerations.

Initially, when electronic data are acquired at one site and stored at another, it is unclear whether these data must be maintained at the transmitting site, the receiving site, or both sites. As discussed previously, the ACR Standard for Teleradiology requires only that data be maintained at the transmitting site. Certainly, any applicable law, regulation, or institutional policy with regard to where data must be maintained should be observed.

The form of stored image data is another consideration. Given the present cost of electronic storage and the amount of that storage necessary to archive medical images, many centers compress data to save resources. If compression is reversible, there is no intrinsic problem. However, when irreversible, "lossy" compression is employed, there is a question of a medical record being altered and clinically relevant data being lost. In the somewhat analogous setting of hardcopy medical records, any alteration may be extremely problematic legally, as it calls into question the validity of the entire record (Andrews 1992). It remains to be seen whether storage with lossy compression practice will become an issue for the courts.

The retention period of medical records is subject to federal, state, and institutional laws and policies. Laws and policies for the jurisdiction where electronic data are being stored should be followed. In addition, the ACR Standard for Teleradiology suggests that teleradiology data being stored at the receiving facility meet the storage standards at the transmitting facility.

This policy is prudent, given the probable applicability of the transmitting state's laws to the teleradiology practitioner.

A final consideration with any stored medical record is confidentiality. Various authorities, the physician/patient privilege, ethical considerations, the constitutional right to privacy, and some state statutory law demand that this confidentiality be maintained (Andrews 1992). Although electronic storage may be a more convenient and accessible format for storing and accessing medical records, this form of record keeping may be more vulnerable to security breaches.

As described in the ACR Teleradiology Standard and the ACR Standard for Digital Image Data Management, security is needed for electronically stored medical records. The ACR standards notwithstanding, there is virtual certainty that the courts would apply the same privacy standards to electronic records that have been applied to traditional medical records (*Alberts v. Devine* 1985). This imposes a duty on the physicians and institutions using teleradiology and PACS to develop policies that assure reasonable patient confidentiality, or face potential liability for breaches of confidentiality.

## CONCLUSION

Teleradiology and PACS technology and application have expanded greatly in the last decade, in many ways leaving behind the laws and policies intended to regulate and control the field. Even where policies have been developed, such as the ACR Standard for Teleradiology and the ACR Standard for Digital Image Data Management, it is unclear what impact these policies will have on the practice of teleradiology and the use of PACS. Many of the legal and policies questions being asked by radiologists and others today will not be answered for years, as legislatures, courts, and professional societies develop approaches to the novel problems posed by the technology. Until that time, physicians using teleradiology and PACS technology should use caution and common sense when confronted with unsettled legal or regulatory questions.

## REFERENCES

42 USCA Section 1320a-7 (1998).

*Alberts v. Devine*, 479 NE2d 113 (Mass 1985), cert. denied, 474 US 1013 (1985).

- American College of Radiology: ACR Standard for Digital Image Data Management. Effective January 1, 1999. Reston, VA (1999a).
- American College of Radiology: ACR Standard for Teleradiology. Revised 1998; effective January 1, 1999. Reston, VA (1999b).
- American Medical Association: House of Delegates Resolution 117 (I-96). Interim Meeting. 1996.
- American Medical Association: status report on telemedicine. Council on Medical Service Report 8-I-97. Chicago 1997.
- American Medical Association: the promotion of quality telemedicine. Joint report of the Council on Medical Education and Council on Medical Service, 1996. In: *American Medical Association: Continuing Medical Education Resource Guide*. Chicago; 1996.
- American Medical Association: the promotion of telemedicine. Quoting AMA Policies 410.973 and 410.987. American Medical Association. Chicago; 1999a.
- American Medical Association: Policy 410:987[3]. American Medical Association. Chicago; 1999b.
- Andrews, BW. Medical records liability. *Health Lawyer*. Summer 1992;6:11.
- Berger S, Cepelewicz BB. Medical legal issues in teleradiology. *AJR* 1996;166: 505–510.
- California Business and Professions Code § 2305 (1998a).
- California Business and Professions Code § 2314, § 2315 (1998b).
- Caryl CJ. Malpractice and other legal issues preventing the development of telemedicine. *J Law Health*. 1998;12:173.
- Classification for five medical image management devices: final rule. 63 *Federal Register* 43242 (April 29, 1998).
- Compuserve, Inc. v. Patterson*, 89 F3d 1257 (6th Cir 1996).
- Cuzmanes PT, Orlando CP. Automation of medical records: the electronic superhighway and its ramifications for healthcare providers. *J Pharm Law* 1997;6:19.
- Dent v. West Virginia*. 129 US 114 (1889).
- Federal Food, Drug and Cosmetic Act, § 201 (1999) (codified at 21 USCA 321).
- Federation of State Medical Boards of the United States, Inc.: a model act to regulate the practice of telemedicine or medicine by other means across state lines, executive summary. Eules, TX; October 31, 1995.
- Goldberg AS, Gordon JF. *Telemedicine: Emerging Legal Issues*. Washington, DC: American Health Lawyers Association; 1998.
- Hagland M. Six opinions on IT security. *Health Manage Technol*. 1998;19(6):24–29.
- Hand v. Tavera*, 864 SW2d 678 (Ct. App Tex 1993).
- HealthCare Financing Administration. internet security policy. Washington, DC; November 24, 1998a.
- HealthCare Financing Administration. Proposed rule: security and electronic signature standards. 63 *Federal Register* 43242 (August 12, 1998b).

*International Shoe v. Washington*, 326 US 310 (1945).

*McKinney v. Schlatter*, 1997 WL 67702 (Ohio App 12 Dist 1997).

NORCAL Mutual Insurance Co. Professional Liability Insurance Policy. Part III  
Exclusions, A-7. San Francisco; 1997.

North Carolina General Statutes § 90-18 (1998).

Telemedicine related activities. Center for Devices and Radiological Health, US  
Food and Drug Administration. Rockville, MD; July 11, 1996.





# **TECHNICAL ISSUES**

THOMAS J. SCHULTZ

# COMPUTER FUNDAMENTALS

---

---

KEITH J. DREYER

**T**he personal computer revolution has dramatically changed our lives forever. Over the past 20 years, with the decrease in hardware costs, increase in computational performance, and a seemingly endless variety of software products, computers have entered nearly every facet of our personal and professional lives. One of the best things about computer hardware and software is that you do not have to know how they work to make them work for you. However, if you can understand just the basics (which is not too hard to do), you can better appreciate how computers can be applied to the field of radiology. With this in mind, the following pages review the basics of modern computers. Mainframes or minicomputers are not discussed. Instead, the goal is to define the fundamental components of a microcomputer—known today as a PC or personal computer—which can easily function as a workstation for primary interpretation of medical images. Important software concepts such as operating systems and programming languages are also discussed.

You can think of computers like humans in that they share 4 of the same basic functions—input, output, memory, and processing (or thought). And since all these functions occur at different locations, there needs to be

a way for them to communicate with one another. While humans use the nervous system for this task, computers use a bus. This digital electronic bus forms the backplane or printed circuit board of today's computers. As a human, there are certain things you know from birth (genetically acquired information) and certain things you learn. Computers are the same (they are just "born" a little differently). The set of preset functions or programs inherent to a computer is known as the operating system. Like humans, computers have a "genetic code." This code can be altered or reprogrammed by viruses, rendering the system useless and often requiring reinstallation of the operating system to regain its usefulness. Finally, those tasks that are "learned" after "birth" (i.e., installed after purchase) are known as software applications, and have made many of their "teachers" (i.e., software manufacturers) very wealthy. And, just as we need to use a language when teaching humans, computers are trained using programming languages.

The remainder of this chapter describes in detail these basic computer functions: input, output, memory, processing, the bus, operating systems, computer languages, and application programs.

## INPUT

Often a primary goal for computer applications is to provide communication of information between computers and their human owners. Therefore, many of the input and output functions of computers are much like those of humans. For humans, input comes from the senses (vision, hearing, touch, smell, and taste). For computers, input comes from input devices (keyboard, mouse, microphone, and digital camera, for example; see Figure 9.1). These devices provide a computer with a rudimentary form of 3 of the 5 human senses. (We still do not have input devices that provide a computer with smell or taste—probably because there has not been a great need for them.)




FIGURE 9.1

Modern personal computer with standard I/O devices.

## OUTPUT

Output for humans comes simply through the movement of muscles. This allows us to grasp and move other objects (such as a pencil or a musical instrument), which we can further use to communicate our output data (such as writing or playing music). We can also move certain muscles to generate sound (phonation) to further communicate our output (speech). So in summary, we humans write and draw, talk and sing, and move a variety of objects to generate output, or more *humanly* stated, express our thoughts. With computers, we have tried to emulate several of these human outputs so our computers can communicate with us. Drawing and writing come in the form of computer monitors and printers (both the film and paper types) (see Figure 9.1). Talking and music generation come through sound cards and speakers connected to the computer, providing computers with the ability to communicate with humans over telephones. Finally, the movement of objects through the use of computers (known as robotics) is well developed for specific applications (e.g., assembly-line automation) but is not generally available to the public yet.

## MEMORY

Human memory is poorly understood and quite variable. Its capacity is tremendous though its accuracy is anything but perfect (you can test this by trying to recite the last paragraph you have just read word for word). Computer memory, on the other hand, is well defined, highly accurate, and available in very large (albeit expensive) quantities. Computer memory serves a variety of purposes and comes in a variety of forms. Memory is needed to store input and output data, as well as data calculated during processing (see Figure 9.2), and to store application programs (that list of instructions that




FIGURE 9.2

Removable (floppy) drive and media.

tells the computer what to do). To serve this variety of purposes best, a wide variety of memory types are available. In general, they come in 3 forms: solid-state—RAM (random access memory), spinning media—HDD (hard disk drive), RAID (redundant array of inexpensive disks), and linear media—optical and magnetic tape storage. These vary in size, speed of access, volatility, and cost.

## PROCESSING

At the core of the computer is the central processing unit, or CPU (“processor” for short) (see Figure 9.3). It is the human brain analog. The CPU interprets a set of instructions (known as machine code) and performs whatever tasks it is told to do. A list of these tasks, or instructions, is known as a computer program. Programs tell the CPU exactly what to do and when to do it (more on programs later).

All modern CPUs are digital, meaning information within them is processed using digital or finite states of electronic voltage (v). Typically, 0 volts (0v) represents the number “zero” while 5 volts (5v) represents the number “one.” All remaining numbers can be represented by a series of 0s and 1s just as we use the series of numerals 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9 to represent all remaining numbers larger than nine. Humans developed using the decimal system (base 10) probably because we have 10 fingers. Computers use the binary system (base 2) because a single digital gate (a digital transistor) has 2 states (on and off, 5v and 0v, or 1 and 0) to process and store its information. Conversion from their binary system to our decimal system is typically performed by the computer before displaying these numbers for human consumption. As an example, binary 10110101 (or  $1 \times$




FIGURE 9.3

Various CPU boards.

$2^0 + 0 \times 2^1 + 1 \times 2^2 + 0 \times 2^3 + 1 \times 2^4 + 1 \times 2^5 + 0 \times 2^6 + 1 \times 2^7$ , or  $1 + 0 + 4 + 0 + 16 + 0 + 64$ ) equals decimal 85 (or  $5 \times 10^0 + 8 \times 10^1$ , or  $5 + 80$ ).

While humans typically place commas between 3 groups of our decimal digits (e.g., 1,500,325), computers group their binary digits (bits) into groups of 8 (bytes). CPUs today process numbers contained in chunks of 1 to 16 bytes (8 to 128 bits) at a time. Modern CPUs use millions of transistors to process these huge binary numbers and billions of transistors to store their results, all in chips about the size of your thumbnail. They can execute instructions at well over 1 GHz (more than a billion instructions per second). In 1985, computer processors were considered screamers if they ran at 1 MHz. Thus, modern CPUs are running a thousand times faster than they did just 20 years ago—and they keep getting faster.

A wide variety of CPUs are available today from a number of vendors such as Intel, AMD, Texas Instruments, Motorola, IBM, and Sony. As long as manufacturers continue to place more and more transistors on a single integrated circuit, CPUs will continue to decrease in price while increasing in performance, thereby making all our applications run faster.

## THE BUS

Inside your computer is a “motherboard” (see Figure 9.4). Basically, this board *is* the computer, to which is attached a power supply, a surrounding metal case, and all input/output devices (such as the keyboard, mouse, and monitor). On the motherboard sit the CPU, the RAM, and connections to the hard disk. These devices are electronically connected to one another on the motherboard via the computer bus. The bus has at its disposal a variety of possible connectors or slots for attaching external devices to your com-



**FIGURE 9.4**

PC motherboard and power supply.

puter as well. The number and type of connectors, which vary between computers, have a large impact on the computer's expandability. One connector, the AGP (accelerated graphics port), is typically used for connection to a video display card, which contains more memory and a special CPU, called a GPU, or graphical processing unit, which manages communication to the computer monitor for image display. Basically, data is transmitted from computer memory to the video display card, where bytes are converted into pixels, generating the digital image.

All this traffic must go back and forth quickly if an image is to make it to the screen in the blink of an eye. Each time the information moves, it takes a finite period of time. The wider the bus, the faster the digital traffic can flow across it. There are many kinds of bus architectures, but for the most part the combination of CPU and operating system functions on a predetermined or specified bus. In other words, once you buy the computer, you cannot change the bus. Because the bus can be a bottleneck to performance, a new bus is designed from time to time to provide greater performance. Recently, USB-2 (Universal Serial Bus 2) has become a popular bus for connecting all sorts of external devices to expand the performance and functionality of today's computers.

## OPERATING SYSTEMS

The operating system (OS) is the software at the heart of the computer. It is simply a program that performs a variety of basic functions for all other programs. (For example, the mouse pointer arrow is actually generated by the OS. Any program that needs a mouse pointer [such as a radiology workstation] simply uses the OS to generate it.) Examples of some operating systems that you might have heard of are CP/M, DOS, LINUX, OS-X, MUMPS, OS/2, SOLARIS, UNIX, VMS, and Windows XP. Without an OS, your computer is merely a lump of silicon and metal. The operating system is what determines what software you can and cannot run on the computer. If you have UNIX, forget about running Microsoft Word for Windows. If you have Windows XP, you cannot run Macintosh OS-based software.

*Historical note:* MUMPS was the first OS designed primarily for medical application programs. It was designed and developed in 1966 by G. Octo Barnett and Neil Pappalardo to facilitate the management of text-based medical information through the use of a hierarchal storage array. It is still used in a number of modern radiology departments even today.

An operating system is usually designed to run on a specific CPU or series of CPUs. That is, you can run DOS on an Intel 286, 386, 486, or Pentium, but it will not run on a Motorola-based Mac. As computers have become more powerful and compact, operating systems have been designed to take advantage of this power, and they have become graphical so that ordinary people can interact with the computer more intuitively. What facilitates their interactions are GUIs (pronounced goo-eez) or graphical user interfaces. Because these GUI-based OSs expose many high-level functions to the end user (e.g., scrolling, drop and drag, wastebaskets), application programs running on them have a common look and feel, regardless of who makes them. In 1984 Apple introduced the Macintosh. Its OS was probably the first widely accepted GUI. Shortly thereafter came the Microsoft GUI (aka Windows). Once the first GUI took hold, more GUIs followed, and programs started to take advantage of their advanced, yet repetitive, functions and features. In fact, you can even find major medical equipment, including magnetic resonance (MR) and computed tomography (CT) scanners and virtually every kind of picture archiving and communication systems (PACS) component, running applications on GUI OSs on PCs.

## COMPUTER LANGUAGES

As mentioned earlier, at the lowest level computers know of only 1s and 0s. If you really want to communicate with a computer in its “native tongue,” you would have to speak binary. Furthermore, the CPU of a computer knows a very few, rudimentary instructions—what is known as machine language or machine code. Very few programmers can “talk” in this native tongue, and it is a very inefficient form of human-computer communication (kind of like getting your point across to someone by tapping on a table). Enter the high-level computer language.

Even though the computer can understand only this machine code, if it had an interpreter it could understand other languages as well. Therefore, if computer scientists could think of a logical way to communicate commands to a computer, they would just have to write a program (known as an interpreter, a compiler, or an assembler) to convert that language to machine language and submit it to the computer. C is an example of just that. It is a high-level language that is much more intuitive for programmers to understand than machine code and therefore makes them much more efficient. A single line of C might convert to thousands of lines of machine language, which the CPU can process in the blink of an eye. Visual C# is a specific



version of the language from Microsoft that uses a GUI programming environment for RAD (rapid application development). With Visual C#, major works can be developed in days.

Other examples of popular high-level languages include C++, Visual Basic, Pascal, and Java. Java, developed by SUN Microsystems, is the newest of these languages, and its popularity is growing rapidly because it can run on a variety of CPUs and OSs without altering any lines of code.

## APPLICATION PROGRAMS

Application programs are those software packages that you purchase and install after you buy your computer. (Some are actually installed before you buy, but after the computer is manufactured, and are called “bundled” software.) Application programs are probably the reason you bought your com-

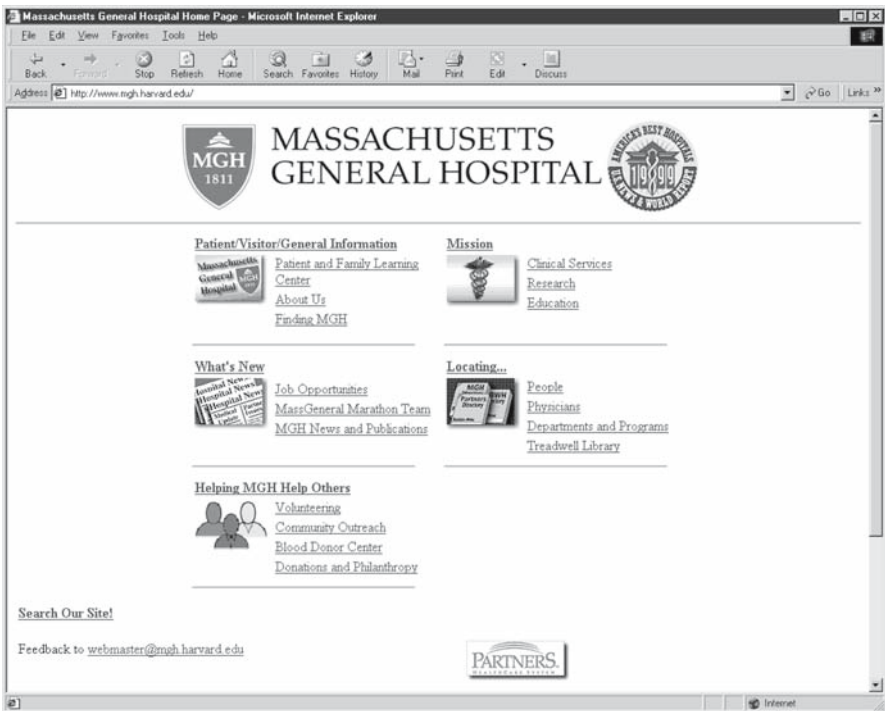


FIGURE 9.5

Sample Web site display using a Web browser.

puter in the first place. Examples include word processors, spreadsheets, and radiology primary interpretation applications.

Sometimes application programs become so widely used that they work their way into the OS. A good example of this is the browser application used to “surf the Web” (see Figure 9.5). Considerable debate has been under way recently as to whether the browser should reside in the OS or should be sold separately. This seemingly simple decision drastically changes the landscape for nearly every software vendor—thus the debate.

## CONCLUSION

We could not sustain medical imaging growth (or civilization, for that matter), as we currently know it, without computers. They continue to improve in performance and usability and decrease in cost and size. For them to become even more widely accepted, however, they will have to communicate better with their human users. This is happening already (see Chapter 25, “Voice Recognition”). It is shortsighted to exclude anything from the range of possibilities of what computers may be able to do in the future. Compared to the millions of years required for nature-guided human evolution, in the past few decades computers have progressed at unprecedented rates. One can only hope that before we make them too intelligent, they will have learned to appreciate their creators. For more information on computer fundamentals, please visit <http://www.MyRadiology.com>.

# DIGITAL IMAGING FUNDAMENTALS

---

---

KEITH J. DREYER • MANNUDEEP K. KALRA

**F**rom its inception over a century ago, radiology has provided a view into human anatomy and pathology via signal detection and image generation. For decades, the only imaging modality available to the science was projection radiography, in which all nonabsorbed signals were converted to an image by a sheet of radiographic film. There was no opportunity for digital image processing because there were no digital images.

Things have changed drastically for radiology in the last half century. New digital methods of capturing anatomy, pathology, and physiology have been invented (e.g., computed tomography [CT], magnetic resonance [MR], ultrasound, and positron emission tomography), and even projection radiography has become digitized (e.g., computed radiography and direct radiography). It is this infiltration of digital imaging into radiology that has allowed digital image processing to become a fundamental part of all modern radiology departments.

## WHY IMAGING?

One may ask, why generate images at all? Many of the native signals detected by radiological modalities today bear little resemblance to an image. Typically, extensive signal processing is necessary to generate even a rudimentary two-dimensional image, let alone the complex images that we are familiar with today. Keeping in mind that humans contain biological signal processors capable of interpreting all sorts of analog data, why not simply present the signal itself in its native form for interpretation? With more than half of the human brain dedicated to image processing, one can quickly see that a visual image is one of the best signals to use when representing data for sophisticated human interpretation.

Radiological modalities that convert their native signal information into visual images are far easier for humans to interpret. The more the resulting images resemble directly observed human anatomy, the easier it is for the observer to create a frame of reference for their subsequent interpretation. The ability to differentiate between normal and abnormal (pathologic states are detected by the identification of abnormal anatomy and/or physiology) is fundamental to the radiological diagnosis of disease. While much of the ability to differentiate these states depends on the observer's experience (knowledge of human pathophysiology and its imaging presentation states), creating an ideal image is paramount to making the correct diagnosis. It is, in large part, the science of digital image processing that will allow us to create an ideal image to enhance the interpretation process.

## DIGITAL IMAGES

Before we begin exploring digital image processing, we must first understand how a digital image is constructed. A digital image is a representation of a two-dimensional image as a finite set of digital values called picture elements or pixels (Figure 10.1). (*Note:* The term “digital image” also applies to data associated with points scattered over a three-dimensional [3-D] region, such as those produced by CT and magnetic resonance imaging (MRI). In that case, each signal sample is called a voxel, or volume element, instead of a pixel. These 3-D image types are described in detail in Chapter 22.) Pixels are organized into rows and columns. In fact, the spatial resolution of an image is determined by the number of rows and columns. Typical resolutions for medical images are  $512 \times 512$ ,  $1024 \times 1024$  and  $2048 \times 2048$ . Pixels of medical images are typically so small and so numerous that, when dis-




FIGURE 10.1


Digital image coordinate system (pixels).

played on a computer monitor, they appear to merge into a smooth continuous image.

## SHADES OF GRAY

Though each pixel can represent a color and intensity of light, for most of radiology the color is limited to gray while the intensity determines the pixel's shade of gray. This is because each pixel typically represents the signal acquired by the modality for that location in space. As the signal intensity increases, the pixel shade becomes more black (or more white, depending on the modality). It is the difference in these shades of gray (contrast resolution) that allows us to visually differentiate between different tissues of the body in an image. Therefore, the greater the signal range of values detectable, the more shades of gray can be displayed, resulting in better image quality. Radiology modalities today can generate pixels with up to 65,536 different shades of gray, or a 16-bit grayscale. However, most computer monitors (and the human visual system) can only differentiate among about 256 different shades of gray, or an 8-bit grayscale. The issue is resolved by the concept of down sampling the higher-resolution grayscale (from 16-bit to 8-bit) by a method called windowing and leveling.

The *window width* determines how many of the original shades will be displayed at once by dividing each pixel by a fixed value (*width*). The *window level* determines where to focus attention in the original grayscale by adding an offset number (*level*) to each pixel (Figure 10.2). These values are then



**FIGURE 10.2**

Visual representation of window level techniques.

normalized down to a resultant 8-bit pixel for subsequent computer monitor display.

## IMAGE PROCESSING

Image processing can be thought of as the digital manipulation of an image that results in a new (and, one hopes, improved) image (Table 10.1). A good example of this is the windowing and leveling algorithm described above. It provides the observer with the ability to enhance different aspects of the original image, producing new images that when collectively displayed provide more information to the observer than the original image did. Other algorithms used for image processing include histogram equalization, geo-

---



---

**TABLE 10.1**  
Digital Imaging Techniques

Technique	In	Out
Image Processing	Image	Image
Image Analysis	Image	Measurements
Image Understanding	Image	High-Level Description

---



---

metric transformations (zoom, rotate, pan), image fusion, noise reduction, edge enhancement, segmentation, frequency transformations, and image compression. Many of these rudimentary image-processing functions are available with picture archiving and communication systems (PACS) and their respective primary interpretation workstations.

## IMAGE ANALYSIS

Where image processing ends, image analysis typically begins. Image analysis is defined as the processing performed on an image that results in measurements or other low-level descriptors. Examples of image analysis are the algorithms used to measure bone density on CT, cardiac index on CT angiography, or tumor volumes on MR. The results are measurements that can be compared to an index and/or previous values to monitor disease progression. Most medical image analysis algorithms today require some manual intervention (semiautomated analysis) using the human to begin the process (seed planting) with the computer providing tireless consistency to the remaining, otherwise subjective, measurement tasks.

## IMAGE UNDERSTANDING

Beyond image analysis is the complex field of image understanding. In medical imaging, image understanding is often referred to as CAD, or computer-aided diagnosis. Most CAD systems of today attempt to find (through feature extraction) an aspect of the medical image data that suggests the presence of disease. Oftentimes these areas of concern are indicated to the radiologist by an overlay (i.e., arrow, circle) produced by the image-understanding algorithm. It remains the responsibility of the radiologist to determine the validity of these automatically identified areas of concern.

## CONCLUSION

The use of digital imaging in medicine has provided a wealth of value to radiology in the healthcare enterprise. It has allowed for the elimination of film and for the digital storage, retrieval, transfer, and display of images anywhere throughout the enterprise, or throughout the world. Though image-processing algorithms are performed daily in hospitals everywhere, image-analysis and image-understanding methods remain elusive. It

is these authors' belief that the next great digital revolution in medical imaging will occur when the secrets to these final challenges start to be revealed. For more information on digital-imaging fundamentals, please visit <http://www.MyRadiology.com>.



# IMAGE ACQUISITION

---

---

KATHERINE P. ANDRIOLE

**D**igital acquisition of data from the various imaging modalities for input to a picture archiving and communication system (PACS) is covered in detail in this chapter. Essential features for successful clinical implementation including conformance with the Digital Imaging and Communications in Medicine (DICOM) standard, radiology information system–hospital information system (RIS-HIS) interfacing, and workflow integration are discussed. Image acquisition from the inherently digital cross-sectional modalities such as computed tomography (CT) and magnetic resonance imaging (MRI) are reviewed, as well as digital acquisition of the conventional projection x-ray utilizing computed radiography (CR), digital radiography (DR), and film digitizers for digital acquisition of images originally recorded on film.

Quality assurance (QA) and quality control (QC) for a PACS are described with emphasis on QA-QC procedures and troubleshooting problems occurring specifically at image acquisition. Future trends in image acquisition for digital radiology and PACS will be introduced including anticipated changes in image datasets (such as increased matrix size, increased spatial resolution, increased slice number and study size, and

improved image quality); changes in the imaging devices themselves (such as smaller footprints and more portability); and image-processing capabilities for softcopy display.

## INTRODUCTION TO IMAGE ACQUISITION

### INTEGRATION WITH PICTURE ARCHIVING AND COMMUNICATION SYSTEMS

Image acquisition is the first point of data entry into a PACS, and as such, errors generated here can propagate throughout the system, adversely affecting clinical operations. General predictors for successful incorporation of image acquisition devices into a digital imaging department include ease of device integration into the established daily workflow routine of the clinical environment, high reliability and fault tolerance of the device, simplicity and intuitiveness of the user interface, and device speed.

### DIGITAL IMAGING AND COMMUNICATIONS IN MEDICINE

Imaging modality conformance with the Digital Imaging and Communications in Medicine (DICOM) standard is critical; only a basic summary is included here. DICOM consists of a standard image format as well as a network communications protocol. Compliance with this standard enables an open architecture for imaging systems, bridging hardware and software entities and allowing interoperability for the transfer of medical images and associated information between disparate systems.

The push by the radiological community for a standard format across imaging devices of different models and makes began in 1982. Collaboration between the American College of Radiology (ACR) and the National Electrical Manufacturers Association (NEMA) produced a standard format (ACR-NEMA 2.0) with which to store an image digitally. It consisted of a file header followed by the image data. The file header contained information relevant to the image, such as matrix size or number of rows and columns, pixel size, and grayscale bit depth, as well as information about the imaging device and technique, (e.g., Brand X CT scanner, acquired with contrast). Patient demographic data such as name, date of birth, and so on, were also included in the image header. The ACR-NEMA 2.0 standard specified exactly where in the header each bit of information was to be stored, such that the standard required image information could be read by any device,

simply by going to the designated location in the header. This standard unified the format of imaging data but functioned only as a point-to-point procedure.

In 1994, at the Radiological Society of North America (RSNA) Meeting, a variety of imaging vendors participated in an impressive demonstration of the new and evolving imaging standard (ACR-NEMA 3.0) or what is currently known as the DICOM standard. Participants attached their devices to a common network and transmitted their images to one another. In addition to the standard image format of ACR-NEMA 2.0, the DICOM standard included a network communications protocol, or a common language for sending and receiving images and relevant data over a network.

The DICOM standard language structure is built on information objects (IO), application entities (AE), and service class users (SCU) and providers (SCP). Information objects include, for example, the image types, such as CT, MRI, CR, and the like. The application entities include the devices, such as a scanner, workstation, or printer. The service classes (SCU, SCP) define an operation on the information object via service object pairs (SOP) of IO, SCU, and SCP. The types of operations performed by an SCU-SCP on an IO include storage; query-retrieve; verification; print; study content notification; and patient, study, and results management.

The DICOM standard is used, for example, to negotiate a transaction between a compliant imaging modality and a compliant PACS workstation. The scanner notifies the workstation, in a language both understand, that it has an image study to send to it. The workstation replies to the modality when it is ready to receive the data. The data is sent in a format known to all, the workstation acknowledges receipt of the image, and then the devices end their negotiation. Figure 11.1 shows the results of an example PACS tool for reading the DICOM header. Shown are elements in Groups 8 and 10 pertaining to image identification parameters (such as study, series, image number) and patient demographics (such as patient name, medical record number, date of birth), respectively.

Prior to DICOM, the acquisition of digital image data and relevant information was extremely difficult, often requiring separate hardware devices and software programs for different vendors' products, and even for different models of devices made by the same manufacturer. Most of the major manufacturers of imaging devices currently comply with the DICOM standard, thus greatly facilitating an open systems architecture consisting of multivendor devices. For many legacy devices purchased prior to the establishment of DICOM, an upgrade path to compliance can be performed. For those few devices that do not yet meet the standard, interface boxes consisting of hardware equipment and software programs that convert the image

Group	Element	Prefix	Suffix	Descr
0010	0010			
0010	1010	0010 0040	0010 0020	
0008	0020			
0008	0030			
0008	0031			
0008	0033			
0008	0050		ACCES#	
0008	0070			
0008	0080			
0008	0090			
0008	1010		SYS#	
0008	1030			
0008	103e			
0010	0010			pati
0010	0020			pati
0010	0030			pati
0010	0040			pati
0010	1010			pati
0010	1030			pati

FIGURE 11.1

The output of an example of a PACS tool for reading the DICOM header. Shown are elements in Groups 8 and 10, pertaining to image identification parameters (such as study, series, image number) and patient demographics (such as patient name, medical record number, date of birth), respectively.

data from the manufacturer's proprietary format to the standard form are available.

#### RADIOLOGY INFORMATION SYSTEM—HOSPITAL INFORMATION SYSTEM INTERFACING FOR DATA VERIFICATION

Equally essential, particularly at acquisition, is integrating the RIS and/or HIS with the PACS. This greatly facilitates input of patient demographics (name, date, time, medical record number [MRN] to uniquely identify a patient, accession number [AccNum] to uniquely identify an imaging examination, exam type, imaging parameters, etc.), and enables automatic PACS data verification, correlation, and error correction with the data recorded in the RIS-HIS. Most imaging modalities are now tightly coupled with the RIS, providing automatic downloading of demographic information from the RIS via barcode readers or directly to the scanner console (via modality worklist

capability) and hence to the DICOM header. This eliminates the highly error-prone manual entry of data at acquisition.

Health Level Seven (HL7) is the RIS-HIS standard, and compliance with it is desirable. RIS-HIS databases are typically patient-centric, enabling query and retrieval of information by the patient, study, series, or image data hierarchy. Integration of RIS-HIS data with the PACS adds intelligence to the system, helping to move data around the system based on “*how* and *what* data should be delivered *where* and *when*,” automating the functions performed traditionally by the film librarian.

## MODALITY WORKLIST

Many vendors now provide the capability to download RIS-HIS schedules and worklists directly to the imaging modality, such as most CT, MRI, digital fluoroscopy (DF), and ultrasound (US) scanners. In these circumstances, the imaging technologist need only choose the appropriate patient’s name from a list on the scanner console monitor (by pointing to it on a touch-screen pad), and the information contained within the RIS-HIS database will be downloaded into the PACS header and associated with the image data for that patient examination.

In the general DICOM model for acquisition of image and relevant data from the imaging modality, the modality device acts as an SCU, and the data is stored to an SCP device such as a PACS acquisition gateway or an image display workstation. In the modality worklist function, however, the image device receives the pertinent patient demographics and image study information from a worklist server, such as a PACS, RIS, or RIS-HIS interfaced device.

There are two modes for accomplishing the RIS-HIS data transfer to the imaging modality. In the first, data is transferred automatically to the modality based on the occurrence of an event trigger, such as a scheduled examination or a patient arrival. The second method involves a query from the modality to the RIS-HIS. This may be initiated by entry of some identifier at the modality, such as bar coding of the study AccNum or the patient MRN from the scheduling card. This initiates a request for the associated RIS-HIS information (patient name, date of birth) to be sent from the worklist server on demand.

The benefits of the DICOM modality worklist cannot be overstated. Incorrectly (manually) entered patient demographic data, such as all the permutations of patient name (e.g., James Jones, J Jones, Jones J) can result in mislabeled image files and incomplete study information; correct demo-

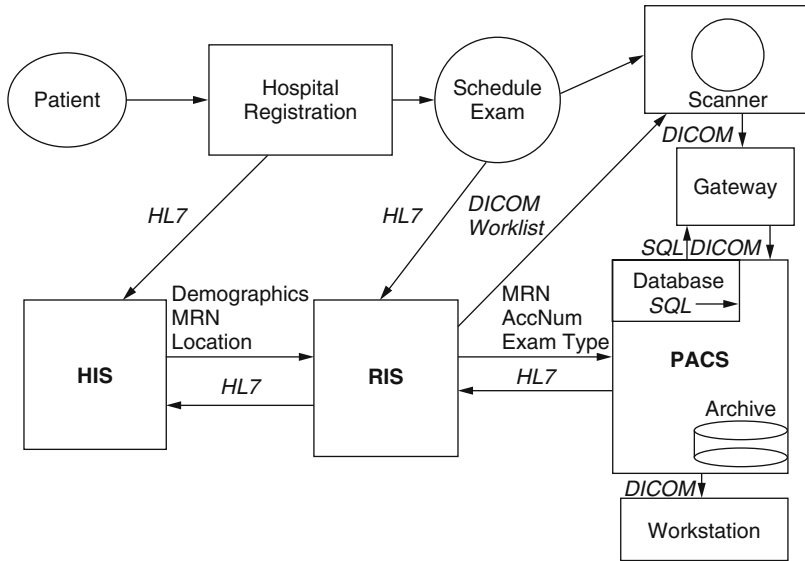


FIGURE 11.2

Diagram of how RIS, HIS, and PACS systems might interact upon scheduling an examination for image acquisition into a PACS.

graphic data are crucial to maintaining the integrity of the PACS database. Furthermore, the improvements in departmental workflow efficiency and device usability are greatly facilitated by modality worklist capabilities. For those few vendors not offering a DICOM modality worklist for their imaging devices, several interface or broker boxes are available that interconnect PACS to RIS-HIS databases translating DICOM to HL7 and vice versa.

Figure 11.2 diagrams an example of how RIS, HIS, and PACS systems might interact upon scheduling an examination for image acquisition into a PACS.

## ACQUISITION OF THE NATIVE DIGITAL CROSS-SECTIONAL MODALITIES

Image acquisition from the inherently digital modalities such as CT, MRI, and US should be a direct digital DICOM capture. Direct digital interfaces allow capture and transmission of image data from the modality at the full spatial resolution and full bit depth of grayscale inherent to the modality,

while analog (video) frame grabbers digitize the video signal voltage output going to an image display, such as a scanner console monitor. In the frame-grabbing method, as in printing an image to film, the image quality is limited by the process to just 8 bits (or 256 gray values) while most modalities have the capability to acquire in 12, 16, or even 32 bits for color data. Capture of only 8 bits may not allow viewing in all the appropriate clinical windows and levels or contrast and brightness settings and is therefore not optimal.

For example, when viewing a CT of the chest, one may wish to view in lung window and level settings and in mediastinal and bone windows and levels. Direct capture of the digital data will allow the viewer to dynamically window and level through each of these settings on the fly (in real time) at the softcopy display station. Whereas, to view all appropriate window and level settings on film, several copies of the study would have to be printed, one at each window and level setting. If one performs the analog acquisition or frame grabbing of the digital data, the viewer can only window and level through the 8 bits captured, which may not be sufficient. Thus, direct capture of digital data from the inherently digital modalities is the preferred method of acquisition. Table 11.1 lists the cross-sectional modalities commonly interfaced to PACS along with their inherent file sizes and bit depths.

---



---

**TABLE 11.1**

The Commonly PACS-Interfaced Cross-Sectional Modalities and Their Inherent File Sizes

<b>Modality</b>	<b>Image Matrix Size</b>	<b>Grayscale Bit Depth</b>
Computed tomography (CT)	512 × 512 pixels	12–16 bits
Digital angiography (RA) and digital fluoroscopy (DF)	512 × 512 pixels or 1024 × 1024 pixels or 2048 × 2048 pixels	8–12 bits
Magnetic resonance imaging (MRI)	256 × 256 pixels	12–16 bits
Nuclear medicine images (NUC)	64 × 64 pixels or 128 × 128 pixels or 256 × 256 pixels	8–32 bits
Ultrasound (US)	64 × 64 pixels or 128 × 128 pixels	16–32 bits

---



---

## ACQUISITION OF PROJECTION RADIOGRAPHY

Methods for digital image acquisition of the conventional projection x-ray include via CR scanners (imaging with photostimulable or storage phosphors), digitization of existing analog film, and DR devices. Digital acquisition of images already on film can be accomplished using a variety of image digitization devices or film scanners. These include the infrequently used analog video cameras with analog-to-digital converters (ADC), digital cameras, charge coupled devices (CCD), and laser scanners.

### FILM DIGITIZERS

Film digitizers will still be necessary even in the all-digital or filmless imaging department, so that film images from outside referrals lacking digital capabilities can be acquired into the system and viewed digitally. Film digitizers convert the continuous optical density values on film into a digital image by sampling at discrete evenly spaced locations and quantizing the transmitted light from a scan of the film into digital numbers. Several types of film digitizers exist today, with some used more frequently than others in PACS and teleradiology applications.

The analog video camera with ADC, or camera on a stick, has been used in low-cost, entry-level teleradiology applications but is infrequently used in PACS applications today because of its manual operation. The analog video camera requires an illumination source and careful attention to lens settings, focus, f-stop, and so forth. In addition, it has a maximum resolution of  $1024 \times 1024 \times 8$  bits (256 grays), thus limiting the range of window and level, or contrast and brightness values the resulting digital image can be displayed in. Digital cameras produce a digital signal output directly from the camera at a maximum resolution of  $2048 \times 2048 \times 12$  bits (4096 grays) but are still infrequently used in PACS due to their high cost.

More commonly used are film scanners such as the CCD and laser scanners, sometimes called flatbed scanners. CCD scanners utilize a row of photocells and uniform bright light illumination to capture the image. A lens focuses the transmitted light from the collimated, diffuse light source onto a linear CCD detector, and the signal is collected and converted to a digital electronic signal via an ADC converter. CCD scanners have a maximum resolution of  $4096 \times 4096 \times 8$  to 12 bits, but have a narrow film optical density range to which they can respond. CCD scanners have been used in high-end teleradiology or entry-level in-house film distribution systems, such as image transmission to the intensive care units (ICUs).



The laser scanner or laser film digitizer uses either a helium-neon (HeNe) gas laser or a solid-state diode laser source. The laser beam is focused by lenses and directed by mirror deflection components, and the light transmitted through the film is collected by a light guide, its intensity detected by a photomultiplier tube, converted to a proportional electronic signal, and digitized in an ADC. Laser scanners use a fine laser beam of generally variable or adjustable spot sizes down to 50 microns (producing an image sharpness of approximately 10 line pairs per millimeter [lp/mm]). They have a maximum spatial resolution of  $4096 \times 5120$  and a grayscale resolution of 12 bits and can accommodate the full optical density range of film. They are semi- or fully automatic in operation and are currently the scanner of choice for PACS applications even though they are often more expensive than CCD scanners.

## COMPUTED RADIOGRAPHY

Computed radiography refers to projection x-ray imaging using photostimulable or storage phosphors as the detector. In this modality, x-rays incident upon a photostimulable phosphor (PSP)-based image sensor or imaging plate (IP) produce a latent image that is stored in the IP until stimulated to luminesce by laser light. This released light energy can be captured and converted to a digital electronic signal for transmission of images to display and archival devices. Unlike conventional screen-film radiography in which the film functions as the imaging sensor or recording medium, as well as the display device and storage media, CR eliminates film from the image-recording step, resulting in a separation of image capture from image display and image storage. This separation of functions potentiates optimization of each of these steps individually. In addition, CR can capitalize on features common to all digital images, namely, electronic transmission, manipulation, display, and storage of radiographs.

Technological advances in CR over time have made this modality widely accepted in digital departments. Hardware and software improvements have been made in the PSP plate, in image reading-scanning devices, and in image-processing algorithms. Overall reduced cost of CR devices as well as a reduction in cost and increased utility of image display devices has contributed to the increased acceptance of CR as a viable digital counterpart to conventional screen-film projection radiography. This section provides an overview of the state-of-the-art in CR systems, including a basic description of the data acquisition process, a review of system specifications, image quality and performance, and advantages and disadvantages inherent in CR.

An explanation of the image-processing algorithms that convert the raw CR image data into useful clinical images will be provided. The image-processing algorithms to be discussed include image segmentation or exposure data recognition and background removal, contrast enhancement, spatial frequency processing including edge enhancement and noise smoothing, dynamic range control (DRC), and multiscale image contrast amplification (MUSICA). Note that the same types of image-processing algorithms utilized for CR may be applied to DR images as well. Examples of several types of artifacts potentially encountered with CR are given along with their causes and methods for correction or minimization of these effects.

## PROCESS DESCRIPTION

A CR system consists of a screen or plate of a stimulable phosphor material that is usually contained in a cassette and is exposed in a manner similar to the traditional screen-film cassette. The PSP in the IP absorbs x-rays that have passed through the patient, “recording” the x-ray image. Like the conventional intensifying screen, CR plates produce light in response to x-rays at the time of exposure. However, storage phosphor plates have the additional property of being capable of storing some of the absorbed x-ray energy as a latent image. Plates are typically made of a europium-doped barium-fluoro-halide-halide crystallized matrix. Electrons from the dopant ion become trapped just below the conduction band when exposed to x-rays. Irradiating the IP at some time after the x-ray exposure with red or near-infrared laser light liberates the electrons into the conduction band, stimulating the phosphor to release some of its stored energy in the form of green, blue, or ultraviolet light, the phenomenon of photostimulable luminescence. The intensity of light emitted is proportional to the amount of x-ray energy absorbed by the storage phosphor.

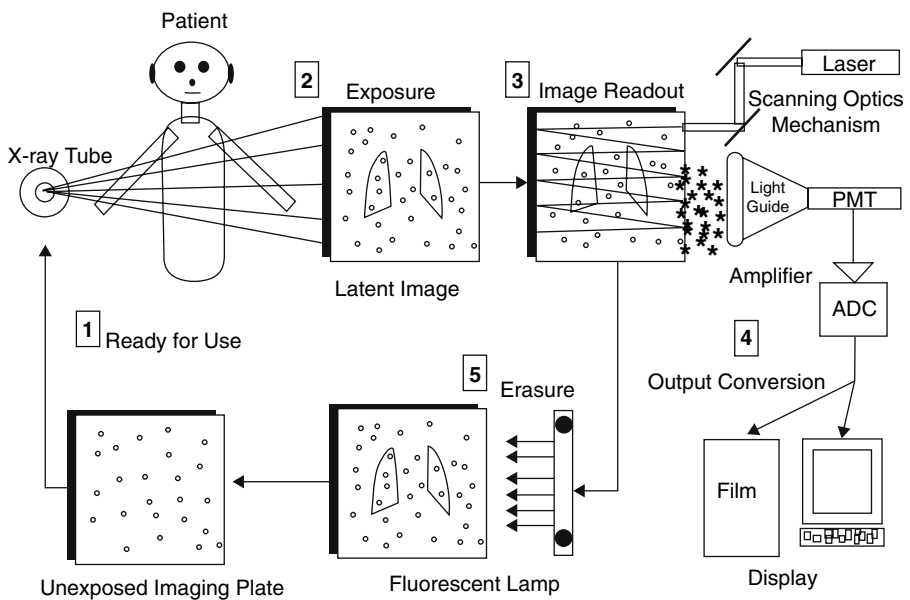
The readout process uses a precision laser spot scanning mechanism in which the laser beam traverses the IP surface in a raster pattern. The stimulated light emitted from the IP is collected and converted into an electrical signal, with optics coupled to a photomultiplier tube (PMT). The PMT converts the collected light from the IP into an electrical signal, which is then amplified, sampled to produce discrete pixels of the digital image, and sent through an ADC to quantize the value of each pixel (i.e., a value between 0 and 1023 for a 10-bit ADC or between 0 and 4095 for a 12-bit ADC).

Not all of the stored energy in the IP is released during the readout process. Thus, to prepare the IP for a new exposure, the IP is briefly flooded with high-intensity (typically fluorescent) light. This erasure step ensures removal of any residual latent image.

A diagram of the process steps involved in a CR system is shown in Figure 11.3. In principle, CR inserts a digital computer between the IP receptor (PSP screen) and the output image. This digital processor can perform a number of image-processing tasks including compensating for exposure errors, applying appropriate contrast characteristics, enhancing image detail, and storing and distributing image information in digital form.

### SYSTEM CHARACTERISTICS

One of the most important differences between CR and screen-film systems is in exposure latitude. The response of a digital imaging system relates the incident x-ray exposure to the resulting pixel value output. System sensitivity is the lowest exposure that will produce a useful pixel value, and the



**FIGURE 11.3**

The image production steps involved in CR. The imaging plate is exposed to x-rays, read out by a laser scanning mechanism, and erased for reuse. A light guide collects the photostimulated luminescence and feeds it to a photomultiplier tube (PMT), which converts the light signal to an electrical signal. Amplification, logarithmic conversion, and analog-to-digital conversion produce the final digital signal that can be displayed on a cathode-ray tube monitor or sent to a laser printer for image reproduction on film.

dynamic range is the ratio of the exposures of the highest and lowest useful pixel values. Storage phosphor systems have extremely wide exposure latitude. The wide latitude of storage phosphor systems and the effectively linear detector characteristic curve allow a wider range of exposure information to be captured in a single image than is possible with any screen-film system. In addition, the wide dynamic range of CR allows it to be used under a broad range of exposure conditions without the need for changing the basic detector. This also makes CR an ideal choice for applications in which exposures are highly variable or difficult to control as in portable or bedside radiography. Through image processing, CR systems can usually create a diagnostic image out of under- or overexposures via appropriate lookup table correction. In the screen-film environment, such under- or overexposures might have necessitated retakes and additional exposure to the patient.

Dose requirements of a medical imaging system depend on the system's ability to detect and convert the incoming signal into a usable output signal. It is important to stress that CR systems are *not* inherently lower-dose systems than screen-film. In fact, several studies have demonstrated a higher required exposure for CR to achieve equivalent optical density on screen-film. However, the wider latitude of storage phosphor systems makes them much more forgiving of under- or overexposure. As in any DR system, when dose is decreased, the noise due to quantum mottle increases. Reader tolerance of this noise tends to be the limiting factor on the lowest acceptable dose.

In some clinical situations, the radiologist may feel comfortable in lowering the exposure technique factor to reduce dose to the patient—such as in pediatric extremity x-ray exams. In other situations, such as imaging the chest of the newborn, one may wish to increase exposure to reduce the more visible mottle (at lower doses) to avoid mistaking the noise over the lungs as indication of pulmonary interstitial emphysema, for example. Computed radiography systems are signal-to-noise-limited (SNR-limited) whereas screen-film systems are contrast-limited.

## IMAGE QUALITY

**DETECTIVE QUANTUM EFFICIENCY** Objective descriptors of digital image quality include detective quantum efficiency (DQE), which is a measure of the fidelity with which a resultant digital image represents the transmitted x-ray fluence pattern (i.e., how efficiently a system converts the x-ray input signal into a useful output image) and includes a measure of the noise added. Also taken into account are the input-output characteris-

tics of the system and the resolution response of unsharpness or blur added during the image capture process. The linear, wide-latitude input-output characteristic of CR systems relative to screen-film leads to wider DQE latitude for CR, which implies that CR has the ability to convert incoming x-ray quanta into “useful” output over a much wider range of exposures than can be accommodated with screen-film systems.

**SPATIAL RESOLUTION** The spatial resolution response or sharpness of an image capture process can be expressed in terms of its modulation transfer function (MTF), which in practice is determined by taking the Fourier Transform of the line spread function (LSF), and relates *input* subject contrast to *imaged* subject contrast as a function of spatial frequency. The ideal image receptor adds no blur or broadening to the input LSF, resulting in an MTF response of 1 at all spatial frequencies. A real image receptor adds blur, typically resulting in a loss of MTF at higher spatial frequencies.

The main factor limiting the spatial resolution in CR, similar to screen-film systems, is x-ray scattering within the phosphor layer. However, it is the scattering of the stimulating beam in CR, rather than the emitted light as in screen-film, that determines system sharpness. Broadening of the laser light spot within the IP phosphor layer spreads with the depth of the plate. Thus, the spatial resolution response of CR is largely dependent on the initial laser beam diameter and on the thickness of the IP detector. The reproducible spatial frequency of CR is also limited by the sampling utilized in the digital readout process. The spatial resolution of CR is less than that of screen-film, with CR ranging from 2.5 to 5 lp/mm using a 200 $\mu$ m laser spot size and a digital matrix size of approximately 2000  $\times$  2500 pixels versus the 5 to 10 lp/mm or higher spatial resolution of screen-film.

Finer spatial resolution can technically be achieved today with the ability to tune laser spot sizes down to 50 $\mu$ m or less. But the image must be sampled more finely (approximately 4000  $\times$  5000 pixels) to achieve 10 lp/mm. Thus there is a tradeoff between the spatial resolution that can technically be achieved and the file size to practically transmit and store. Most general CR examinations are acquired using a 200 $\mu$ m laser spot size and a sampling of 2000  $\times$  2500 pixels. For examinations requiring very fine detail resolution such as in mammography, images are acquired with a 50 $\mu$ m laser spot size and sampled at 4000  $\times$  5000 pixels.

**CONTRAST RESOLUTION** The contrast or grayscale resolution for CR is much greater than that for screen-film. Note that since overall image quality resolution is a combination of spatial and grayscale resolution, the superior contrast resolution of CR can often compensate for its lack of inher-

ent spatial resolution. By manipulating the image contrast and brightness, or window and level values, respectively, small features often become more readily apparent in the image. This is analogous to “bright-lighting” or “hot-lighting” a bone film, for example, when looking for a small fracture. The overall impression is that the spatial resolution of the image has been improved, when in fact, it has not changed—only the contrast resolution has been manipulated. More work needs to be done to determine the most appropriate window and level settings for initial display of a CR image. Lacking the optimum default settings, it is often useful to “dynamically” view CR softcopy images with a variety of window and level settings.

**NOISE** The types of noise affecting CR images include x-ray dose-dependent noise and fixed noise (independent of x-ray dose). The dose-dependent noise components can be classified into x-ray quantum noise, or mottle, and light photon noise. The quantum mottle inherent in the input x-ray beam is the limiting noise factor, and it arises in the process of absorption by the IP, with noise being inversely proportional to the detector x-ray dose absorption. Light photon noise arises in the process of photoelectric transmission of the photostimulable luminescence light at the surface of the PMT.

Fixed noise sources in CR systems include IP structural noise (the predominant factor), noise in the electronics chain, laser power fluctuations, quantization noise in the analog-to-digital conversion process, and so on. Imaging plate structural noise arises from the nonuniformity of phosphor particle distribution, with finer particles providing noise improvement. Note that for CR systems, it is the *noise* sources that limit the DQE system latitude, whereas in conventional x-ray systems, the DQE latitude is limited by the narrower *exposure response* of screen-film.

**COMPARISON WITH SCREEN-FILM** The extremely large latitude of CR systems makes CR more forgiving in difficult imaging situations such as portable examinations and enables decreased retake rates for improper exposure technique, as compared to screen-film. The superior contrast resolution of CR can compensate in many cases for its lesser spatial resolution. Cost savings and improved radiology department workflow can be realized with CR and the elimination of film for projection radiographs.

## AVAILABLE COMPUTED RADIOGRAPHY SYSTEMS

**HISTORICAL PERSPECTIVE** Most of the progress in storage phosphor imaging has been made post–World War II. In 1975, Eastman Kodak

Company (Rochester, NY) patented an apparatus using infrared-stimulable phosphors or thermoluminescent materials to store an image. In 1980, Fuji Photo Film (Tokyo, Japan) patented a process in which PSPs were used to record and reproduce an image by absorbing radiation and then releasing the stored energy as light when stimulated by a helium-neon laser. The emitted phosphor luminescence was detected by a PMT, and the electronic signal produced reconstructed the image.

Fuji was the first to commercialize a storage phosphor-based CR system in 1983 (as the FCR 101) and published the first technical paper (in *Radiology*) describing CR for acquiring clinical digital x-ray images. The central-processing-type second-generation scanners (FCR 201) were marketed in 1985. Third-generation Fuji systems marketed in 1989 included distributed processing (FCR 7000) and stand-alone (AC-1) types. Fuji systems in the FCR 9000 series are improved, higher-speed, higher-performance, third-generation scanners. Current Fuji systems include upright chest units, CR detectors in wall and table buckeyes, multiplate autoloaders, and more compact stand-alone units.

In 1992, Kodak installed its first commercial storage phosphor reader (Model 3110). Later models include autoloader devices. In 1994, Agfa-Gevaert NV (Belgium) debuted its own CR system design (the ADC 70). In 1997, Agfa showed its ADC Compact with greatly reduced footprint. Agfa also introduced a low-cost, entry-level single plate reader (the ADC Solo) in 1998, appropriate for distributed CR environments such as clinics, trauma centers, ICUs, and the like. In 1998, Lumisys presented its low-cost, desktop CR unit (the ACR 2000) with manual-feed, single-plate reading. Numerous desktop units have been introduced including the Orex CR. Konica Corp debuted its own device (XPress) in 2002 and later (the Regius) upright unit, both of which have relatively fast scan times (at 40 and 16s cycle times, respectively). Many companies have been involved in CR research and development, including NA Philips Corp, EI DuPont de Nemours & Co, 3M Co, Hitachi, Ltd, Siemens AG, Toshiba Corp, General Electric Corp, Kasei Optonix, Ltd, Mitsubishi Chemical Industries, Ltd, Nichia Corp, GTE Products Co, and DigiRad Corp.

## TECHNOLOGICAL ADVANCES

Major improvements in the overall CR system design and performance characteristics include a reduction in the physical size of the reading/scanning units, increased plate-reading capacity per unit time, and better image quality. These advances have been achieved through a combination of

changes in the IPs themselves, in the image reader or scanning devices, and in the application of image-processing algorithms to affect image output.

The newer IPs developed for the latest CR devices have higher image quality (increased sharpness) and improved fading and residual image characteristics. Higher image quality has resulted from several modifications in the IP phosphor and layer thickness. Smaller phosphor grain size in the IP (down to approximately  $4\mu\text{m}$ ) diminishes fixed noise of the IP, while increased packing density of phosphor particles counteracts a concomitant decrease in photostimulable luminescence. A thinner protective layer is utilized in the plates, tending to reduce x-ray quantum noise. In and of itself, this would improve the spatial resolution response characteristics of the plates as a result of diminished beam scattering. However, in the newest IPs, the quantity of phosphor coated onto the plate is increased for durability purposes, resulting in the same response characteristic of previous IPs.

A historical review of CR scanning units chronicles improved compactness and increased processing speed. The first Fuji unit (FCR 101) from 1983 required roughly  $6\text{m}^2$  of floor space to house the reader and could only process about 45 plates per hour, while today's Fuji models as well as other vendors' devices occupy less than  $1\text{m}^2$  and can process more than 110 plates per hour. This is a decrease in apparatus size by a factor of approximately one sixth and an increase in processing capacity of roughly 2.5 times. Desktop models reduce the physical device footprint even further.

Computed radiography IP sizes, pixel resolutions, and their associated digital file sizes are roughly the same across manufacturers for the various cassette sizes offered. For example, the  $14\text{in} \times 17\text{in}$  ( $35\text{cm} \times 43\text{cm}$ ) plates are read with a sampling rate of 5 to 5.81 pixels per mm, at a digital image matrix size of roughly  $2000 \times 2000$  pixels ( $1760 \times 2140$  pixels for Fuji and  $2048 \times 2508$  pixels for Agfa and Kodak). Images are typically quantized to 12 bits (for 4096 gray levels). Thus, total image file sizes range from roughly 8MB to 11.5 MB. The smaller plates are scanned at the same laser spot size ( $100\mu\text{m}$ ), and the digitization rate does not change; therefore, the pixel size is smaller. The  $10\text{in} \times 12\text{in}$  ( $24\text{cm} \times 30\text{cm}$ ) plates are typically read at a sampling rate of 6.7 to 9 pixels per mm, and the  $8\text{in} \times 10\text{in}$  ( $18\text{cm} \times 24\text{cm}$ ) plates are read at 10 pixels per mm.

Cassetteless CR devices have been introduced in which the detector is incorporated into a chest unit, wall, or table buckey to speed throughput and facilitate workflow much as DR devices do. Dual-sided signal collection capability is available from Fuji, increasing overall signal-to-noise. Agfa has shown a product in development (ScanHead CR) that stimulates and reads out the IP line by line as opposed to the point-by-point scanning that occurs



in most CR devices today. Increased speed (5 s scan time) and higher DQE have been demonstrated. In addition, needle-phosphors have been explored as a possible replacement to powder-phosphors, having shown improved spatial resolution and DQE.

## IMAGE-PROCESSING ALGORITHMS

Image processing is performed to optimize the radiograph for output display. Each manufacturer has a set of proprietary algorithms that can be applied to the image for printing on laser film or for display, initially only on the manufacturer's own proprietary workstations. Prior to the DICOM standard, only the raw data could be directly acquired digitally. Therefore, to attain the same image appearance on other display stations, the appropriate image-processing algorithms (if known) had to be implemented somewhere along the chain from acquisition to display. Now image-processing parameters can be passed in the DICOM header, and algorithms can be applied to CR images displayed on generic workstations, though advanced real-time manipulation of images can typically only be done on each manufacturer's specific processing station. In general, the digital image processing applied to CR consists of a recognition or analysis phase, followed by contrast enhancement and/or frequency processing. Note that the same general types of image processing applied to CR can also be applied to DR images.

**IMAGE SEGMENTATION** In the image recognition stage, the region of exposure is detected (i.e., the collimation edges are detected), a histogram analysis of the pixel gray values in the image is performed to assess the actual exposure to the plate, and the appropriate lookup table specific to the region of anatomy imaged and chosen by the x-ray technologist at the time of patient demographic information input is selected. Proper recognition of the exposed region of interest is extremely important as it affects future processing applied to the image data. For example, if the bright white area of the image caused by collimation at the time of exposure is not detected properly, its very high gray values will be taken into account during histogram analysis, increasing the "window" of values to be accommodated by a given display device (softcopy or hardcopy). The effect would be to decrease the overall contrast in the image.

Some segmentation algorithms, in addition to detection of collimation edges in the image, enable users to blacken the region outside these edges, in the final image if so desired. This tends to improve image contrast appearance by removing this bright white background in images of small body parts

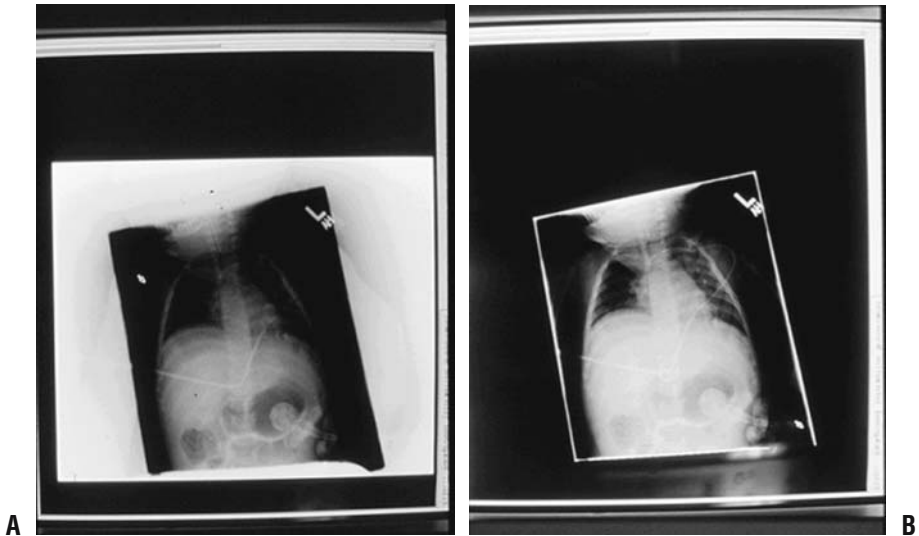



FIGURE 11.4

Example of image segmentation algorithm detection of (white) colimation edges of exposure region in (A), with “blackened surround” applied in (B). Note the improved overall contrast in (B).

or pediatric patients. Figure 11.4B demonstrates this feature of “blackened surround,” as applied to the image in Figure 11.4A.


**CONTRAST ENHANCEMENT** Conventional contrast enhancement, also called gradation processing, tone scaling, and latitude reduction, is performed next. This processing amounts to choosing the best characteristic curve (usually a nonlinear transformation of x-ray exposure to image density) to apply to the image data. These algorithms are quite flexible and can be tuned to satisfy a particular user’s preferences for a given “look” of the image. Lookup tables are specific to the region of anatomy imaged. Figure 11.5 shows an example of the default adult chest lookup table (A) applied to an image and the (B) same image with high-contrast processing. A reverse contrast scale or “black bone” technique, in what was originally black in the image becomes white, and what was originally white in the image becomes black, is sometimes felt to be beneficial for identifying and locating tubes and lines. An example is shown in Figure 11.6 where the contrast

**FIGURE 11.5**

Chest image processed with (A) default mode and (B) high-contrast algorithm applied.

reversal algorithm has been applied to the image in Figure 11.6A, resulting in the image in Figure 11.6B.

**SPATIAL FREQUENCY PROCESSING** The next type of image processing usually performed is spatial frequency processing, sometimes called edge enhancement. These algorithms adjust the frequency response characteristics of the CR systems, essentially implementing a high or band pass filter operation to enhance the high spatial frequency content contained in edge

**FIGURE 11.6**

Chest image processed with (A) default mode and (B) blackbone or contrast reversal algorithm applied.

information. Unfortunately, noise also contains high spatial frequency information and can be exacerbated by edge enhancement techniques. To lessen this problem, a nonlinear unsharp masking technique is typically implemented, serving to suppress noise via a smoothing process. Unsharp masking is an averaging technique that, via summation, tends to blur the image. When this is subtracted from the original image data, the effect is one of noise suppression. Specific spatial frequencies can be preferentially selected and emphasized by changing the mask size and weighting parameters. For example, low spatial frequency information in the image can be augmented by using a relatively large mask, while high spatial frequency or edge information can be enhanced by using a small mask size.

**DYNAMIC RANGE CONTROL** An advanced algorithm by Fuji, for selective compression or emphasis of low-density regions in an image, independent of contrast and spatial frequency, is known as dynamic range control (DRC) processing. The algorithm consists of performing an unsharp mask for suppression of high spatial frequency information, then applying a specific lookup table, mapping to selected regions (i.e., low-density areas). This mask is then added back to the original data with the overall result being improved contrast in poorly penetrated regions, without loss of high frequency and contrast emphasis. In a clinical evaluation of the algorithm for processing of adult portable chest exams, DRC was found to be preferred by five thoracic radiologists in a side-by-side comparison, providing improved visibility of mediastinal details and enhanced subdiaphragmatic regions.

**MULTISCALE IMAGE CONTRAST AMPLIFICATION** Multiscale image contrast amplification (MUSICA) is a very flexible advanced image-processing algorithm developed by Agfa. MUSICA is a local contrast enhancement technique based on the principle of detail amplitude or strength and the notion that image features can be striking or subtle, large in size or small. MUSICA processing is independent of the size or diameter of the object with the feature to be enhanced. The method is carried out by decomposing the original image into a set of detail images, where each detail image represents an image feature of a specific scale. This set of detail images, or basis functions, completely describes the original image. Each detail image representation and the image background are contrast equalized separately; some details can be enhanced and others attenuated as desired. All the separate detail images are recombined into a single image, and the result is diminished differences in contrast between features regardless of size, such that all image features become more visible.


## IMAGE ARTIFACTS

The appearance and causes of image artifacts that can occur with CR systems should be recognized and corrected. Artifacts can arise from a variety of sources including those related to the IPs themselves, to image readers, and to image processing. Several types of artifacts potentially encountered with CR have been minimized with the latest technology improvements but may still be seen in older systems.

Lead backing added to the aluminum-framed, carbon-fiber cassettes has eliminated the so-called lightbulb effect, darkened outer portions of a film due to backscattered radiation. High sensitivity of the CR plates renders them extremely susceptible to scattered radiation or inadvertent exposure, thus routine erasure of all CR plates on the day of use is recommended, as is the storing of IPs on end, rather than stacking of cassettes one on top of another. The occurrence of persistent latent images after high exposures or after prolonged intervals between plate erasure and reuse has been lessened by the improved efficiency of the two-stage erasure procedure utilized in the latest CR systems. Improved recognition of the collimation pattern employed for a given image allows varied (including off-angle) collimation fields and, in turn, improves histogram analysis and subsequent processing of the imaged region, although these algorithms can fail in some instances. Plate cracking from wear-and-tear can create troublesome artifacts.

Inadvertent double exposures can occur with the present CR systems, potentially masking low-density findings such as regions of parenchymal consolidation, or leading to errors in interpreting line positions. Such artifacts are more difficult to detect than with screen-film systems because of CR's linear frequency processing response, optimizing image intensity over a wide range of exposures (i.e., due to its wide dynamic range). Figure 11.7 shows an example of the double exposure artifact. Laser scanning artifacts can still occur with current CR readers and are seen as a linear artifact across the image, caused by dust on the light source. Proper and frequent cleaning of the laser and light guide apparatus as well as the IPs themselves can prevent such artifacts.

The ability of CR to produce clinically diagnostic images over a wide range of exposures depends on the effectiveness of the image analysis algorithms applied to each dataset. The specific processing parameters used are based on standards tuned to the anatomic region under examination. Incorrect selection of diagnostic specifier or inappropriate anatomic region can result in an image of unacceptable quality. Understanding the causes of some of the CR imaging artifacts described here as well as maintaining formal,



**FIGURE 11.7**

Example of inadvertent double exposure.

routine QA procedures can help to recognize, correct, and avoid future difficulties.

#### SUMMARY OF COMPUTED RADIOGRAPHY

Computed radiography can be utilized for the digital image acquisition of projection radiography examinations into a PACS. As a result of its wide exposure latitude and relative forgiveness of exposure technique, CR can improve the quality of images in difficult imaging situations, such as in portable or bedside examinations of critically ill or hospitalized patients. As such, CR systems have been successfully utilized in the ICU setting, in the emergency room (ER) or trauma center, as well as in the operating room (OR). Computed radiography can also be cost effective for a high-volume clinic setting or in a low-volume site as input to a teleradiology service and has successfully reduced retake rates for portable and other examinations.

Technological advances in CR hardware and software have contributed to the increased acceptance of CR as a counterpart to conventional screen-film projection radiography, making the use of this modality for clinical purposes more widespread. Computed radiography is compatible with existing x-ray equipment, yet separates out the functions of image acquisition or capture, image display, and image archival versus traditional screen-film, in which film serves as the image detector, display, and storage medium. This separation in image capture, display, and storage functions by CR enables optimization of each of these steps individually. Potential expected benefits are improved diagnostic capability (via the wide dynamic range of CR and the ability to manipulate the exam through image processing) and enhanced radiology department productivity (via networking capabilities for transmission of images to remotely located digital softcopy displays and for storage and retrieval of the digital data).

## DIGITAL RADIOGRAPHY

In addition to CR devices for digital image acquisition of projection x-rays, there are the maturing direct digital detectors falling under the general heading of digital radiography (DR). Unlike conventional screen-film radiography in which the film functions as the imaging sensor, or recording medium, as well as the display and storage media, DR, like CR, eliminates film from the image-recording step, resulting in a separation of image capture from image display and image storage. This separation of functions potentiates optimization of each of these steps individually. In addition, DR, like CR, can capitalize on features common to digital or filmless imaging, namely, the ability to acquire, transmit, display, manipulate, and archive data electronically, overcoming some of the limitations of conventional screen-film radiography. Digital imaging benefits include remote access to images and clinical information by multiple users simultaneously, permanent storage and subsequent retrieval of image data, expedient information delivery to those who need it, and efficient, cost-effective workflow with elimination of film.

In this chapter, DR refers to devices in which the digitization of the x-ray signal takes place within the detector itself, providing an immediate full-fidelity image on a softcopy display monitor. Compare this with CR, which utilizes a PSP IP detector in a cassette design that must be processed in a CR reader following x-ray exposure, for conversion to a digital image. Digital radiography devices may be classified as direct or indirect based on their detector design and conversion of absorbed x-rays into an image. Note

that the acronym “DR” may be used by some to refer to direct radiography, also called direct digital radiography (DDR), as the subset of digital radiography in which x-ray absorption within the detector is converted into a proportional electric charge without an intermediate light conversion step.

This section provides an overview of the state-of-the-art in digital radiography systems. A basic description of the data acquisition process will be given, followed by a review of system specifications, image quality, and performance, including signal-to-noise, contrast, and spatial resolution characteristics. Advantages and disadvantages inherent in CR and DR, and a comparison with screen-film radiography will be given with respect to system performance, image quality, workflow, and cost.

## PROCESS DESCRIPTION

**INDIRECT VERSUS DIRECT CONVERSION** DR refers to devices for direct digital acquisition of projection radiographs in which the digitization of the x-ray signal takes place within the detector. Digital radiography devices, also called flat-panel detectors, include two types, indirect conversion devices in which light is first generated using a scintillator or phosphor and then detected by a CCD or a thin-film-transistor (TFT) array in conjunction with photodiodes; and DDR devices, which consist of a top electrode, dielectric layer, selenium x-ray photoconductor, and thin-film pixel array. Figure 11.8 shows a comparison of the direct and indirect energy conversion steps in the production of a digital x-ray image. DDR devices offer direct energy conversion of x-ray for immediate readout without the intermediate light conversion step.

The basis of DR devices is the large area TFT active matrix array, or flat panel, in which each pixel consists of a signal collection area or charge collection electrode, a storage capacitor, and an amorphous silicon field-effect transistor (FET) switch that allows the active readout of the charge stored in the capacitor. Arrays of individual detector areas are addressed by orthogonally arranged gate switches and data lines to read the signal generated by the absorption of x-rays in the detector. The TFT arrays are used in conjunction with a direct x-ray photoconductor layer or an indirect x-ray-sensitive phosphor-coated light-sensitive detector or photodiode array.

An example DDR device, diagrammed in cross section in Figure 11.9, uses a multilayer detector in a cassette design, in which the x-ray energy is converted directly to electron-hole pairs in an amorphous selenium (Se) photoconductive conversion layer. Charge pairs are separated in a bias field such that the holes are collected in the storage capacitors and the electrons



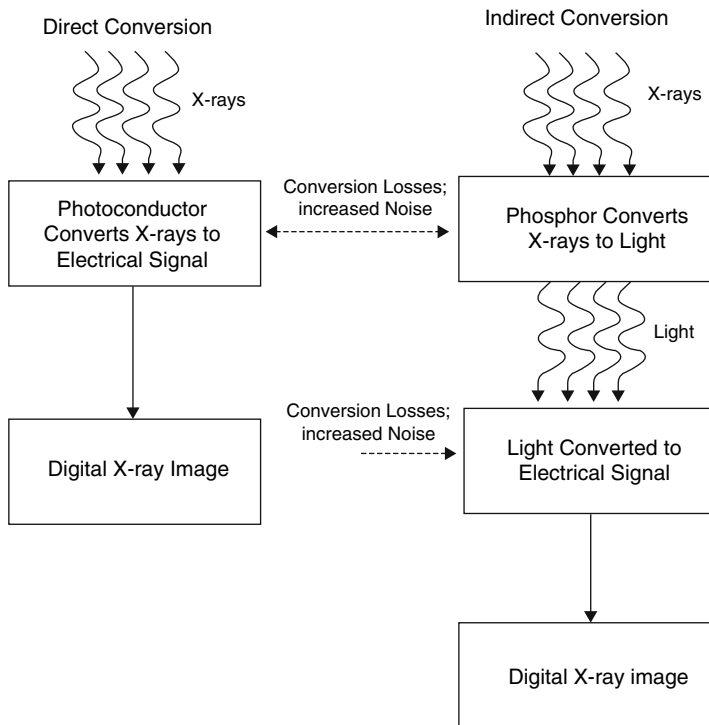


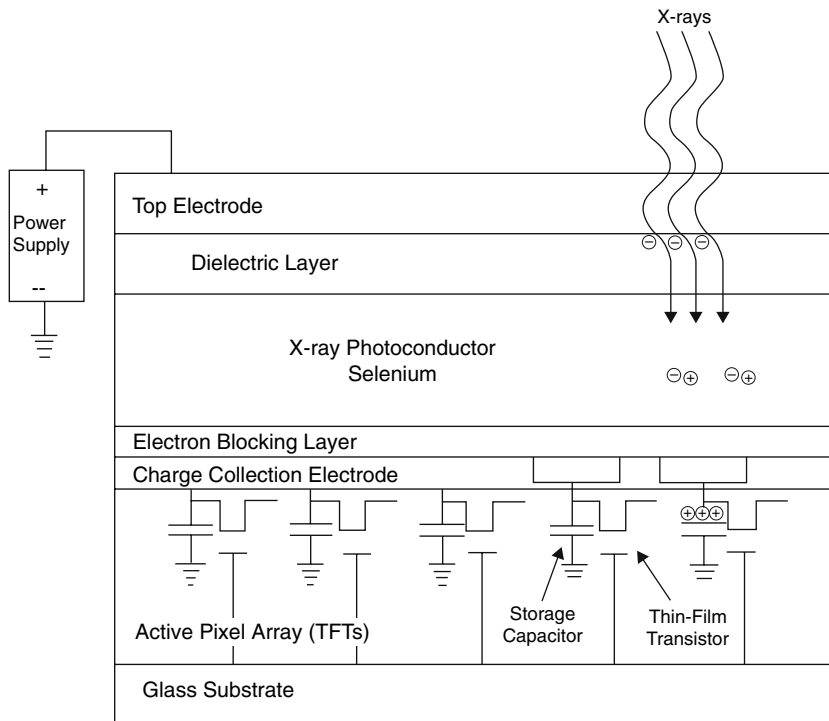
FIGURE 11.8

The image production steps involved in direct and indirect digital radiography detectors.

drift toward the Se-dielectric interface. At the end of exposure, the image resides in the pixel matrix in the form of charges, with the charge proportional to the absorbed radiation. At the end of readout, the charges are erased to prepare for another detection cycle.

An example indirect DR device uses an x-ray-sensitive phosphor coating on top of a light-sensitive flat panel amorphous silicon (Am-Si) detector TFT array. The x-rays are first converted to light and then to a proportional charge in the photodiode (typically a cesium iodide [CsI] scintillator), which is then stored in the TFT array where the image signal is recorded.

**SYSTEM CHARACTERISTICS** DR detectors have high efficiency, low noise, and good spatial resolution; wide latitude; and all the benefits of digital or filmless imaging. Similarly, DR has a very wide dynamic range of quan-



**FIGURE 11.9**

Cross-sectional view of an example of a direct digital radiography (DDR) detector panel.

tization to thousands of gray levels. These devices are becoming more widely used clinically and are available in table buckeyes as well as chest units. Digital radiography units have superior workflow and increased patient throughput due to the elimination of cassette handling.

The short imaging cycle time of DR may lend itself to combined static radiographic and dynamic fluoroscopic uses in future applications. This is true especially for the indirect devices. The direct Se detector, for example, has a ghosting problem due to charge trapping, which introduces a lag time at the end of each cycle, lengthening the time to readiness for the next exposure.

The cost of sensor production is still high so that the overall price of devices has not dropped appreciably. Digital radiography is sometimes referred to as a one-room-at-a-time technology since the detectors are built into the room and matched to the x-ray source. Detector fragility and poor

portability make DR difficult to use in the bedside x-ray environment, but some portable devices are now being introduced.

**IMAGE QUALITY** DR image quality is comparable to that of CR. However, DR devices have higher DQEs than CR, capturing roughly 80% absorption of the x-ray energy at optimum exposures. Thus, DR is a very high-efficiency, low-noise detector, converting much of the incoming x-ray signal into useful output. Several recent studies have demonstrated high image quality at lower radiation dose to the patient. The ability to lower exposure would be a significant advantage for DR. A factor limiting DR efficiency involves the packing fraction or the ratio of active detector area to dead space taken up by the data readout devices (transistors, data lines, capacitors, etc.). Because the physical size of the data readout components is currently fixed, the smaller the pixel size, the smaller the packing fraction, with a larger proportion of dead area overwhelming the active area, in some cases reducing the active area to 30% or less. The overall effect is a reduction in geometric and quantum efficiency.

The spatial resolution of DR is comparable to that of CR, which is still less than that for analog x-ray. Typical matrix sizes are on the order of 2000 to 2500 pixels  $\times$  2000 to 2500 pixels. The pixel size of the TFT array detector is the limiting factor for spatial resolution, with the direct Se detector yielding a better inherent spatial resolution than do indirect detectors. This can lead to better signal modulation and superior contrast.

DR design presents a delicate trade-off between detector efficiency, inversely proportional to pixel size, and spatial resolution, affected directly by pixel size. Typically, DR devices are specified for higher detection efficiency at a cost of much less spatial resolution than screen-film, with compensation by a wide dynamic range or high contrast resolution. Design complexities requiring further development include wiring configurations to minimize dead space and maximize the detector packing fraction, fast and robust signal readout methods, and better error correction matrices for more accurate signal readout.

## COMPARISON OF COMPUTED RADIOGRAPHY AND DIGITAL RADIOGRAPHY

Table 11.2 lists the advantages of CR and DR, including all the benefits of digital images, which can be electronically processed, manipulated, distributed, displayed, and archived. The superior contrast resolution of the digital modalities can compensate in many cases for the lesser spatial resolution as

**TABLE 11.2**

## Summary of Advantages of CR and DR Systems

---

---

Produce digital images capable of being electronically processed, manipulated, distributed, displayed, and archived
Large latitude systems allowing excellent visualization of both soft tissue and bone in the same exposure image
Superior contrast resolution potentially compensating for lack of spatial resolution
Decreased retake rates
Potential cost savings if film is eliminated
Improved radiology department workflow with elimination of film handling routines

---

---

compared with screen-film. Both CR and DR can be utilized for the digital image acquisition of projection radiography examinations into a PACS.

As for any digital image acquisition device, CR and/or DR would be the first point of entry into a PACS. Errors may propagate from here, with the quality of the PACS output being directly dependent on the quality of the signal in. In addition to image quality, essential features for successful clinical implementation of CR or DR systems for a PACS include the following: DICOM conformance of the modality is essential and includes compliance with the image data and header format, as well as the DICOM communication protocol. Equally critical is interfacing to the RIS-HIS. Integration of the CR/DR system with the RIS-HIS can reduce human errors on patient demographic information input and improve efficiency. Ease of integration of the device into the daily workflow routine, and simplicity and robustness of the user interface are very important. Reliability, fault tolerance, and capabilities for error tracking are also major issues to consider, as are device speed and performance.

As a result of CR's convenient workflow and portability, as well as its wide exposure latitude and relative forgiveness of exposure technique, CR can improve the quality of images in difficult imaging situations, such as in portable or bedside examinations of critically ill or hospitalized patients, and enable decreased retake rates for improper exposure technique. As such, CR systems have been successfully utilized in the ICU setting, in the ER or trauma center, as well as in the OR. Computed radiography can also be cost effective for a high-volume clinic setting, or in a low-volume site as input to a teleradiology service. Cost savings and improved radiology departmental workflow can be realized with CR and the elimination of film.

Technological advances in CR hardware and software have contributed to the increased acceptance of CR as the current counterpart to conventional screen-film projection radiography, making the use of this modality for clinical purposes more widespread. Computed radiography is compatible with existing x-ray equipment, yet separates out the functions of image acquisition or capture, image display, and image archival versus traditional screen-film, in which film serves as the image detector, display, and storage medium. This separation in image capture, display, and storage functions by CR enables optimization of each of these steps individually. Potential expected benefits are improved diagnostic capability (via the wide dynamic range of CR and the ability to manipulate the data through image processing), and enhanced radiology department productivity (via networking capabilities for transmission of images to remotely located digital softcopy displays and for storage and retrieval of the digital data).

DR devices have more efficient detectors, offering direct energy conversion of x-ray for immediate readout. The higher DQE may enable DR to produce high-quality images at a lower radiation dose to the patient. These detectors have low noise and good spatial resolution, wide latitude, and all the benefits of digital or filmless imaging. But cost is still high since detector production is difficult and expensive, and DR is a one-room-at-a-time detector. Digital radiography may be cost effective in high volume settings with constant high patient throughput. However, meeting the cost competitiveness of screen-film systems is difficult unless film printing is eliminated from the cost equation. Digital radiography may be preferable for imaging examinations requiring very high quality, such as in mammography, upright chest exams, and bone work.

Future improvements in image-processing algorithms, with a better understanding of optimum display settings for softcopy viewing, have the potential to greatly facilitate and standardize softcopy reading of digital projection radiographs and further the acceptance of CR and DR in the clinical arena. It is likely that CR and DR devices will coexist for some time.

## **QUALITY ASSURANCE AND QUALITY CONTROL FOR ACQUISITION**

The current trend for radiology departments and medical imaging within healthcare enterprises is an increasing move toward the all-digital or filmless medical image management system or PACS. Operational concerns with PACS implementations can arise at all stages of the process, from the design

specifications to installation, training, and acceptance. Quality control procedures necessarily become modified in the filmless radiology department, and new processes must be put in place to better prepare for the total digital clinical department.

Forming and maintaining a continuing quality improvement (CQI) committee may facilitate PACS installation and training periods, workflow modifications, quality assurance, and clinical acceptance. This committee should include radiologists at all levels (resident, fellow, attending), radiology technologists, film library personnel, ED and ICU clinician end users, and PACS team members. The CQI committee may assist in the creation of new management procedures, provide a means for user feedback and education, and contribute to the overall acceptance of and user satisfaction with the system.

## PROBLEMS OCCURRING AT ACQUISITION

The imaging modality is the first entry point into the PACS, and any errors in data input here can propagate throughout the system. Thus, interfacing of a PACS to the RIS-HIS and, better yet, DICOM modality worklist capability at the imaging device, are essential. When a PACS is properly interfaced to the RIS-HIS, input data can be verified by comparison of pertinent demographic data (name, date, time, MRN, AccNum, and exam type) at the PACS acquisition gateway with the data recorded in the RIS. Thus, any imaging exam entering the PACS will be RIS-verified prior to archival and distribution, maintaining the data integrity of the system.

Most imaging modalities are now tightly coupled with the RIS and provide automatic downloading of demographic information from the RIS, via barcode readers, to the modality, and hence the DICOM header. This eliminates the highly error-prone manual entry of data at acquisition. Unfortunately for manual data-entry devices, any errors in data may result in the image data being held in a queue pending manual (human) inspection and resolution. Continuous feedback should be given to technologists making repeated errors in data entry.

The well-designed PACS holds newly acquired studies in a restricted area (fix-queue or “penalty box”) until the demographic data in the header is matched to a pending exam request from the RIS-HIS. If any failure occurs, such as an incorrect MRN or DOB (date of birth), the new exam will not pass automatically into the system to be archived (although it may be displayable) until the discrepancy has been resolved by human intervention. However, the inverse test has not been implemented. Pending exam orders

held in the RIS-HIS that do not relate to any incoming PACS image data within a certain time frame should be flagged. Full PACS acquisition QA requires this bidirectional process monitor to ensure that data in the PACS is valid and verified with data in the RIS-HIS, and all data in the RIS-HIS is acquired into the PACS. This may also assist QA procedures in determining which of the studies that have been ordered and completed have no associated report and therefore may not have been read.

Some DICOM transfer of imaging exams (i.e., from CT and MR scanners) to the PACS requires autosend-networking pathways to be enabled at the scanner. Unfortunately, these features can easily be turned off (frequently by the service manufacturer), resulting in missed real-time transfer of images to the PACS. Stressing the importance of having the autosend enabled at the time the examination is performed to the imaging technologists as well as the manufacturer's service personnel can reduce this problem.

Although many digital angiographic/fluorographic systems are DICOM compliant, few have been integrated into a PACS. The large volume of data typically generated by angiographic procedures is one reason for this. One way to reduce the data volume, and perhaps facilitate connection to a PACS, might be to store only key images of an angiographic run, much the way they are filmed. Some manufacturers allow operators to create summary series of the examinations, which could then be transmitted to the PACS for viewing on display workstations. A second problem for incorporation of angiographic images into a PACS arises from the inability of most PACS to do subtraction, pixel shifting, and rapid mask selection—features utilized in most angiographic examinations.

## QUALITY CONTROL PROCEDURES/ TROUBLESHOOTING

During the PACS planning, specifications, and lab testing phases it can be beneficial to involve anticipated users of the system. User awareness of the goals of a PACS implementation and its system features prior to clinical installment can affect the overall success of the system. Installation of system components will be most successful when scheduled during low-volume periods and when all affected users are notified well in advance of the install date. Backup contingency plans must be in place prior to going live in the clinical environment.

Among the many roles of the medical physicist and/or PACS engineer in incorporating an imaging modality into the diagnostic imaging department are acceptance testing of the device and QA-QC. The medical physi-

cist should be involved in the siting and planning of the imaging system, as well as the installation, testing and tuning, and training. Formal QA-QC procedures are still evolving, particularly for the newer modalities such as CR. In fact, substantial efforts have been under way to standardize CR QA-QC, such as the American Association of Physicists in Medicine (AAPM) Task Group #10 draft document “Computed Radiography Acceptance Testing and Quality Control.”

In spite of the fact that CR is more forgiving of a broad range of exposures, the use of this modality is not an excuse to employ poor radiographic technique. Computed radiography exposures as well as image quality should be routinely monitored on a per-examination-type basis. Understanding the causes of possible CR imaging artifacts as well as maintaining formal, routine quality assurance procedures can help to recognize, correct for, and avoid future difficulties with this relatively new modality, just as with the proven modalities. Formal QA-QC procedures should be put in place and diligently adhered to.

Maintenance, QA-QC, and workflow procedural modifications continue to be developed as incidents occur and system troubleshooting is carried out, as more radiology departments implement PACS technologies. Documentation of events and CQI committee review and analysis of system functioning in conjunction with review of user comments and suggestions is extremely important to the successful clinical operation of the PACS.

## **FUTURE TRENDS IN IMAGE ACQUISITION**

Although the types of imaging modalities will probably not change all that much in the next several years, the anticipated future trends in image acquisition for digital radiology and PACS include changes in the image dataset sizes, changes in the imaging devices themselves, and improvement in image processing for softcopy display of digital images.

### **IMAGE DATASETS**

No new types of imaging modalities are foreseen for the near future. However, it is anticipated that the image data sets acquired from the existing modalities will increase in overall study file size, in some cases dramatically; to a certain extent this has already begun. For example, many radiology departments have begun installing multiple detector array or multislice CT scanners, which tend to generate a greater number of individual images than



do the single detector array scanners. This is due to the slice thickness in helical acquisition (~5 mm) versus the single detector arrays (~7 to 10 mm) and the clinical imaging protocols used, as well as the increasing clinical utility of three-dimensional image display representations.

Image matrix sizes for the digital projection radiography devices (CR and DR) have gone up from roughly from  $1000 \times 2000$  square matrices to  $4000 \times 5000$  pixels squared for mammography applications. The increased sampling was done to improve the spatial resolution. Most laser film digitizers can now vary their spot sizes from  $200\mu\text{m}$  down to  $50\mu\text{m}$ , greatly improving the inherent spatial resolution of the resulting images of the scanned analog film, with a concomitant increase in file size.

The bit depth representation of grayscale pixel values has also increased from 8 bits to 10, 12, and 16 bits, and color images are stored as 32-bit- or 4-byte-per-pixel data files. Furthermore, the addition of post-processing results or slice reconstructions, and cinegraphic sequences to the image dataset, while improving the overall quality of the image, may greatly increase the amount of data to be acquired into a PACS.

## DEVICES

While image datasets and file sizes are getting larger, the imaging devices themselves will continue to get smaller in physical footprint. This has been seen most dramatically with the CR devices, going from requiring roughly  $36\text{m}^2$  of floor space and special electrical power and cooling to desktop devices that can be placed in most any location. CT and MRI devices too are becoming smaller in size, more portable, and more robust. Hopefully these devices will continue to become less expensive.

## IMAGE PROCESSING

An important area of increased attention continues to be image-processing capabilities for softcopy image display. Future processing techniques will most likely go above and beyond the simple window and level (or contrast and brightness) manipulation techniques. These post-processing algorithms are currently available and tunable at the imaging modality or accompanying modality acquisition workstation but may in time be manipulable in real-time at the display station.

Image compression is currently being debated but may eventually be available at the modality to reduce image transmission time and archival

space. Some techniques, such as the wavelet transform, may become more widely utilized not only as compression techniques, but also for image enhancement at the imaging devices.

In time, it is anticipated that the percentage of all imaging devices utilized by healthcare enterprises that are digital in nature will increase greatly. Further, the percentage of digital image acquisition from the devices that are capable should increase, decreasing the amount of film used as an acquisition, display, and archival medium.

## CONCLUSION

Image acquisition is the first point of data entry into a PACS. Errors generated here can propagate throughout the system, adversely affecting clinical operations. General predictors for successful incorporation of image acquisition devices into a digital-imaging department include:

- ▶ Ease of device integration into the established daily workflow routine of the clinical environment
- ▶ High reliability and fault tolerance of the device
- ▶ Simplicity and intuitiveness of the user interface
- ▶ Device speed

Imaging modality conformance with the DICOM standard is critical. DICOM consists of a standard image format as well as a network communications protocol. Compliance with this standard enables an open architecture for imaging systems, bridging hardware and software entities and allowing interoperability for the transfer of medical images and associated information between disparate systems.

- ▶ Compliance with the DICOM standard has greatly facilitated image acquisition for PACS and digital radiology departments.
- ▶ DICOM compliance should be required from the modality and PACS vendors.
- ▶ Most modalities today do comply with the DICOM standard.
- ▶ Interface boxes are available to convert legacy devices that do not comply with the DICOM standard.

Equally essential, particularly at acquisition, is interfacing the RIS-HIS with the PACS. This greatly facilitates input of patient demographics (name,

date, time, MRN, AccNum, exam type) and imaging parameters and enables automatic PACS data verification, correlation, and error correction with the data recorded in the RIS-HIS. Most imaging modalities are now tightly coupled with the RIS, providing automatic downloading of demographic information from the RIS, via barcode readers, to the modality, and hence to the DICOM header. Additionally, many vendors comply with the DICOM modality worklist function, providing the capability to download RIS-HIS schedules and worklists directly to the imaging modality. Both of these features eliminate the highly error-prone manual entry of data at acquisition.

- ▶ RIS-HIS-PACS database integration is essential for a clinically functioning digital department and is intelligence added.
- ▶ HL7 compliance is desirable.
- ▶ RIS-HIS databases are typically patient-centric, enabling query and retrieval of information by the patient, study, series, or image data hierarchy.
- ▶ Modality worklist capability at the imaging device can greatly improve departmental workflow and efficiency.

Image acquisition from the inherently digital modalities such as CT and MRI should be performed as a direct digital DICOM capture, as opposed to frame grabbing of the data. Direct digital interfaces allow capture and transmission of image data from the modality at the full spatial resolution and full bit depth or grayscale inherent in the modality, while analog (video) frame grabbers digitize the video signal voltage output going to an image display, such as a scanner console monitor, and image quality is limited by the process to only 8 bits (or 256 gray values). Direct capture of digital data from the inherently digital modalities is the preferred method of acquisition.

Digital acquisition of images originally recorded on film can be accomplished using a variety of image digitization devices. These include the infrequently used analog video cameras with ADCs, digital cameras, CCDs, and laser scanners. Laser scanners use a fine laser beam of generally variable or adjustable spot sizes down to 50  $\mu\text{m}$ . They have maximum resolution of  $4096 \times 5120 \times 12$  bits, can accommodate the full optical density range of film, and are semiautomatic or fully automatic in operation. Laser scanners are the film digitizer of choice for most PACS needs and have an adjustable spot size to meet the application.

Computed radiography and DR are modalities that digitally acquire projection radiographs without the use of film. Computed radiography uti-

lizes a PSP IP in a cassette design similar to that of screen-film cassettes. The PSP IPs have the property that, when exposed to x-ray, a latent image is formed in the plate until it is subsequently stimulated to luminesce by a laser. The emitted light is then captured and converted to a digital image. Digital radiography uses conversion of x-ray energy to electrical charge (often via an intermediate light conversion step). The electrical charge is collected by a TFT pixel array, offering electronic readout initiated immediately at the time of exposure. Both CR and DR have:

- ▶ All the benefits of digital (filmless) images including the ability to acquire, transmit, display, manipulate, and store data digitally
- ▶ Wide latitude response detectors, potentially reducing retakes resulting from poor exposure as in the screen-film based environment
- ▶ Equivalent spatial resolution capabilities

Computed radiography has superior procedural flexibility for some examinations with difficult positioning, can accommodate portable bedside examinations, and is a proven clinical modality. Digital radiography has a higher efficiency detector with potential for dose reduction and immediate image readout with higher patient throughput, particularly in settings with a steady flow of high volume.

Some current DR problem areas include portability issues, the requirement for a single device per radiographic room (“one-room-at-a-time technology”), and its high production cost. It is likely that CR and DR will coexist as digital radiographic devices for some time. Quality assurance and quality control procedures become necessarily modified in the filmless radiology department, and new processes must be put in place. Forming and maintaining a CQI committee may facilitate:

- ▶ PACS installation and training
- ▶ Workflow management
- ▶ Quality assurance
- ▶ Clinical acceptance

The major problem occurring at acquisition is inaccurate data input at the modality. This source of (human) input error can be greatly reduced by integrating the RIS-HIS with the PACS. This allows RIS verification of PACS data prior to archival, and an image audit monitoring of what is in the PACS as compared with what is in the RIS, and vice versa. Quality control for the imaging modalities is just as essential in the digital world as it is in

**TABLE 11.3**  
Summary of Future Trends in Image Acquisition

Image matrix size	↑
Image quality	↑↑
Spatial resolution	↑
Number of image slices	↑↑↑
Size of imaging examinations	↑↑↑
Size of devices	↓↓
Portability of devices	↑
Cost of devices	↓
Percentage of image devices that are digital	↑↑↑
Percentage of image acquisition that is digital (elimination of film)	↑↑

the film-based arena. Proper planning and communication with end users early and often as to system status can reduce dissatisfaction with the system.

Future trends anticipated in image acquisition for digital radiology and PACS are summarized in Table 11.3 and include changes in image datasets such as increased matrix sizes, increased inherent spatial resolution, increased slice numbers and study sizes, and improved image quality. Imaging devices themselves will continue to become smaller in physical footprint, more portable and robust, and possibly less costly. The image-processing capabilities for softcopy display will be an important area receiving increased attention in the future. Many of these capabilities will be developed by the image acquisition modality vendors and by research teams for application to the optimal softcopy display of digital images for diagnosis and review. In time, the percentage of images acquired digitally will increase, and less film will be printed for medical imaging.

## REFERENCES

- Agfa. The highest productivity in computed radiography. *Agfa-Gevaert NV Report*. Belgium; 1994.
- Andriole KP. Anatomy of picture archiving and communication systems: nuts and bolts—image acquisition: getting digital images for imaging modalities. *J Digit Imaging*. May 1999a;12(2 suppl 1):216–217.

- Andriole KP. Computed radiography overview. In Seibert JA, Filipow LJ, Andriole KP, eds. *Practical Digital Imaging and PACS*. Medical Physics Publishing; 1999b:135–155.
- Andriole KP. Productivity and cost assessment of CR, DR and screen-film for outpatient chest examinations. *J Digit Imaging*. January 2003;15(3):161–169.
- Andriole KP, Avrin DE, Yin L, Gould RG, Arenson RL. PACS databases and enrichment of the folder manager concept. *J Digit Imaging*. February 2000;13(1):3–12.
- Andriole KP, Gooding CA, Gould RG, Huang HK. Analysis of a high-resolution computed radiography imaging plate versus conventional screen-film radiography for neonatal intensive care unit applications. *SPIE Phys Med Imaging*. 1994;2163:80–97.
- Andriole KP, Gould RG, Avrin DE, Bazzill TM, Yin L, Arenson RL. Continuing quality improvement procedures for a clinical PACS. *J Digit Imaging*. August 1998;11(3 suppl 1):111–114.
- Barnes GT. Digital x-ray image capture with image intensifier and storage phosphor plates: imaging principles, performance and limitations. *Proceedings of the AAPM 1993 Summer School: Digital Imaging*. Monograph 22. Charlottesville, VA: University of Virginia; 23–48.
- Berg GE, Kaiser HF. The x-ray storage properties of the infra-red storage phosphor and application to radiography. *J Appl Phys*. 1947;18:343–347.
- Bogucki TM, Trauernicht DP, Kocher TE. Characteristics of a storage phosphor system for medical imaging. *Kodak Health Sciences Technical and Scientific Monograph 6*. New York: Eastman Kodak Co; July 1995.
- Gringold EL, Tucker DM, Barnes GT. Computed radiography: user-programmable features and capabilities. *J Digit Imaging*. 1994;7(3):113–122.
- Honeyman JC, Staab EV. Operational concerns with PACS implementations. *Appl Radiol*. August 1997a;13–16.
- Honeyman JC, Jones D, Frost MM, et al. PACS quality control and automatic problem notifier. *SPIE Medical Imaging 1997: PACS Design and Evaluation 1997b*;3035:396–404.
- Ishida M. *Fuji Computed Radiography Technical Review*, No. 1. Tokyo: Fuji Photo Film Co, Ltd; 1993.
- Kodak. Digital radiography using storage phosphors. *Kodak Health Sciences Technical and Scientific Monograph*. New York: Eastman Kodak Co; April 1992.
- Kodak. Optimizing CR images with image processing: segmentation, tone scaling, edge enhancement. *Kodak Health Sciences Technical and Scientific Manuscript*. New York: Eastman Kodak; March 1994.
- Kotera N, Eguchi S, Miyahara J, Matsumoto S, Kato H. Method and apparatus for recording and reproducing a radiation image. US Patent 4,236,078 (1980).
- Lee DL, Cheung LK, Jeromin LS. A new digital detector for projection radiography. *Proc SPIE Phys Med Imaging*. February 26–27, 1995;2432:237–249.

- Luckey G. Apparatus and methods for producing images corresponding to patterns of high energy radiation. US Patent 3,859,527 (June 7, 1975). Revised No. 31847 (March 12, 1985).
- Matsuda T, Arakawa S, Kohda K, Torii S, Nakajima N. *Fuji Computed Radiography Technical Review*, No. 2. Tokyo: Fuji Photo Film Co, Ltd; 1993.
- Oestman JW, Prokop M, Schaefer CM, Galanski M. Hardware and software artifacts in storage phosphor radiography. *Radiographics*. 1991;11:795–805.
- Ogawa E, Arakawa S, Ishida M, Kato H. Quantitative analysis of imaging performance for computed radiography systems. *SPIE Phys Med Imaging*. 1995;2432:421–431.
- Seibert, JA. Photostimulable phosphor system acceptance testing. In: Seibert JA, Barnes GT, Gould RG, eds. *Specification, Acceptance Testing and Quality Control of Diagnostic X-ray Imaging Equipment*. Medical Physics Monograph No. 20. Woodbury, NY: AAPM; 1994:771–800.
- Solomon SL, Jost RG, Glazer HS, Sagel SS, Anderson DJ, Molina PL. Artifacts in computed radiography. *AJR Am J Roentgenol*. 1991;157:181–185.
- Sonoda M, Takano M, Miyahara J, Kato H. Computed radiography utilizing scanning laser stimulated luminescence. *Radiology*. 1983;148:833–838.
- Storto ML, Andriole KP, Kee ST, Webb WR, Gamsu G. Portable chest imaging: clinical evaluation of a new processing algorithm in digital radiography. 81st Scientific Assembly and Annual Meeting of the Radiological Society of North America; November 26–December 1, 1995; Chicago.
- Volpe JP, Storto ML, Andriole KP, Gamsu G. Artifacts in chest radiography with a third-generation computed radiography system. *AJR Am J Roentgenol*. 1996;166:653–657.
- Vuytsteke P, Dewaele P, Schoeters E. Optimizing radiography imaging performance. *Proceedings of the 1997 AAPM Summer School*, 107–151.
- Willis CE, Leckie RG, Carter J, Williamson MP, Scotti SD, Norton G. Objective measures of quality assurance in a computed radiography-based radiology department. *Proc SPIE*. 1995;2432:588–599.
- Wilson AJ, West OC. Single-exposure conventional and computed radiography: the hybrid cassette revisited. *Invest Radiol*. 1993;28(5):409–412.

# IMAGE COMPRESSION

---

---

BRADLEY J. ERICKSON

Implementation of any picture archiving and communication system (PACS) requires a discussion of how to deal with the large quantities of data that must be transmitted and stored. Even before PACS, clever engineers devised encoding schemes that reduced the apparent size of images to reduce the demands placed on transmission and storage devices. Encoding images in order to reduce storage and transmission demands is called “image compression.” Compression of medical images is controversial—some argue that it should not be used, because soon storage and networks will be cheap enough that compression is not necessary. This argument has been made from the early days of PACS and continues to be made. I suspect this issue will continue to be controversial until humans no longer view images.

Those faced with the need to implement systems in the present point to the benefits that can be achieved. Sometimes the argument for using compression is purely financial: using compression tips the scale from a losing proposition to a winning one. It can allow one to use less expensive networking technologies and requires less (potentially much less) storage. These can be major components of the capital acquisition as well as operating costs.



In other cases, compression can make the difference between a project that is feasible and one that cannot work. When a specific turnaround time is required to achieve satisfactory service and the network capacity is limited, compression can be the only solution. This is most often the case when remote facilities are an integral part of the PACS equation.

Whether cost savings or turnaround time is the issue, the counter-argument is frequently made that if one delays the project, the advance of technology can obviate the need for compression. This is an appealing argument, for Moore's law seems to apply not only to processor speed but also to network and storage devices. Unfortunately, this oversimplifies the situation. The primary reason this argument fails is that the demands on the imaging system are not static. The standards for PACS performance continue to rise in step with Moore's law. That is because the same advances in processing power that make a PACS cheaper are also making imaging devices that produce higher image data volumes. A few years back, we reported that over a long period of time, the daily image data volume for computed tomography (CT) and magnetic resonance (MR) scanners was paralleling the increases in network speed and storage density. Figure 12.1 shows that the relationship has held true for 30 years. The conclusion is that one should not delay a project in expectation that technology advances will eliminate the need for compression. *If a PACS implementation is not viable without compression today, it likely will not be viable in the future. Similarly, if compression is*

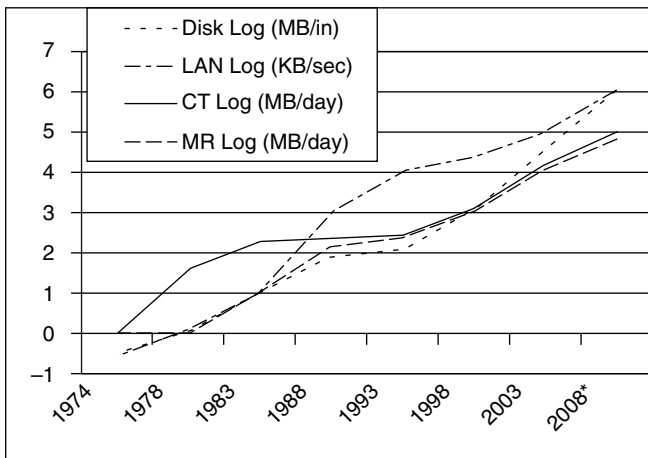


FIGURE 12.1

Graph of network speed, storage density, and daily CT and MR data volumes from 1974 to 2004. \* = projected.

*necessary at the start of a project, one should plan that it will continue to be necessary for its lifetime.*

## IMAGE COMPRESSION TECHNOLOGIES

There are two main categories of image compression: reversible, or lossless, and irreversible, or lossy. Reversible compression will exactly reproduce the original image after compression and decompression. There is little controversy surrounding the use of lossless compression—the major question is whether the computational effort is worth the modest gain. For most medical images, the gain is on the order of 2:1 to 3:1. A 2:1 ratio means that the storage or network demands are 50% of what the original image would require; 3:1 would require just 33%. Many systems implement this type of image compression because the computation cost is small compared to the cost savings in storage and networking. In some cases, this compression occurs only when the image is saved to the long-term archive, saving only storage costs. Some storage devices actually have hardware that performs this type of compression.

There are 2 main ways that lossless compression algorithms work. The first type is the group that utilizes repetition within an image to compress the image. For instance, if there were several consecutive pixels with the same value, it is faster to say “5 pixels of value 312” than to say “312 312 312 312 312.” Other than some outer regions of images (e.g., the pixels that fill the circle of a CT to make a square) or three-dimensional (3-D) renderings, exact repetition of pixel values is rather rare for medical images. However, an adaptation is to use the prior value as a starting point and just send the difference. Other than where there are strong edges, small numbers (which take less space) added to the prior value will correctly represent the pixel value.

An alternative to this strategy is to look for patterns within the image. If one were to look through this text, there would be a fairly high frequency of “the” and “to,” and one would find that “q” was always followed by “u.” If a single-letter code were substituted for these patterns, the storage space required for the text could be significantly reduced. This is a dictionary-style compression method and is commonly used for highly structured information, such as text. It is used in most of the “zip” programs that are available to compress files.

The lines that follow this paragraph show examples of lossless compression methods. The original pixel values are shown in [a]. A run-length encoding of the pixels could be that shown in [b]. Note that there are more

values sent in line [b] than line [a], though the value encoding how many of a kind may not take as much space as the value. Line [c] shows a differential encoding of the pixels. The number of values is the same, but if the change values are small, they will take less space to transmit. The next line, [d], has two parts: a dictionary and the data. The data values are used to look up the actual pixel values for the output. The first entry in the dictionary is the most common pattern and uses the fewest bits possible, with less frequently used patterns consuming more bits. Note that in all of these examples, the compressed data is larger than the original. This is because in such a small dataset there is little redundancy to take advantage of and because different data sizes (bits) are not represented.

```
[a] 1  4  4  4  5  6  7  8  8  4  5  4  5  1  4
[b] 1  1) (3  4) (1  5) (1  6) (1  7) (2  8) (1  4)
    (1  5) (1  4) (1  5) (1  1) (1  4)
[c] 1  3  0  0  1  1  1  1  0 -4  1 -1  1 -4  3
[d] dictionary:
    (4  5) (4  4) (1) (6) (7) (8  8) (4)
    data:
    3  2  1  4  5  6  1  1  3  7
```


While there is some structure to medical images, the gain achieved by compression of images is not as great as that achieved by compression of text. Compared with the repetition strategy described above, pattern-based compression typically will achieve higher compression ratios but requires more central processing unit (CPU) power and memory to compress and decompress. However, one can devise methods in which most of the CPU power is used for the compression task (which is done just once) and less for decompression (which is done every time the image is viewed). Unlike lossless compression, where the compression ratio is the same for a given image, the implementation of lossy compression requires decisions about how much compression to apply. Depending on the algorithm, one may specify a “quality factor,” which roughly correlates with perceived image quality after compression, or the algorithm may require a specific compression ratio. It would seem that the quality factor approach is more desirable, as each image would then be optimally compressed. In practice, however, the correlation with perceived image quality is variable. This will be discussed in greater depth later in this chapter.

Lossy compression begins with a very different step from lossless compression that is not very intuitive. That first step is a transformation, and the

type of transformation is often used to characterize the type of compression. An early popular transformation was the discrete cosine transform (DCT). In some cases, the entire image had this transform applied. However, this is computationally intensive, particularly for the computers of the 1980s, when the technique gained popularity. Therefore, an industry group known as the Joint Photographic Experts Group (JPEG) created a processing algorithm for images that would break the image into  $8 \times 8$  pixel blocks and apply the DCT to each block. These smaller blocks were more amenable to computers of the day and permitted dedicated hardware to be built for this task. In either case, the DCT would create an image (with the same matrix size as the original image/block) in which the low-frequency components were concentrated in the lower left corner, with the highest-frequency elements in the upper right. The values in this transformed image matrix are sometimes referred to as coefficients because they are coefficients for a mathematical formula that can re-create the image. More recently, the various forms of the wavelet transformation have been applied for image compression. Such wavelet transforms produce multiple images at different scales, where the next-higher-resolution image has the higher-frequency changes from its lower-resolution predecessor. Figure 12.2 shows an example of a 4-level wavelet transformed image.

After the transformation of choice has been performed, the image has been altered, but there has been no loss of information—one could invert the transformation and exactly reproduce the original image. The next step after transformation is quantization, which is the step of deciding which parts of the transformed image are important and which are not. One could select a low-frequency rectangle and compute the inverse transform of it to re-create an image that would be somewhat blurrier than the original. How blurry the resulting image is depends on how effectively the transformation concentrated information into the selected parts of the transformed image.

There are 2 primary ways one can save storage/transmission space: by assuming that coefficients not sent are 0 (zero-filling) or by reducing the accuracy of the coefficients that are sent (e.g., converting from 80-bit floating point values to 20-bit scaled integers). In some schemes, including JPEG, one can also preferentially save accuracy on those coefficients that represent frequencies that the human eye is sensitive to, while saving bits on those components to which the eye is not sensitive. The multiresolution nature of wavelets also permits more sophisticated selection of parts of the image than just a rectangle. Some algorithms, such as Set Partitioning in Hierarchical Trees (SPIHT), efficiently represent areas of images that have much activity while saving bits in uninteresting areas. This is complementary to



**FIGURE 12.2**

Example of a wavelet transformed image. The lowest-frequency components are in the upper left corner. The original image can be exactly reconstructed from this image.

preferring spatial frequencies to which the eye is sensitive, as described above.

In some cases, it is also desirable to sequence the coefficients so that those that are most important to recreating a faithful image are first in the stream of data, rather than ordered by an X, Y position. This can allow the user to get a quick view while the image continues to improve as more information is sent. This is called progressive streaming (referring to the order of data transmission) or successive refinement (referring to the actual displaying of improved images).

After the coefficients have been quantized, there is often still some redundancy, and so the third step in lossy compression is to use one of the lossless techniques described above (typically a dictionary-type method) to further compress the quantized coefficients.

Unlike lossless methods, in which the compression ratio is fully determined by the algorithm selected, lossy methods allow the user to select either the compression ratio or a quality factor. In either case, the user's choice

affects how much the coefficients are quantized, meaning that a judgment is made about acceptable compression level. Unfortunately, in most cases it is not feasible for the user to interactively select the appropriate compression for each and every image, and so general rules for compression must be used.

## APPLICATION TO MEDICAL IMAGING

The description of image compression provided above is not unique to medical images—these image compression algorithms are used for images on the Internet or from your digital camera. However, some unique aspects of medical images deserve comment. The first is that most medical images are grayscale, with more gray levels than most commercial/consumer images. Radiographic images are 10 to 12 bits, while CT and MR images are 12 to 16 bits. Since computers prefer to deal with pixels that are multiples of 8 bits, most medical images are stored using 16 bits per pixel. One exception is ultrasound. In some cases (particularly Doppler), ultrasound images use 24-bit color, which is like that used in commercial/consumer imaging. Grayscale ultrasound images are variable and may be either 8-bit or 16-bit. The significance is that JPEG was created without attention to medical imaging, and the extensions to JPEG to allow it to support 12 or 16 bits are not always compatible.

Another difference is that CT and MR images are often from spatially adjacent locations, meaning there is significant redundancy between 2 or more images. This is not considered in any of the (two-dimensional [2-D]) algorithms described above. However, an update of the JPEG standard (referred to as JPEG2000) does have components that support the notion of a 3-D stack of slices. While the JPEG2000 standard is still advancing, the basic 2-D methods have been ratified, and there are implementations that have been tested on medical images. Some of those results will be discussed later in this chapter. Because compression has been considered important to the success of PACS by many of its leaders, there was a concerted effort to participate in the JPEG2000 standard in order to make it better support medical images—in particular, by addressing the 16-bit issue. The Digital Imaging and Communications in Medicine (DICOM) standard has included JPEG-based compression in its transfer syntaxes for some time. The approved part of JPEG2000, which focuses on 2-D image compression using a wavelet transformation, has already been adopted as a DICOM standard (Supplement 61, November 2001). It includes support for 16-bit images, progressive encoding, and region-of-interest functions (improved quality

in specific regions of the image). It is expected that as other parts of JPEG2000 are ratified, such as the volumetric image functions (JP3D or JPEG2000 Part 10) and Motion JPEG for endoscopy, DICOM will add them to its standard.

## EVALUATION IN MEDICAL IMAGING

### IS IT WORTH IT?

There is controversy over whether there is any acceptable level of lossy compression for medical images or if there should be no allowance for alteration of pixel data. I believe that lossy compression can be acceptable and justified. Many major medical centers have come to the same conclusion. There are several bases for this conclusion.

First, as with nearly all parts of medical practice, one must exercise judgment about the trade-off of benefits and risks. Higher-quality radiographs can be obtained if higher-exposure images are obtained. Whether this would improve diagnosis or simply reduce quantum noise is the judgment we make when an exposure/dosage decision is made. In the case of MR and CT, one could reconstruct many more images than is typically required using interpolation. In the case of MR, one could use many more signal averages to reduce noise. But in order to conserve resources (reduce MR examination time or reduce total images per examination) we make judgments about how much is enough to be diagnostic. The same is true with image compression. As long as there is a thorough understanding of the trade-offs, and confidence that diagnostic accuracy is not significantly impacted, compression is no more “dangerous” than other decisions in medical imaging.

Second, the pixel data altered by compression has nearly always been altered even before it is compressed. CT uses reconstruction filters that may blur to reduce noise. Computed radiographs have lookup tables applied to emphasize certain regions of the histogram, while suppressing others, and may also have spatial filtering. These are also carefully considered alterations of pixel data in order to improve radiologist efficiency. Perhaps the greatest loss of data occurs in fluoroscopy, in which only a few spot images are taken from several minutes of an examination.

Third, it has been shown that at modest levels, lossy compression actually improves diagnostic performance because it preferentially loses the noise. This has been shown to improve detection of breast cancers as well

as lung pathologies. From this perspective, compression could be viewed as a valuable processing step. This also raises issues about implementation. If lossy compression makes a lesion more conspicuous, compressing an image only after interpretation could cause one to miss a lesion that is more apparent with compression. If the lossy compressed version is the only version archived, it also could lead to greater difficulties if there should be any medico-legal issues about lesion perception/diagnosis. For this reason, when our department had the opportunity to implement lossy compression prior to diagnostic review, we did so. In that case, no one ever saw the original images—the same lossy compressed image is interpreted by the radiologist, distributed to the referring physician, and archived for any later review. Of course, this was done only after a large-scale review of the compression algorithm that was implemented.

In these days when containment of medical costs is receiving much attention, we should not be reluctant to make a carefully considered, thoroughly studied decision to reduce costs by implementing lossy compression.

#### SETTING THE STANDARD

If one agrees that using lossy compression can be an acceptable decision, the next question is how to determine what is good enough. As noted in the technical description of lossy algorithms, nearly all allow the user to select at least the range of compression that is achieved. That is really the challenge of compression: how to maximize the compression ratio without degrading the image. At least 4 methodologies have been used to answer this question.

The first method is to use human visual ratings. In this case, one simply asks radiologists either: (1) Are the compressed images good enough? or (2) Is there any difference between compressed and uncompressed images? The first question is a bit more difficult to justify, as it involves a judgment about what is required for diagnosis. In many cases, the abnormalities are so obvious that even marked image degradation could still permit a diagnosis. This is usually how x-ray exposure is selected. An alternative that seems more acceptable is to see whether there is any perceptible difference between compressed and uncompressed images. If one cannot tell the difference, it is unlikely to make a diagnostic difference. This type of study is easy to execute, since one can select randomly from all images and then present pairs of images (the original and the compressed/decompressed) to a blinded rater. If the rater prefers the compressed images at the same rate (or higher rate due to de-noising effect), then the compression ratio is acceptable.



The second method seeks to prove that compressed images are diagnostically equivalent to uncompressed. To do this, one must have proof of what the diagnosis is for each case. Furthermore, one should really have a range of diseases that would be expected to be diagnosed with that type of image. Ideally, the gold standard for knowing that disease is present should be established by an independent method, since we know that compression can improve performance. If an independent standard were not used, all the correct diagnoses made possible by improved signal-to-noise ratio would be counted as misses. And finally, you must have enough of each type of disease to prove a difference. Note that failing to find a difference is not the same as proving equivalence. The latter is much more demanding statistically, while failing to find a difference is a certainty if you have only a few cases.

One example of this methodology is a study we did with chest radiographs, in which we selected 2 common pathologies: nodules and fibrosis. We selected only 2 diseases to keep the receiver operating characteristic (ROC) design simple. We also selected them because they had very different properties—uncalcified nodules are low-frequency findings, while fibrosis is a high-frequency finding. Chest CT was used as the gold standard. Mammography can employ the same technique—masses versus microcalcifications, with biopsy or clinical follow-up being the gold standard.

The third method is to use computer-aided diagnosis (CAD) output as a detector for diagnostic degradation. This has the advantage of allowing many cases to be assessed with less effort. In addition, the variability of a human observer is removed. The major concern is that the features important to a CAD algorithm may not be the same for a human. Furthermore, CAD algorithms are generally focused on detecting a single disease (e.g., cancer). In cases like mammography, where cancer is nearly the entire disease range, this is not a problem, but for nearly all other modalities, preservation of CAD performance ensures only that features important to that CAD algorithm are preserved, not necessarily all the features necessary for a radiologic interpretation.

The fourth evaluation method is also computer-based but uses a model of the human visual system to predict when lossy compression has produced alterations that would be visible to a human. Like method 3, this method has the advantage of allowing many cases to be evaluated with little human effort. It has the additional advantage over the CAD method in that it should be valid for any type of image or disease. However, this method is complex to implement—one must have a valid model of the human visual system, which, in turn, requires knowledge of the output device. The computations

for this are also nontrivial. Nevertheless, there are now some reports that are using this method.

## COMPRESSION APPLICATIONS IN MEDICAL IMAGING

Clearly, the goal of medical image compression is to reduce the amount of data that must be stored or transmitted. Depending on the application, either lossy or lossless compression may be the best choice. Rather than review applications using this division, it is probably more appropriate to evaluate compression for each modality. This is because the compressibility of images is best correlated with modality.

### MAMMOGRAPHY

Mammography is an area that has received much attention from the compression industry. That is because the data sets are large (40 megabytes [MB] per image  $\times$  4 images per screening exam), there is a large exam volume because mammography is used for screening a large at-risk population, and expertise is often not distributed to areas that acquire the images. This combination of factors means that moving the images from the acquisition location to a distant location where there is expertise in interpreting mammograms is often necessary.

One factor that has made analysis of mammogram compression easier is that there is really only one disease process of interest: breast cancer. However, breast cancer can have any of 3 types of appearance. The first is that of a small area of tiny calcifications, also known as microcalcifications. The second appearance is of a mass, which will often have irregular borders. The third is distortion of the architecture of the tissue within the breast. Because these 3 appearances can be fairly easily characterized, the analysis of compression effects is simpler than for other image types.

Early mammography compression studies used digitized films and either JPEG or a wavelet variant. Because early algorithms tended to lose high-frequency information first, most attention was paid to the effect on microcalcifications. Studies from the 1990s showed that JPEG compression could be applied at levels up to 15:1 without producing perceptible changes in the images, as assessed by image-processing experts. One should note that this study did not use radiologists to evaluate the images. This is often a problem—to do a good study, one must have many cases with subtle disease

with an independent gold standard, and these many cases should be evaluated by radiologists. Just compiling a large set of subtle cases with proof of diagnosis is difficult. Getting several experts to review these cases is more difficult. And since the cases selected are subtle, there will be significant variability in the ratings—variability that will obscure any subtle difference caused by compression.

For these reasons, some have begun using computer-aided diagnosis (CAD) as a measure of acceptable compression. This is an elegant way to avoid the variability inherent in human observers, and to evaluate the large number of cases required to detect small differences.

Early studies, of necessity, used digitized film mammograms. Digital detectors are beginning to appear on the market. It is possible that the image properties are different, and so it is necessary to perform separate studies on these images to determine the correct compression ratio. But at the same time, compression algorithms, such as JPEG2000, are advancing. JPEG2000 is receiving much attention because it has been adopted into the DICOM standard and because it holds promise for better compression performance than standard JPEG. One recent study using an alternative forced-choice method found that digital mammograms compressed with JPEG2000 at ratios up to 20:1 were indistinguishable from the originals, which is similar to results for digitized films.

## COMPUTED RADIOGRAPHY/DIGITAL RADIOGRAPHY

Compression of radiographic images also has a long history, beginning with study of digitized radiographs but now focusing on computed radiography (CR) and digital radiography (DR) images. It is encouraging that the results seem similar: digitized radiographs seem to be about as compressible as CR and DR. For chest radiographs, ratios in the range of 20:1 seem to produce no visible or diagnostic degradation. Slone found that with very close inspection/magnification, images compressed at 10:1 using JPEG are still indistinguishable from originals, and at normal viewing conditions, 20:1 is equivalent. Compression at ratios of up to 32:1 (either JPEG or wavelet) did not degrade detection of simulated nodules on chest phantom images. We found that for either nodule diagnosis or interstitial disease detection, wavelet compression of digitized chest X-rays (CXRs) at up to 40:1 was not significantly different for nodules and for interstitial disease, and performance at 10:1 showed a trend to be superior to original images. There are fewer studies of musculoskeletal radiographs, but those that exist also show that compression in the range of 20:1 does not alter visual appearance.

## COMPUTED TOMOGRAPHY AND MAGNETIC RESONANCE IMAGES

Images with a smaller matrix, such as CT and MRI, might be expected to be less compressible because there is less potential for redundancy; this seems to be the case. Whereas radiography seems to tolerate ratios in the 20:1 range, acceptable compression ratios for CT and MR seem to be closer to 10:1. There is also very little difference between compression methods, likely because there is less redundancy for full-frame methods such as JPEG2000 and wavelet to leverage. In a study of wavelet compression, ratios of 8:1 did not affect accuracy of diagnosis of acute appendicitis, but 16:1 and 24:1 did show decreased sensitivity. Another study showed no change in nodule detection on low-dose chest CT at 10:1 utilizing wavelet compression; 10:1 was also found equivalent for both JPEG and wavelet for detection rate for lung cancers on low-dose chest CT. Brain MRIs compressed using a wavelet algorithm at up to 20:1 showed no difference in ROC value for variety of lesions for a  $512 \times 256$  matrix and 10:1 for a  $256 \times 256$  matrix. This reflects rather nicely that the higher matrix had little additional information. JPEG compression of head CTs at up to 20:1 did not degrade ROC performance for detecting infarction, though infarction is a low-frequency finding that may be more compression tolerant.

While it would be nice to have a single ratio for a modality, this is not the case. Conventional chest CTs compressed at 6:1 (JPEG) were considered acceptable, but only 4:1 was acceptable for thin section (2 mm) CT. This is an interesting finding that despite the fact that compression seems to discard noise preferentially, noisy images are less compressible. It demonstrates the importance of redundancy and texture versus information in visual appearance. This is an important problem because of the rapid expansion of multidetector CT. It is fairly simple to create very thin images as well as thicker images for a given body part. Studies comprising thousands of images could become routine. The fact that the greater noise in these thin sections dominates the redundancy is significant—not only are there more slices to store and transmit, but they are also less compressible. Since potentially multiple datasets are derived from the same projection data, some have proposed that it may be more efficient to compress the projection images than to compress the reconstructed sections.

## FLUOROSCOPY

The largest systematic study of fluoroscopic compression, and probably of any medical imaging modality, is for cardiac angiograms. This large, multi-

site trial, which focused on the visibility and appearance of coronary artery stenosis, showed that JPEG compression at 6:1 was equivalent to the original images, with some degradation of quality and performance at 10:1. It is reasonable to expect that other fluoroscopic images will exhibit similar compressibility.

## NUCLEAR MEDICINE

Nuclear medicine images also comprise a wide range of image types. Some are very low resolution ( $64 \times 64$ ), while others are similar to CT. We have found the small matrix images to be rather incompressible, and in those cases, we simply use lossless compression. It is similar to the text example above—there are so few samples to work with that finding redundancy is difficult. For 256 and up matrix images, we have found compression ratios similar to CT and MRI (10:1) to be acceptable.

## ULTRASOUND (STATIC AND VIDEO)

Diagnostic ultrasound actually produces images of several types: static grayscale images that may be captured video with 256 gray levels, static grayscale images with 1000 to 4000 gray levels, static color images with 24 bits of color, and real-time or video signals. Despite this wide variety of image types, all seem to allow compression on the order of 10:1, and this also seems to transcend compression algorithm. Some have also demonstrated no alteration in automated intimal wall thickness from ultrasound images after compression.

## THE FUTURE

Information theory can measure how much unique information is present in an image. That should be the upper limit on lossless image compression ratios. However, there are (fortunately) many clever people who find that those limits can be broken by using additional information about the images to cheat the limits. One familiar example from CT is that the reconstruction typically is circular, while the image is square. Rather than encoding the information outside the circle, one could (and some scanner manufacturers did) never send any information for pixels outside the circle, thus improving the compression ratio. This is obviously an extreme case, but other infor-

mation can be used to effectively reduce the data size without altering pixels of interest. For instance, CT data is viewed at narrow window widths over only a limited range of window levels. By altering the histogram, one can improve compression ratios without producing any perceptible image degradation *as normally viewed*.

## THREE-DIMENSIONAL (3-D) IMAGING

Three-dimensional imaging modalities, such as helical CT, MR, and ultrasound (US), are those that are experiencing the greatest growth in data volumes. Since they can also produce 3-D data, improved compression rates may be achievable by utilizing methods that take advantage of the coherence that is through-plane, as well as in-plane. While early results with 3-D compression were disappointing, more recent attempts with thinner section data show more promise. Figure 12.3 shows a comparison of 2-D versus a videolike 3-D compressor versus a true 3-D spatial compressor. It is likely that this improved performance for 3-D in the later study is because the thinner sections possible with multidetector computed tomography (MDCT) allow more coherence of data between slices. However, the results continue to fall below expectations, likely because these thinner slices also have more noise, which decreases compressibility.




FIGURE 12.3

Comparison of JPEG2000 compression methods. The original image is the left-most panel. The next image is after 2-D JPEG2000 compression has been applied. The third uses a videolike compressor, and the fourth is 3-D JPEG2000 compression. All compressed images have a final compression ratio of 100:1. Note the marked degradation of the 2-D compressed image and the near-perfect 3-D compressed image. (Images courtesy of Michael Marcellin, University of Arizona.)

## RAW DATA COMPRESSION

For modalities like CT and MR, the images that are reconstructed are only a small subset of the possible images that could be produced. Given this flexibility and the different personal preferences in how to view these increasingly complex image sets, it may be most effective simply to compress and store the raw data and the reconstruction parameters than to compress and store all the image products derived from the raw data. In the case of CT, this would mean compressing the projection data, while in MRI, it would mean compressing the k-space data. While there was some early investigation of these methods, recent trends may revive these techniques. While these examples are for CT and MRI, it is possible and even likely that multiple-image products will be created from a single source/raw dataset. For that reason, being able to compress the raw data, and describe (compactly) how the derivatives were made, will likely be more effective than storing all the derivatives.

## HUMAN VISUAL SYSTEM–BASED COMPRESSION

A crucial element in the decision to use lossy compression is that there are no perceptible changes in the image, or that the changes are so small that they are of no diagnostic significance. Several studies have been done using the latter criteria, but they are somewhat less satisfying because there is always the nagging doubt of whether a particularly challenging case just might be compromised by compression. So-called visually lossless compression is more appealing—if the radiologist cannot see the difference, it is hard to imagine that a difference in diagnostic performance could exist. Operating on that assumption, there are some efforts under way to use models of the human visual system (HVS) to determine the threshold at which compression-related image alterations become perceptible. This means that each image is examined and its optimal compression ratio is computed and applied. This is probably more valid than general modality rules for acceptable compression ratios. Anyone who has applied lossy compression knows that there are always exceptional images that are very incompressible, for which general rules do not apply. These exceptional images would not be overcompressed with an HVS-based compression system. While an HVS system would be more expensive because of the need to compute the optimal ratio for each image, it might also save money by allowing a higher average compression ratio. Today, conservative ratios are applied to make sure most images are good enough. With HVS, *every* image would be good enough,

and those that are very compressible would not be *undercompressed*. Work on HVS-based compression is still preliminary, but such a system is a reasonable possibility and might address concerns that exist within the imaging community.

This chapter introduced compression by pointing out that it was controversial. But once computer algorithms are used to interpret images, there will be a quantitative measure of required image characteristics that will permit precise optimization of compression. This is directly analogous to modeling the HVS to determine acceptable compression. The difference is that HVS models are based on estimates of humans and specific viewing conditions—these may not match actual viewers and viewing conditions. With CAD, such variability is not present, enabling more confident statements that compression is not affecting image interpretation. Some may still argue that at some point, computers and networks will be inexpensive enough to render compression needless. The fact that this argument has been made for decades, and since Moore's law specifies an 18-month doubling time, I doubt this will be true in the foreseeable future. Indeed, the forces driving computer technology also drive medical imaging devices. Compression is a technology that we will have to grapple with for the foreseeable future. Much like x-rays themselves, compression can be applied to advantage as long as those employing it understand its strengths and weaknesses.

## REFERENCES

- Bae KT, Whiting BR. CT data storage reduction by means of compressing projection data instead of images: feasibility study. *Radiology*. 2001;219(3):850–855.
- Brennecke R, et al. American College of Cardiology/European Society of Cardiology international study of angiographic data compression phase III. Measurement of image quality differences at varying levels of data compression. *Eur Heart J*. 2000;21(8):687–696.
- Chan H, et al. Image compression in digital mammography: effects on computerized detection of subtle microcalcifications. *Med Phys*. 1996;23(8):1325–1336.
- Good W, Maitz G, Gur D. Joint photographic experts group (JPEG) compatible data compression of mammograms. *J Digit Imaging*. 1994;7(3):123–132.
- Erickson BJ, et al. Requirements for an enterprise digital image archive. *J Digit Imaging*. 2001;14(2):72–82.
- Hangiandreou NJ, et al. The effects of irreversible JPEG compression on an automated algorithm for measuring carotid artery intima-media thickness from ultrasound images. *J Digit Imaging*. 2002;15(suppl 1):258–260.



- Huang H. Progress in image processing technology related to radiological sciences: a five-year review. *Comput Methods Programs Biomed.* 1987;25(2):143–156.
- Kerensky RA, et al. American College of Cardiology/European Society of Cardiology international study of angiographic data compression phase I. The effects of lossy data compression on recognition of diagnostic features in digital coronary angiography. *Eur Heart J.* 2000;21(8):668–678.
- Ko JP, et al. Wavelet compression of low-dose chest CT data: effect on lung nodule detection. *Radiology.* 2003;228(1):70–75.
- Kotter E, et al. Evaluation of lossy data compression of chest X-rays: a receiver operating characteristic study. *Invest Radiol.* 2003;38(5):243–249.
- Li F, et al. Effects of JPEG and wavelet compression of spiral low-dose CT images on detection of small lung cancers. *Acta Radiol.* 2001;42(2):156–160.
- Manduca A, et al. Histogram transformation for improved compression of CT images. In *Medical Imaging 1997*. Newport Beach, CA: SPIE; 1997.
- Manduca A, et al. 3-D compression of medical images with Set Partitioning in Hierarchical Trees. In *RSNA*. Chicago: RSNA; 1997.
- Megibow AJ, et al. Computed tomography diagnosis utilizing compressed image data: an ROC analysis using acute appendicitis as a model. *J Digit Imaging.* 2002;15(2):84–90.
- Ohgiya Y, et al. Acute cerebral infarction: effect of JPEG compression on detection at CT. *Radiology.* 2003;227(1):124–127.
- Persons KR, et al. Evaluation of irreversible JPEG compression for a clinical ultrasound practice. *J Digit Imaging.* 2002;15(1):15–21.
- Ritenour ER. Lossy compression should not be used in certain imaging applications such as chest radiography. For the proposition. *Med Phys.* 1999;26(9):1773–1774.
- Said A, Pearlman W. A new fast and efficient codec based on set partitioning in hierarchical trees. *IEEE Trans Trans Circuits and Systems for Video Technology.* 1996; 6:243–250.
- Savcenko V, et al. Detection of subtle abnormalities on chest radiographs after irreversible compression. *Radiology.* 1998;206(3):609–616.
- Siddiqui K, et al. Improved image compression at various slice thicknesses for multi-slice CT using 3-D JPEG2000 (part 2) in comparison with conventional 2-D compression. In *Soc Computer Applications in Radiology*. Vancouver, BC: SCAR; 2004.
- Slone RM, Muka E, Pilgram TK. Irreversible JPEG compression of digital chest radiographs for primary interpretation: assessment of visually lossless threshold. *Radiology.* 2003;228(2):425–429.
- Sung MM, et al. Clinical evaluation of compression ratios using JPEG2000 on computed radiography chest images. *J Digit Imaging.* 2002;15(2):78–83.
- Suryanarayanan S, et al. A perceptual evaluation of JPEG 2000 image compression for digital mammography: contrast-detail characteristics. *J Digit Imaging.* 2004; 17(1):64–70.

- Tuinenburg JC, et al. American College of Cardiology/European Society of Cardiology international study of angiographic data compression phase II. The effects of varying JPEG data compression levels on the quantitative assessment of the degree of stenosis in digital coronary angiography. *Eur Heart J*. 2000;21(8):679–686.
- Yamamoto S, et al. Evaluation of compressed lung CT image quality using quantitative analysis. *Radiat Med*. 2001;19(6):321–329.
- Zhang Y, Pham B, Eckstein M. Automated optimization of JPEG 2000 encoder options based on model observer performance for detecting variable signals in X-ray coronary angiograms. *IEEE Trans Med Imaging*. 2004;23(4):459–474.
- Zheng B, et al. Applying computer-assisted detection schemes to digitized mammograms after JPEG data compression: an assessment. *Acad Radiol*. 2000;7(8):595–602.

# PACS ARCHITECTURE

---

---

KENNETH HECKMAN • THOMAS J. SCHULTZ

**O**ver the past several years, picture archiving and communication system (PACS) architecture has evolved to incorporate the advances that have occurred in computer, imaging, and networking technologies. These technological advances have resulted in increased server processing speeds, faster networks, decreased costs of storage and hardware, increased reliability of hardware, and an increase in image size as well as in the number of images in a study. Because of all these factors, it is increasingly important for PACS professionals to have an understanding of the evolving technologies and the impact they have on the PACS environment.

In this chapter we will discuss PACS architecture, storage models, historic and current trends, disaster recovery, and medical-legal archiving. The medical-legal archive topic alone could consume an entire book. Legal issues surrounding the archive will not be addressed; instead this section will discuss some of the most frequently employed technologies and strategies.

## WHAT IS PACS ARCHITECTURE?

Architecture is a general term that can be applied to a wide variety of situations. In general, it is used to define the construction of an entity. In computer systems it is a logical and functional description of the hardware configuration of the system. It is very important to carefully plan the PACS core architecture; both the current and foreseeable future goals of the institution must be accounted for. This is often difficult because both the organization and the technology are never in a static state.

### KEY TERM DEFINITIONS

To discuss the architecture, it is crucial to use common definitions of the PACS components. The following section describes and defines the most relevant PACS components (Figure 13.1).

#### CORE SERVERS

Core servers (A through G in Figure 13.1 plus the database) are the computers that are integral to the functioning of the PACS. Depending on the architecture and software, the core could consist of one or more servers. The server's roles include, but are not limited to, database, Digital Imaging and Communications in Medicine (DICOM; gateway) exam import/export, radiology information system (RIS) interfacing, storage systems, Web servers, and other interface/image distribution servers.

#### DATABASE

The database is the component of the PACS system that is used to store all the information required for the PACS to operate. It contains all the specific study information, that is, patient demographics, prior reports, annotations and notes, archive (exam storage) locations, exam routing rules, and similar types of information.

#### HEALTH LEVEL SEVEN SERVER

The Health Level Seven (HL7) server (A in Figure 13.1) is responsible for receiving the data from the (RIS) and populating this data into the PACS

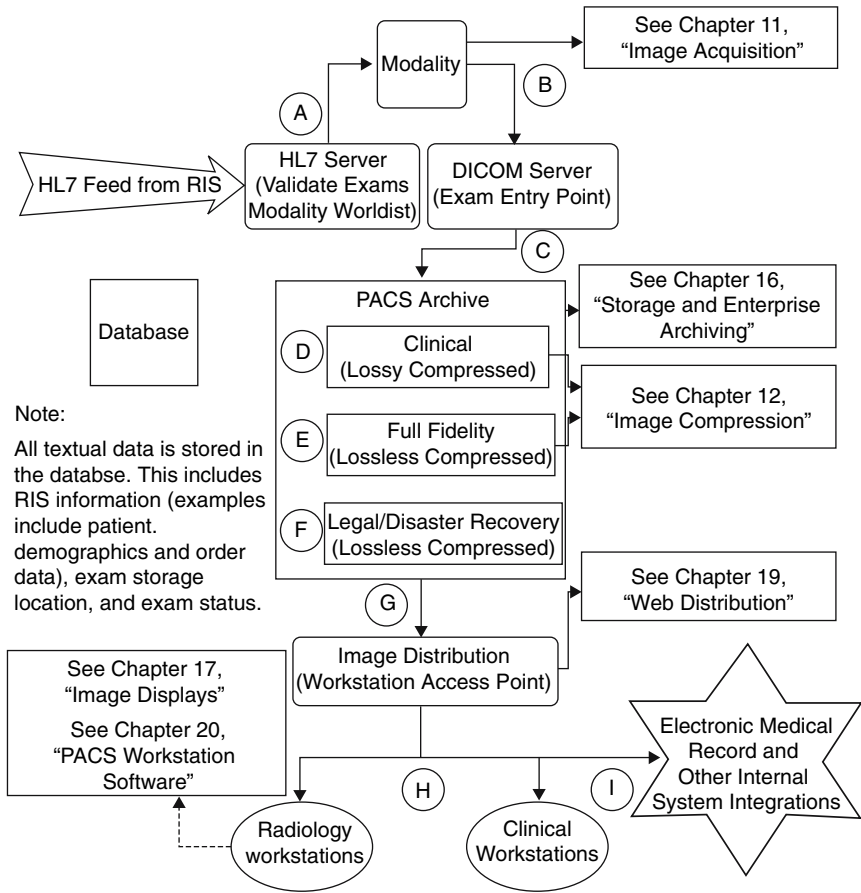


FIGURE 13.1

Basic PACS components.

database. HL7 is the interface standard for data transfer between hospital information systems.

### MODALITY WORKLIST

The modality worklist (MWL) function (shown as A and B in Figure 13.1 and defined in the DICOM standard) allows the technologist at modalities to query the PACS or workflow server. The returned data automatically populates the appropriate fields at the modality. Because the data is electroni-

cally entered, the integrity of the demographic data transferred with the images is significantly higher than that of data entered manually.

#### DIGITAL IMAGING AND COMMUNICATIONS IN MEDICINE SERVER

The role of the DICOM server (C in Figure 13.1) is to accept DICOM-compliant images sent from the modalities. The process usually begins by validating site-specific key demographics contained in the DICOM header with the HL7 populated data in the PACS database (recall that the RIS is continually transmitting information to PACS via HL7). Once the study has been validated (demographic data contained in the image data match those of the RIS), the textual data (modality-specific data) are imported into the database and the images are compressed (in lossy and/or lossless fashion, depending on policies) and stored in the archive(s). If routing (automatic forwarding to another DICOM system) is implemented, the exam is then forwarded to the foreign host.

#### PACS ARCHIVE

Archiving is discussed in depth in Chapter 16, but a general understanding of the terminology is essential for a thorough discussion of PACS architecture. The archive (D through F in Figure 13.1) is defined as the long-term storage system for PACS. The archive is generally split into “full fidelity” and “clinical” or “lossy compressed,” also known as the clinical and disaster recovery (DR) archives. The DR archive usually employs lossless compression, which has very low compression ratios (typically 2:1), but preserves all the image quality (no image quality lost due to compression). The clinical archive usually offers significantly higher compression ratios (10:1 or more) and is generally used for clinical review. Lossy compression is typically achieved by processing the images through a spatial to frequency domain conversion (JPG2000). During the conversion process, an algorithm determines which data are most insignificant and throws those data away—the higher the compression ratio, the greater the loss of image quality. A highly compressed clinical archive allows fast retrieval of prior exams (smaller files lessen the bandwidth demands on the network). Even more beneficial, the amount of storage (redundant array of inexpensive disks [RAID]) required to keep all priors online is magnitudes less than the amount that would

otherwise be required. It is also important to note that there are important legal implications for storing medical exams that vary from state to state.

## IMAGE DISTRIBUTION SERVERS

The image distribution servers (G in Figure 13.1) are the servers that interface the workstations and other display devices with the PACS. This is accomplished either by routing (automatically pushing exams to the workstations via polices) or by the workstation pulling the images on demand. The on-demand method is usually implemented via Web servers and is the preferred option for most new PACS offerings.

## WORKSTATIONS

Workstations (H and I in Figure 13.1) are the computers with which end users (radiologists, technologists, and referring providers) view and interpret radiographic images. These workstations generally provide different functionality depending on the need of the user.

In general, workstation software falls into two categories: thin client and thick client. The thin client is usually a Web (standards-based) application that has minimal hardware and software requirements and offers centralized display station installation/upgrade management. The result is workstation software that any user can install without the assistance of the PACS vendor or the institution's PACS team. The thick client generally requires a larger amount of workstation power, storage, and configuration management. The thick client display station software is usually built on proprietary software, and industry standards are not leveraged. This results in complex installation/maintenance procedures often requiring that the PACS vendor install and maintain every workstation. (In some instances the institution's PACS team can be trained by the PACS vendor to install/maintain the workstations, but this is usually not the case.)

## STORAGE

Three models are generally used for online spinning storage (RAID) (D through F in Figure 13.1). They are direct attached storage (DAS), storage area network (SAN), and network attached storage (NAS). The details of the storage types are thoroughly discussed in Chapter 16.

The intention of the following is to give a basic understanding of the 3 storage strategies and how they relate to PACS architecture.

**DIRECT ATTACHED STORAGE** Direct attached storage is a storage device that is directly attached to a server. The server that is attached to the storage device is the only server that can directly utilize that storage.

**STORAGE AREA NETWORK** Storage area network is a high-speed, dedicated (usually fiber-optic) network of shared storage devices. The SAN architecture works in a way that makes all storage available to all servers via a dedicated network—not the same network used for e-mail and Web browsing. Scaling storage in a SAN is accomplished by simply adding more storage devices to the storage network. Standards are emerging, but most solutions are usually proprietary.

**NETWORK ATTACHED STORAGE** Network attached storage is a dedicated computer that provides file (disk) sharing and is often referred to as a NAS head. A key differentiator between NAS and SAN is that the NAS head is shared and accessed across the same network as the institution's e-mail and business applications; a dedicated storage network is usually not required. The underlying storage infrastructure behind the NAS head (computer connected to the storage offering file sharing) may be DAS or SAN (the NAS heads may participate in a storage network). Scalability is achieved either by adding more NAS heads or by connecting more storage to existing NAS heads. In this environment, commodity and proprietary solutions can coexist.

## HIERARCHICAL STORAGE MANAGEMENT OR LIFE CYCLE MANAGEMENT

Hierarchical storage management (HSM), also referred to as life cycle management, is an application that automatically moves data among different storage systems (E and F in Figure 13.3). The concept is to move infrequently used data to slower (cheaper) tiers of storage. Some examples include undictated exams on fast disk (RAID), relevant prior exams on slower disk (RAID), and exams not reviewed in 18 months on tape (CD/DVD). Some solutions move data among geographical locations but keep all data on spinning disk (life cycle strategies are discussed later in this chapter under “Medical-Legal Archive”). There are only two practical tiers of storage specific to PACS: disk (RAID) and removable media. Contrary to



traditional belief that storage tiers should also exist for disk (fast, expensive disk for unread studies and slow, inexpensive disk for prior exams), our experience has shown that having tiers of RAID (fast and slow) yields immeasurable performance differences. Regardless of the selected approach, the role of the HSM is to prevent catastrophic data loss due to hardware, software, or human error.

## CLUSTERING

Clustering is the binding of two or more computers together in such a way that the application behaves as if it were installed on a single computer. The primary reasons for clustering are load balancing and fault tolerance. Load balancing spreads the requests for application services across all participating servers. For example, if a researcher submits a request for 1000 exams, 1 server could satisfy the request while another server is refreshing a workflow screen for a radiologist. Fault tolerance refers to the ability of a system to handle hardware or software failure while continuing to behave as if there has been no failure from the perspective of the end user. There are multiple degrees of fault tolerance, which can be as simple as planning for power interruptions or as involved as being able to sustain a multiserver failure. A common fault-tolerant implementation of clustering is referred to as fail-over. A fail-over cluster refers to the practice of mirroring a system so that every function that a server performs can be performed on a duplicate (backup) system. Therefore, when the primary system fails, the backup system automatically takes control and satisfies all incoming requests.

## ELEMENTS OF THE ARCHITECTURE

Several key topics need to be evaluated to determine an appropriate architecture for an institution. It is important that a multidisciplinary team be involved in designing the technical and operational components of the PACS. Depending on the organizational structure, the team should include the following: general information system (IS)/information technology (IT) staff, networking specialists, radiologists, and radiology clinical operations staff. Once the team is assembled, at least the following should be considered: operational fit (workflow), versatility, scalability, robustness, and fault tolerance.

## OPERATIONAL FIT

Operational fit refers to how the workflow that is built into the PACS product will work with the current workflow requirements of the institution. A PACS implementation will, by its very nature, require many workflow changes by the radiology department. However, each vendor has a different way to deal with these changes, and the impact on radiologist, technologists, and administrative/support staff needs to be very carefully evaluated. An excellent PACS product may be worthless if implemented in an institution whose workflow cannot not be modified to fit how the selected PACS product operates.

## VERSATILITY

Versatility, in this context, defines a system's ability to incorporate new technologies and adapt to internal institutional changes. Technology changes could be anything from the addition of a new storage medium to the implementation of (perhaps fewer) more powerful servers. Operationally, changes could range from new modalities or new modality imaging protocols to a rollout of clinical review stations to referring physicians.

## SCALABILITY

Scalability, in the PACS world, refers to the ability of the PACS software to grow with the institution's goals. This statement refers to software, not hardware. The limitation of a system's ability to grow in terms of volume, depth of storage, or number of supported users and modalities is completely dependent on the PACS provider's design/implementation. Hardware is subservient to the software. For example, if a PACS provider's platform supports 1 server (all PACS operations operate on 1 computer), then the system can only scale by purchasing a faster, more powerful server. If the PACS provider supports the spreading of the PACS processes over multiple servers, then the system can be scaled by adding more servers.

## ROBUSTNESS/FAULT TOLERANCE

Robustness applies to several important elements in the system that allow administrators to provide high-quality services (acceptable uptime from the end user's perspective). This includes items such as self-monitoring capabil-

ities (to e-mailing or paging the PACS administrator when a monitored event occurs), ability to interact with third-party software (antivirus and system backup applications), and resilience to hardware failure (clustering) and virus attacks. A robust system will also allow users to be able to work with the system without understanding the architecture or even knowing when a system fault has occurred.

## COMPARISON OF PACS ARCHITECTURES

PACS architectures typically fall into 2 major categories: distributed and centralized. Older PACS are usually implemented on the distributed model, and newer PACS usually on the centralized model. Each approach has specific strengths and weaknesses, depending on how the PACS vendor implements a design. It is critical that the design philosophy of a selected PACS vendor is fully understood and accepted. Once a PACS is installed and it is discovered that the PACS vendor's and the institution's visions are not aligned, it is usually too late to make a change. Below is a discussion of the earlier design philosophies and how the architecture has evolved.

### HISTORIC DISTRIBUTED MODEL

One of the first successful designs utilized in PACS was the distributed model (Figure 13.2). This model was based on the client-server computing approach in the late 1980s. There were several other factors that influenced this model. At the time network speeds were about 1/100 of what the standard network speed is today, and disk storage was slower and far more expensive. The highest density of an individual disk was about 36 gigabytes (GB); today it is 400 GB and climbing. It is also worth noting that lossy compression was generally not accepted for any purpose. The lack of compression exacerbated the cost of storage and stressed the network. These and other issues resulted in a PACS design that utilized transient storage at the workstations to avoid network bottlenecks and relied excessively on the medical-legal archive because of insufficient disk space within the PACS core. Exams were only temporally cached on disk, resulting in constant retrieves from near-line storage (usually magnetic optical disk [MOD] or digital linear tape [DLT], tape technology).

Generally, the PACS data flow was as follows: the RIS interfaced with the PACS via HL7. The process that arbitrated the communications was often called a broker (A in Figure 13.2). Modalities would interact with the broker system to facilitate modality worklists. Upon receiving images from the modality, the DICOM server would interact with the broker to validate

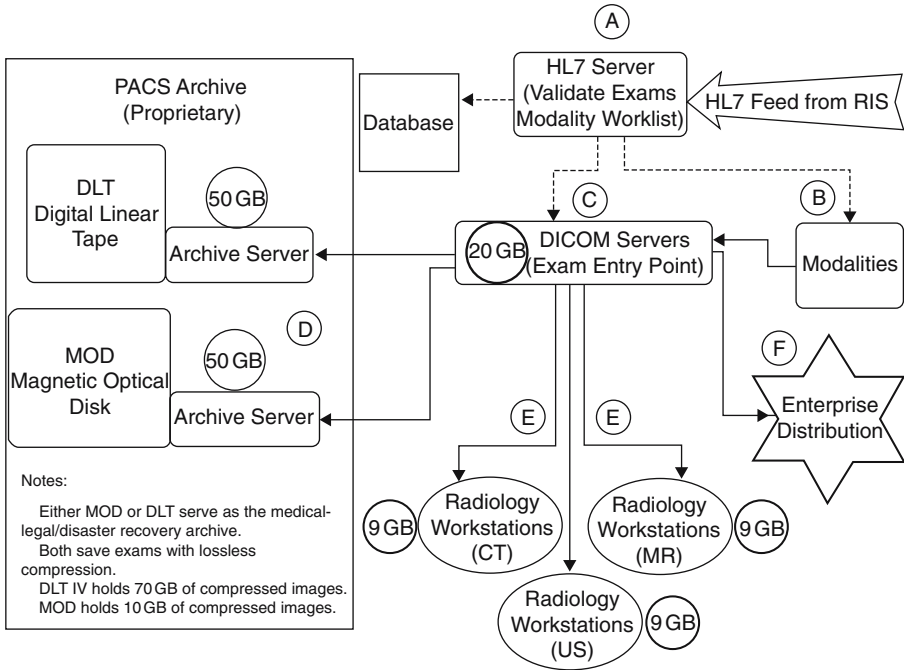


FIGURE 13.2

PACS architecture circa 1995.

the incoming exam's demographic data to that stored in the PACS database (B and C in Figure 13.2). If entries did not match, the exam was made unavailable to the radiologist until a human rectified the differences. Once validated, the exams were then routed to the assigned archives, enterprise distribution systems, and groups of workstations (policies were based on modality/body part) (D, E, and F in Figure 13.2).

Because of the routing model, if a radiologist in the CT/chest group wanted to view an exam that was an MR/head, the radiologist would have had to retrieve the exam to the local workstation because it would have been prerouted only to the CT/chest workstation. The same would be true if the radiologist wanted to view a prior exam that was not considered a relevant prior (previous radiographic exams related to the same body part) and the exam existed only in the archive. It could take minutes (in some cases hours) before the exam was routed to the local workstation. Further, because the workstations were required to store all the images that may be viewed, there was a requirement for the workstation to have an adequate amount of storage

(simple rules flushed exams when the disk began to fill). This often frustrated radiologists who would retrieve a study, view the exam, then, wanting to rereview the exam, had to reretrieve the exam from the archive because the exam no longer resided on the local workstation's disk.

Operationally, the workflow of a department needed to be designed in such a way that organized the radiologist workload by subspecialty. Therefore, because of the requirement of the distributed model to route exams, routing/relevant prior policies were constantly being tweaked. Perfect policies were rarely attained.

## CURRENT CENTRALIZED MODEL

This section will show how the current trend of leveraging standards leads to a simpler, more supportable PACS (Figure 13.3). Many of the constraints (routing, immediate access to all exams, client administration, proprietary standards, and others) have been alleviated and in many cases completely eliminated.

As PACS systems have matured, the goal has become to make all images readily available to all users regardless of what workstation configuration is available (uncontrolled hardware). This has caused a shift in the industry to a more centralized architecture utilizing Internet standards-based technologies.

In a centralized system (Internet standards-based), there are two distinct layers, the core and the workstation. Figure 13.3 shows a generic centralized PACS. In this figure, A through F are all considered part of the core, while G and H are workstations. Usually, the client workstations are based on thin clients leveraging standards and technology available through the operating system (file sharing) and Web browsers (HTML). One of the crucial elements for a centralized PACS is adequate network speed and bandwidth to accommodate the movement of images quickly from one location (archive) to another (workstations). Due to advancements in network technologies during the past 10 years—there have been substantial performance gains and significant reductions in costs—it is very affordable to install 100 megabyte (MB) or even 1000 MB switched network infrastructure (see Chapter 15 for more network information). The same can be said of storage, especially given the growing acceptance of lossy compressed priors. The cost per terabyte (TB) of storage has fallen from as much as \$100,000 per TB to as little as \$4,000 per TB. Consider the actual cost when 10:1 compression is factored in. When RAID costs less and less RAID is required, the institution can realize significant savings.

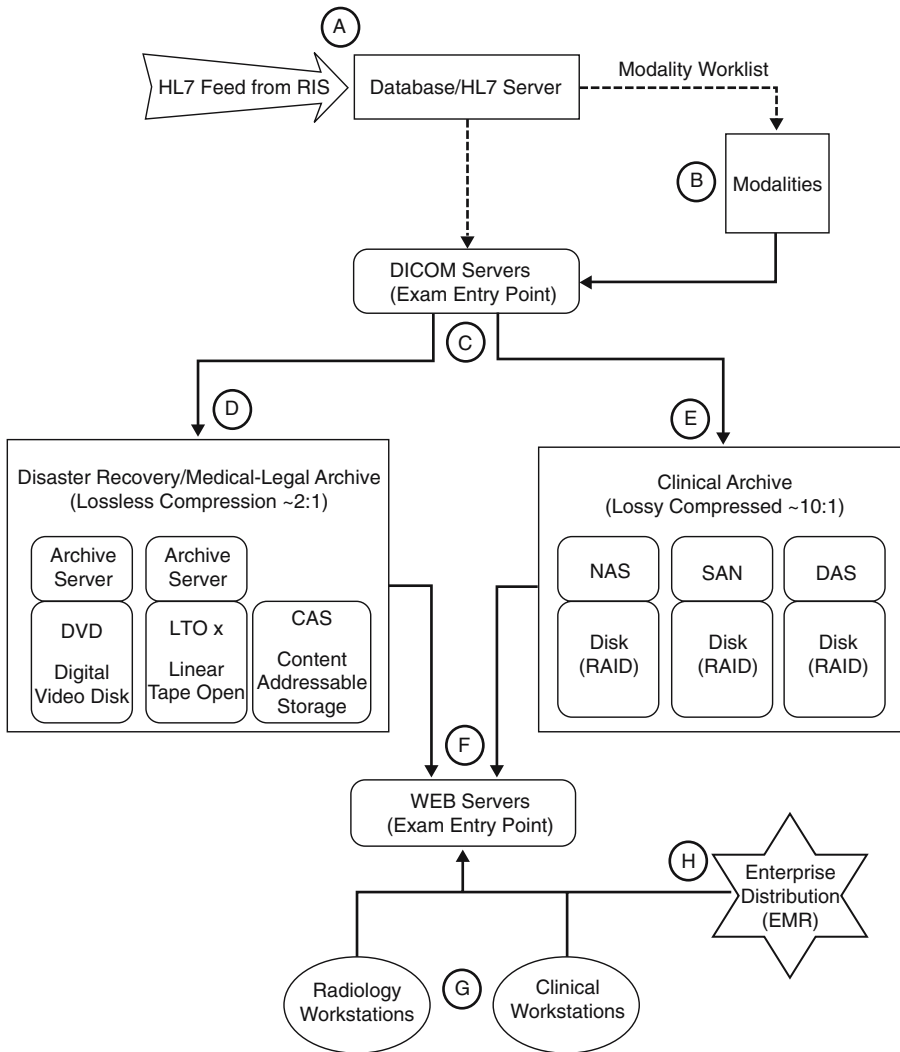


FIGURE 13.3

PACS architecture circa 2005.

In this approach, the data flow is as follows: Demographic data is transferred to the PACS from the RIS via HL7 (A in Figure 13.3). The HL7 process (in most cases) is now considered an integral part of the PACS; it is not referred to (conceptually or physically) as a separate “broker” process.

The imported RIS data is then used for modality worklists and exam validation. After the images are acquired at the modality, the images are transferred via DICOM to the PACS DICOM/image server (C in Figure 13.3). Once the data are received and validated by the DICOM/image server, the images are sent to the archive, where multiple copies, both lossless and lossy, are stored (D and E in Figure 13.3). At that point the images are available via the image distribution servers (F in Figure 13.3) to the various types of PACS workstations and enterprise distribution systems (G and H in Figure 13.3).

Integration of PACS into other hospital applications becomes much simpler when a standards-based PACS is installed. It can be as simple as embedding a Web link in the hospital's electronic medical records system (EMR). The link could work in a process like the following: The user navigates the hospital medical record system to the radiology section and selects an exam to review. Within the report text, the familiar underlined blue text representing a Web link appears. When a user clicks on the link, a browser window opens up, launches the PACS client, and displays the images of the exam being reviewed in the hospital's EMR application.

In this model, no autorouting or pre-fetching of images is required. This completely eliminates the administrative time required to build and maintain routing rules, and the frustration felt by radiologists unable to access any exam at any workstation at any time is also eliminated. Another significant benefit is that the scalability of the system is such that growth (in exam volume and new modalities) can be attained and new technologies added (see Chapters 15 and 16 on storage and server technologies) with no or little impact to the radiologist or clinician. This truly opens the possibility of viewing the PACS provider as a software company sticking with its core competency and staying out of the hardware and storage business.

Still at the center of newer PACS designs is the database. The role of the database has not substantially changed. It is still responsible for storing a wide variety of information about the exams and exam locations.

## COMBINED DISTRIBUTED AND CENTRALIZED MODEL: TRANSITIONAL TECHNOLOGY

Currently several PACS products in the market incorporate elements of both models into their architecture. These products will typically allow for pre-fetching and routing, as well as fast retrieval of images from any location. They also tend to have multiple versions of the workstation that range from a thin client version for use by clinicians outside of radiology and a thick

client version for primary interpretation. While these products address the issues of both distributed and centralized models, they also require the support that is required by both models. They tend to utilize newer technologies that are available and may have a more complex architecture than a pure centralized or decentralized model. These complexities may or may not be more difficult to support and need to be closely evaluated.

## MEDICAL-LEGAL ARCHIVE

The discussions in this chapter, to this point, have focused on the PACS architecture design considerations and definitions and explanations of PACS core components. It is also important for institutions to understand that the installation of PACS is no guarantee that images cannot be lost. In fact, digital archiving solutions are burdened with far greater expectations for data retention and retrieval compared to the film analog world. This section will compare and contrast the operational and technical options available for legal archiving of medical images.

## DISASTER RECOVERY

PACS disasters can and do appear in a variety of forms. To a system administrator, PACS disasters may be thought of as a result of technical failures that impact overall performance. To the radiologist or clinician, a disaster may be perceived as a single exam that cannot be immediately viewed, the technical reasons being irrelevant; in fact, it may simply be due to a poor PACS architecture or implementation strategy. The PACS could merely be fraught with insufficient clinical storage, poor archival policies, or limited RIS integration. The focus of this section will be the prevention of PACS disasters caused by the actual loss of data.

There are two basic strategies for storing exams in “safe” locations: disk-to-disk and removable media. The term “safe” refers to a remote/second location that is unlikely to be affected by the same event(s) that led to the data loss at the primary site.

## DISK TO DISK

“Disk-to-disk” simply refers to the replication of exam data to other disk storage devices ideally located at remote location. Because of the significant



price decreases in disk-based storage devices over the past few years, many disk-only disaster recovery solutions have been penetrating the PACS market. Currently there are 3 general approaches to applying disk-to-disk disaster recovery (Figure 13.4).

**STORAGE AREA NETWORK** In the storage area network (SAN) approach, a storage device is installed at all remote sites where copies of the exams are desired. This also necessitates reliable (usually dedicated) network fiber connectivity between the storage sites (not the same network as the intuition would use for Web/e-mail and other day-to-day business activities). Iron Mountain and other physical document storage vendors are now providing SAN storage devices as part of their offerings. The replication/storage management is provided by the storage devices; external computers are not required to manage the data replication process. Some storage companies offering this type of solution include EMC, IBM, Winchester Systems, and HP. See Chapter 16 for more information on SAN.

**NETWORK ATTACHED STORAGE** The network attached storage (NAS) approach is similar to the SAN approach. The primary difference is that the

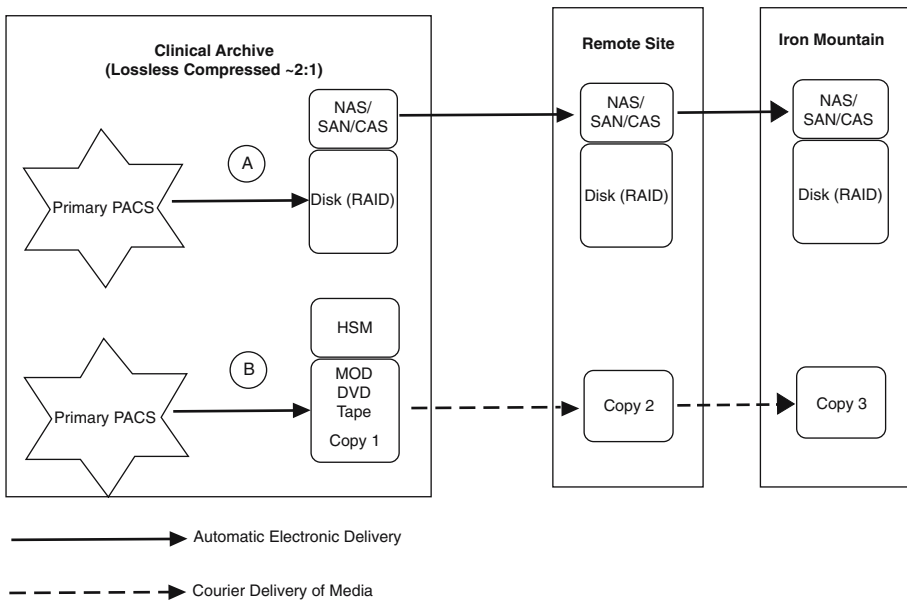


FIGURE 13.4

PACS medical-legal archive examples.

replication process is managed by software running on computers connected to the storage, and the network connecting the devices is not required to be dedicated. The one layer of abstraction of replicating the data at the operating-system level makes a multivendor solution more attainable. For example, if the primary PACS utilizes Windows 2003 on a Dell server connected to 2 TB of disk storage and the remote location has an IBM server running Linux connected to 2 TB of storage, the exam data could be replicated between the two locations via standard file-sharing protocols available to both operating systems. The replication software will leverage these standard protocols to synchronize the two locations. Nearly all storage companies provide NAS solutions. Two companies that specialize in NAS are Network Appliance and Quantum's Snappy Appliances. See Chapter 16 for more information on NAS.

**CONTENT ADDRESSABLE STORAGE** Content addressable storage (CAS) is a scalable storage/computing grid. The idea is to provide scalability and performance by adding self-contained units that act together as a hive. The units can be defined as a small server containing at least 4 disks (~1 TB raw). All the units run the same software, which allows them to act as one. The aggregate storage is the sum of all the units. Some manufacturers use elaborate algorithms to provide a unique key that allows the verification of file integrity and retrieval. This key is often called a token, and the key/integrity generation algorithm commonly employed is message digest 5 (MD5). Some manufacturers hide the token complexities behind a file system, effectively making the grid device look like a NAS. Companies offering grid storage include EMC (Centera) (Hopkinton, MA) and ExaGrid (Westborough, MA). See Chapter 16 for more information on CAS.

## REMOVABLE MEDIA

In this approach, the primary PACS site hosts a computer running software that manages the process of moving exams from disk to removable media and removable media back to disk. The solution usually contains a robotic library that allows for near-line access to large numbers of media. The genre of software that provides this service is called HSM (see "Hierarchical Storage Management or Life Cycle Management" earlier in this chapter). The term is being replaced with "life cycle management". The storage polices available allow for the creation of two or more copies of the media. One copy can stay local to the intuition, another could be sent to a remote site, and a third copy could be sent to a document storage company such as

Iron Mountain. Removable media solutions are the least expensive and easiest to implement. Companies providing this type of solution include StorageTek (ASM) (Louisville, CO), EMC (Legatto), and IBM (Tivoli) (Armonk, NY). See Chapter 16 for more information on removable media.

Table 13.1 is a quick summary of key considerations and how they relate to common archival strategies.

## TECHNIQUE CONSIDERATIONS

We have defined and described basic strategies for archiving, but how does one select the best approach? This is a difficult question to answer. Institutional experience, IT resources, politics, and vendor relationships all play a factor in selecting an archival solution. Below are some general guidelines to consider that are independent of the internal considerations.

**TABLE 13.1**  
Comparison of Approaches for Medical-Legal Archiving

Category	Storage	Network	Content	Removable
	Area	Attached	Addressable	
	Network	Storage	Storage	
Suitable for PACS	*****	*****	*****	*****
Resistant to virus attack	***	**	***	*
Resistant to human activity	***	***	****	*****
Proprietary	*****	**	*****	***
Scalability	*****	*****	*****	*****
Fast to recover exams	*****	****	***	*
Requires human intervention	**	***	**	*****
Multivendor friendly	*	*****	*	*****
Technically simple	**	****	****	*****
Reasonable initial cost	***	****	****	*****
Reasonable maintenance costs	***	***	****	*****

Number of asterisks indicates degree to which solution agrees with category attribute.

## SOLUTION REQUIREMENTS

- ▶ Technically simple
- ▶ Provide copies of exams in multiple locations
- ▶ Low cost
- ▶ Accept some manual processes
- ▶ Accept slow retrieval from off-site locations (physical transportation of media)
- ▶ Provide enough storage on the clinical RAID; if insufficient storage exists on the clinical archive, significant/continuous retrieves of prior exams from the DR archive can cause media failures and affect the radiologist workflow

In this scenario, a removable media solution would probably be the leading candidate. Once the equipment is installed and configured, a human would need only to insert media as required and export filled media for transport to one (or more) safe locations. The cost for multisite storage is only the price of media and transport.

## SOLUTION REQUIREMENTS

- ▶ No manual interaction—complete automation
- ▶ Large budget needed
- ▶ Provide copies of exams in multiple locations
- ▶ Allow fast retrieval from any location (electronic transportation)

This strategy requires deploying a RAID in each location where it is desired to store a copy of each exam. All copies occur automatically with no human involvement once the solution is configured. To sustain this approach, each year disk storage will need to be purchased in equal increments at all locations. Any budget cuts could seriously jeopardize the integrity of the archive; there may be enough funds to store exams at only 1 site instead of the planned 2 or 3. There is also an increased risk that an engineer performing maintenance on a storage device could cause accidental data loss that could be replicated to all other sites. Even worse, it may be possible for a virus attack to affect all copies of the exams because all copies are always online.

## USING DISASTER RECOVERY TO SWITCH VENDORS?

Although the current relationship with the PACS provider may be satisfactory, what would happen if that were to change? What if a new PACS was released to the market that had desired capabilities that the current vendor would not be able to match for years? This topic could consume an entire article—perhaps even an entire book. What is important to consider during the development of an archive strategy is how to prepare for the possible switch to another PACS provider. The answer may be to store all exam data in a nonproprietary format (DICOM standard part 10) on open standard storage systems. Many PACS storage products store data in a format that only their software can retrieve; consider products that are open. This will ensure that your new PACS vendor can read the stored exams.

## RECOVERY AND EXAM PRIORITY

In the end, whatever strategy is selected must provide a means to quickly recover (business continuance) from a local disaster. For DR it is important to consider the pattern of usage in radiology; data (exams) that are stored in PACS are less likely to be retrieved (reviewed) once dictated. This is an important consideration when a disaster hits. For example, assume a water pipe broke near the computer room and RAID devices were lost that contained 4 years' worth of exams. Conventional approaches would be to repopulate the RAID devices from the archive file by file, usually in reverse chronological order. The problem with this approach is that the usage of the exam data is random, not sequential. Therefore, if a radiologist is waiting to display an exam that is 2 years old, the wait could be days or longer. A better approach is to employ a system that can automatically repopulate the PACS by demand. In this system a request to display a missing exam would result in triggering automatic processes that would reload the exam back into PACS. This method would import the most relevant (requested) exams as a priority and populate the unrequested exams as resources permit.

# NETWORKING FUNDAMENTALS

---

---

SCOTT M. ROGALA


**T**he world of networking and the world of picture archiving and communications systems (PACS) are two different environments that converge in such a way as to require a very special skill set. This chapter gives you the basics of computer networking and shows you how they apply to a PACS environment.

It is important to appreciate the role the network plays in a PACS implementation. Our goal is not to make you an expert network engineer, but rather to give you enough information so you can navigate the often confusing, cluttered world of computer networking. With the right kind of information you will feel comfortable enough with the terminology to understand what your vendor(s) are providing, on both the PACS and network sides. We also want to impress on you the importance of good network design and implementation; these are integral parts of the PACS system. Failure to create a strong, robust network infrastructure will result in unhappy users, finger-pointing, and loss of confidence in PACS. If the network is designed and implemented correctly, it can contribute immensely to a successful PACS implementation.

We will start with some basic concepts, intended not for the veteran network engineer but for those who have had no exposure to computer networking. Through extensive use of analogies, most of you will grasp the concepts well enough to see how much thought is needed in the design of networks, and why the necessary investment of time and capital must be made to achieve a successful implementation. We hope that by explaining in familiar terms what is considered wizardry and hocus-pocus will help bridge the gap between network engineers and radiologists.


## FOR REFERENCE: THE INTERNATIONAL ORGANIZATION FOR STANDARDIZATION MODEL

The International Organization for Standardization (ISO) model (see Figures 14.1 and 14.2) was set down by the ISO as a framework to make it easier to construct computer networks from the application (as one views an image) all the way down to the physical layer (i.e., the wires). It defined how networks should interoperate. Note that the ISO model serves only as a guideline, and no network, to our knowledge, is set up exactly to the ISO definition.



**FIGURE 14.1**

The ISO model.



**FIGURE 14.2**


The ISO model with Ethernet and TCIP/IP.

## A SIMPLE NETWORK

Let us start by setting up a very simple network of 2 computers, a server and a client, to illustrate many of the concepts.

In our example, the server machine is an image archive in a small PACS system, and the client machine represents a primary interpretation workstation. For the PACS to work, these 2 computers exchange data with each other, for instance, radiology images. This simple network would look something like Figure 14.3.

What are the components of this architecture? First we will work from the top down, and then we will explain in more detail from the bottom up.



**FIGURE 14.3**

Two computers connected.




## THE ARCHITECTURE AND COMPONENTS AND THE INTERNATIONAL ORGANIZATION FOR STANDARDIZATION MODEL

The archive machine hosts the images the clinical display needs. An application on the archive knows how to answer image requests from clinical displays. In the ISO model this portion of communication happens at the higher levels. The application needs to transfer this information from itself to the clinical display requesting the information. The overall picture would look something like Figure 14.4. Figure 14.4 is an extremely simplified way of looking at the ISO model. It is displayed in this way to emphasize the fundamental components that we will need to understand to effectively make PACS networking decisions.

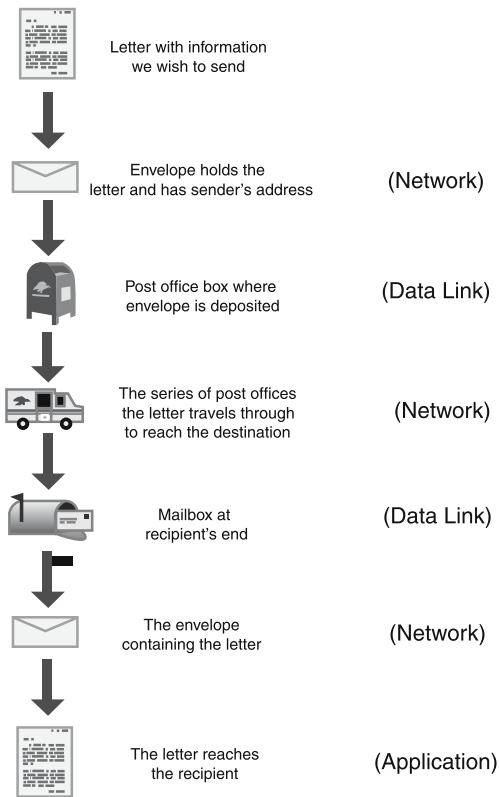
Each layer is interested only in the exchange of information between the layer directly above and that directly below. (For example, the hardware layer generally does not care how the protocol layers pass information to the application: it is concerned only with passing the information up to the protocol layer.) In this way the application communicates with the protocol stack, which in turn hands the information over to the hardware layer (the network interface card, or NIC). Next, the hardware layer puts the information out onto the network.

Each layer has a different and well-defined task. The application layer knows the data it wants to transmit, and it knows which machine it wants to transmit it to. The protocol stack knows how to find the computer in question (or how to find a device that knows how to locate it). The network layer knows how to transmit the data over the network, and the network knows



**FIGURE 14.4**

A simplified ISO model.



**FIGURE 14.5**

Postal system analogy.

how to transmit the data to its destination. On the other end, the information works its way back up until it ultimately reaches the application and the image is viewable.


In many ways, this is analogous to the way the U.S. postal system works. In our scenario it would look something like Figure 14.5.

In this example it is easy to see how each layer is independent of the others. It is not important to the envelope what information is contained in it, only that the information fits. The same goes for the post office box, which does not care what size envelope or even small box is put in it, only that it is not a koala or something else that does not belong in post office boxes. Moreover, the post office box could not care less about the contents of the envelope. As we will see later, the post office system uses the address on the envelope to move the envelope throughout the postal system.

Now we will work our way back up through the layers, starting with the physical layer. The physical layer consists of the wires that make up the network. It may also include the NIC in the computers attached to it. In general, the physical layer can be thought of as the “plumbing” or highway of a network. The quality and width of the pipe can determine the speed at which the data can be transmitted, the reliability of the network, and the distance at which the data can be transmitted.

Most of you have probably seen the cable coming out of the wall plate and connecting to your computer’s NIC. These cables are commonly known as unshielded twisted pair (UTP) copper cabling, also referred to as Category 5 (CAT5) or Category 3 (CAT3), depending on the exact quality. Networks can also be made up of telephone wires, coax cable (otherwise known as thinnet or 10Base5), and other types of wires. One advance that has changed networks dramatically in recent years has been the use of fiber-optic cabling, which can transmit more data over longer distances than conventional cabling by using light or lasers instead of electrical signals. The relatively high cost of fiber cabling has relegated it generally to the core of networks where bandwidth is needed most. Of late, fiber cabling has become the only type of cabling that can support faster transmission rates, and it is finding its way to the desktop as it becomes more popular and less expensive. Later in this chapter, we will discuss which types of cabling are generally used where.

Let us return to our example in Figure 14.3. Our hypothetical network will use CAT5 cabling between the 2 workstations for now. The workstations could be directly connected or, using a device discussed earlier, we could use a hub, to which both devices can be connected, as shown in Figure 14.6. Just like the hub of a wheel, a hub in networking terms is a device



**FIGURE 14.6**

Two computers connected with a hub.

connecting multiple computers. Hubs can have from as few as 4 to as many as 24 or 48 connections for computers.

One layer higher, the hardware (data link) layer consists of the NIC, as well as the software drivers installed on the computer that allow the computer to communicate with the hardware. The most common types of cards are Ethernet cards, which are designed to communicate on Ethernet networks. There are also ATM cards (asynchronous transfer mode, discussed later), token ring cards, and a variety of others. As with all the other layers, direct communication occurs between cards of the same type (i.e., Ethernet to Ethernet, or ATM to ATM). As we get into more complicated network topologies, we will see that networks can become very diverse, consisting of computers with Ethernet cards, ATM cards, and token ring cards, all communicating via the use of various internetworking devices.

For now, to keep it simple, let us say that both computers have Ethernet cards in them. Both computers also have the appropriate software drivers installed.

As we continue to move up our simplified ISO model, we need to discuss protocols. Loosely defined, protocols are a set of predetermined rules that 2 or more parties follow in communication or other interactions. In computer networks, protocols are the set of rules, or languages, that enable computers to communicate over the underlying network infrastructure. Protocols and how they are used are very much akin to languages in the real world: just as humans speak languages such as English, French, and Swahili, networks use languages such as TCP/IP, IPX, and Appletalk. For 2 computers to communicate, they must be speaking the same language. In our scenario we will say that our 2 computers are speaking TCP/IP with each other, thus using a common language to communicate over an Ethernet network on CAT5 cabling.

Last but not least is the application that uses this language to communicate. In this example, that application is PACS.

## COMMUNICATION AND CONVERSATIONS

It is time to delve more deeply into exactly how conversations occur. Once we establish that 2 computers want to exchange data, they must use their shared set of rules (protocols) to do so. Further, these protocols need to communicate over a network. So far we have talked about Ethernet, ATM, and token ring. These are sometimes called network topologies, and they all occur at the data link layer. Certain rules must be followed on these topologies for them to work correctly. We have also identified a few of the more common protocols (or languages), such as TCP/IP and IPX. These

computer communication languages operate at the network layer. Any protocol that is routable is using layer 3, the network layer, to make decisions. Protocols that are not routable are called bridged protocols and are aware of only the data link layer. We discuss this further when we discuss routing and bridging.

To better explain the relationship between a topology, such as Ethernet, and a protocol, such as TCP/IP, we again use the analogy of human communication. When we communicate, we can do so via a wide variety of media: the air (speech), paper (the written word), or hands (sign language). What language we use is completely independent of the medium we use. Similarly, computers can communicate via TCP/IP or IPX and do so over Ethernet or copper wire or token ring or fiber-optic cabling. (Later we discuss devices that do the functional equivalent of taking written English and verbalizing it, or vice versa. You are much less likely to find a network device that does the functional equivalent of translating from English to French, however; such devices are commonly referred to as gateways.)

Since protocols are much like languages, it is useful to use the conversation model to explain how computers communicate. The primary obstacle in using this analogy is that communication for us is second nature; we talk and write without being aware that we are in fact following a set of rules. In the world of computers, where only logic exists and everything is taken literally, conversations are actually very complicated to carry out.

As an example, in human behavior, there are protocols for starting a conversation, such as saying “Hello, how are you?” and responding “Good, and yourself?” These sorts of things occur naturally (assuming everyone is civil!). Similarly, ending a conversation has sets of rules concerning good-byes, and so on. Even during the conversation, it is important not to interrupt the person speaking, and to acknowledge what the person is saying with a nod. If you do not understand what the person is saying, you ask the person to repeat himself or herself. Very similar things happen on computer networks, as we will explain.

## AN EXAMPLE OF A CLASSIC COMPUTER NETWORK—SHARED ETHERNET

To get a better understanding of how computers handle this, let us look at a conversation between our archive computer (server) and our clinical display (client) and postulate how they might communicate. The client knows it wants certain information from the server. Its request makes its way down through the protocol stack using a language that it knows the archive

understands. This message makes its way out onto the wire, heading for the destination machine. The destination, in this case the archive, is listening for anybody looking to talk to it, and sees the packet (piece of data) on the network. The 2 machines begin with their “Hello, how are you?” routine and agree to communicate with certain rules.

They agree, for instance, on how much data the 2 will transmit at any given time, what to call their conversation, and a great number of other things. All the while, since there is a single shared network between the 2 devices, only 1 of the 2 computers can be speaking at any given time. If the client and archive attempt to send information at the same time, both of their transmissions are lost, and they need to start their sentences over. When such a situation arises, it is termed a “collision,” because the electrical signals effectively collide with one another on the network of shared cabling. In these instances, algorithms in the program determine when it is safe to start trying to repeat the message while decreasing the likelihood of another collision.

When only 2 computers are in the equation, the likelihood of a collision is smaller than you might imagine. However, as more and more computers compete for the same network, the likelihood of 2 or more computers trying to transmit at once increases greatly. At some point, when you are building networks with hundreds of computers on them or networks in which computers are constantly trying to communicate, collisions become the norm rather than the exception, and the network cannot bear the burden of all those sentences in the conversation. The sentences, in network terms, are called packets.

What we are describing is called shared Ethernet, in which all the computers share the same network. Even today it is probably the most popular type of network in the world, although newer technologies are becoming more prevalent. Only 1 computer can be transmitting its electrical impulses onto the network at any given time. We like to use the example of a conference room to make the idea of a shared medium a little easier to understand (see Figure 14.7).

Let us say that our conference room holds 200 people. First, we set up a simple rule: anybody can talk to anybody else, but only 1 person in the entire room can be talking at a time. If somebody tries to talk while somebody else is in midsentence, they both have to stop and start their sentences again. One can see very quickly how, in a room full of people operating under these rules, very little is going to get said, especially if a lot of conversations are in progress. In this example, the conference room is analogous to a network segment, the air is the medium (Ethernet), the people are the computers, and the information they are trying to transmit could be PACS



layer

- 4 PACS, Baseball
- 3 English or TCP/IP
- 2 The People
- 1 The Air the Sound Travels Over

**FIGURE 14.7**

The conference room.

knowledge, baseball news, or the latest juicy gossip. Some could be speaking English or French, but only those who understand French are able to communicate with others who understand French. Similarly, on our network, computers could be talking TCP/IP to one another while others are speaking IPX. They all share the same network segment, but they communicate only with those speaking the same language.

## ENTER BRIDGING

You can see in this scenario how inefficient such communications can be. The same was true for early networks. The first attempt to tackle the problem of collision was with a concept called bridging. A bridge is a device that connects 2 or more networks together. (As we see later, a bridge understands only the 2 lower layers of our simplified ISO model. Next we talk about a device that understands the lower 3 layers, a router.)

To explain exactly what a bridge does, let us take the conference room and split it in two. We assign each person a unique number, known as his or her MAC address, or media access control address. The MAC address is unique not only in the conference room, but also throughout the world. In computing, the MAC address is derived from the network interface card. It is how each computer (or in our example, each person) is identified.

In Figure 14.8 you see that the conference room is split in two, with people numbered 1, 2, and 3 on the left side and 4, 5, and 6 on the right side. In the middle of the conference room is a dividing wall, which stops all noise unless the bridge, B, allows it through. It is the bridge's job to keep

track of who is where, and to bridge the traffic through the dividing wall as necessary.

For the bridge to do its job effectively, it listens to all network traffic. It does this by listening for broadcasts. In a network, a broadcast is a method a computer uses to locate another computer when it does not know exactly where to find that computer. It is almost like yelling out, “Joe, where are you?” In the case of a simple shared network segment, the client looking for the archive for the first time sends out a broadcast requesting the MAC address of the archive. The archive answers this request, and the conversation between these 2 computers ensues. In the case of our conference room without the dividing wall, these broadcasts are just like any other packet—they can cause a collision.

A bridge sits between 2 network segments (i.e., in the middle of the conference room) and keeps track of who is on which half. Initially the bridge has no idea who is on which side, so it begins to listen to traffic and build a list. When device 1 makes a request to find device 2, the bridge records that 1 is on its left side and that 2 responded from the left side. This is illustrated in Figure 14.9, and the conference room analogy is carried forward in Figure 14.10.

From this point on, any traffic between 1 and 2 is isolated from the right half, effectively creating 2 different collision domains. Devices 1 and 2 can communicate directly with each other without affecting the right side. Similarly, devices 3 and 4 on the right side can function the same way. When the bridge sees device 4 broadcast to locate device 1, the bridge knows that device 1 is on its left side, and will pass the traffic. In performing this sort of task on what once had been 1 large, congested segment, the bridge can greatly decrease the number of collisions, provided that it is properly located (that is, that proper network design and traffic analysis have been performed). We look more closely at those functions as we design our PACS network.




FIGURE 14.8

The conference room split by a bridge.






FIGURE 14.9

A segmented network with a bridge and the bridge's table.

## ROUTING

Bridging solved some, but not all, of the problems computer networks faced with respect to traffic. Devices still needed to broadcast in order to find their destinations, and in larger networks with many segments bridged together, networks could still collapse under the sheer weight. On networks that comprised hundreds and even thousands of computers, bridging just was not up to the task. Also, although we will not go into detail here, another problem was that bridging did not scale particularly well, meaning that large, complex networks were susceptible to prolonged outages that could be very difficult to troubleshoot.




FIGURE 14.10

The equivalent conference room.

The challenge facing programmers was to create a more intelligent networking device that could limit broadcasts and more efficiently use the bandwidth. Routing is a technology developed in the mid-1980s to address these problems. Routing made use of protocols that had more intelligence built into them so they could pass traffic more effectively throughout the network. Protocols that came along before routing existed had no concept of layer 3 of the ISO model; they simply made use of the MAC address and used broadcasts to find their destinations. Protocols like that are called bridged protocols, and examples are LAT and Netbeui. When routing was introduced, protocols came along such as TCP/IP and IPX, which sent out information in packets that routers could use to more intelligently handle the traffic. These are commonly referred to as routable protocols. Such protocols could limit broadcasts and make it much easier to find devices on a large network.

Again we will use the conference room example, except this time we will expand our world to include the hallway outside the conference room, and add other conference rooms off this hallway (see Figure 14.11).

In this example, each room has a letter assigned to it: A through D. There are folks in each room and now they can be addressed in two ways. MAC addresses still need to be unique, but the protocol address needs to be unique only within that conference room. Addressing at the routable protocol level would look something like Table 14.1.

The benefit is realized when we have many conference rooms, as in Table 14.1. When device A.a wants to communicate with device D.c, it knows that (1) device D.c is not in the same conference room as itself (otherwise it would have the same room letter) and (2) the hallway knows how to find room D. So device A.a sends its information out into the hallway (the router, sometimes referred to as its “default gateway”), and the router relays the conversation to room D, where device D.c receives the information.

Rm 1	Rm 2	Rm 3	Rm 4
a (1) b (2)	a (5) b (6)	a (9) b (10)	a (13) b (14)
c (3) d (4)	c (7) d (8)	c (11) d (12)	c (15) d (16)

In the user 1. a protocol (network) layer address, the “1” is the data link layer unique address.

FIGURE 14.11

Four conference rooms (A through D) with devices in each.

---

---

**TABLE 14.1**  
Conference Rooms and Addresses of Their Devices

<b>Room</b>	<b>Devices' Addresses</b>
A	A.a, A.b, A.c, A.d
B	B.a, B.b, B.c, B.d
C	C.a, C.b, C.c, C.d
D	D.a, D.b, D.c, D.d

---

---

Earlier we mentioned that there are devices that can take the spoken word and convert it to written, and the reverse, as well. Both routers and bridges are capable of doing this, and commonly do. For instance, room C could be a room in which all the devices are communicating via the written word, while rooms A, B, and D are communicating via oral speech. The router would know how to verbalize everything coming out of room C and going into the other rooms, and know how to do the reverse for traffic going into room C.

Routing has been tremendously beneficial in large, complicated environments. Although it adds another layer of complexity and requires significantly more engineering to set up, it goes a long way toward simplifying, troubleshooting, and preventing a network from collapsing during heavy traffic. Unfortunately, routing can slow traffic down, because it has to perform some intensive tasks in order to determine exactly where the packet needs to travel. Keeping with our analogy, the time spent using the hallway to get to another conference room slows the conversation. In most cases the slowdown is negligible, but as we discuss in designing PACS networks, it can be detrimental to PACS performance.

A router is slow because it has to examine more of the packet to make routing decisions. Bridges work much more quickly because they have fewer decisions to make. The nature of networks is such that they are never set up as efficiently as possible—the server you need so you get your information inevitably has 2 or 3 routers in the way. In addition, as bandwidth demands increase, segments need to become smaller, which means more routers, which means slower response time to the server. When routing became too slow, another technology was introduced. Enter switching.

## SWITCHING

Switching, as simple as it is, revolutionized the networking industry. Many networks were collapsing under the sheer weight of routers, and engineers were unable to break up their segments any more than they already had without creating hopelessly complex architectures. The idea of switching is to step back and look at only as much of the packet as a bridge looks at, the data link layer. This means that a switch is essentially a bridge. To a lot of people this is confusing, so let us look at it this way—a switch is basically a bridge with many, many interfaces, and sometimes each interface has only 1 device connected to it. Switches were developed because it became cost-effective to build bridges with numerous interfaces on them. Switches are essentially bridges that are applied in a different manner than original bridges because they are more powerful. As we begin to look at network design and talk about bandwidth, we will see how the idea of switching (and the fact that it came about as higher-speed networks were being developed), when incorporated with routing, helps make network bottlenecks easier to remove.

Switches are used in a variety of ways. Some switches are used in the core of the network (more on this later), whereas other switches are used to replace hubs. When they replace hubs, they have the same effect as changing the rules in the conference room. Recall our first rule of the conference room: anybody can talk to anybody else, but only 1 person in the room can be talking at a time. If more than 1 person speaks we have a collision. To help alleviate the problem, we first installed a bridge in the conference room; then we added conference rooms with a hallway as the router. But that scheme can be applied only so far. At some point we cannot take everybody out of the conference room, and 50 chatty people are going to have some serious collision problems. It does not seem productive to install 25 bridges in the conference room. But when switching came along, it greatly simplified the situation by doing just that. In effect, switching looked at the rules we established concerning talking and changed 1 rule in the conference room—that which was the most inefficient. Switching allows us to say that anybody can talk to anybody else, as long as each person is talking to only 1 other person. Sometimes when a switch is used to separate everybody from everybody else in the conference room, it is referred to as a switching hub.

But what happens when the users in the conference room do not want to talk to each other, but instead want to talk to somebody in another room, that is, through the router? Then we need to talk about bandwidth.

## BANDWIDTH

Bandwidth, simply stated, is the rate at which information can be transmitted in a specific time interval, usually seconds. In computer networks, the rate at which data can be transmitted is measured in bits per second. A bit, or binary digit, is either a 0 or a 1. A bit can be considered the atom of the computer world. Eight bits are called a byte. The reason 8 bits are significant is that eight 1s in binary form (11111111) are equivalent to 255 in decimal form, which is enough to represent all variations in the English alphabet of upper- and lower-case letters, numbers, and other characters.

One million bytes are commonly referred to as 1 megabyte, or 1 MB. This is how most computer users understand file sizes, hard drive sizes, and the like. In the data networking world, however, transmission rates are measured in megabits, not megabytes. One megabit is denoted as 1 Mb to distinguish it from a megabyte. In this way, data transmission rates can be deceiving because a network that can transmit data at 10 Mb per second is really transmitting 1.25 MB per second (one eighth the speed). Why megabytes are used when referring to file sizes and megabits are used when referring to data transfer rates is a mystery. It could have been a marketing ploy by network vendors to make their networks sound faster than they really are.

To most people, megabits mean very little. It is like telling people your car has 282 horsepower when they have not experienced 150, 250, or 300 horsepower. To put megabits into perspective, let us use the example shown in Table 14.2.

Many different types of common networks are available. Along with the shared Ethernet mentioned in our example network, others are listed in Table 14.3 in relative order of performance (though bear in mind that certain

---

---

**TABLE 14.2**  
Putting Megabits into Perspective

The typical individual reads at approximately 600 bits/second.

The image transmitted to your television screen represents about 3,000,000 bits/second.

Most modems connect to the Internet at 56,000 bits/second.

The most common type of network in existence today is shared Ethernet, which operates at 10,000,000 bits/second.

---

---

**TABLE 14.3**

Different Types of Networks Available and Their Speeds

---

---

4 Mb/s token ring
10 Mb/s Ethernet
16 Mb/s token ring
100 Mb/s Ethernet (also known as Fast Ethernet)
155 Mb/s ATM (also known as OC3)
622 Mb/s ATM (also known as OC12)
1000 Mb/s Ethernet (also known as Gigabit Ethernet)
2488 Mb/s ATM (also known as OC48)

---

---

traffic conditions may alter the order somewhat, meaning that mileage may vary).

As data transmission rates increase, so does the quality of the cabling required to support those rates. For instance, there is no standard to transmit anything faster than 100 Mb/s Ethernet on copper cabling. All speeds higher than 100 Mb/s require some type of fiber-optic cabling. The industry continues to work on the standard for Gigabit Ethernet that will support copper, but the distance will likely be very short (25 m).

Note that all network transmission rates are theoretical speeds. To say that 10 Mb/s Ethernet can transmit a 1.25 MB file in 1 second would be fallacious. It is important to make the distinction between theoretical and actual bandwidth and understand why the difference (sometimes called overhead) exists.

Many factors contribute to the difference. To understand overhead, consider how, in shared Ethernet, there is the possibility of a collision. This is one form of overhead. Anytime there is a collision, the data that is being transmitted is discarded and must be retransmitted. In a shared Ethernet segment, this affects the actual bandwidth. Other network types, such as token ring, ATM, and to a lesser extent switched Ethernet, do not share the problems of collisions, but all network types must deal with some form of overhead that prevents them from reaching their theoretical bandwidth.

To understand some of the other factors, we need to examine conversations again, and realize just how little of the conversation in a data network


has to do with the actual data being transmitted. Each half of spoken dialogue is made of sentences. In a network, each of these sentences is called a packet. Packets are the elements composed of bits that carry data between the 2 computers on the network.

Let us look at the example of the letter and the envelope to see how a packet is formed.

In Figure 14.12, the letter represents the data that needs to be sent to the recipient. Realize that this letter is only a small part of what is actually getting sent through the postal system. The letter requires the envelope, the recipient's address, and the return address. All of this contributes to overhead. In some cases, the actual data could comprise as little as 33% of the packet. On top of that, each packet must be checked to make sure it made it to its destination without any errors (from electrical interference, for instance, which may cause a 0 to be interpreted as a 1 or vice versa). There is overhead in doing the checking, as well. As you might imagine, all of this extra information can quickly contribute to overhead on a network, gobbling up bandwidth you thought you had all to yourself. In some protocols, packets also have to be acknowledged by the receiver, or else the sender resends the packet. Imagine the overhead involved in sending a message that simply says, "Yes, I got your letter." There is very little data in that packet compared to the actual bandwidth it consumes.

(Bear in mind that the packet in Figure 14.13 is an example of a packet, and not to be taken too literally. Each packet is actually made up of other information, depending on the exact protocol and data link layer for which it is designed. Our example packet has been simplified dramatically.)

All of this goes a long way toward dragging the actual bandwidth of a network down. Most engineers agree that between collisions and overhead,



**FIGURE 14.12**

A letter and envelope with address are like a packet.




FIGURE 14.13

A simplified packet.


actual bandwidth on a shared Ethernet segment is somewhere between 3 and 4 Mb/s. Any time you try to transmit more than that you will experience the law of diminishing returns in which collisions become the norm rather than the exception, and very little gets transmitted. The same holds true for 100 Mb/s shared Ethernet, which typically maximizes at somewhere between 30 and 40 Mb/s. Switching can go a long way toward bringing the theoretical and actual closer together by eliminating collisions on a segment. For example, on 100 Mb Ethernet segments (known as Fast Ethernet) that have implemented switching, the maximum bandwidths are closer to 75%, or 75 Mb/s. For reasons we discuss later, ATM has different theoretical figures based on a variety of assumptions and other factors.

## FACTORS TO CONSIDER IN BUILDING A NETWORK

Up to this point we have discussed all the various working parts of a network. We have discussed hubs, routers, bridges, and switches. We have examined different types of cabling used in building networks, and we have discussed how data actually is transmitted between computers. Given the vast array of protocols, hardware, and topologies at our disposal, how is one supposed to determine exactly how to go about building a network? Our original example network is quite simple: 2 computers, a server and a client connected to a hub using CAT5 cabling to transmit image information. Depending on how important it is to get the information from the archive to the client, we could upgrade them to a faster network speed, such as 100 Mb/s Ethernet. What happens as we start to add workstations? (See Figure 14.14.)

If all the devices are connected to the same hub, they share the same collision domain. Because all the devices are relatively close together (within 300 ft), we continue to use CAT5 cabling. Each device is configured to communicate on the network at 10 Mb/s. In this scenario, each of the clinical



**FIGURE 14.14**

Adding workstations to our network.

displays is attempting to access information on the archive computer. The clinical display can theoretically access the information at 10Mb/s, and the fastest the archive can respond with information collectively is 10Mb/s. Even if we were to implement a switched architecture here, that would only make the situation worse, as the clinical displays currently do not have to compete with one another to send out their packets. They have to compete only for the archive's 10Mb/s of bandwidth. Here we can see a bottleneck.

Another way to look at bottlenecks is to look at the example of a human conversation. Let us say that with speaking and listening you can process 200 words per minute in a dialogue hypothetically. This means that you can listen at the rate of 200 words per minute or you can speak at the rate of 200 words per minute. Any more than that and you might become confused and request that the other person slow down in speaking, or you have to slow yourself down in speaking. For argument's sake, let us say you are pretty good at multitasking and can actually be having 2 conversations at once; you can be discussing PACS networking with 1 individual and the Red Sox with another. Nevertheless, you can discuss these 2 subjects collectively at a total of 200 words per minute. If 2 more people try to talk to you, those 200 words need to be divided among all 4 parties. At some point you become the bottleneck, as each of the other 4 parties is capable of handling 200 words per minute but you are able to give each party only 50 words apiece. Either they have to get their information from another source, or you have got to find a way to speak and listen more rapidly. Similarly, we will see how we can intelligently upgrade our network in the proper places to get the most for our investment. In our example, we could upgrade the server to "speak" at

100Mb/s Ethernet, while keeping the clinical displays at 10Mb/s Ethernet. Upgrading all the clinical displays to 100Mb/s Ethernet while the archive processes at 100Mb/s will accomplish little except in an environment where only 1 or 2 clinical displays are requesting information at once, in which case optimum transmission rates can still be realized. If multiple workstations are operating at 100Mb/s and requesting data from the archive, the archive could again become the bottleneck, at which time we could upgrade it to Gigabit Ethernet or ATM of one type or another.

It is clear there are multiple ways to handle this arms race. In the next example, we begin to integrate the various devices and architectures into a coherent PACS solution.

## A REAL-WORLD EXAMPLE

It would be helpful to use a real-world example of how our institutions used to use switching and routing together to increase productivity for several thousand users on a network experiencing high utilization, collisions, and very poor performance. In one case, 1200 users in a particular building were using one 10Mb/s Ethernet segment sharing the same collision and broadcast domain; that is to say, they were all in the same conference room. Each computer in the room was a client, all accessing servers off their network segment (outside the conference room). Simply installing a bridge would have accomplished little in this case, because all the traffic was attempting to use the same 10Mb/s router interface. We could not take half the users and put them on a different routed segment, because that would have meant readdressing 600 computers by hand to tell them they were on a different network segment (in a different room). We needed to somehow keep all the users on the same segment, while decreasing the number of collisions. Installing a switch for all these users would help some, as they would still be competing for the same 10Mb/s interface on the router. We needed to do a combination that allowed faster access out of the conference room (the router interface) and used switching within the room so all 10Mb/s users would suffer from fewer collisions. New equipment enabled us to install a switch and give the router a 100Mb/s interface into the room, while allowing all users to communicate on the network at 10Mb/s in a switched architecture. In doing so, we dramatically decreased the number of collisions and circumvented the bottleneck that would have resulted from upgrading to 100Mb/s the router's interface onto the segment (see Figures 14.15 and 14.16).




FIGURE 14.15

Network before upgrade.




FIGURE 14.16

Network after upgrade.

## PACS NETWORKS

As you probably see by now, networks can be incredibly diverse. There are a number of factors to take into consideration when building a network, and we have many tools available to us in the form of routers, switches, Ethernet, ATM, and copper and fiber cabling. How effectively we put this together depends on the amount of knowledge we have of the tools available to us, as well as an intimate understanding of exactly how traffic on the network flows. We also need to take cost into consideration. Among the many important factors, 4 stand out (Table 14.4).

The decision-making process is similar to that used for choosing a car. Potentially, one could purchase a high-performance car for a midrange price, but in the long run there may be reliability issues with it, and maintenance may be expensive. On the other hand, a more reliable car might cost the same amount but lack the performance we might require on the highway. Getting both reliability and performance costs more money.

PACS networks are different from most networks. Typical networks are intended to transmit e-mail, share small files, share printers, and handle other processes. If it takes 2 minutes or 3 minutes to transmit an e-mail, does it really matter? If it takes 10 seconds to open that Word document rather than 8, who is going to notice? On the other hand, PACS networks are transmitting large amounts of data, and that data needs to get where it is going quickly and reliably. PACS networks require both performance, because of the sheer amount of data being transmitted, and reliability, because of the data's clinical nature. The reason PACS is becoming so popular is that only recently have fast, reliable technologies become available to make PACS implementations practical and cost-effective. In the following examples, you will see networks on 3 different scales and learn the rationale behind their design.

---

---

**TABLE 14.4**  
Four Important Factors in Building a Network

Performance—How quickly the network can transmit data.

Reliability—How reliable the network is. How easy is it to maintain?

Upgradability—How easy is it to upgrade?

Price—How much does the network cost?

---

---

## ASYNCHRONOUS TRANSFER MODE

First let us look at few technologies that we will use in building this network. One such technology is ATM, or asynchronous transfer mode. We will explain the differences between ATM and Ethernet by explaining how they were derived.

Ethernet, in its purest form, is a “dumb” medium. It does not know what traffic is passing over it, nor does it need to. It operates under simple rules and principles that govern how traffic will be transported. Your e-mail about the Red Sox has no higher or lower priority than a PACS image or a telephone call. In this sense, Ethernet is said to have no concept of quality of service, or QoS.

ATM, on the other hand, was designed with the intent of allowing the network administrator to classify different types of traffic with different priorities. In this sense, an ATM network offers different levels of QoS depending on the application or workstation.

ATM is able to do this so well in large part because it is a much more predictive architecture. In an Ethernet network, packets can vary in size, reaching a maximum of 1500 bytes. However, some packets can be 300 bytes, some 50 bytes, and some 1500 bytes. Thus, it is very difficult for an Ethernet network to be able to predict traffic patterns. On the other hand, ATM packets are called “cells” and are always 53 bytes in length. Because of the constant size of the cells, ATM is predictive: a packet of streaming video will not get “stuck” behind a 1500-byte e-mail packet in a router while it tries to reach its destination.

ATM is often compared to a telephone network because of the high predictability it ensures. ATM networks can be built so that voice cells will have a high priority and always get through. ATM does this in much the same way the telephone system does. When you pick up the telephone and attempt to make a call, the ATM or telephone network checks to see whether it can complete the call with the parameters you require. In the case of ATM and computer networks, you may request a certain amount of bandwidth. If the ATM network cannot guarantee that amount of bandwidth, the computer and the network can negotiate a lower bandwidth rate, or not start the conversation at all. This is very much like a telephone call, especially a cell phone call, where you occasionally get a “fast busy” because the voice network cannot guarantee enough bandwidth to complete the call.

That said, ATM has had a long, troubled journey toward universal acceptance. It has taken years from its inception to make it to a finished product. Along the way various vendors have released products based

only loosely on not-yet-agreed-upon standards that make vendor interoperability very unlikely. Because of this, when you choose an ATM vendor, understand the high probability of the product's not working with other ATM hardware.


Another consideration in choosing to use ATM in networks is the work a router or switch needs to perform in converting packets into cells and back. This conversion contributes a certain overhead, which is often viewed as detrimental. Strong consideration should be made when mixing ATM with other network types, such as Ethernet. However, those considerations go beyond the scope of this book.

There are other downsides to ATM. The number of ATM installations in most LAN environments is rather small when compared to those of Ethernet. Hence, it can be harder to find engineers who have expertise in the area. What is more, the price per port of ATM is significantly higher than that of Ethernet. Add to this the fact that Gigabit Ethernet is beginning to gain steam and provide high bandwidth at a lower price, so an investment in ATM may not be a good one.

## WIDE AREA NETWORKING

All the technologies that we have talked about so far (ATM, Ethernet, token ring) have been designed for local campus settings; that is, for connections on which we can run either fiber or copper cabling. When we leave the campus setting, however, we no longer have the luxury of being able to run our own cable. Whether trying to go across the street, across town, or across the world, we need an external provider to connect sites. Telephone companies (also known as Regional Bell Operating Companies, or RBOCs), such as Bell Atlantic or AT&T, have massive networks designed for handling telephone calls all over the world. They can use these networks to provide data transmission as well.

RBOCs lease "circuits" to companies over which to transmit data. These circuits are generally point-to-point connections between 2 sites. Generally speaking, the circuits they provide offer bandwidth from 56 kilobits (Kb)/s all the way up to 155 Mb/s. Unfortunately, charges are recurring (usually monthly), and the price per megabit is much higher than on the local area network (LAN) for a connection of the same speed. For instance, a standard type of circuit, a T1, operates at 1.554 Mb/s, and can cost several hundred dollars per month. The farther the circuit extends, the more the RBOC charges. Several factors in the industry have kept these costs fairly stable for many years, but technologies to get around these charges are




**FIGURE 14.17**

Illustration of how leased circuits from RBOCs can be integrated into LANs to allow connectivity between 2 LANs.

becoming more popular, and competition is starting to drive down prices (see Figure 14.17). We discuss some of these technologies and the competition factors later in this chapter.

## A SIMPLE PACS NETWORK

Figure 14.18 shows a relatively simple PACS network. As you can see, there are 2 modalities, 3 clinical and diagnostic displays, and an archive server. All



**FIGURE 14.18**

A simple PACS network.

these devices are connected to a central Ethernet switch, probably using CAT5 cabling. This network assumes that all devices connecting to the central switch are within cable specifications for connecting at speeds of 100Mb/s.

In this diagram the archive server is connected at 100Mb/s, as are the diagnostic and clinical displays. The modalities are connected at 100Mb/s. As most traffic is being sent to the archive server and retrieved from the archive server, we see our first potential bottleneck, but in a network of this size, it is very unlikely that the 100Mb/s connection will be swamped. In some cases, a 1000Mb/s (Gigabit Ethernet) connection could be substituted for the 100Mb/s connection to the server, but the likelihood that it would enhance performance is minimal.

## MAKING LIFE MORE COMPLICATED

Now we will complicate our PACS network somewhat. First, we will add some modalities and displays. Let us say that these additions are located at the other end of the building and thus the distance exceeds our maximum cable requirements. We will need to connect them to a second Ethernet switch, and then connect the Ethernet switches together, probably with fiber-optic cabling because the 2 switches are 750 feet apart.

Adding this additional closet (switch) causes us to look at our network in a whole new light. Now we have to consider how much bandwidth we need between the 2 closets (1, in Figure 14.19). This second switch has 3 modalities and 3 display stations and could conceivably saturate 330Mb/s of bandwidth if all are transmitting at the same time. This is very unlikely, and Fast Ethernet is likely to be enough bandwidth for quite some time.

Still, just as in the simple PACS network, all the traffic is going to the archive. So it is even more conceivable that the connection between the archive and the switch will be the most likely bottleneck. Given the technologies available these days, both links (2 and 3 in Figure 14.19) could be 155 Mb/s ATM, although that is not much of an improvement over 100Mb/s Ethernet. We also need to consider the cell conversion if the rest of the devices are still Ethernet and if we want to add ATM to an all-Ethernet environment. Another option is 622 Mb/s (OC-12) ATM, which offers significantly better performance than 155 Mb/s ATM. Gigabit Ethernet is yet another option, and would probably come in at the same price point or better than 622 Mb/s ATM. These options offer a perfect example of how you must work with your PACS vendor and network vendor to determine which connection type best fits your situation.



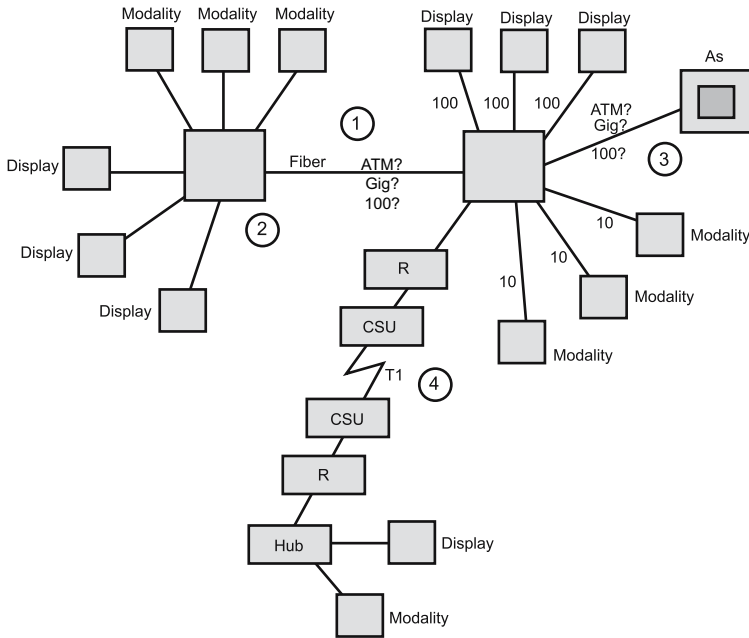


FIGURE 14.19

A more complicated network.

In our network example let us say we decide to add a remote site across town. We have routers on each end (1 at our local site and 1 at the remote site). Our RBOC is providing a T1 (1.554Mb/s), a relatively slow connection speed for PACS. With the modality at the remote site, images could still be stored on the archive at a reasonable rate, meaning that compressed images could be read at the remote site. If we need to view uncompressed images in a more timely manner we probably need a faster connection type, perhaps multiple T1s or a T3 (which an RBOC offers at 45 Mb/s) (Table 14.5).

## COMPLEX PACS NETWORK

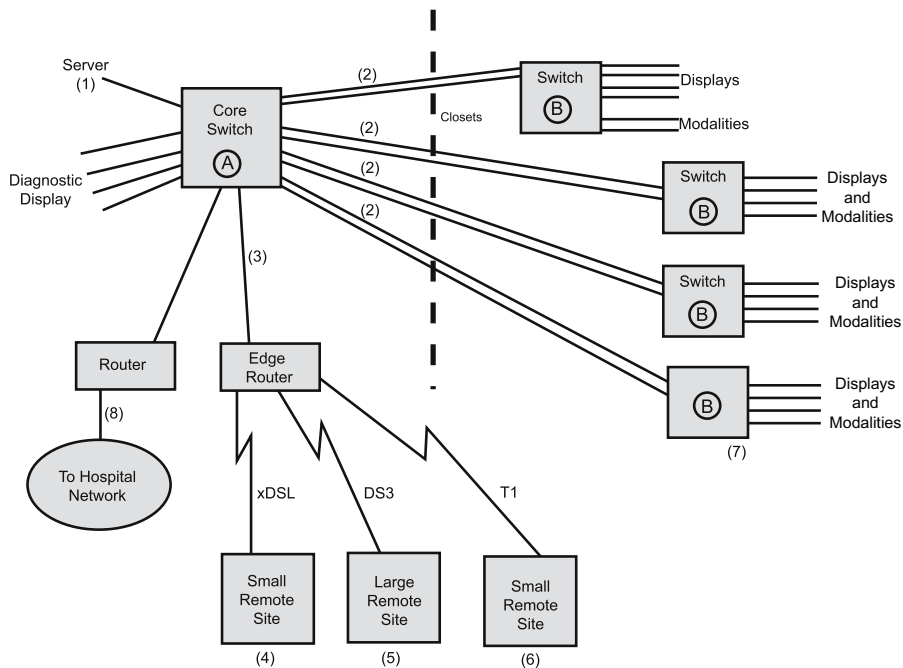
Our complex network incorporates a very large number of devices. Including what is shown in Figure 14.20, 50 or so modalities could be spread around our hypothetical hospital, as well as 50 or more displays for reading the images. We will need a fast, robust network to address bandwidth and reliability concerns. Although the scale has changed dramatically, many of

**TABLE 14.5**  
RBOC Circuit Types and Bandwidth

Circuit Type	Speed
Fractional T1	56 Kb/s–1.554 Mb/s
T1	1.554 Mb/s
T3/DS 3	45Mb/s
OC3	155 Mb/s
OC12	622Mb/s

the same technologies are at our disposal. It is a question of how we use them.

We have chosen what is commonly referred to as a “star topology” for this network. In this architecture each switch B in Figure 14.20 connects via



**FIGURE 14.20**

Complex PACS network.

fiber to the core switch A in Figure 14.20. Modalities and displays can be connected to either the closet switch closest to them or to the core switch, if it is closer. We chose this design in our example because the archive server was centrally located to all of them, thus cutting down on wiring costs.

As for the actual connections between the core switch A and the closet switches B, we have a number of technologies. Again, your PACS vendor may have preferences. In this case, we chose Gigabit Ethernet for the connections (uplinks) between the core and closet switches (2 in Figure 14.20). You may notice that we include 2 fiber-optic uplinks from the closets to the core switch. This allows greater reliability in the event that 1 of the pairs of fiber-optic connections experiences a failure, or if the fiber-optic connection on either the core switch or the closet fails.

Although not illustrated here, an even more resilient design would include 2 core switches A, with the archive server having a connection to both core switches, and each closet switch having a connection to both core switches. This type of connection would create a somewhat more complicated design to troubleshoot in the event of a failure, and it would certainly cost more, but it provides a high level of redundancy (meaning cable failures and hardware failures have less of an impact on the network) and ensures a high degree of availability of the archive server.

You also should notice a connection, 3, in Figure 14.20 to an edge router. This router acts just the same as the router in our previous PACS design, except that multiple sites of varying speed connect to it. In the current example, we are connecting to a remote site (4) using xDSL at 144Kb/s. This speed is too slow for sending or retrieving compressed images, but fast enough for a remote site that needs only to view compressed images. The site connected with the DS3 (5) acts just like the site in the previous network example (Figure 14.19).

The site with the DS3 connection can communicate back to our main network at 45 Mb/s. This allows ample bandwidth for the large remote site to host a few modalities and display computers, although at a premium in terms of price. If bandwidth becomes an issue over this link we could upgrade to an OC3 connection providing 155 Mb/s, but the cost would be great.

Finally, 8 in Figure 14.20 shows a router that is providing connectivity to the rest of the hospital network, including its clinical systems. Note that this is the only router we are using on the LAN, and that in none of our previous examples did we include a router, as it can slow down performance in moving large images. However, in interfacing with the rest of a hospital network, a router should provide an efficient way of allowing us to move patient information between the 2 networks, as well as allowing cli-

nicians to view compressed images throughout the hospital on lower-priced workstations.

## FUTURE TRENDS

The networking industry as a whole experiences changes at an extremely rapid pace. The rate of change is dizzying. There are several changes in the industry that will affect how PACS is implemented in the near future.

### GIGABIT ETHERNET

Gigabit Ethernet has gained acceptance. Now Gigabit Ethernet is offered by nearly all vendors at reasonable prices. For the first time on the LAN, the ability for the network to provide inexpensive bandwidth has outpaced the computer's ability to transmit data. This means that since it is difficult for a Sparc workstation to transmit much more than 200Mb/s, a Gigabit Ethernet connection offers plenty of bandwidth. Before Gigabit Ethernet, only expensive OC-12 ATM could make such a promise, while most workstations could saturate a 100Mb/s Ethernet pipe. Of course, as computers get faster, their ability to transmit data will also increase. At this point, however, Gigabit Ethernet provides a lot of room for growth. Even so, it is possible for computers to collectively saturate Gigabit Ethernet backbones and connections, so careful network design is still required. Further, there are groups currently working on terabit Ethernet speed solutions.

### RESOURCE RESERVATION PROTOCOL

Earlier, when explaining the advantages of ATM, we said that ATM had the ability to distinguish between different traffic types, thus allowing it to efficiently pass video and voice traffic while still moving e-mail along at a reasonable rate. Ethernet vendors, not wanting to be left in the dust, recognized this distinguishing capability as an important feature and have been working toward implementing some of those features into the Ethernet architecture. Called RSVP (resource reservation protocol), this feature will give Ethernet some ATM-like capabilities by allowing it to offer a lesser degree of QoS (almost a poor man's QoS). When combined with the brute force of Gigabit Ethernet, it offers a challenge to the future of the more elegant ATM in the LAN environment.

## LAYER 3 SWITCHING

As detailed at length earlier in this chapter, switching occurs at the ISO model layer 2 and routing occurs at layer 3. The lines are becoming blurred, however, as vendors seek to offer the speed and flexibility of switching and the intelligence and efficiency of routing. Vendors are calling this layer 3 switching, although different vendors have widely varying ideas of how to implement it and what it means. In effect, layer 3 switching should combine the best of both worlds to such an extent that the idea of using the corridor in the conference room model will no longer be seen as an impediment to fast, efficient communication between different segments. In a large network, where PACS seeks to peacefully coexist with any number of applications, layer 3 switching offers potentially dramatic benefits by getting large images and studies out to a vast clientele.

## VIRTUAL PRIVATE NETWORKS

Virtual private networks, or VPNs, have constituted a major buzzword in the industry. The goal of this technology is to make the Internet a safe medium for the secure transfer of private information, such as clinical data or images. It does this by encrypting the information at the source and decrypting it at the destination. In doing so a private virtual network is set up over the public, unsecure Internet. There is hope that companies can use this technology to decrease costs by using the Internet to exchange information that they normally would be able to share only over private circuits leased from RBOCs, thus saving money. This is also an ideal technology to use for remote access to the corporate network for home users or in virtual office environments. In the near future, as this technology matures and vendors are able to produce scalable solutions, VPNs will make a great impact on how the Internet is used.

## LOWER COSTS

Competitiveness among vendors also means lower costs in the networking industry, just as in the computer industry. In the first year that Gigabit Ethernet became available, prices dropped nearly in half. The price of NICs continues to drop, as well as the price of routers and switches. On the WAN front, RBOCs are seeing increasing competition from one another, from cable providers, and from other third-party providers who are seeking to take

advantage of the tremendous hunger for bandwidth on the WAN. The cost and speed of the connection to the home or remote office, often called the “last mile,” is quickly reaching attractive price points for many applications, especially PACS. RBOCs will be forced to lower prices of T1s and T3s as new technologies such as cable modems, xDSL, and wireless become more common. Similarly, use of the Internet and technologies such as VPN will also allow for significant price reductions for remote access.

# SERVERS AND OPERATING SYSTEMS

---

---

KENNETH M. NESBITT • THOMAS J. SCHULTZ •  
ROBERTO DASILVA

**S**ometimes policies and procedures for all areas of system (picture archiving and communication system [PACS]) operations are overlooked, or it is assumed that the PACS vendor will provide them. Although every vendor provides a standard build document, it usually lacks procedures for contingencies such as server deployment, maintenance, change management, virus prevention and recovery, and server recovery from hardware and software failures. Therefore, it is important that the institution understand the available technologies and, jointly with the selected PACS provider, develop a PACS server strategy that reflects the goals of the institution. This chapter has 3 sections that discuss technologies and issues to consider during the deployment process: “The PACS Core (Servers),” “Operating System Security and Policies,” and “Viruses, Trojans, and Worms.”

## THE PACS CORE (SERVERS)

There are numerous server deployment scenarios from which to choose. Most strategies are based on one or more of the following techniques: single server, multiserver, load balancing, and server clustering.

### SINGLE SERVER

The single server model consists of one server handling all PACS functions. These functions include the importation of exams from DICOM modalities, the processing of the database requests, and the processing of image distribution requests from the radiologists' interpretation workstations. Workstations can access exams via standard Web protocols (the same Web protocols used to book your vacation on the Internet) or through a proprietary application program interface (API) built into the PACS vendor's software. The single server scenario is often ideal for an imaging center or small community hospital where the size and volume of exams are relatively small. Thus, the need for additional servers to provide the distributed processor power, memory, and hard disk space is not as great as it would be in a larger hospital, where higher volumes and exam sizes necessitate more resources. Figure 15.1 illustrates a single server model for PACS using Web technology as a means of distribution, with a backup server waiting in the wings.

### MULTISERVER

In a multiserver model, a separate server is responsible for each component of the PACS. One server handles DICOM importing, another server maintains the database, and one handles image distribution requests. This design affords the luxury of having processor power, memory, and hard disk space dedicated to each PACS component on an individual basis rather than having the components competing for the resources of one server. For example, this proves useful in institutions where large exam sizes and the number of studies performed can cause the database to grow to considerable size (200 gigabytes [GB] or more). Figure 15.2 shows a multiserver model for PACS using Web technology as a means of image distribution. All requests for information or images are funneled through the Web server. The Web server acts like a traffic cop, directing information requests to the proper servers. Image requests are serviced from the archive(s), and textual requests are serviced through the database.




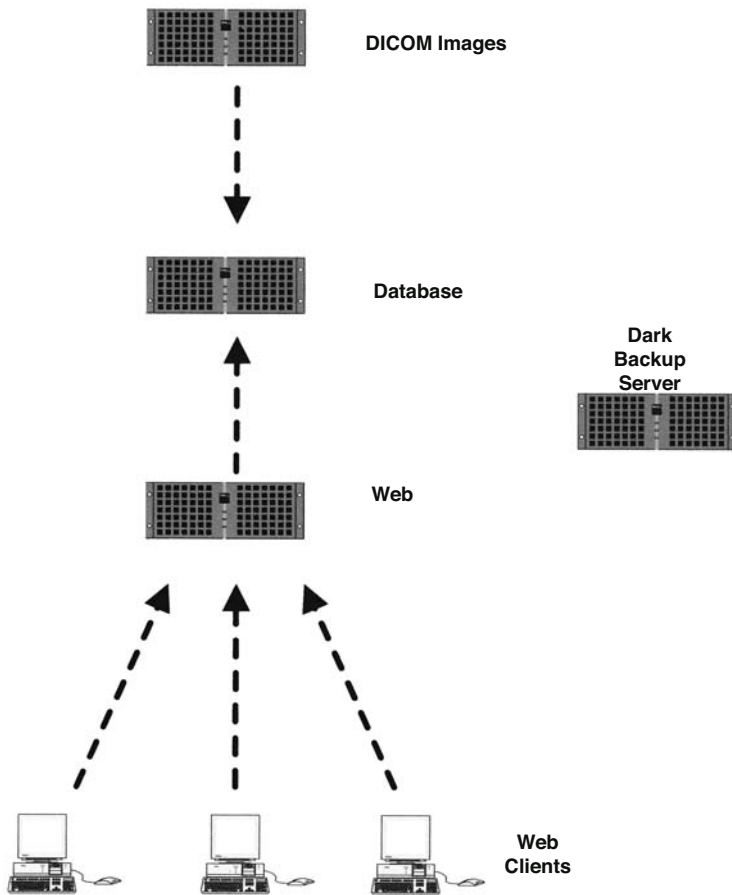


FIGURE 15.1

A single server model for PACS using Web technology as a means of distribution, with a backup server available.

## LOAD BALANCING

The load-balancing model is based on the theory of distributing tasks evenly so that no one system is overwhelmed; each server assigned to a task shares that task evenly. In PACS this technique works well with DICOM image servers and image distribution servers. For example, if a PACS had two DICOM image servers and two image distribution servers, the request to send images into the PACS would be intelligently routed to the DICOM server with the least amount of activity. Requests by the PACS workstations for images would be satisfied by the distribution server that is the least busy. The intelligence required to load balance is not provided by the PACS vendors but is provided by third-party technology vendors. Load balancing can be implemented at the hardware level (through content switch, a device that monitors network/server activity and makes decisions on what server should service a given request) or the software level (through Microsoft; the software runs on the servers instead of on a dedicated device). It is worth noting that database servers are costly to load balance (See “Clustered Failover” following this section) because transactions need to be coordinated among all participating servers.

**FIGURE 15.2**

A multiserver model for PACS using Web technology for image distribution. All requests for information or images are funneled through the Web server, which directs information requests to the proper servers.

Figure 15.3 illustrates a typical load-balancing model for PACS, using the Web as a means of image distribution. Following the diagram, both DICOM image servers forward their respective information to the database server, while each Web server makes respective calls to the database server to fulfill the client requests.

## CLUSTERED FAILOVER

The clustered failover model provides a high level of hardware redundancy protection for PACS servers (this approach is often referred to as active-passive clustering). Clustered failover means that one server controls the task being performed while the second server passively waits for the first one to fail. In this model, the DICOM image servers are clustered together in pairs, as are the image distribution servers, providing simple failover redundancy.

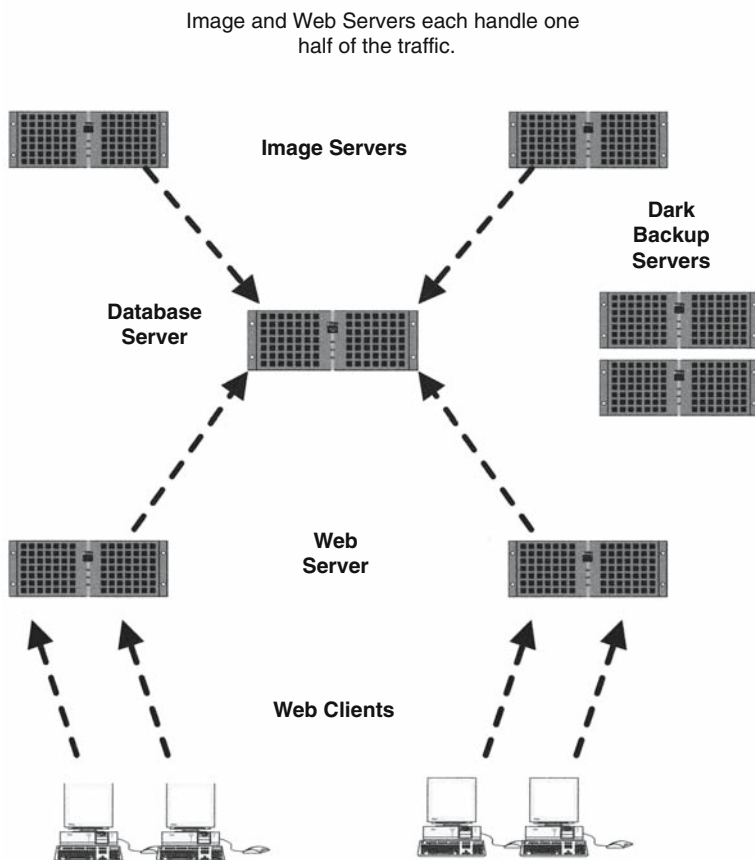


FIGURE 15.3

A typical load-balancing model for PACS, using the Web as a means of image distribution. Both DICOM image servers forward their information to the database server, while each Web server makes calls to the database server to fulfill the client requests.

The failover between servers is managed through software, either with an enterprise-class operating system such as UNIX/Solaris, Windows 2000 Advanced Server, or Windows Server 03 Enterprise edition, or through commercially available cross-platform products such as Veritas. A caveat of this approach is that if the PACS software caused the failure of the first server, the second server will also fail because it is running the same software as the first.

It is also possible to combine techniques (often called active-active clustering). For example, if the database (IBM DB2, Oracle, or Microsoft SQL Server) were clustered and load balanced, then during normal operations all data requests would be evenly distributed between both physical servers. If one failed, all operations would be funneled to the remaining server. Figure 15.4 shows a clustered failover (load-balanced) model for PACS using Windows 2000 Advanced Server as the operating system. This

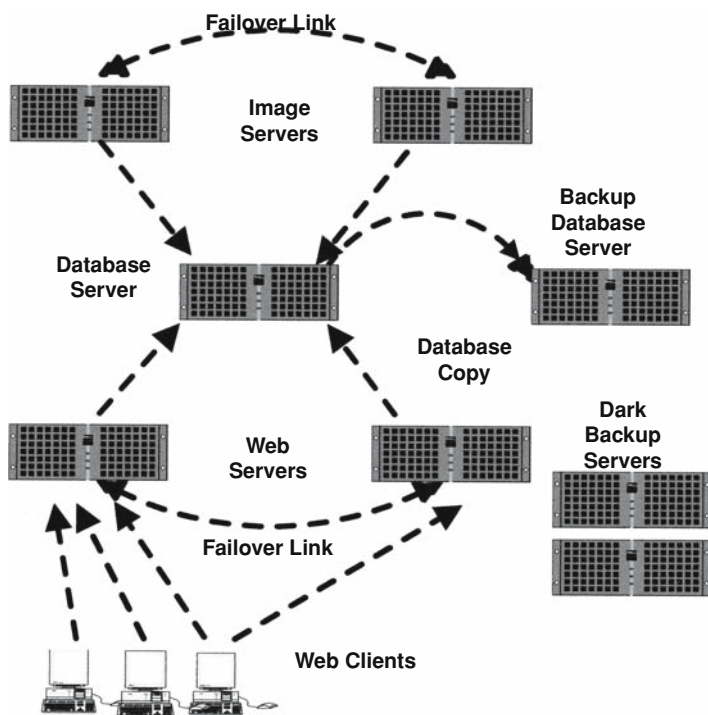


FIGURE 15.4

A clustered failover (load-balanced) model for PACS; Windows 2000 Advanced Server is the operating system.

**TABLE 15.1**  
Sample Server Hardware

<b>Vendor Model</b>	<b>Processor Speed</b>	<b>Max Processors</b>	<b>RAID Level</b>	<b>No. Internal Disks (Max Size)</b>	<b>Memory (Max)</b>	<b>Redundant Components</b>
DELL 2850	3.6 GHz	2	0,1,5,10	6–300 GB	12 GB	Yes
DELL 6600	3.0 GHz	4	0,1,5	8–300 GB	32 GB	Yes
HP DL380	3.6 GHz	2	0,1,5	6–146 GB	8 GB	Yes
HP DL580	3.0 GHz	4	0,1,5	4–146 GB	32 GB	Yes
IBM X345	2.67 GHz	2	0,1,5	4–146 GB	8 GB	Yes
Sun V440	1.28 GHz	4	0,1,5	4–73 GB	16 GB	Yes

Note: Usually for PACS, having more than 2 processors in a server yields little gain.

approach is usually very costly and offers little advantage over active-passive clustering.

## SERVER HARDWARE AND OPERATING SYSTEMS

There are a number of manufacturers and configurations to choose from when deciding on server hardware. Companies such as Dell ([www.dell.com](http://www.dell.com)) (Austin, TX), Hewlett-Packard ([www.hp.com](http://www.hp.com)) (Palo Alto, CA), IBM ([www.ibm.com](http://www.ibm.com)) (Armonk, NY), and Sun ([www.sun.com](http://www.sun.com)) (Palo Alto, CA) all offer quality entry level to high-end servers depending on an organization's needs and the chosen PACS solution. If the PACS solution is Windows based, then DELL, HP, and IBM are all excellent choices at very competitive prices. If the PACS vendor offers a UNIX/Solaris-based solution, then Sun is the only choice, as the hardware and operating system are proprietary (Linux may be an option for some implementations in the near future). Tables 15.1 and 15.2 show a quick comparison of some server product features and major operating system features. No matter which PACS solution an organization implements, there are powerful, high-quality server hardware and operating systems available.

**TABLE 15.2**  
Sample Server Operating Systems

<b>OS</b>	<b>Supported Processors</b>	<b>Supported Memory</b>	<b>Cluster Capable</b>
Windows 2000 Server	4	8 GB	No
Windows 2000 Enterprise	8	32 GB	Yes
Windows 2003 Server	4	8 GB	No
Windows 2003 Enterprise	8	32 GB	Yes
Linux	64	64 GB	Yes
Solaris	64	64 GB	Yes

## DISASTER RECOVERY

Disaster recovery (DR) planning is equally important as choosing the right PACS solution. Without thorough and concise planning, an institution may find itself severely impacted if a major failure occurs on one of the PACS core components. System downtime could be counted in weeks if proper strategies are not developed. This section discusses different aspects of DR, such as uninterruptible power supply (UPS) and emergency power, server operating system (OS) imaging, archiving, and database backups.

### UNINTERRUPTIBLE POWER SUPPLY

Having a UPS and emergency power installed in the room where the PACS servers will be housed is critical to surviving anything from power blips to complete power outages. An enterprise-class UPS consists of several large batteries, built-in power filtering, and system-monitoring capabilities. The ideal UPS will provide enough battery power to allow servers to shut down gracefully in the event of primary power failure. Companies such as MGE ([www.mge.com](http://www.mge.com)) (Saint Ismier, France) and APC ([www.apc.com](http://www.apc.com)) (W. Kingston, RI) offer entry level to enterprise-class UPSs capable of accommodating any PACS scenario.

Most hospitals and trauma centers supply backup emergency power to key areas such as the emergency ward and intensive care units. Upon the installation of PACS, it is advisable to have the PACS servers integrated into the same emergency backup power system. Emergency power can keep PACS

servers operational through power failures and enable an organization to continue operating until normal power has been restored. The role of a UPS in this scenario is to provide enough power to bridge the gap between the initial power failure and the institution's switch over to the backup power system.

## OPERATING SYSTEM IMAGING

Server OS imaging involves using third-party software to take a snapshot (state in time) or image of the OS volume (or any volume) that encompasses all the applications and server patches installed. Images of other volumes, such as the database or log volumes, may also be acquired. By implementing a strategy of making OS images of servers before making changes such as installing security patches, tweaking performance, or upgrading applications, extensive downtime can be avoided. If a system change causes issues, it can take less than 20 minutes to reset the affected server back to the last known running state. (Consider the peace of mind provided by knowing that it will only take 20 minutes to restore a server to a known working state after a virus attack.) This is significantly faster than having to go through the process of reinstalling the OS and other applications, then possibly having to re-create the database from tape backups and database log files. Two vendors of server-imaging software are Acronis ([www.acronis.com](http://www.acronis.com)) (Herndon, VA) and Powerquest ([www.powerquest.com](http://www.powerquest.com)) (Orem, UT). Both offer excellent products to help protect your server environment and reduce the time required to recover from a disaster. Both products offer the capability to restore system images from remote network locations (NAS) or from CD or DVD.

## REMOVABLE MEDIA

Archiving to removable media provides the institution with the option of storing the removed media in a secure offline location. (Storage devices are covered in detail in Chapter 16.) There are many choices for removable media, such as CD, DVD, DLT (digital linear tape, proprietary technology developed by Quantum), Super DLT 1, Super DLT 2, LTO (linear tape open) 1, LTO 2, and the recently released LTO 3. For smaller institutions, a DVD archive would probably suffice. With DVD media capable of writing to both sides, the ability to store 9.4GB on one disc is appealing, and the new 25 GB DVD release is imminent. One company that offers this solution is DAX Archiving Solutions ([www.daxarchiving.com](http://www.daxarchiving.com)) (The Netherlands). The DLT solution has been around for a while and is still a viable option for most

institutions. It provides up to 80 GB of storage using data compression. Super DLT 1 takes things up a notch with capacities of 160 GB raw and 320 GB compressed. Super DLT 2 goes even further, with capacities of 300 GB raw and 600 GB compressed. Linear tape open is an open standards-based technology supported by companies such as Hewlett Packard, IBM, and Seagate (Scotts Valley, CA). Linear tape open 1 tapes can store 100 GB raw and 200 GB compressed. The second-generation LTO 2 drives can store 200 GB raw and 400 GB compressed. Linear tape open 3 became available in the fourth quarter of 2004 and offers capacities of 400 GB raw and 800 GB compressed. Two companies that offer complete DLT and LTO solutions are Storage ([www.storageetek.com](http://www.storageetek.com)) (Louisville, CO) and IBM ([www.ibm.com](http://www.ibm.com)).

## DATABASE BACKUPS

Database backups are an essential of everyday life in the world of PACS. They are a vital piece of the DR puzzle that must be completed to ensure a minimal loss of data (PACS and server configurations) in the event of a substantial failure or environmental disaster. In conjunction with server imaging, the database backup can be used to restore the database to a state that would be less than 24 hours old. It is highly recommended that database backups be performed regularly (every 24 hours). Occasionally restoring a backup to a test system to ensure the integrity of the backup is good idea. Backing up to a high-density tape such as Super DLT or LTO enables retaining many days' worth of backups, which is useful in the event a single backup copy is bad. It is important to be cognizant of the fact that database corruption or hardware failure can occur at any time, so keep several backups to prevent major headaches.

## OPERATING SYSTEM SECURITY AND POLICIES

To maintain a secure network environment and ensure system functionality, it is crucial that server security policies be clearly developed and outlined. The following is a general outline of recommended server policies.

### OPERATING SYSTEM PATCHES

Operating systems (defined as the first program that runs when a computer is turned on; this program manages all other programs, including applica-



tions such as e-mail or PACS) all have flaws or bugs. The OS provider continually offers updates (patches) to fix problems that can cause crashes or to seal security flaws to protect the server from hackers. It is very important to routinely and diligently monitor the OS provider's Web site for available patches. Usually a security hole is discovered and fixed before hackers can take advantage of the flaw. How? Hackers reverse engineer a security patch to learn of the OS flaw and then create malicious programs (viruses) to exploit the flaw. Their intention is to affect systems that are not patched (protected from the flaw).

## VIRUS PROTECTION

Operating system patches are the first line of defense against malicious (virus) software (see "Viruses, Trojans, and Worms" in the next section). The second line of defense is a genre of software (obviously named) antivirus. Antivirus software runs on each server and workstation and continually monitors files and running programs for signature behaviors. The antivirus software provider continually researches the Web and other sources for virus sightings. Once a new virus is discovered, it is then studied and reverse engineered to determine what security flaws it takes advantage of, how it replicates itself, and whatever else it tries to do (e.g., steal data, clog the network). Once the virus is understood, this information is electronically sent to all customers subscribing to the antivirus provider's service. The updated information is often called "virus signatures." Once the antivirus software is updated with the new signatures, it will then scan the server to determine whether it has been compromised. If it has been, the antivirus software will remove the virus, cleansing the system. Sometimes the system is so severely compromised that it cannot be repaired. In this scenario the system will need to be rebuilt from a known clean OS image or by reinstalling the OS and applications. Providers of antivirus software include McAfee ([www.mcafee.com](http://www.mcafee.com)) and Norton ([www.symantec.com](http://www.symantec.com)).

## PASSWORDS

Servers should have individual user accounts and fairly complicated passwords to prevent a security breach by a hacker. Most hackers will write a simple algorithm looking for easy access through weak passwords or, better yet, no password at all. Examples of weak passwords are "123" or "abc."

## PASSWORD CONSTRUCTION GUIDELINES

- ▶ Include both upper- and lowercase characters (e.g., a-z, A-Z).
- ▶ Have digits and punctuation characters as well as letters (e.g., 0-9, !@#\$%^&\*()\_+|~=-\`{}:”;’<>?,./).
- ▶ Have at least 8 alphanumeric characters.

### LIMIT THE NUMBER OF UNNECESSARY ACCOUNTS

Avoid creating duplicate user accounts, test accounts, shared accounts, general department accounts, and the like. Use group policies to assign permissions as needed, and audit your accounts regularly.

### NO UNENCRYPTED AUTHENTICATION

It is advisable that user authentication (logging in) be encrypted (an encoding that makes a password difficult to reverse engineer). For example: “baby” could be encoded to “%%\$FFDFDR” during transmission over the network, but when the server receives the login data it converts it back to “baby.” This is important to consider because hackers have tools that allow them to monitor network traffic and snatch passwords that are transmitted in clear text (“baby”).

### NO UNAUTHORIZED E-MAIL DELIVERIES

All operating systems have the ability to send mail messages to other servers. Malicious code (viruses) has the ability to leverage this process and cause havoc throughout the institution (spam and other mail-clogging problems). To help avoid this situation, it is advisable to turn off the SMTP (simple mail transfer protocol) program unless it is absolutely needed.

### PHYSICAL SECURITY

All servers and related equipment should be located in a secure room with limited access to only those individuals supporting the system. This is an

obvious security measure to take, but many institutions have had issues with missing/modified equipment due to poor access polices.

## SECURE BACKUP MEDIA

All backups and emergency repair disks should be kept in a secure, fireproof location in a building other than the one where the server resides.

## SECURITY GUIDELINES FOR WINDOWS (NT, 2000, 2003, AND XP)

Following are basic guidelines unique to Windows-based systems.

### ADMINISTRATOR ACCOUNTS

The administrator account in Windows grants the highest level of access. Any system change can be performed with this account. Therefore, the administrator account (in domains, the domain administrator) should be restricted to senior members of the information technology (IT) staff. Each IT staff member requiring administrative privileges should have his or her own account created and placed within the administrator's group. This will allow for clear tracking of who is doing what and when. For situations in which IT staff require fewer privileges, such as reading e-mail, a second account should be created with a lower level of access. Following this practice will prevent accidental changes from being made to critical systems.

### REPLACE THE "EVERYONE" GROUP WITH "AUTHENTICATED USERS" ON FILE SHARES

"Everyone" in the context of Windows security means that anyone who has access to your server can access the data contained within those shares. The "Everyone" group is assigned by default when a share is created and thus should be replaced with the "Authenticated Users" group.

### PASSWORD PROTECT THE SCREENSAVER

A simple method of preventing access to a system in an unsecured location is to password protect the screensaver. This is a basic feature built into all Windows operating systems. For details, visit [www.microsoft.com](http://www.microsoft.com).

## USE NEW TECHNOLOGY FILE SYSTEM ON ALL PARTITIONS

New technology file system (NTFS) is the only file system type that implements strong security and encryption features. It also contains autocorrection capabilities when corruption occurs. For these and other reasons, it is recommended that this file system be implemented for most applications. Visit [www.microsoft.com](http://www.microsoft.com) for more information.

## RUN WINDOWS UPDATE

Microsoft provides a mechanism for automatic notification and installation of critical OS patches. It is strongly recommended that you take advantage of this service. On servers it is usually the best practice to set the service to download all patches automatically, but to install them manually. This approach allows for review of how the patch will affect the PACS applications. It is possible that some patches may prevent PACS from running correctly.

## DISABLE THE DEFAULT SHARES

Windows NT, Windows 2000, Windows Server 03, and Windows XP open hidden shares on each installation for use by the system account. You can disable the default administrative shares in two ways. One is to stop or disable the server service, which removes the ability to share folders on the server. The other way is by editing the registry (see [www.microsoft.com](http://www.microsoft.com) for details). Keep in mind that disabling these shares provides an extra measure of security but may cause problems with applications. Test your changes in a lab before disabling these shares in a production environment.

## **VIRUSES, TROJANS, AND WORMS**

More than 30 new viruses appear every day, and the roster of active viruses increases by more than 10,000 per year.

## DEFINITION OF A VIRUS

A virus is defined as a piece of code that replicates by attaching itself to another object, usually without the user's knowledge or permission. Viruses

can infect program files, documents, or low-level disk and file-system structures such as the boot sector and partition table. Viruses can run when an infected program file runs and can also reside in memory and infect files as the user opens, saves, or creates the files. When a computer virus infects a computer running Windows, it can change values in the registry, replace system files, and take over e-mail programs in its attempt to replicate itself.

## DEFINITION OF A WORM

Worms are a type of virus that replicate by copying themselves from one computer to another, usually over a network. The general goal of a worm is to replicate itself to as many systems as possible. Often the worm acts as a transport mechanism for delivering yet another virus.

## DEFINITION OF A TROJAN HORSE

A Trojan horse is a virus that masquerades as a useful program. Sometimes it actually performs a useful purpose, but behind the scenes it releases a virus. A Trojan horse program can come from an e-mail attachment or a download from a Web site, usually disguised as a joke program or a software utility of some sort.

## EDUCATION AND TRAINING

The operations and IT team should work together to develop a sound education and training program for the PACS users. The more users understand about passwords, viruses, and safe computing, the less likely it is that they will be the ones to allow a compromise of the PACS system.

## GLOSSARY

**CORE** The central foundation of the PACS, consisting of one or more servers.

**SERVER** Powerful, robust computers that possess multiple-gigahertz-speed processors, several gigabytes of memory, and several gigabytes of hard-drive space (RAID).

**PROCESSOR** The component of a server that does most of the data processing.

**MEMORY** The component of a server that assists the processor with data processing and preserves data for retrieval.

**HARD DRIVE** The component of a server that holds and spins a magnetic disk and reads/writes information on it. It is also referred to as a hard disk.

**GIGABYTE (GB)** A unit of computer memory or hard drive capacity equal to 1,024 megabytes.

**GIGAHERTZ (GHZ)** A unit of frequency used to measure modern processor clock speed. One GHz is equal to 1 billion clock cycles per second.

**SYSTEM ARCHITECTURE** Deployment scenario/server model: the topology of the CORE systems, workstations, and archive, meaning the number of servers used and their respective roles.

**PRIMARY SERVER** The server that performs most or all of the PACS functions on a regular basis.

**BACKUP SERVER** A server that is attached to the network and configured to deploy in the event of a failure or virus attack on the primary server.

**DARK SERVER** A server that is detached from the network to prevent infection from virus attacks. This provides a preconfigured, virus-free server to deploy in the event that the primary server requires downtime for repairs.

# STORAGE AND ENTERPRISE ARCHIVING

---

---

PAUL G. NAGY • THOMAS J. SCHULTZ

**A** study conducted at the Berkeley School of Information Management and Systems concluded that the world generated more than 5 exabytes (5 billion gigabytes [GB]) of recorded information in 2002. This represents a 30% growth per year for the past 5 years. More information has been generated in the past 3 years alone than in the previous 40,000 years of civilization combined. The entire printed Library of Congress holds 17 million books, or 136 terabytes (TB) of data.

Hospitals have seen a parallel data explosion, especially in the imaging modalities. A single 2000-slice computed tomography (CT) procedure captures 1 GB of information. Advances in molecular imaging, multispectral imaging, and physiologic imaging suggest that the rate of growth in storage requirements will only accelerate in the future. We have clearly entered the era of information overload, and our success now depends on how well we can capture, retrieve, and synthesize this avalanche of information.

**TABLE 16.1**  
Storage Terminology: Powers of 10

Kilo	$10^3$	1 kilobyte (KB) = 1000 bytes
Mega	$10^6$	1 megabyte (MB) = 1000 KB
Giga	$10^9$	1 gigabyte (GB) = 1000 MB
Tera	$10^{12}$	1 terabyte (TB) = 1000 GB
Peta	$10^{15}$	1 petabyte (PB) = 1000 TB
Exa	$10^{18}$	1 exabyte (EB) = 1000 PB
Zetta	$10^{21}$	1 zettabyte (ZB) = 1000 EB
Yotta	$10^{24}$	1 yotta (YB) = 1000 ZB
Google	$10^{100}$	

The storage industry has been providing innovations at a breakneck pace to attempt to manage this data deluge. Expanding storage and the escalating pace of capacity and performance are among the true marvels of the computer age. Historians like to study the storage industry in particular because it is somewhat like studying the evolution of fruit flies, with companies innovating and the market reforming itself every few years to constantly embrace disruptive changes. Gordon Moore, the cofounder of Intel, predicted in 1965 that the number of transistors on a microprocessor would continue to double every 18 months for the same cost. Now known as Moore's law, this observation applies not only to microprocessors but to almost every facet of computing, especially storage. Over the last 50 years, the storage areal density (number of bits that can be squeezed into 1 square inch) has increased by a factor of 17 million at the same time that the vocabulary of storage terminology has expanded to accommodate previously undreamed of capacities (Table 16.1). Widespread agreement on the accuracy of Moore's law has led to an entirely new philosophy of equipment buying: if everything will be better and cheaper next year, then buy only what you need and make sure your system is designed to incorporate new technology that has not yet been invented.

Archiving technology is described as either online, near-line, or offline. Hard drives are online technology, because they are instantly accessible on a moment's notice. Tape and optical media are considered near-line because



they may need to be automatically retrieved from a robotic jukebox and mounted before being accessible. This generally takes 10 to 30 seconds if drives are available in the jukebox.

Offline storage requires manual intervention to load the media off a shelf and into the reader. As storage technology is becoming dramatically less expensive, offline storage should be discouraged. The true cost of offline storage is higher than most users perceive, because it includes the delay in time to retrieve data, the operator's time, and, most important, the much higher probability of losing data by mislabeling, misplacing, or damaging the media through improper storage.


## HARD DRIVES

The hard drive was invented at the research laboratories of IBM in 1952. The random access method of accounting control (RAMAC) stored its data in 50 stacked, 2-ft-wide aluminum platters. The platters rotated at 1200 revolutions per minute (rpm), and the system weighed more than 1 ton. At the time, RAMAC was considered a marvel because it could hold 5 megabytes (MB) of data—the storage equivalent of 50,000 punch cards.

### DRIVE MECHANICS

A hard drive consists of cylindrical platters, with thousands of single-bit-wide tracks of data (Figure 16.1). The surface of the platter is layered with a ferromagnetic material and polished so that the surface is extremely uniform. To read information off the surface of the hard drive, an electromagnet on a moving arm skims the surface of the platter at a height of less than 0.10mm. When the head is in read mode, it will induce a current on the head from the magnetic surface of the hard drive and thus be able to read information. When the hard drive is writing to the disk, the head is energized and changes the polarity of the magnetic surface of the platter. When the ferromagnetic material is exposed to an external field, it is permanently magnetized.

The hard drive surface is broken into rings, also known as tracks. Each track on the hard drive is broken into sectors. A platter can have upwards of 30,000 tracks and between 100 and 500 sectors per track. When a computer requests a file from the hard drive, the file is indexed by the track and sectors in which it is located. Sectors can be of various sizes, but, for large

**FIGURE 16.1**

Terminology and internal mechanisms of a hard drive.

file applications like a picture archiving and communications system (PACS), a typical sector size is 64 kilobytes (KB). The hard drive spins the platter underneath the head and rotates at a rate of 5400rpm for relatively slow drives and up to 15,000rpm for high-speed drives. The two key performance metrics for hard drives are defined as the seek time and the data transfer rate.

The seek time is the time that it takes to access the beginning piece of data requested. On average, this is the time it takes for half of a rotation to position the sector requested underneath the head (a lag known as rotational latency). This is dependent on the speed at which the hard drive is spinning. For a 5400-rpm hard drive, the access time would be calculated as:

$$\frac{1 \text{ minute}}{5400 \text{ rotations}} \times 0.5 \text{ rotation} = 5.6 \text{ milliseconds}$$

For a 15,000-rpm hard drive, the access time is significantly reduced:

$$\frac{1 \text{ minute}}{15,000 \text{ rotations}} \times 0.5 \text{ rotation} = 2.0 \text{ milliseconds}$$

A high-speed hard drive is ideal for applications such as databases and e-mail servers that can process more than 300 in-and-out requests per second from

a storage system. For storing large image files, however, it is arguable that a high overall transfer and throughput rate is more important than saving a few milliseconds on access time.

Data transfer rate is not quite as simple to calculate because it is dependent on the geometry of the hard drive and the location of the data. Data transfer rates are twice as great around the outer radius of hard drives because the circumference at the outer radius is twice that of the tracks around the inner radius. On average, high-rpm drives transfer data at 40 to 60 MB per second, and slower-rpm drives transfer data at 25 to 40 MB per second.

The highest-capacity hard drive released to date is the 400-GB hard drive by Hitachi (Tokyo, Japan). The 400-GB hard drive has 5 platters and 10 heads that read both sides of the platters. The Hitachi drive rotates at 7200 rpm, has an access time of 8.5 milliseconds, and is rated at an average transfer rate of 45 MB per second.

## DISK ARRAYS AND INTERFACES

Hard drive performance is often limited by the way in which the drive is connected to the computer. The hard drive is defined by this connector, called the interface. Many types of interfaces are available, but the principal ones for enterprise storage applications are serial ATA, small computer system interface (SCSI), and Fibre Channel. Parallel ATA, or AT attached, was originally named for the IBM AT computer and is the most common type of hard drive in the commercial PC market.

Serial ATA is an evolved version of parallel ATA that provides enterprise performance and reliability at commercial pricing. ATA drives are included in the majority of the PC market, with a volume of 10 times that of SCSI, which allows them a much lower cost per gigabyte of storage.

SCSI has historically been the hard drive style of choice for enterprise applications. Multiple hard drives can be daisy-chained off a single SCSI connection. Fibre Channel is an outgrowth of SCSI. In fact, it uses the same SCSI command set but transports it over a serial connection, whereas SCSI uses a parallel bus. Fibre Channel is an unfortunate misnomer in that it is a protocol and does not depend on fiber-optic connections at the physical layer. Although Fibre Channel can be run over copper wire, it is usually run over fiber optics because of the extremely fast speed that optical switching provides. Fibre Channel hard drives have the fastest interface and are also the most expensive type of hard drive on the market. Table 16.2 summarizes hard drive interface performance.

**TABLE 16.2**  
Hard Drive Interface Performance

<b>Interface</b>	<b>Theoretical Performance (MB/s)</b>
USB 1.0/1.1	1.5
Ultra ATA/33	33
IEEE 1394/Firewire	50
USB 2.0	60
Ultra ATA/66	66
Ultra ATA/100	100
Ultra ATA/133	133
Serial ATA/150	150
Ultra 160/SCSI	160
Serial ATA II/300	300
Ultra 320/SCSI	320
Fibre Channel	400+

*Source:* Connolly C. Highpoint 1520 RocketRAID Serial ATA/150 controller. Game PC. July 27, 2002. Available at: [www.gamepc.com/labs/view\\_content.asp?id=hpsataraid&page=2](http://www.gamepc.com/labs/view_content.asp?id=hpsataraid&page=2). Accessed September 15, 2004.

## **RELIABILITY AND CALCULATING MEAN TIME TO FAILURE**

Failure of computer components is a given and must be taken into account when designing an archive. Computer components with moving parts, such as the hard drive, have a higher probability of failure than solid-state components, such as the central processing unit and motherboard. Hard drive bearings can seize, heads can crash into platters, or the actuator arm can lock up. Most enterprise hard drives are rated with a mean life-to-failure of more than 1 million hours of operation. This does not mean that the average hard drive will actually last for 1 million hours. This model assumes a 5-year lifespan that is covered in the warranty and that all drives are replaced every 5 years. This is an important difference and means that a migration plan should be in place to prevent going over the warranty. Over the 5 years (43,800

hours) of warranty for the hard drive, there is a 4% failure rate. In other words, for every 22 hard drives running over the 5 years, expect at least 1 hard drive to fail.

## REDUNDANT ARRAY OF INEXPENSIVE DISKS

Redundant array of inexpensive disks (RAID) is one of the core concepts of implementing a high-performance, high-availability large storage solution. RAID allows multiple hard drives to work in concert and appear to the computer as a single storage device. Several RAID configurations are available and provide different combinations of redundancy and performance characteristics. To orchestrate the hard drives, a RAID controller is utilized between the hard drives and the computer. The RAID controls the hard drives directly and provides blocks of storage to the operating system.

RAID 0, striping, splits data to 2 hard drives simultaneously. The benefit to striping is that the storage system can be twice as fast at reading and writing data because it utilizes 2 drives. The problem with striping, however, is that there is no data redundancy. If either drive fails, all data on both drives are lost.

RAID 1, mirroring, copies the data to a shadow hard drive. RAID 1 provides complete redundancy in the event that 1 of the hard drives fails but does not enjoy any performance benefits because the user is still reading from a single drive. The other challenge to RAID 1 is effectively the loss of the capacity of the second drive.

RAID 10 is called a nested strategy and combines RAID 1 and RAID 0 levels. Two hard drives are striped, and both of those hard drives have a mirror copy in case of failure. This simple technique provides the performance benefits of 2 drives while allowing for failure. The drawback of RAID 10 is that half the storage is used making a backup of the data.

RAID 3 requires a minimum of 3 hard drives and is a combination of striping and a checksum. As data are split among multiple hard drives, a checksum, or parity bit, is calculated. If any of the drives fails, the information can be reconstructed onto a new hard drive from this checksum. RAID 3 uses a fixed parity drive with synchronized disk rotation and calculates checksums at the bit level. The challenge to parity drives is that the capacity of 1 drive is lost and that writing data to the disk array can be slower because it is calculating the parity bit. The read speed from RAID 3 is quite fast because the data is striped over all the drives in the array.

RAID 4 is a variant of RAID 3. RAID 4 calculates the parity bit on a dedicated drive at the block level instead of at the lower bit level. This

improves the random access performance of the disk array compared with RAID 3.

RAID 5 is a variant of RAID 4 that also calculates the parity at the block level. However, RAID 5 does not use a dedicated drive to store the checksum but, instead, distributes the parity information among all the drives; this improves random access write performance. RAID 5 has relatively poor write performance but excellent read performance that matches well to the clinical requirements of PACS. Performance truly matters with PACS when a user selects a patient name on a list and wants instantly to read large image studies. PACS is considered a read-intensive application in that every study that is written to the system is read 5 to 10 times in its lifetime for diagnosis, relevant priors, enterprise distribution, archiving, and migration.

A new technique not yet in common use is RAID 6. RAID 6 is similar to RAID 5 but calculates 2 parity checks that allow the system to suffer any 2 hard drive failures without compromise. The catch is the loss of capacity of 2 drives on the system.

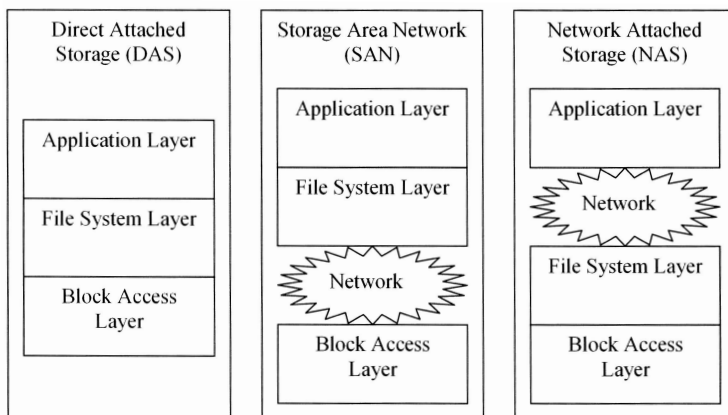
In clinical practice, RAID 5 is the predominant technique for retaining medical images on storage arrays. RAID 5 is considered the ideal configuration for the PACS industry because it provides good performance through striping across multiple drives and at the same time provides redundancy and can survive losing the capacity of 1 drive in the array. With 1 drive out, however, the system must run in compromised condition in which another drive failure would result in the loss of the entire RAID array. That is why it is always good policy to have an extra hard drive sitting on a RAID array to act as a hot spare. During the event of a failed drive, the system will activate the spare and start rebuilding the spare drive. The time it takes to rebuild the spare hard drive and restore the RAID is called the mean time to recovery (MTTR). This recovery rebuild rate is usually in the range of 20 to 50GB/hour. For a 200 to 300GB drive, this exposed time could be several hours. A RAID 5 system is fully operational to external requests during the rebuild time, although its performance will be degraded because it is rebuilding a drive in the background. The probability of another failure during the rebuild time is extremely remote. When a system seems to go against the odds by having more failures than are statistically probable, environmental issues, such as high internal temperature or excessive vibration in the system rack, may be the cause. One common myth is that hard drives and system components fail without any warning signs. A study performed in 1999 tracked the system log of 368 hard drives and showed 4 hard drive failures over an 18-month period. The drives started warning the system between 5 and 186 hours before each failure, generating on average more

than 1000 messages. It is important to have a systems management tool in place and adequately configured so that the administrator can be alerted to the problem and be proactive in finding a solution.

The most common systems management tool is the simple network management protocol (SNMP). This protocol allows health information about a server to be communicated between systems to provide a holistic view of the status of the entire system. SNMP tools can track fans, hard drives, and power supplies for failure and can e-mail administrators in the event of failure. These tools also can routinely monitor the temperature of these components, as well as the drawing power used. Limits can be set for identifying impending problems. The Storage Management Initiative Specification (SMI-S), an open standard developed by the Storage Network Industry Association (SNIA), has been recently introduced. Whereas SNMP is developed around a passive mode for catching error messages, SMI-S is designed to be more active, allowing administrators to control and reconfigure enterprise storage devices.

## ROLE OF THE FILE SYSTEM

Now that all your hard drives are sitting in a nice disk drive rack mount array, how does the PACS application connect to and make use of them? Sitting between an application and the low-level block access storage system lies the file system, as shown in Figure 16.2 The file or filing system is the



**FIGURE 16.2**

File system.

card catalog of a storage system and knows where files are stored and presents a hierarchical view that helps users locate them. When searching through the treelike structure of your file system in an application such as Windows Explorer, you are viewing the file system and not actually scanning through the entire hard drive. When you defragment a hard drive, the file system reorganizes the hard drive to put files into contiguous sectors for faster reading. The common file systems for Windows are FAT16, FAT32, and NTFS. NTFS is the newest file system from Microsoft and has the best features for PACS applications. NTFS can address up to 256 TB of data using a 64-KB cluster size and can support up to 16-TB files. FAT32 can address up to only 32 GB, and the largest file it can support is 2 GB.

## **DIRECT ATTACHED STORAGE, STORAGE AREA NETWORKS, AND NETWORK ATTACHED STORAGE**

Direct attached storage (DAS) is using the hard drives that come with the server (Figure 16.3). Storage area networks (SANs) and network attached storage (NAS) are the only 2 methods for scaling storage by inserting a network topology at either the file level (NAS) or the block level (SAN) between the storage system and the application.


### **DIRECT ATTACHED STORAGE**

A 5U server might have the capacity for 10 to 12 drives to be added, typically yielding a capacity of around 2 to 3 TB. This is the least expensive approach for small storage requirements but also provides the least amount of redundancy. If any component fails on the server and stops the application, access to the data will also be lost. Another challenge of DAS is the likelihood of running out of space. A direct attached solution does not scale well.

### **STORAGE AREA NETWORK**

Storage area network refers to a network dedicated to storage access at the block level (Figure 16.3). Networking typically involves Ethernet and



**FIGURE 16.3**

Simple example of direct attached storage (DAS), storage area network (SAN), and network attached storage (NAS).

TCP/IP standards, but a SAN uses neither. A SAN uses the Fibre Channel communications protocol to directly access storage devices. The file system remains on the computer, but low-level block-level access is routed through an external network of storage devices. These external storage devices can be daisy-chained together into an arbitrary loop or can be driven from a star topology using a SAN switch.

A SAN is perfectly suited for a busy information services department running dozens or even hundreds of servers. A SAN can be used to handle the storage needs of servers running different operating systems, although each system sees only its own blocks of data. Instead of having to monitor and manage the storage needs on each of those computers, storage management can be consolidated into a SAN and managed centrally. A SAN puts the storage in one place to manage and can automatically grant computers more storage on demand when it sees them approaching their storage limits. This provides significant economies in the operating costs for the servers. A SAN can provide logic at the block level to copy data to different storage devices and provide a transparent level of fault tolerance to the computers as well. Another unique advantage of SAN is that it can move and back up data from 1 node of the SAN (such as a RAID array) to another node (such as a tape archive) without drawing any resources from the host computers. This is called a clientless backup.

A SAN is the most frequently used storage method for PACS database redundancy. With fault-tolerant clustering, 2 servers share the same SAN partition, although 1 server sits passive. The 2 servers are linked via a heart-beat connection that monitors the active server and, in the event of failure, alerts the passive server to assume the network name of the server and take over the shared storage partition and database processes.

SAN is the highest performance type of scalable storage architecture because it provides such low-level access—but it is also the most expensive. One reason is that optical connections are more costly than copper ones. In addition, the storage network industry is orders of magnitude smaller than the Ethernet network industry and so cannot provide buyers the benefits of economies of scale. Because it is not used as widely, the Fibre Channel standard is not as plug-and-play as networking standards. Some incompatibilities are often intentional. SAN providers often include proprietary code for the following: performance/value-added features (such as multihost bus adapter [HBA], the card that connects the server to the SAN), unique software for caching read/write requests in fast memory, replication (backup to another location) technology, fault monitoring, and SAN management. A SAN can become an engineering exercise and require complex installation

and ongoing maintenance. Expertise in Fibre Channel is not as common in information technology organizations as networking expertise and is therefore more costly. Figure 16.4 is a PACS example of a multivendor SAN implementation.

To alleviate these problems, an exciting and potentially disruptive technology has been developed: Internet SCSI or iSCSI. Internet SCSI is a communications protocol that sends and receives low-level SCSI

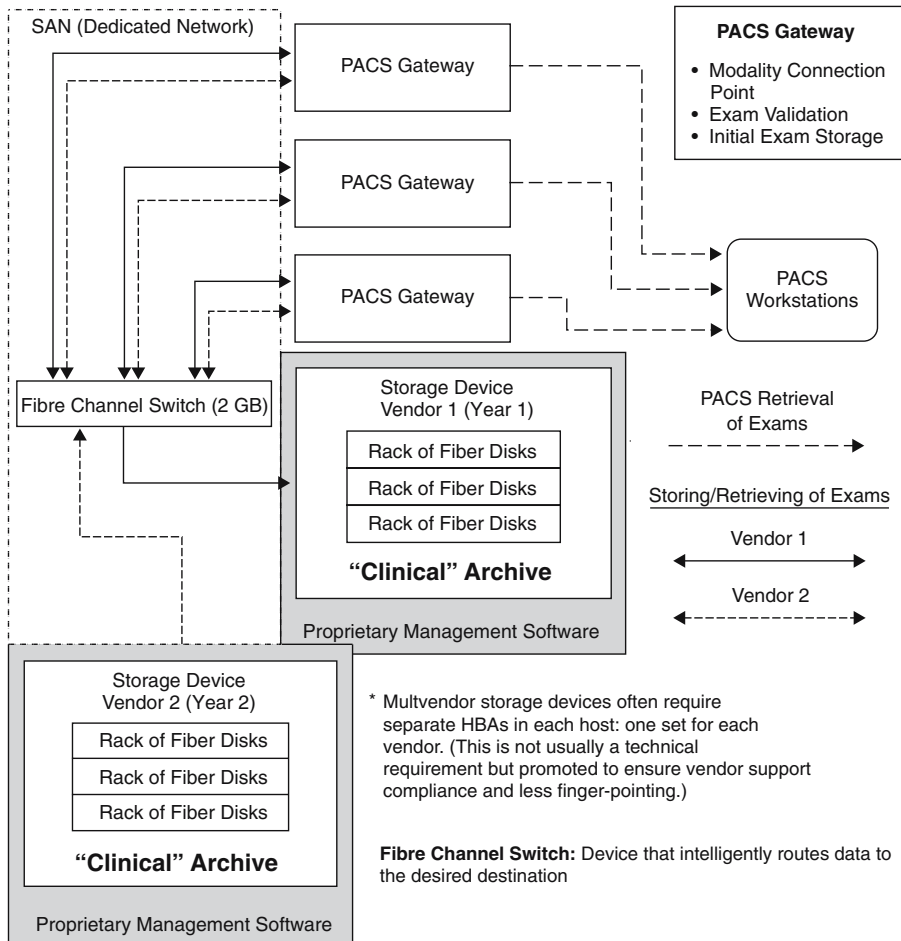


FIGURE 16.4

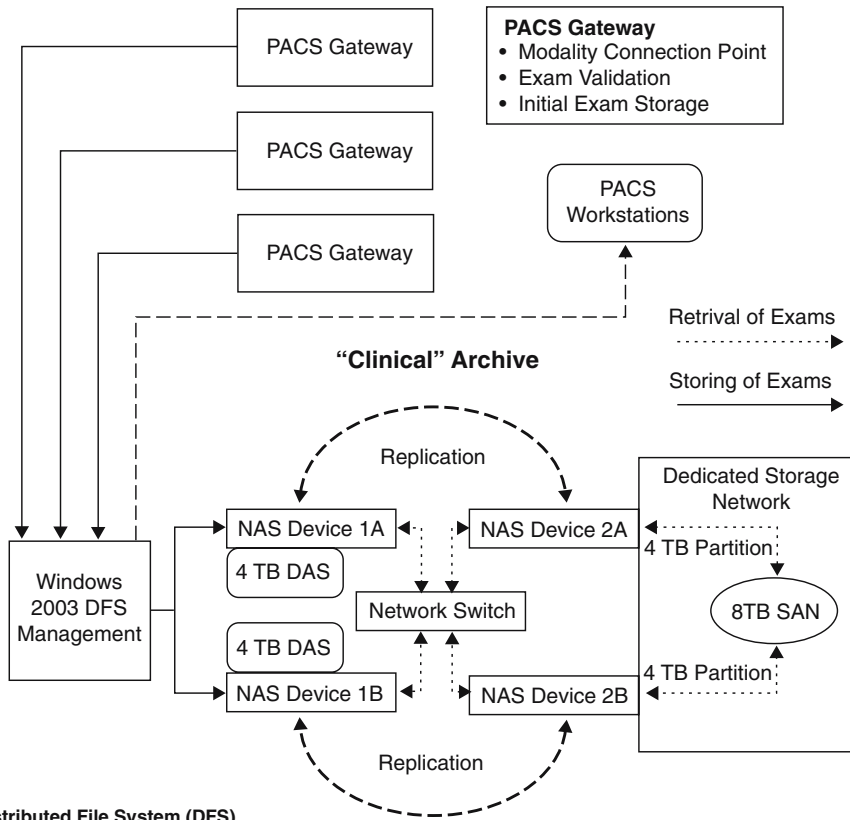
Sample PACS (SAN) implementation.

commands on a copper Category 5 Ethernet cable sitting on top of TCP/IP routing protocols. Internet SCSI can use the same inexpensive cabling as Ethernet networking and even sit on an Ethernet network. Internet SCSI is giving rise to what is being referred to as IP SANs, and it is predicted that this strategy will drastically reduce the cost of SANs installation and support with the introduction of plug-and-play (multivendor) components.

## NETWORK ATTACHED STORAGE

The third technique for adding storage is by working at the file level and using NAS (Figure 16.5). NAS servers, also known as filers, are dedicated to providing remote file-level access. PACS applications can access storage sitting on NAS servers using a high-level file-level access protocol called the common Internet file system (CIFS; Microsoft Windows-type file sharing). The benefit of using CIFS is that it is a widely used standard in the industry and is platform independent. Linux and Unix servers work extremely well as CIFS file servers using the powerful and open-source Samba server. CIFS uses the universal naming convention (UNC) for addressing files, such as: `\\servername\share\path\filename`. The first portion of the UNC is the server name or network address at which the NAS server resides.

NAS is a relatively easy way to scale a storage system when the PACS application supports UNC. The database keeps a UNC pointer to the locations at which the files reside. One potential drawback of the NAS architecture is that data are sent on the same network twice, with one hop sending data from the NAS to the server and the second hop sending data from the server to the end-client workstation. This contrasts with a SAN, which typically uses a dedicated Fibre Channel network to afford fast, direct access to the storage devices. Some PACS applications avoid this potential problem by passing the UNC address to the end workstation to directly access the files from the NAS server. This architecture scales well, because it avoids having all system file access go through a central point that becomes a potential bottleneck. As an example, suppose there are 5 NAS devices, each containing 1 TB of storage on a 1 GB network, and the PACS software is architected to spread the storage load among the 5 NAS devices. The aggregate network bandwidth would be  $5 \times 1 \text{ GB}$  or 5 GB, and the aggregate storage would be 5 TB. Both the amount of storage and the pipeline to the storage increase proportionally as (vendor-agnostic “inexpensive”) storage devices are added.



**Distributed File System (DFS)**

DFS redirects requests for files to assigned NAS devices. The primary advantage of redirecting client requests is to take advantage of replication. If two or more NAS devices are replicated and the primary NAS device fails, DFS will automatically direct all future client requests to one of the still functioning backup NAS devices.

**Process Overview**

In this example there are 4 NAS access points. Because the NAS devices present the underlining storage as a file system, the technology utilized by physical storage devices is irrelevant. To illustrate this point, this example utilizes 2 distinctly separate technologies (DAS and SAN). A simplistic way to look at this is that the operating system presents all disk storage as a file system. Storage vendors write plug-in software that interfaces the storage device to the operating system. This software is usually referred to as "device drivers." Therefore, if the operating system can mount (use) the storage device, it can then, in turn, offer the file system to other users on the network (in Windows-speak, file share).

**FIGURE 16.5**

Sample PACS (NAS) implementation.

Another exciting NAS-related technology is contained in new features offered by the distributed file system (DFS) that is standard with the Microsoft Windows 2003 operating system. A distributed file system can make several disparate NAS servers appear logically grouped and present them to the PACS application as a single source. A distributed file system provides replication between NAS servers that can act as a fault-tolerant backup to switch over to in the event 1 of the NAS servers fails. These features were once available only from high-end storage vendors.

Challenges involved in supporting a large NAS environment include the installation/support of replication/backup software, monitoring tools, antivirus measures (including daily signature updates), and, of course, those pesky service patches (because a NAS server is running an operating system with a file system, both will need to be maintained on all the NAS servers to avoid security vulnerabilities and file corruption).

## GRID STORAGE

Grid storage creates fault-tolerant storage pools based on commodity “inexpensive” servers, often referred to as grid servers (Figure 16.6). Each grid server usually contains 4 internal serial ATA disk drives, 1 processor, and a modest amount of memory. Each of the servers participating in the grid runs vendor proprietary software that allows the devices to act as one. At least 1 server acts as the interface between the grid and external application (PACS) servers. This grid’s interface server appears to the PACS servers as a NAS device presenting a standard file system. The interface server also hosts hierarchical storage manager (HSM)-like software that manages the location of files based on user configurable policies. These policies allow for defining 1 –  $n$  copies of a file, with each copy stored on a different grid server. This is how redundancy/fault tolerance is achieved. If 1 or more of the grid servers fails, the data can still be retrieved to the interface node (standard file system) from a copy located on a functioning grid server.

## CONTENT ADDRESSABLE STORAGE

Content addressable storage (CAS) is functionally the same as grid storage. However, instead of presenting files to the application (PACS) servers via NAS, a proprietary interface is utilized (Figure 16.6). Moreover, the integrity of the files is guaranteed by elaborate algorithms that provide a unique value (token) based on the data in the file. The algorithm commonly employed is

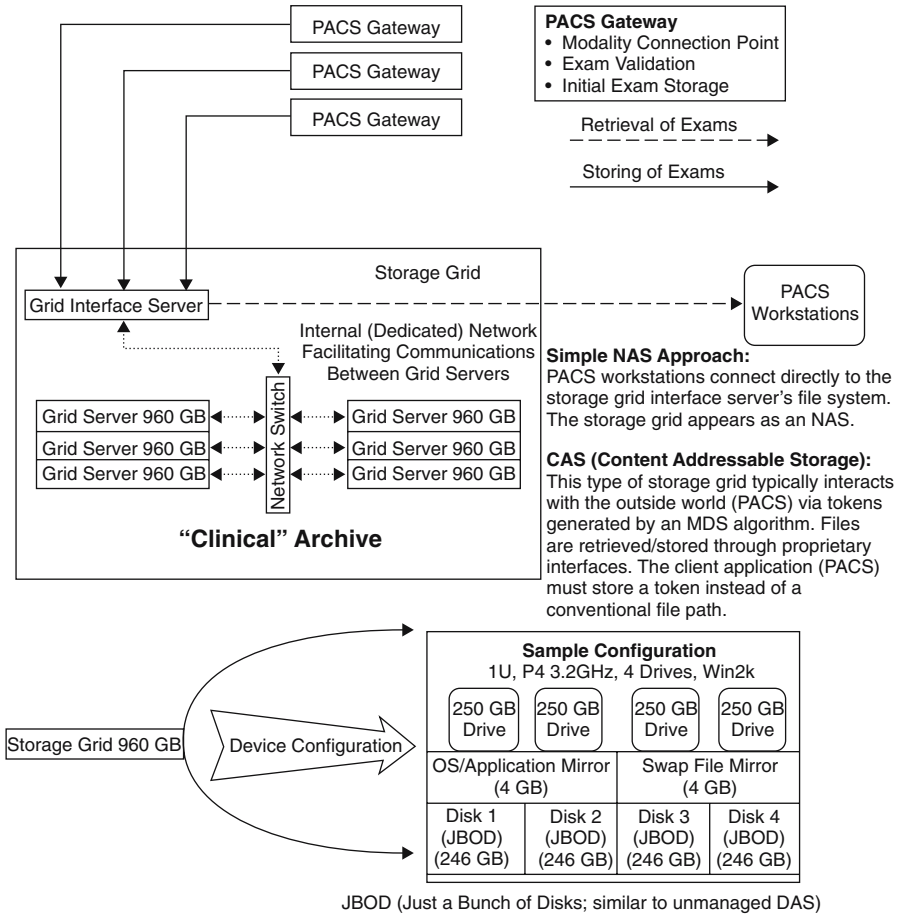


FIGURE 16.6

Grid storage.

MD5 (message-digest algorithm). The generated value acts as a unique key for retrieving the file and must be stored by the (PACS) application. If the data on the storage grid somehow changed, the file's unique value would no longer match the token stored in the (PACS) application database, resulting in an error being sent from the storage grid back to the requesting (PACS) application. This approach also supports the ability to logically group files together as related objects and store metadata (data about the data; for example, to store date, last access date, or information relating to the sources of the data.)

## FOOD AND DRUG ADMINISTRATION AND STORAGE

A common myth is that storage is a U.S. Food and Drug Administration (FDA)-approved device and hence can be purchased only through a PACS vendor. The FDA has clarified this issue by stating that devices that need clearance are those that alter image data, such as image-processing algorithms. The bottom line is that storage does not have to be purchased from PACS vendors, especially when markup costs are large. This freedom to purchase is particularly useful in a NAS architecture where the PACS application is encapsulated from the storage by the CIFS standard. This should shorten the validation cycle time for testing new storage solutions and allow customers more flexibility in enjoying continued innovation from the storage industry.

## NEAR-LINE TECHNOLOGY

### OPTICAL STORAGE

Optical storage can also be used as an archive device for storing data. The compact disc (CD) is an aluminum disc encased in a transparent plastic layer, with binary information encoded as a series of bumps that start at the center of the disc and spiral out to the ends. A laser is focused onto the spiral path as the CD spins. In the absence of a bump, the laser is reflected back onto a photodetector that records a hit. Bumps on the disc cause the laser to deflect, and the absence of a hit on the photo detector is recorded as a miss (or 0). These series of hits and misses become 1s and 0s that form a massive data stream. The spiral path is a mere 50 $\mu$ m wide and is separated from adjacent paths by just 1.6 $\mu$ m. The linear length of a standard 700-MB, 12-cm CD path is more than 3.5 mi long. Commercial CD burning technology can make CDs by using an aluminum disc that is covered by a layer of photosensitive material. A write laser can stimulate the photosensitive material, which becomes opacified. Thus, when a read laser hits that area, there is no reflection from the underlying aluminum plate and a miss is recorded.

Digital versatile discs (DVDs) use the same underlying technology as CDs. The difference is that a DVD uses a higher-frequency laser that provides a smaller area on which to focus. This means that the bumps can be



smaller and packed more closely together. A DVD has a spiral path that is a mere 320nm wide, separated from adjacent tracks by 740nm. A DVD can store 4.7GB of data, 7 times the capacity of CDs. DVDs can be written on both sides, providing twice the capacity (9.4GB). A newer technology allows a DVD to store 2 layers on the same side by having a semitransparent layer made from gold overlaid on top of the base aluminum layer. A laser can focus either on the gold layer or through it on the aluminum layer. This technique requires slightly longer bump sizes to ensure accuracy, so the total capacity is slightly less than twice that of a regular DVD. A double-sided, double-layer DVD can hold 15.9GB of data.

An interesting hybrid between magnetic and optical technology is magneto-optical (MO) drives. An MO drive writes to the media using a magnetic field that thermally changes the properties of the material. When polarized light is shone on the material, the material shifts the frequency of the reflected light. This gives MO faster write speeds than typical optical drives, up to 4MB/s. MO disks are encased in a hard envelope and can be selected as rewriteable or “write once read many” (WORM). The WORM option provides an exceptionally long lifetime for the media, projected at more than 60 years. Magneto-optical media come in 5.2- and 9.1-GB capacities but cost several times as much per gigabyte as DVD media.

If optical media are selected for archive technology, it is likely that a jukebox will be part of the solution. A jukebox system uses a robotic arm to load and unload optical media from readers and stores the media in slots. Optical media are sensitive to extremes in temperature, humidity, and direct exposure to sunlight. A jukebox should be used, especially for exposed media such as CDs and DVDs that do not have enclosed hard cases and so are susceptible to scratches from manual handling. Today, large optical jukebox systems can hold up to 2000 slots and as much as 20TB of storage using double-sided DVD or MO technology.

## TAPE

Tape is based on the same magnetic technology used for hard drives. The difference is that tape encases the magnetic material in a flexible plastic strip that is 8 to 12mm wide and holds multiple tracks of data. Tape is a sequential access device, and hard drives and optical media are random access devices. High-capacity cartridges hold more than 2000 linear feet of tape. A tape must wind to (seek) the location at which data are stored. This can take on the order of 30 to 60 seconds for larger-capacity tapes,

**TABLE 16.3**  
Tape Capacity and Performance Matrix

<b>Tape Type</b>	<b>Uncompressed Capacity (GB)</b>	<b>Transfer Rate (MB/s)</b>	<b>Transfer Rate (GB/h)</b>	<b>Average File Access Time (Seconds)</b>
DLT	40	3	10.8	68
SDLT 320	160	16	57.6	70
SDLT 600	300	36	129.6	79
AIT-1	35	4	14.4	27
AIT-2	50	6	21.6	27
AIT-3	100	12	43.2	27
SAIT	500	30	108.0	70
LTO-1	100	15	54.0	52
LTO-2	200	35	126.0	52
SLR50	25	2	7.2	55
SLR100	50	5	18.0	58

whereas hard drives and optical drives can do this in well under a second (Table 16.3). The engineering trick is to wind the tape quickly with a minimal amount of tension that might cause wear and eventual failure. In a jukebox, an additional 10 to 30 seconds is usually needed to pull a tape from a slot and mount it into a reader. Tape jukebox systems have a MTBF (mean time between failures) of 200,000 hours but, in many cases, remain operational during component failures. Tape heads are the most common component to fail in a MTBF with approximately 40,000 hours of operation. In the event of a tape head failure, another drive unit can simply read the tape media.

Many different formats of tape are offered by vendors and/or consortiums of vendors. Formats are usually backwards-compatible over 2 or 3 previous generations. Be aware that several tape formats use the compressed capacity as the name of the format. For example, SDLT 320 holds only 160GB of uncompressed data and 320GB of data when compressed losslessly by a factor of 2.

## ROLE OF THE HIERARCHICAL STORAGE MANAGER BETWEEN ONLINE AND NEAR-LINE STORAGE

The hierarchical storage manager software sits on servers, manages storage located in various storage devices, and provides a single view to the end user. The HSM has traditionally worked in PACS with a small portion of RAID and a large tape/optical jukebox. The HSM manages the process of synchronizing files on the RAID to 1 or more cartridges of removable media in the jukebox. As the RAID fills up past a point, known as a high-level watermark, the HSM will then flush files (a process where all but metadata about the file are deleted—enough information is retained on the RAID to know the location on a given cartridge where the file can be retrieved) from the RAID onto the near-line storage device.

Requests from the PACS for exams are intercepted by the HSM software and, if the exam has been flushed off the RAID, the HSM will restore the exam back to the RAID and then satisfy the request by allowing the requesting client to access the file. To help eliminate latency for prior exam requests, PACS pre-fetching algorithms use the scheduled procedure list for the next day to select and restore those patients' prior studies back to the RAID well before the newly scheduled procedure. For PACS applications, there is an 80% probability that a relevant prior is less than 6 months old. Physicians unlucky enough to request a patient's study that is not on the RAID typically have to wait between 40 seconds and several minutes to gain access to those patients' files.

There are 2 challenges to this multitier architecture. The first lies in understanding the trade-off that underlies the assumption behind HSMs: hard drives are fast but expensive, and the jukebox is slow but less expensive. This assumption is no longer compelling, as online storage devices continue to drop in price at a faster rate than enterprise jukebox technology. The second challenge is that predicting clinical need is not a perfect science. Waiting several minutes for a study to load from a jukebox when a clinician's expectation is that it should take only a few seconds might create a number problem reports to the PACS administrators.

As PACS vendors move toward storing all exams on spinning-disk solutions, the HSM is being reborn under a new name with a slightly different function. The new buzz names for HSM are life cycle storage management (LSM) or storage virtualization. This new approach to HSM helps manage the life cycle of the storage architecture and protect and maintain continuous availability of data. When new storage is added to the system, the storage

manager(SM) can have that additional storage appear seamlessly to the PACS application as 1 pool. An SM can also copy data from 1 disk array onto a new storage device in the background. This way, storage devices can be retired seamlessly, rebuilding and renaming new devices to take over their predecessors' roles. This can be very useful in ensuring that storage devices remain serviceable by the storage provider and that the department does not get caught with antiquated, unsupported hardware. The HSM storage manager can be utilized to duplicate copies of exams stored in separate geographic locations to provide disaster recovery as well as better local performance. The logic of the SM can be placed at any level of storage architecture. It can reside as part of a SAN (integrated into the storage device level directly) or it could reside in the operating system at the file level. The SM could also reside as part of the PACS application. Since the functions of storage virtualization extend well beyond the PACS arena, this logic should ideally be processed at the file or block level and not by the application. This gives the PACS vendor one less piece of the puzzle to figure out and also allows customers to enjoy the economy-of-scale benefits of these applications that are used well beyond the medical community. This also gives more freedom to customers to select the best storage approach for their needs.

## CONCLUSION

Storage is a fundamental enabling technology of the PACS industry. The only way to position ourselves to survive the onslaught of storage consumption is to be able to embrace the advances being made in the storage industry. Look for the continued trend toward commoditization through the use of well-accepted and proven industry standards such as copper Category 5 Ethernet network cable for storage networks in IP-based SANS. Expect PACS vendors to provide better encapsulation of their application from the storage architecture to allow customers to better enjoy the technological and economic benefits of this competitive market.

## ABBREVIATIONS

**AIT**    Advanced intelligent tape

**ATA**    AT advanced

**BIT** Binary representation (0 or 1)

**BYTE** 8 bits

**CAS** Content addressable storage

**CD** Compact disk

**CIFS** Common Internet file system

**DAS** Direct attached storage

**DAT** Digital audio tape

**DFS** Distributed file system

**DLT** Digital linear tape

**DVD** Digital video disc

**FC** Fibre Channel

**FCIP** Fibre Channel over IP

**Gigabyte 1000 megabytes**

**HBA** Host bus adapter

**HSM** Hierarchical storage manager

**IDE** Integrated drive electronics

**ISCSI** Internet SCSI

**JBOD** Just a bunch of disks

**LAN** Local area network

**LTO** Linear tape open

**MTTF** Mean time to failure

- NAS** Network attached storage
- RAID** Redundant array of inexpensive disks
- SAIT** Super advanced intelligent tape
- SAN** Storage area network
- SAS** Serially attached SCSI
- SATA** Serial ATA
- SCSI** Small computer system interface
- SMB** Server message block
- SMI-S** System management interface
- SNMP** Simple network monitoring protocol
- WAN** Wide area network

## ► REFERENCES

- Christensen C. *The Innovator's Dilemma*. New York; HarperCollins; 2003:3–68.
- Coufal H, Grochowski E. The future of data storage, principles, potential, and problems. FAST Conference on File and Storage Technologies. January 2002. Available at: [www.usenix.org/publications/library/proceedings/fast02/coufal.pdf](http://www.usenix.org/publications/library/proceedings/fast02/coufal.pdf).
- Farley M. *Building Storage Networks*. 2nd ed. New York: McGraw-Hill; 2001: 160–207.
- Gardner S, Hughes K. FDA ruling opens door to new PACS storage options. Auntminnie.com. July 2003. Available at: [www.auntminnie.com/default.asp?Sec=sup&Sub=pac&Pag=dis&ItemId=58787&d=1](http://www.auntminnie.com/default.asp?Sec=sup&Sub=pac&Pag=dis&ItemId=58787&d=1). Accessed September 15, 2004.
- Hennessy J, Patterson D. *Computer Architecture: A Quantitative Approach*. 3rd ed. New York: Morgan Kaufman Publishers; 2003:676–785.
- Hitachi Global Storage 2004. Available at: [www.hitachigst.com/hdd/support/7k400/7k400.htm](http://www.hitachigst.com/hdd/support/7k400/7k400.htm). Accessed September 15, 2004.
- Lyman P, Varian H. How Much Information? 2003. Available at: [www.sims.berkeley.edu/research/projects/how-much-info-2003/index.htm](http://www.sims.berkeley.edu/research/projects/how-much-info-2003/index.htm). Accessed August 4, 2004.

- Murray W, Maekawa K. Reliability evaluation of 3M magneto-optic media. Proceedings of Magneto-Optical Recording International Symposium. *J Magn Soc Jpn.* 1996;20(suppl S1):309–314.
- Nagy P, Farmer J. Demystifying data storage: archiving options for PACS. *Appl Radiol.* May 2004:18–22.
- Nice K. How DVDs work. Available at: <http://electronics.howstuffworks.com/dvd.htm>. Accessed September 15, 2004.
- Webster J. Everything you need to know about SMI-S. *Infostor.* 2004;8:28–32.

## ► SUGGESTED READING

- Anderson M. The toll of downtime. *Healthcare Informatics.* April 2002. Available at: [www.healthcare-informatics.com/issues/2002/04\\_02/leading.htm](http://www.healthcare-informatics.com/issues/2002/04_02/leading.htm). Accessed September 15, 2004.
- Andriole KP, Avrin DE, Yin L, Gould RG, Luth DM, Arenson RL. Relevant priors prefetching algorithm performance for a picture archiving and communication system. *J Digit Imaging.* 2000;13(2 suppl 1):73–75.
- Avrin DE, Andriole KP, Yin L, Gould R, Arenson RL. Simulation of disaster recovery of a picture archiving and communications system using off-site hierarchal storage management. *J Digit Imaging.* 2000;13(2 suppl 1):168–170.
- Avrin DE, Andriole KP, Yin L, Gould RG, Arenson RL. A hierarchical storage management (HSM) scheme for cost-effective on-line archival using lossy compression. *J Digit Imaging.* 2001;14:18–23.
- Battaglia G, Maroldi R, Chiesa A. Evaluating the digital storage requirements for a partial picture archive and communication system. *J Digit Imaging.* 1990;3:54–59.
- Bauman RA, Gell G, Dwyer SJ III. Large picture archiving and communication systems of the world. Part 2. *J Digit Imaging.* 1996;9:172–177.
- Behlen F, Maughan J. Migrating data to the new PACS. *Health Imaging & IT.* May 2004. Available at: [www.healthimaging.com/archives/HIIT/HIIT\\_2004/HIIT0205/HIIT020510.htm](http://www.healthimaging.com/archives/HIIT/HIIT_2004/HIIT0205/HIIT020510.htm). Accessed September 15, 2004.
- Benner A. *Fibre Channel for SANS.* New York: McGraw-Hill; 2001:5–12.
- Blado ME. Management of the picture archiving and communications system archive at Texas Children's Hospital. *J Digit Imaging.* 2001;14(2 suppl 1):84–88.
- Blado ME, Tomlinson A. Monitoring the accuracy of a PACS image database. *J Digit Imaging.* 2002;15(suppl 1):87–95.
- Bui AA, McNitt-Gray MF, Goldin JG, Cardenas AF, Aberle DR. Problem-oriented prefetching for an integrated clinical imaging workstation. *J Am Med Inform Assoc.* 2001;8:242–253.

- Chesson E. Long-term image archiving. *Health Imaging & IT*. June 2004. Available at: [www.healthimaging.com/archives/HIIT/HIIT\\_2004/HIIT0206/HIIT020607.htm](http://www.healthimaging.com/archives/HIIT/HIIT_2004/HIIT0206/HIIT020607.htm). Accessed September 15, 2004.
- Channin D, Parisot C, Wanchoo V, Leontiev A, Siegel E. Integrating the Health-care Enterprise: a primer: part 3. What does IHE do for ME? *Radiographics*. 2001;21:1351–1358.
- D'Arcy S, D'Arcy J. Securing data from disaster. *Decis Imaging Economics*. June 2003. Available at: [www.imagingeconomics.com/library/200306-02.asp](http://www.imagingeconomics.com/library/200306-02.asp). Accessed September 15, 2004.
- Dumery B. Digital image archiving: challenges and choices. *Radiol Manage*. June 2002. Available at: [www.ahra.com/AHRAArticles/AHRAArticles.dll/Show?ID=328](http://www.ahra.com/AHRAArticles/AHRAArticles.dll/Show?ID=328). Accessed September 15, 2004.
- Erickson BJ, Persons KR, Hangiandreou NJ, James EM, Hanna CJ, Gehring DG. Requirements for an enterprise digital image archive. *J Digit Imaging*. 2001; 14:72–82.
- Henderson M, Behlen F, Parisot C, Siegel E, Channin D. Integrating the Health-care Enterprise: a primer: part 4. The role of existing standards in IHE. *Radiographics*. 2001;21:1597–1603.
- Honeyman JC, Huda W, Frost MM, Palmer CK, Staab EV. Picture archiving and communication system bandwidth and storage requirements. *J Digit Imaging*. 1996;9:60–66.
- McEnery KW, Suitor CT, Thompson SK, Shepard JS, Murphy WA. Enterprise utilization of “always on-line” diagnostic study archive. *J Digit Imaging*. 2002; 15(suppl 1):81–86.
- Mendenhall R, Dewey M, Soutar I. Designing fault-tolerant distributed archives for picture archiving and communication systems. *J Digit Imaging*. 2001;14(2 suppl 1):80–83.
- Nagy P. The ABC's of storage technology. *SCAR University: Educating Healthcare Professionals for Tomorrow's Technology*. Great Falls, VA: Society for Computer Applications in Radiology; 2004:123–127.
- Nagy P, Warnock M. Predicting PACS loading and performance metrics using Monte Carlo and queuing methods. In: *Proceedings of the International Society for Photo-Optical Engineering*. San Diego, CA: SPIE; 2003:46–55.
- Nagy P, Daly M, Warnock M, Ehlers K. PACSPulse: A Web-based DICOM network traffic monitor and analysis tool. *Radiographics*. 2003;23:795–801.
- National Electrical Manufacturers Association. DICOM resources 2004. Available at: <http://medical.nema.org>. Accessed September 15, 2004.
- National Institute of Standards and Technology. Prefixes for binary multiples. Available at: <http://physics.nist.gov/cuu/Units/binary.html>. Accessed September 15, 2004.
- Oosterwijk H. *PACS Fundamentals*. Aubrey, TX: Otech, Inc; 2004:53.



- Pavlicek W, Zavalkovskiy B, Eversman WG. Performance and function of a high-speed multiple star topology image management system at Mayo Clinic Scottsdale. *J Digit Imaging*. 1999;12(2 suppl 1):168-174.
- Piedad F, Hawkins M. *Highb Availability: Design, Techniques, and Processes*. Upper Saddle River, NJ: Prentice Hall; 2001:13-29.
- Qi H, Snyder WE. Content-based image retrieval in picture archiving and communications systems. *J Digit Imaging*. 1999;12(2 suppl 1):81-83.
- Shelton PD, Lyche DK, Norton GS, et al. Benchmark testing the Digital Imaging Network-Picture Archiving and Communications System proposal of the Department of Defense. *J Digit Imaging*. 1999;12:94-98.
- Siegel E, Channin D. Integrating the Healthcare Enterprise: a primer: part 1. Introduction. *Radiographics*. 2001;21:1339-1341.
- Smith EM, Wandtke J, Robinson A. The strategic and operational characteristics of a distributed phased archive for a multivendor incremental implementation of picture archiving and communications systems. *J Digit Imaging*. 1999;12(2 suppl 1):71-74.
- Smith EM, Wright J, Fontaine MT, Robinson A. Archive selection for the MICAS, a multi-vendor incremental approach to PACS. *J Digit Imaging*. 1998;11(3 suppl 1):32-34.
- Smith R. Rethinking disaster recovery. *Imaging Econ*. December 2001. Available at: [www.imagingeconomics.com/library/200112-06.asp](http://www.imagingeconomics.com/library/200112-06.asp). Accessed September 15, 2004.
- Tamm E, Thompson S, Venable S, McEnery K. Impact of multislice CT on PACS resources. *J Digit Imaging*. 2002;15(suppl 1):96-101.
- Wong AW, Huang HK. Subsystem throughputs of a clinical picture archiving and communications system. *J Digit Imaging*. 1992;5:252-261.
- Wong AW, Huang HK, Arenson RL, Lee JK. Digital archive system for radiologic images. *Radiographics*. 1994;14:1119-1126.

# IMAGE DISPLAYS

---

---

DAVID S. HIRSCHORN

## HISTORY

Until the advent of picture archiving and communication systems (PACS) in the early 1990s, film was the primary medium for radiologic image display. At the end of the 19th century, the first radiographs were acquired on a piece of film. Radiography was the first modality of radiology, and even after a hundred years it remains the mainstay of the field, comprising about 70% of the study volume worldwide. Until the 1980s, the only way to acquire a radiograph was on film. Even when computed radiography was introduced in the 1980s, its images were initially printed to film for interpretation. Likewise, as other modalities such as ultrasound, computed tomography (CT), magnetic resonance imaging (MRI), digital fluoroscopy, and nuclear medicine were developed during the 20th century, their inherently digital images were typically printed to film for viewing. PACS did not become a reality until computers and networks became fast and robust enough to handle digital image transfers and storage.

But what about image display? Could the standard computer displays of the early 1990s provide adequate image quality? It was quickly realized that the answer was no and that a specialized display would need to be developed to support medical imaging. The primary deficiencies in the typical color cathode-ray tube (CRT) displays of that time were in the areas of luminance, spatial resolution, and contrast resolution, not to mention issues of geometric and other distortions. The quality of such displays was just not comparable to that of radiographic film. Thus, a whole industry was started to produce displays that were optimized to the requirements of radiologic image display of radiographs.

Monochrome CRTs were created with special phosphors to produce high luminance levels with less distortion and at higher spatial resolutions. These displays met the requirements for radiography. Because all the other modalities seemed to have less stringent requirements, there was a basic assumption that what was good enough for radiography was good enough for everything else. While this was true, it was seldom considered whether such displays were really necessary for nonradiography images. In fact, ultrasound images are typically acquired on equipment with color CRT displays. Furthermore, nuclear medicine workstations almost always had color CRTs because color was so necessary in the diagnostic viewing of these images.

In the mid-1990s, liquid crystal display (LCD) displays were introduced to the general consumer market. Medical display manufacturers quickly adapted this technology as they did for the CRT; they produced a monochrome version with higher luminance and higher spatial resolutions. While these displays clearly outperformed their CRT predecessors, they made the question of nonradiography image viewing more acute. The old, commercial-grade, color CRTs did not meet minimum luminance requirements, but the newer color, commercial-grade LCD displays did. Many radiologists began to wonder whether they really needed the more expensive monochrome displays to see what they needed to see, especially for nonradiography.

There is little reason left to buy CRTs for medical image viewing today. They are heavier, bulkier, hotter, and less durable than LCDs and provide much less sharpness, especially over time. Rather, more attention is being paid to deciding which type of LCD to choose—monochrome or color, what size, and so forth. This chapter explores the basic features of medical imaging displays and their clinical relevance.

## HOW MANY MONITORS?

For many years films were hung on film alternators in a 4-over-4 configuration so as to maximize the number of images that could be viewed at once. Changing the view meant having the machine mechanically rotate the panels, a process that could take up to 15 seconds per change. Thus, radiologists grew accustomed to looking at up to 8 films at once. The first PACS workstations offered up to 4 monitors, and many departments purchased these configurations because they were the closest to what radiologists had grown used to seeing, albeit extremely expensive. However, over time, most radiologists noticed that their viewing habits changed considerably with PACS. Once images could be scrolled rapidly in “cine” mode, they found that it was no longer necessary and in fact less desirable to view multiple panels of images in the tile mode they had used for many years.

Siegel and Reiner examined radiologist productivity versus the number of monitors. They found that while there was a significant gain in having 2 imaging monitors instead of just 1, having more than 2 did not yield further gain. Furthermore, radiologists usually do not compare more than 2 images at a time; it is usually just a frontal view to a lateral one, a current image to a prior one. It is usually desirable to view several pulse sequences at one time with MRI images, but in those cases, since the native resolution of these images is so small, the screen is often split into 4 or more sections, obviating the need for more than 2 monitors.

One exception to this rule may be mammography. Currently there is an influx of softcopy reading systems for digital mammography on the market. Since a screening mammogram has 4 views and comparison to the prior mammogram is always performed when available, there may be justification for an increased number of imaging displays in this case. Time will tell as they undergo more clinical usage.

The PACS workstation has evolved to the extent that the term is probably outmoded. People do not refer to their desktop computers as word-processing workstations just because they use them for word processing. So, too, PACS viewing software is beginning to take its place as just another application on the radiologic interpretation computer. While it may be the primary application for image interpretation, the same computer may be used for access to the radiology information system (RIS), the hospital information system (HIS), the electronic medical record (EMR), the Internet, e-mail, and even a speech recognition dictation system. As such, today's systems usually include one additional low-cost color LCD display for these other functions, and sometimes even two. Even the exam selection portion

of the PACS viewing software itself is often better viewed this way, as it permits the graphical user interface of this part of the system to employ color.

## PHYSICAL SIZE

Image displays that have been optimized for medical image viewing are typically 21 in. diagonally. Displays larger than 25 in. can not be comfortably viewed at the typical viewing distance of 18 in. without having to move one's head around, thereby causing neck strain. Consumer displays range in size from 17 up to 21 in. diagonally. The size of LCD displays usually corresponds to the resolution given in Table 17.1.

This measurement does not include the size of the bezel, which is the frame around the display. Bezels range in thickness from 0.5 to 2 in., are usually made from plastic, and may have several buttons to control the display and some indicator lights. Some displays include audio speakers in the bezel, which increases the size of the bezel. There has been a trend since

**TABLE 17.1**  
Resolution and Pixel Pitch of Various Displays

Display Size and Type	Width × Height (pixels)	Resolution		Pixel Pitch (microns)
		Megapixels		
		Actual	Commonly Called	
15-in color	1024 × 768	0.79	1	297
17-in color	1280 × 1024	1.3	1	264
18-in color	1280 × 1024	1.3	1	281
19-in color	1280 × 1024	1.3	1	294
20-in color	1600 × 1200	1.9	2	255
21-in color	1600 × 1200	1.9	2	270
21-in monochrome 3MP	1536 × 2048	3.1	3	207
21-in monochrome 5MP	2048 × 2560	5.2	5	165
22-in color 9MP	3840 × 2400	9.2	9	125

2003 to produce displays with thinner bezels, as this allows multiple displays to be more closely placed side by side. For medical imaging, the ideal bezel is thin and black, as this de-emphasizes the bezel and allows the displayed image to dominate.

Liquid crystal displays are much thinner than their CRT counterparts. Cathode-ray tube displays are usually somewhat cubical; their thickness roughly equals their height and width. Liquid crystal displays are usually about 2 to 3 in. thick, regardless of height and width. Accordingly, LCD displays are lighter and easier to move around.

## RESOLUTION AND PIXEL PITCH

### DEFINITIONS

The actual display area of a display is divided into a matrix of picture elements called “pixels” (pix for “picture,” els for “elements”). The term “resolution” refers to the size of that matrix—either its length and width, or the total number of pixels, which is obtained simply by multiplying the two dimensions. As demonstrated in Table 17.1, these range from 1 megapixel (MP) up to 9 MP. Displays with 2000 pixels on one side are sometimes called “2K monitors.” However, since this term could mean either a 3 MP or a 5 MP display, it is less specific.

It should be apparent from Table 17.1 that as the display size in inches increases, so does the display resolution. However, this is not a 1:1 relationship; 1.3 MP displays come in both 18- and 19-in. sizes, and 1.9 MP come in 20- and 21-in. sizes. Furthermore, the relationship is not linear; the number of pixels increases faster than the display area.

This relationship between physical size of the display area and the resolution is called the “pixel pitch.” It can be obtained by simply dividing the physical width in microns by the width in pixels. Since pixels are typically square, the same should be obtained using the physical height and the height in pixels.

### REQUIREMENTS

One question that is still being explored is, Just how much resolution does one need for diagnostic viewing of radiologic images? To address this issue, a distinction needs to be made between two categories of images: radiography and nonradiography. Radiography includes the standard chest x-ray,

abdominal radiographs, bone radiographs, and so on. The actual acquisition modality can be either computed radiography (CR) or digital radiography (DR); it does not matter. It also includes mammography, but that will be discussed separately later in this chapter. Nonradiography is everything else: CT, MRI, ultrasound, nuclear medicine, fluoroscopy, and angiography. There are at least two major differences between these categories that affect display requirements. One is the inherent contrast of the features normally found in the images; this will be discussed in the next section about luminance. The other, which is germane to this section, is the native size and resolution of the images.

The typical chest x-ray has been obtained on a 14 × 17-in. film cassette for decades, and this cassette size has been perpetuated in CR and DR machines. However, even displays that are specifically designed for medical imaging have a viewable image size of only 12.5 × 16.6 in. This means that chest x-rays must be reduced by at least 20% to fit on the screen, and this assumes that the orientation of the image and the display match (e.g., both are in portrait) and does not take into account the presence of buttons, labels, and a taskbar, which occupy screen space and further shrink the image. Such displays are usually placed in portrait orientation, which matches that of posterior anterior (PA) and lateral exams but does not match the orientation of the ubiquitous portable chest x-ray, which is typically obtained in landscape orientation. The point is that in order to be displayed in its entirety, a radiograph is often shrunk to fit into the physical viewable area of the display. This is because *the physical size of the display was not designed around the size of the acquired image but rather around the human visual system*. Larger displays are not only more costly to produce but also more difficult to view at a normal viewing distance of 18 in. Consider a picture of a house. No one would want the picture to be the same size as the house! Likewise, radiographs are shrunk somewhat to make them easier to view, and of course the radiologist can and does zoom and pan the image to examine different parts more closely as needed.

Similarly, the pixel pitch of a display should match the limits of the human visual system at a typical viewing distance of 18 in. rather than the pixel pitch of the original image. *Making the pixel pitch smaller than this threshold will not convey more information to the radiologist*. Cramming more pixels onto the screen may display more information, but radiologists will not see it because it is below their visual threshold. One might argue that radiologists will see more if they move their faces closer to the screen, which is what radiologists and indeed people in general do in order to see things better—they move their eyes closer to the object. However, the same effect can be accomplished on the computer display by simply zooming and panning the

image, and this involves far less neck strain, is more ergonomically sound, and is a lot less costly. Exactly what the perception threshold is depends on the visual acuity of the radiologist and is a matter of debate. At the very least, it is unclear how many radiologists who view a 5 MP display with 165 micron pixels at a viewing distance of 18 in. are actually seeing 5 MP.

One rationale for the development of the 5-MP display may have been the size of radiographs. According to American College of Radiology recommendations, radiographs (CR or DR) should be acquired at a resolution of at least 2.5 line pairs per millimeter (lp/mm). For a 14×17-in. cassette, this translates into about 4MP, and in practice the images are rendered at 5 MP. However, as mentioned, images are often shrunk anyway due to elements of the graphical user interface such as buttons and a taskbar as well as mismatches in image versus display orientation. Furthermore, it is more logical to match the resolution of the display to the limits of the human visual system than to the resolution of the acquired image. Film mammography provides a good insight into this principle. Mammographers do move their eyes very close to the film to see those microcalcifications, but they know that even that effort is not enough. They must also use a magnifying glass to be sure they are not missing any such findings. In other words, the resolution of the acquired image is higher than the limits of the human visual system, and important information may be missed if the mammographer does not use a magnifying glass. The magnifying glass brings down the higher resolution of the acquired image into the human viewable range. The advantage of a computer display system is that, unlike film, it has a magnification system built in. Thus, it is not necessary for the display system to have the same number of pixels as the acquired image, so long as all the pixels of the acquired image can be viewed by zooming and panning the displayed image. This last point is essentially what is stated in the American College of Radiology guidelines (ACR Technical Standard for Digital Image Data Management, 2001 Revision).

Much smaller than radiographs, nonradiography images (CT, MRI, ultrasound, etc.) have different resolution requirements. The largest of these are digital fluoroscopy images, which can be as large as 1 MP. But CT images are  $\frac{1}{4}$ MP and MRI images are usually even smaller. When a radiologist is viewing these images in stack mode, a 1-MP display has more than enough resolution. Thus, PACS workstations that are primarily dedicated to nonradiography diagnosis do not need the higher resolutions offered by displays designed around radiography. In fact, *higher resolution displays can actually slow down performance when viewing nonradiography images in stack mode.* Consider trying to scroll through a CT exam with 100 images, each  $\frac{1}{4}$ MP in size, on a 1-MP display. Each image needs to be digitally magnified roughly by a factor of 4 before it can be displayed in order to make it fill most of the



screen. Now consider the same scenario with a 3-MP display—each image needs to be magnified by a factor of 12 instead. At current processor speeds, this creates a very noticeable difference in the speed with which one can scroll through the images. In sum, for scrolling through CT images, a 3-MP display is overkill and can significantly slow down the process. For this reason, most radiologists at Massachusetts General Hospital (MGH) have learned to avoid using the higher-resolution displays for CT and MRI viewing.

As mentioned above, there is another difference between radiography and nonradiography with regard to luminance requirements.

## LUMINANCE

Luminance is the brightness of the display, expressed in foot-lamberts, candela per square meter, or nits ( $1 \text{ cd/m}^2 = 1 \text{ nit} = 0.29 \text{ foot-lambert}$ ). Grayscale radiologic images usually contain 256 shades of gray in any given presentation state. What the ACR recommendations endeavor to do is to make sure that the radiologist can see all 256 of those shades and perceive them in a linear fashion (e.g., a pixel with a value of 200 appears twice as bright as one with a value of 100). To that end, they issued 2 recommendations. The first is that the maximum luminance of the display should be at least 50 foot-lamberts. A display dimmer than may be able to produce 256 distinct shades of gray but not ones that the human eye can distinguish. The difference between each shade of gray and the next one must be at least 1 JND (just noticeable difference). For this, at least 170 nits (50 foot-lamberts) are required. Most color CRTs never reached this brightness, and as such were not suitable for medical imaging. This was one of the factors that spawned the development of the “medical-grade” display. Monochrome CRTs have a luminance of about 400 nits, and their LCD counterparts go up to 700 nits. However, today’s color LCDs are typically 250 nits in brightness; in August 2004, a 400-nit color LCD was introduced to the general consumer market.

Is brighter better? To some degree it probably is. The brighter the display, the greater the step change in luminance from one shade of gray to the next; brighter displays have 2 to 3 JNDs from one pixel value to the next. Subtle findings tend to be more conspicuous on brighter displays. However, there is a point of diminishing returns. Too much light from large white areas (such as the diaphragm on a chest x-ray) on a very bright display can blind the radiologist to subtle findings in the darker areas of the image.

The second recommendation of the ACR is in regard to the perceptual linearization of the shades of gray. This is discussed in the next section on calibration.

## CALIBRATION

Again, radiologic images typically contain 256 shades of gray in any given presentation state. In the case of MRI and CT, this means at any given window/level setting, but the same applies to any image with more than 8 bits of pixel depth. At first glance it might seem trivial to display an image with 256 shades of gray, as that is exactly the maximum amount possible on standard consumer-grade display systems. Each pixel is comprised of 3 sub-pixels—1 red, 1 green, and 1 blue. Each subpixel can be modulated on a scale of 0 to 255, which yields  $256^3$  (about 16 million) possible color combinations. In order to produce a shade of gray, all one need do is set all 3 sub-pixels to the same value, hence 256 possible shades of gray.

There are 2 problems with this approach. The first is that it assumes a uniform increase in luminance output per change in digital driving level (DDL), called a display function. That is, if a DDL of 50 yields 45 nits of light, then a DDL of 100 should yield 90 nits of light. It turns out that many displays have wildly varying display functions. What one might assume to be a fairly linear function is often skewed and sometimes erratic due to a host of factors in the manufacturing process. The second problem is that a linear function is not even what is desired. Why not? Because it turns out that although 90 nits are indeed twice as bright as 45 nits, that difference in brightness is not so perceived by the human eye. Human perception of brightness is not linear. A photometer will correctly measure the difference, but radiographs are still interpreted by humans at this point in time (this is shifting a bit; see Chapter 21, “Breast Imaging, Computer-Aided Detection, and Computer-Assisted Classification”). In a very rough sense, it takes a 1% change in luminance for the eye to detect the change. Reconsider the above example. Let us say that the DDL of 50 yields 45 nits of light. DDL 51 need only yield 0.45 nits more to produce 1 JND between the 2 DDLs. Let us also say that a DDL of 200 yields 130 nits. DDL 201 will have to yield 1.3 nits more in order to produce 1 JND between the 2 DDLs. This is a simplification of the Barten model of the human visual system that was adopted by the ACR but illustrates the point that the light output of a display for medical imaging is supposed to follow a specific curve that is spelled out in the DICOM Part 14 Grayscale Display Function (GSDF).

This curve is a mathematical function that depends on the minimum and maximum light output of the display (called the black level and white level) and whose precision is infinite. The goal is to produce 256 discrete shades of gray that approximate this curve with the smallest possible deviations. The best way to do this is to produce a large pool of potential shades of gray from which to pick the 256 best ones. The greater the number of shades of gray a display system can produce, the closer the 256 best ones will be to the curve. The number of potential shades of gray a display can produce depends on the limitations of the display itself and the video card that controls it. In the days of the medical-grade CRT, calibration precision was actually a simpler matter. As an analog device, the CRT could produce an almost limitless number of shades of gray; it all depended on the precision of the voltages fed to it by the video card. The finer the card could modulate its output voltages, the greater the number of intermediary shades of gray. Analog medical-grade video cards produced 1000 shades of gray or more, from which the best 256 were chosen during the calibration process to be mapped to the 256 DDLs.

Liquid crystal display technology presented a challenge to display calibration, as LCDs are inherently digital devices. Consider an analog radio dial. With patience it can be carefully and slowly adjusted to tune in a given radio station, say 100.1, for example. A digital radio tuner can be adjusted to reach the same station much more quickly. A favorite station can even be saved and reached with the press of a button. But what happens if the digital tuner loses its calibration over time and reads 100.1 when it is actually tuned to 100.2? The digital scale only moves in increments of 0.2, so you can only go down to 99.9, which is actually 100.0, but you cannot reach 100.1. In a digital device, there is no “in between,” whereas with the analog dial there is. You would not even have noticed if, over time, you needed to turn the analog dial slightly more or less to reach 100.1. Medical-grade LCDs are made from the same basic digital material as consumer-grade LCDs, and these panels only accept 256 DDLs per subpixel. One saving grace is that medical-grade displays have the color filters left off the subpixels during the manufacturing process (a change in the standard manufacturing process that is one of the factors that drives up the price for these displays). The color filters are left off because they attenuate the light coming through from the back of the panel, and medical displays are designed to be brighter than typical consumer displays. But a side benefit of this is that now all 3 subpixels produce white light instead of colors, and as such the total number of shades of gray that can be produced is  $256 \times 3 = 768$ . The medical-grade display manufacturers did not stop there, though. They employ temporal modulation as well to achieve a bigger pool of shades of gray to choose from.

This involves leaving some subpixels on at a higher level for only a fraction of a screen refresh cycle. Theoretically, those subpixels are flashing higher and lower rapidly, and one might think that this would produce a flicker effect. But since it takes place so quickly, well above the human eye's threshold for flicker, the eye perceives it as a slight static increase in brightness. The net result is once again a few thousand shades of gray from which to choose the 256 best ones.

What about consumer-grade LCDs? At present they have no such special circuitry for luminance calibration. It was long thought that the only way to calibrate them was to accept that the 256 shades of gray would have to be chosen from the 256 shades produced by red = green = blue—essentially no choice at all. In fact, what actually happened was that some of the shades of gray were thrown away to minimize the overall errors of deviation from the ideal curve, so only 230 or so were used, clearly not an optimal solution. Fortunately, in 2002 a technique was developed by Image Smiths, Inc., to calibrate almost any consumer-grade color LCD with any standard video card and achieve stunning precision, on par with that of the medical-grade display systems. Image Smiths basically challenged the assumption that one must make red = green = blue in order to produce a shade of gray in a pixel. Large deviations will of course produce hues, but what Image Smiths showed was that very small deviations from that formula can produce changes in perception of luminance *without changes in perception of color*. In other words, the rods in the human eye are more sensitive than the cones; therefore, the light output of a color display can be manipulated to produce thousands of shades of gray without the perception of color in them. Again, a photometer would sense the subtle hues, but the human eye will not. Thus, even standard consumer-grade displays and video cards can be calibrated to the DICOM Part 14 GSDF to a very fine precision.

Medical-grade displays come calibrated from the factory and come with circuitry to maintain that calibration as long as possible. Some will even indicate when the display needs to be serviced. The main rate-limiting factor in the durability of a display is its backlight. The basic assumption is that so long as the backlight's light output can be kept constant, then the factory calibration can be considered intact. The medical-grade displays have built-in sensors to monitor the light output and adjust it as necessary to keep it constant. When the backlight's output drops to a level at which maximum compensation (usually increasing the voltage) is inadequate, the backlight needs to be replaced by the manufacturer, at which point it will be recalibrated. Medical-grade LCD displays have not been around long enough to generate significant durability statistics, but they are thought to well outlast their CRT predecessors. They should last about 5 years instead of the typical

1.5 to 2 years of CRTs. Some manufacturers advertise extra-long-life backlights, but of course only time will tell if they really last.

## MOUNT

Most LCD displays come with a standard Video Electronics Standards Association (VESA) 100-mm mount in the back. This is desirable because it permits flexibility in placement of the display. If desired, such a display can easily be connected to a variety of wall, ceiling, and cart mounts and other types of mechanical arms. It is also desirable for the mount to permit pivoting the display between portrait and landscape modes.

## GRAPHICS CARDS

The graphics card is what connects the display to the computer. Other names for it include graphics board, video board, and video card. Several factors affect the suitability of video cards for viewing radiologic images. They include the slot type, video random access memory (VRAM), lookup tables, ability to pivot, maximum resolution, and number and types of outputs.

### SLOT TYPE

Also called the bus interface, the slot type specifies into which kind of slot the card is inserted in the computer. This matters because not every card can fit into every computer. For the past 8 years or so, there have been 2 main types of expansion slots in a computer—the advanced graphics port (AGP) and peripheral component interconnect (PCI). By design there is only 1 AGP slot that can tightly connect only 1 video card to the CPU. Additional video cards, and all other cards such as network interface cards, modems, and sound cards, get plugged into slower PCI slots, of which there can be up to 5 in a single computer. Thus, if a computer already has 1 AGP video card, additional cards must be PCI. Medical-grade video cards have largely remained designed for the PCI slot because there was little to be gained by accommodating the AGP slot and it would limit usability in multiple-monitor setups. Furthermore, some computers with a small form factor chassis can only take a half height AGP card, also called low profile. Such

computers also tend to have fewer PCI slots, some only 2, which gives less flexibility for multimonitor setups. The speed differential between PCI and AGP is imperceptible for the static- and stack-mode image viewing of most PACS workstations. However, it may be noticeable during three-dimensional (3-D) animations and large cine sequences of ultrasound or cardiac angiography. A new slot type was introduced to the market in the summer of 2004 called the PCI Express slot. It comes in 4 speeds: 1×, 4×, 8×, and 16×. The 16× PCI Express slot replaces the AGP slot on newer systems, and the 1× slots are beginning to replace some of the PCI slots. One advantage of this design is that a smaller card (e.g., 1×) can be used in a larger slot (e.g., 16×) if so desired, whereas the AGP and PCI standards are mutually exclusive.

## VIDEO RANDOM ACCESS MEMORY

Most video cards come with plenty of VRAM to support 24-bit color depth at their full resolution. Thirty-two megabytes of RAM is plenty to drive a 2-MP display, and many cards come with 64MB or more. Additional VRAM only becomes an issue with complex 3-D rendering.

## LOOKUP TABLES

The lookup table, also called a gamma correction table, is where the correction factors from the DICOM calibration are stored. Since this is really a function of the display, the best place for it is in the display itself. While this may change in the future, most displays can not accommodate the lookup table, so it is stored in the video card instead, and the video card applies the corrections before it sends its signal out to the display. The important thing to understand about lookup tables is that without them, DICOM calibration cannot be performed. Fortunately, all video cards today have a lookup table. However, one is needed for each display attached to the card, and some dual-headed cards have only one lookup table. Therefore, only dual-headed cards with 2 lookup tables—one for each head, should be used for radiologic image viewing. Furthermore, if a multimonitor workstation is calibrated, and then the cables are disconnected in order to move the workstation, care must be taken to make sure the cables get reconnected to their original video card heads. Otherwise, the calibrations will be applied to the wrong displays.

## PIVOTING

In order for a display to pivot between portrait and landscape modes, it must (1) have a base that pivots and (2) have a means of rotating the display contents 90 degrees. Many consumer-grade manufacturers include software that will perform the rotation. While this may work fine for a single low-resolution display, it is inadvisable for higher-resolution and multimonitor setups, as it places a considerable burden on the computer's central processing unit (CPU) and will likely slow down system performance. It is preferable that the video card itself perform the rotation. Some do and some do not; this function needs to be checked before a card for such a purpose is purchased.

## RESOLUTION AND OUTPUTS

For many years the standard video graphics adapter (VGA) analog video output remained unchanged. So long as CRTs were used, there was no need to change it. But the introduction of flat-panel LCD displays brought a new type of interface between the computer and the display called digital visual interface (DVI). Bearing many more pins, the DVI connection can carry 1 analog and up to 2 digital signals to a display. The intent is for the analog signal method to be phased out over time. Today, the overwhelming majority of video cards in use have analog outputs only, so most users of digital flat panels send them an analog signal. However, this involves taking the digital signal from the computer's CPU, converting it to an analog signal in the video card, and then converting it back to a digital signal in the flat-panel display. Clearly, this is illogical and is highly prone to artifacts from the two conversions. This is especially noticeable in locations with significant electrical interference, which is quite common in hospitals. Therefore, whenever possible, *a digital DVI signal should always be connected to a flat panel for radiologic viewing*. However, DVI outputs have some limitations that do not exist with analog outputs. Most analog outputs can drive up to 3 MPs at a full 60 Hz refresh rate and 24-bit color. In contrast, many DVI outputs are limited to only 1.3 MPs (1280 × 1024) at 24-bit color. This is changing, and there are several DVI cards available today that can support 2 MP (1600 × 1200) at 24-bit color.

There are very few color displays that run above 2 MP today; those that do utilize multiple DVI connections to support their needs. It should be pointed out that medical-grade displays typically come with their own custom-designed video cards, which will provide the best performance from

these displays. Dual-headed display cards are often the best way to support multiple monitors, rather than using multiple single-headed cards. First, they tend to be cheaper; second, they conserve PCI slots. Some quad-headed cards exist, too.

## GRAYSCALE VERSUS COLOR AND MEDICAL GRADE VERSUS CONSUMER GRADE

Why have medical-grade displays traditionally been grayscale? The main reason is luminance. What was true in the CRT world is true in the LCD world, albeit for different reasons: a monochrome display is brighter than its color counterpart. The way LCD displays produce colors is by having each pixel consist of 3 subpixels, each covered with a color filter: red, green, or blue. The way these filters produce their respective colors is by filtering out the other two colors. The blue filter, for example, filters out red and green. The net impact is a considerable attenuation of the backlight in order to produce these colors. Grayscale displays just have the filters left off during the manufacturing process, which lets more light shine through. The result is a display with a maximum luminance of about 700 nits instead of the typical 250 nits of consumer displays. While this may be desirable for radiography, it prevents viewing of color Doppler ultrasound signals, nuclear medicine color scale images, and positron emission tomography (PET)/CT fusion images. Display of color is also becoming increasingly important as functional MRI (fMRI) and computer-aided diagnosis (CAD) move further into the mainstream. Therefore, most PACS installations include a variety of monitor types, depending on the study mix at the institution and which exams tend to get read where. Further flexibility is obtained by using color displays alongside grayscale displays, giving the radiologist a choice depending on the type of exam he or she wants to view. As of this writing, a consumer-grade 1.3-MP display was recently introduced to the market with a luminance of 400 nits. It remains to be seen whether this will mark the beginning of a trend toward brighter consumer-grade displays.

In sum, medical-grade higher-resolution grayscale displays have been optimized for radiography and are probably best suited for routine primary diagnosis of radiography. But for other modalities and for non-routine viewing, consumer-grade, color flat-panel displays are becoming a viable and sometimes better alternative. As these displays evolve, the distinction between the medical-grade and consumer-grade display will continue to blur.



## OPERATING ROOM DISPLAYS

Displaying images in the operating room (OR) presents a somewhat different set of challenges. Primary diagnosis is seldom performed here except for checking intraoperative images. Furthermore, the ambient light is often significant and cannot be eliminated. Some surgeons prefer regular color flat panels. These can be mounted on the wall, on a boom suspended from the ceiling, or on a mobile cart. Some may insist on medical-grade grayscale displays. Another possibility to consider is the large 30- and 40-in. LCD displays. These displays are large in physical size, but only 1 MP in resolution ( $1280 \times 768$ ). Accordingly, their pixel pitch is on the order of  $600 \mu\text{m}$ —huge. These displays look good only at a minimum of 6 ft away; they are meant only for distance viewing, like television screens. But they are brighter, too, at about 350 nits, to accommodate this distance. These displays have been installed at both MGH and Staten Island University Hospital and have met with measured success.

## DISPLAY ERGONOMICS

Ergonomics is an important topic about which most radiologists are unfortunately uninformed. There are a few basic rules to follow to avoid injury and strain from long hours reading at a diagnostic workstation. First, the chair should be adjusted so that the radiologist's feet are resting flat on the floor (not dangling), with the knees bent at 90 degrees. Next, the top of the displays should be level with the eyes, so that the displays lie in the natural viewing angle for the neck. The keyboard and mouse should be level with the arms, which should be resting comfortably on armrests. Therefore, the chair should have an adjustable height and armrests. Also, the height of the monitors and the height of the keyboard and mouse should be independently adjustable. Such tables are available from at least 2 companies, and are targeted for the radiology profession.

# DIGITAL MAMMOGRAPHY

---

---

MARTIN J. YAFFE

**T**he development of digital mammography was motivated by the potential to improve sensitivity and specificity in detection of breast cancer. In screen-film mammography, the processes of image acquisition, storage, and display are all tightly linked, which necessitates important compromises. For example, as phosphor screens are made thicker to absorb x-rays more efficiently, the increased diffusion of light reduces spatial resolution. Film adds random fluctuation to the image due to its granularity, reducing the image signal-to-noise ratio (SNR). Furthermore, the response of the screen-film combination is highly nonlinear with x-ray exposure, causing both image contrast and SNR to fall off sharply outside a relatively narrow intermediate range of exposures. Consequently, image quality deteriorates in areas of the image depicting the more radiolucent and radio-opaque parts of the breast. The storage and retrieval of film mammograms is inefficient; films are occasionally lost and frequently misplaced. Because there is only a single instance of the image, consultation can only occur if the radiologist, the film, and the other discussant are all in the same room. Finally, it is difficult to extract quantitative information from a film image and the quality of that information is limited because of the SNR problems already discussed.

In digital mammography, the processes of image acquisition, storage/retrieval, and display are separated, allowing optimization of each. Acquisition is performed using an electronic x-ray detector whose output is digitized. The image is stored on disk and displayed either on a high-resolution softcopy device or on a film by laser printing. The initially continuous x-ray image is spatially sampled, and the signal level is also discretely quantized. Spatial sampling is performed either by using an x-ray detector that is pixelated, that is, divided into a matrix of discrete sensing elements, or by mechanical scanning of either the entire detector or the detector readout system. Sampling of signal level occurs in the analog-to-digital conversion process. The number of bits digitization must be adequate to represent subtle difference in x-ray attenuation by tissue over a wide dynamic range of x-ray exposure.

To be useful, the digital detector must image as much of the breast tissue as possible, and the spatial resolution must be high enough to allow adequate visualization of the fine detail structures in the breast, for example, morphology and clustering characteristics of microcalcifications and spiculations radiating from tumors. To attain high spatial resolution, the image element or pixel must be smaller ( $100\mu\text{m}$  or less) than that used for general radiographic applications. The number of image elements required is just the ratio between the overall image area and the area of the individual image element.

## IMAGE ACQUISITION

In screen-film mammography, the overall limiting spatial resolution is typically at least 16 line pairs per millimeter (lp/mm). Sampling theory implies that to provide such resolution, the required dimension of the image element must be  $0.031\text{ mm}$  ( $31\mu\text{m}$ ) or less. In many cases, it is not practical with the technologies used to produce electronic detectors to fabricate devices with elements that small. Currently, digital mammography detectors with square detector elements (dels) of dimensions 50, 70, 85, and  $100\mu\text{m}$  are commercially available. One manufacturer of digital mammography equipment provides the capability of producing a partial image with dels of  $25\mu\text{m}$ .

The requirements for quantization of signal level are dictated by the maximum attenuation factor of the breast and the required precision in the part of the image corresponding to the most attenuating region (thickest, densest part of the breast). If the exposure factors are properly chosen, it is reasonable that the attenuation factor can be restricted to at most 70 or so. For 1% precision, this implies that about  $70/0.01 = 7000$  discrete signal levels

are required. This scale can be represented by 13 to 14 bits of digitization. It is also possible to digitize the transmitted x-ray signal logarithmically. This is done in the case of systems using photostimulable phosphor plates as detectors. This reduces the requirement to accommodate the dynamic range to 10 bits. In either case, however, 2 bytes of storage are required for each image element. This results in unused bits, “an overhead” of between 13% and 38%. The currently available digital mammography systems are summarized in Table 18.1 in terms of the number of bits used, the del size, the matrix dimensions, and the consequent size of the images.

In mammography, typically 4 views are acquired per bilateral examination. The daily volume of data,  $V$ , in megabytes (MB), that must be stored from all the digital mammography units in a department is:

$$V = S * 4 * n_d * n_r$$

where  $S$  is the size (in MB) of each digital mammogram,  $n_d$  is the number of bilateral patient exams per day, and  $n_r$  is the number of mammography rooms. Table 18.2 gives the total annual data volumes produced from digital mammography under a variety of operational scenarios.

Some mammography units provide output images at 2 levels. The raw image is essentially the image from the detector on which only some

**TABLE 18.1**

Detector Element Size, Matrix Dimensions, Bit Depth, and Image Size(s) in MB for Current Digital Mammography Systems

<b>System</b>	<b>Nominal Detector Element Size (<math>\mu\text{m}</math>)</b>	<b>Matrix Dimensions</b>	<b>Bit Depth</b>	<b>MB per Image</b>
Fischer Senoscan	54	4095 $\times$ 5625	12	43.9
Fuji	50	4728 $\times$ 5928 (3540 $\times$ 4740)	10 (log)	53.5 (32)
GE Senographe 2000DS	100	1914 $\times$ 2294	14	8.4
Lorad Selenia	70	3328 $\times$ 4096 (2560 $\times$ 3328)	14	26 (16.3)
Sectra	50	4608 $\times$ 5200	14	45.7
XCounter	60	4000 $\times$ 5000	17	57.2

**TABLE 18.2**  
Annual Data Production in TB Based on 250 Operational Days per Year

System	1 Unit			3 Units		
	Exams/Unit/Day			Exams/Unit/Day		
	15	25	40	15	25	40
Fischer Senoscan	0.63	1.1	1.7	1.9	3.1	5
Fuji (S)	0.46	0.76	1.2	1.4	2.3	3.6
Fuji (L)	0.76	1.3	2	2.3	3.9	6.1
GE Senographe DS	0.13	0.2	0.33	0.36	0.6	0.96
Lorad Selenia (S)	0.23	0.39	0.62	0.7	1.2	1.9
Lorad Selenia (L)	0.37	0.62	0.99	1.1	1.9	3
Sectra	0.65	1.09	1.7	2	3.3	5.2
XCounter	0.82	1.36	2.18	2.45	4.1	6.5

elementary operations, such as correction for spatial nonuniformity of the detector response, have been applied. In most cases, the signal level is linear, with transmitted x-ray exposure through the breast. A processed image is also produced. This is the one that is generally viewed by the radiologist and has often been transformed to incorporate a flattening or peripheral equalization to reduce the dynamic range requirements of the display device and by application of contrast or sharpness enhancement algorithms. There is a tendency to archive the processed image as it is the one on which the radiology report was made. If, on the other hand, it is desirable to carry out subsequent quantitative analysis of the images or to compare images with ones that will be produced in the future, when processing algorithms may be different, it may be useful to archive the raw image.

The Digital Imaging and Communications in Medicine (DICOM) Working Group DICOM 3.3-2003 has set standards for the contents of the headers and the information stored in digital mammograms and structured reports from mammographic computer-assisted detection (CAD) systems. Images can be stored as 2 types: For Processing, which is essentially an unprocessed image except for spatial nonuniformity corrections, and For Presentation, which is an image along with enough information for it to be displayed and automatically viewed in a predetermined manner.

The DICOM Grayscale Softcopy Presentation State is a separate DICOM object containing a group of tags identifying the image and various display parameters. It gives the ability to convey an intended presentation state or record a desired presentation state. It includes capabilities for specifying:

- ▶ Output grayscale space in presentation values (P-Values)
- ▶ Grayscale contrast transformation
- ▶ Mask subtraction for multiframe images
- ▶ Selection of the area of the image to display and whether to rotate or flip it
- ▶ Image and display relative annotations, including graphics, text, and overlays

Softcopy Presentation State refers to the grayscale image transformations that are to be applied in an explicitly defined manner to convert the stored image pixel data values in a Composite Image Storage Instance to P-Values when an image is displayed on a softcopy device. The P-Values are in a device-independent, perceptually linear space that is formally defined in DICOM PS 3.14 Grayscale Standard Display Function. At the present time, not all vendors of digital mammography systems support the DICOM Presentation State or comply with the Standard DICOM Grayscale Display Function.

## IMAGE RETRIEVAL

As digital mammograms are very large and, in many cases, may have to be downloaded from a central PACS archive to a specialized high-resolution workstation for review, it is essential that communication bandwidth be sufficient to avoid bottlenecks. Note that for image element sizes of  $50\mu\text{m}$  and  $100\mu\text{m}$ , transmission of a 4-view examination at standard modem rates (56Kb/s) would require about 7 hours and 1.8 hours, respectively. With full 100Mb/s bandwidth, these times would be reduced to 14.4 and 3.8 seconds, although, in practice, bandwidth generally must be shared with other applications. As mammography examinations are normally scheduled in advance of the visit, most PACS systems use the worklist to pre-fetch, or download images from previous examinations during low-demand periods, such as at night, so that they are readily available when required for comparison. Images are often sent to the PACS system, as well as the diagnostic work-

station used for that modality, so that they will be available for immediate viewing.

Data compression can significantly reduce the size of images. JPEG Process 14 lossless compression, which is endorsed in the DICOM Standard can provide a reduction in size by a factor of about 7 times for mammograms in which there is a saturated background. This is accomplished by removing the blank bits in each pixel value and using a technique such as run length encoding to represent the area outside the breast more efficiently. A PACS system should be able to store and expand images provided in compression formats meeting DICOM specifications. Primarily for medico-legal reasons lossy compression is not currently considered acceptable for mammographic imaging.

Despite the claims of compliance with the DICOM standard made by the manufacturers of various digital mammography units and PACS systems, it is still quite common to encounter compatibility problems when storing DICOM images to and retrieving them from PACS systems. This is particularly true when using technology from different vendors. The person(s) responsible for integrating the systems must carefully read the DICOM conformance statements of the different components. They should also be prepared to spend considerable time and effort working with the vendors to ensure that images can be satisfactorily stored and retrieved.

## IMAGE DISPLAY

As has been shown, digital mammograms are very large. Displaying them effectively is challenging. The range of signal levels exceeds that of laser film and also that which can be displayed simultaneously on either cathode-ray tube (CRT) or liquid crystal display (LCD) flat-panel monitors. For hard-copy printing, it is therefore necessary to apply a lookup table to reduce the dynamic range, while rendering the most relevant features in the mammogram with good contrast. Generally some type of peripheral equalization is automatically applied to the image data. Linear or nonlinear transformations can then be performed. This can either be done automatically, generally based on the image histogram of signal levels, or manually by the radiologist or a technologist prior to printing.

For softcopy reading, similar transformations are applied, but they are usually done interactively as the image is viewed. This is effective but inefficient; it can take up a significant amount of a radiologist's time as there is virtually an infinite range of possible settings of the display lookup table. On

most systems, an attempt is made to arrive at initial “semi-optimal” display settings automatically and quickly. The radiologist may then choose to make small improvements in the settings as deemed necessary.

A softcopy workstation should allow rapid adjustment of the position of individual images so that corresponding views can be matched to evaluate asymmetries. Additionally, it should be possible to adjust the contrast of either individual images or all images together as required.

Though laser printers have the spatial resolution capability to represent the full number of acquired image elements in a digital mammogram (Table 18.1), except for the case of the smallest (GE) image, even the highest-resolution softcopy displays do not. For this reason, some facilities choose to print images and view them on a standard mammographic transilluminator. This also allows multiple views of the current and a previous examination to be simultaneously displayed in a convenient manner. Even in an all-digital department, there needs to be a high-quality mammographic view box in order to review previous films from patients and outside films.

On the other hand, once an image is printed, its display characteristics are fixed and if it is desired to explore the signal level space to investigate subtle structural findings, it would be necessary to readjust the image on soft copy and then reprint it. This is time consuming and detracts from the potential value of having a digital image.

If digital mammograms are to be displayed in soft copy, it is necessary to use specialized display strategies to represent all the acquired information. Current high-resolution displays have formats of  $2048 \times 2500$  (5 megapixel [MP] display) or  $3840 \times 2400$  (9.2 MP display). For  $100\mu\text{m}$  image elements, it is possible to represent each element by a pixel and display 1 image at full acquired resolution on a monitor. For smaller elements or for multi-image display per monitor, this is not possible. To visualize the entire breast, images can be down-sampled by averaging adjacent elements, linearly interpolating values or simply selecting appropriately spaced pixels to reduce the size of the image. This, of course, results in reduced spatial resolution. This can be done to the point where 2 views can be displayed on the monitor back to back for correlation or comparison. These might be corresponding views of each breast, craniocaudal versus mediolateral oblique view, or corresponding current and previous views. This overview allows the radiologist to look globally for signs of disease.

If suspicious areas are identified, then the image must be zoomed (a reversal of the averaging process) to recover the full acquired resolution in viewing a region of interest in the breast. Of course, it is critical that signs of cancer must not be missed in viewing the initial, low-resolution version of the image. In our own research, we are exploring one approach to avoid-



ing such a problem by applying a CAD algorithm on the underlying acquired image data to indicate to the viewer the possible presence of suspicious abnormalities.

Recent studies on the efficiency of digital mammography have shown that the examination throughput is improved with respect to the image acquisition by the technologist. Because of the manipulations described above, the time spent by the radiologist on image interpretation is longer with digital compared to screen-film mammography. Some possibilities for automated image adjustment have already been suggested. Additional strategies to streamline image viewing and navigation would be very desirable in terms of making digital mammography more user friendly.

## WORKSTATION DESIGN

The demands of softcopy display of digital mammography preclude the use of monitors with less than 5 MP capacity for image interpretation. It is not suitable to employ standard PC monitors for this purpose. In order to have reasonable efficiency of image reading, at least a 2-monitor workstation is necessary. The system should allow rapid image “hanging” in different configurations and standard sequences that can be customized to suit the preferred reading protocol of individual radiologists.

An effective means of comparing current digital and previous film mammograms is very unlikely. Disparities in physical size of images and the need to locate a bright illuminator near the workstation where stray light can shine on the surface of the digital monitor and reduce image quality are real problems. One solution is to have the ability to digitize the old mammograms easily so that they could be presented alongside the current images.

Ideally, it should be possible to be able to display 8 digital mammographic images at full resolution simultaneously—2 views of each breast for the current and previous examinations. Because of the cost, limitations of monitors and viewing angle considerations, this is not practical, and it is necessary to move electronic images onto the display for making comparisons.

Throughput is a critical factor in mammography, and any design features contributing to improved efficiency of image interpretation would be valuable. Such features might include integrated presentation of CAD marks that could be toggled on and off and superimposed directly on the mammogram as well as voice recognition software that generates a textual radiologist’s report in a structured format.

Another opportunity to facilitate image interpretation is to bring sources of complementary information together on the same screen. For

example, the high spatial resolution anatomical information from digital mammography can be fused with breast ultrasound and functional image information on angiogenesis, water content, or metabolism from magnetic resonance imaging, positron-emission tomography, and the like. Cursors that relate spatial locations between images would also speed the interaction.

## INTERACTION BETWEEN MULTIPLE SYSTEMS

One of the impediments of digital mammography in its current implementation is the tight linkage of workstations from 1 vendor to the image acquisition system for that vendor. To some extent this is historical and simply reflects the necessity of communicating certain key pieces of information between the systems to ensure the most effective display of images. If systems conform in spirit with the DICOM standard, it should be possible for images acquired on any system to be displayed on any workstation meeting certain specifications. In fact, future FDA approval of digital mammography systems is likely to separate the acquisition and display systems in terms of how they are handled. A possible requirement for primary display stations may be conformance with the DICOM Grayscale Display Standard, support for the DICOM Presentation States, and American Association of Physicists in Medicine (AAPM) Task Group 18 quality standards. It is probable that facilities will purchase digital mammography systems and other complementary breast imaging systems from a range of manufacturers, so it is important that industry leaders work to facilitate interoperability.

## REFERENCES

- Digital Imaging and Communications in Medicine (DICOM) PS 3.4-2003 National Electrical Manufacturers Association 1300 N. 17th Street Rosslyn, Virginia 22209 USA DICOM PS 3.14 Grayscale Standard Display Function.
- Digital Imaging and Communications in Medicine (DICOM) PS 3.5-2004 National Electrical Manufacturers Association 1300 N. 17th Street Rosslyn, Virginia 22209 USA Part 5 Data Structures and Encoding, Section 8.2.1, p48 and Annex A.4-3, p68.
- Yaffe MJ. Detectors for digital mammography In: Pisano ED, Yaffe MJ, and Kuzmiac CM, eds. Digital Mammography. Lippincott Williams and Wilkins, Philadelphia 2004.

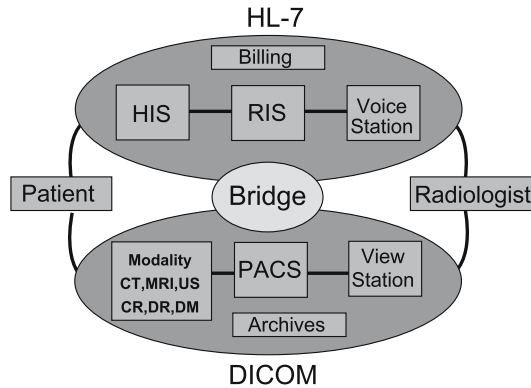
# WEB DISTRIBUTION

---

---

KEITH J. DREYER

**A**t the heart of the digital radiology department are two main computational systems, the radiology information system (RIS) and the picture archiving and communication system (PACS) (Figure 19.1). Although these technologies continue to converge, their functions within an enterprise remain unchanged. The RIS encompasses many text-based computing functions including transcription, reporting, ordering, scheduling, tracking, and billing. The PACS deals with image-based computing functions such as acquisition, interpretation, storage, and local image distribution. The proper use of these automated systems dramatically reduces the use of film and paper within a radiology department. However, removing film and paper removes the conventional method for distributing radiology information throughout the hospital. Because the goal of any radiology department is to deliver timely and accurate interpretations to requesting clinicians, the digital department needs a digital method to deliver its results. Enter the Internet.



**FIGURE 19.1**

The digital radiology department.

## THE INTERNET

An introduction to the emerging world of enterprise image distribution requires a basic understanding of the terminology and jargon associated with the Internet and the World Wide Web. For some, this is new and fundamental information; for others, a review.

The Internet (Figure 19.2) is a collection of computers (i.e., a network) communicating over a variety of transmission lines throughout the world, using a single common protocol known as transfer communication proto-

- **Internet**
  - Inexpensive
  - Ubiquitous
  - You Don't Own It
  - QOS (Quality of Service)
    - ADSL (Asynchronous Digital Subscriber Line)
    - RSVP (Resource Reservation Protocol)
  - Security Issues
    - Authentication
    - Encryption

**FIGURE 19.2**

Advantages of the Internet.

col/Internet protocol, or TCP/IP. This low-level protocol essentially allows all the computers in the world to communicate. It allows computer users and programs to communicate with each other using higher-level protocols such as SMTP (simple mail transport protocol, i.e., e-mail) (Figure 19.3). Hypertext transport protocol (HTTP) allows the distribution of text and images (and a variety of other media types) and is commonly referred to as the World Wide Web (or the Web for short). It is easy to see why the idea of using a universal way to transmit and receive text and images—the Web—to distribute radiologic information is becoming so generally accepted. Essentially, it is the process of turning your digital radiology department into a Web site. And, as with all successful Web businesses, the Web does not change your business (radiology) or your customers (clinicians), it just changes the way you deliver your business to your customers. If done correctly, Web distribution can offer great efficiencies to you and your clinicians.

Of note, on top of the HTTP protocol sit several languages that define the details of the Web pages. Hypertext markup language (HTML) is the Web's basic language and has been in place since 1994. Extensible markup language (XML) is a newer language gaining steam in the Web world due to its ability to separate data from display parameters, thus preserving the structure of all data fields. What does this mean? It means, for example, that even if a patient's medical record number is deeply buried in the middle of a Web page written in XML, a computer program reading that page (yes, they can read Web pages, too) could find it easily and use it to access more information on the patient from a different Web site. Extensible markup language is becoming the de facto standard for electronic data interconnect

- **E-Mail**
  - SMTP—Simple Mail Transport Protocol
- **File Transfer**
  - FTP—File Transport Protocol
- **The World Wide Web**
  - HTTP—Hypertext Transport Prot Protocol
    - HTML—Hypertext Markup Language
    - XML—Extensible Markup Language

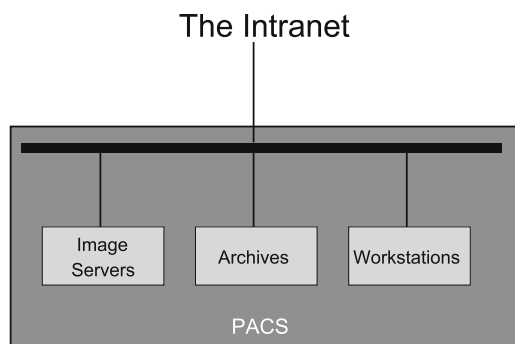
**FIGURE 19.3**

Common protocols available over the Internet.

(EDI), a way for computer programs (running on different computers) to reference each other through Web services, such as a radiology Web image server connecting to a hospital EMR (electronic medical record).

## HOSPITAL-WIDE IMAGE DISTRIBUTION

With the installation of PACS at many institutions around the United States, the task of distributing images to referring clinicians becomes a challenge. In the legacy system, physicians who wish to view their images rely on obtaining their films from the radiology department's film library. Often, these films are used to communicate with patients, family members, and consultants for patient education and patient management. The advent of PACS eliminates this classic film-based workflow. The use of preexisting intranet to distribute images is a well-accepted practice within hospitals due to its ubiquity, portability, and reasonable cost (Figure 19.4). (The intranet is the part of the Internet that is behind your hospital's firewall. See Chapter 14, "Networking Fundamentals," for more detail.) The Internet, and associated hospital intranets, has increasingly become the technological basis for both image management within the radiology department and image distribution to the enterprise (Figure 19.5). For hospitals with legacy PACS installations that have not been designed around Web technology, vendors are able to add Web-based solutions with relative ease. Central to the appeal of Web-based distribution is the ability for any physician, anywhere, whether at home



**FIGURE 19.4**

Typical components of a PACS intranet.


- **Intranet**
  - You Own It
  - Variety of Pipes with Guaranteed Speed
    - Modems, T1, ATM
    - 20 Kb/s, 2 Mb/s, 2 Gb/s
  - Good for Short Runs
  - EXPENSIVE
  - Firewall Security
    - Hard Shell, Soft Center

**FIGURE 19.5**

Advantages and disadvantages of an intranet.

or within the clinical setting, to use a personal computer as a virtual light box and view radiologic images and reports.

Early in the process, visionaries in the PACS arena predicted that the problem of image management and distribution should be approached from an enterprise point of view—creating a digital imaging architecture in radiology that would be integrated into the electronic medical record. The guiding principle in system design should be that images go wherever alphanumeric medical information goes, a task best achieved by a flexible, integrated Web-based solution. The various protocols available allow this flow of information with guaranteed connectivity (Figure 19.6).

**FIGURE 19.6**

Common hospital intranet protocols.

## COSTS

Many PACS vendors still charge high prices for their primary interpretation workstations and thus do not offer a viable alternative for referring physicians who need to review their ordered examinations. The current ubiquity of the desktop PC is a resource that obviates the need for software-only primary interpretation clients deployable via the Web. For example, Partners HealthCare System Inc., Boston, the parent corporation of Massachusetts General Hospital, currently has more than 32,000 PCs on its intranet. Utilizing this resource for image distribution adds essentially no additional cost, provided there is a software-only site-licensed client available over the intranet for primary interpretation.

There remains a confusion of pricing of Web-based PACS solutions within the industry. The alternatives include cost per user, cost per view (or click), cost per maximum users, cost per concurrent users, cost per studies stored, and a one-time cost in annual licensing. This range of pricing structures often creates confusion when comparing products using a price performance analysis. In their selection process, hospitals should always require vendors to converge on the same pricing model. It is the opinion of the author that the ideal pricing model is one that results in a one-time capital expense for a perpetual software-only license with a 10% to 15% recurring annual software maintenance fee that secures all future upgrades. All hardware for the Web-based PACS, archive, and enterprise distribution system should be purchased directly from the hardware vendor to achieve maximum savings and direct customer support.

## INFRASTRUCTURE

Inherent in standard Web protocol is a client-server relationship. Whereas the Web server is a computer that contains the images to be distributed for viewing, the client, in this example, is basically the clinician's PC. The following paragraphs describe the requirements for both the server and client systems of a typical radiology distribution system.

### SERVER

A Web server used for radiology image distribution needs to perform several functions, as follows:



- ▶ Accept images from a variety of DICOM sources.
- ▶ Compress images in either a lossy or lossless format (or both).
- ▶ Store compressed images in a local, fast access device (e.g., redundant array of inexpensive disks [RAID]).
- ▶ Archive all image data to a removable media archive (e.g., linear tape open [LTO])
- ▶ Act as an HTTP server waiting for Web requests over TCP/IP.

Typically, Web servers for radiology are Wintel-based PC servers, available from a variety of vendors, with a range of performance specifications. Sizing the system for this function requires several factors to be evaluated to provide optimum performance throughout the enterprise. These factors include cost, speed, number of concurrent users, reliability, and support. If the hospital distribution workflow calls for more than one Web server, it should be determined whether each server should have its own archive or if one archive should be shared among all servers. In most circumstances, a single archive should be shared, with each Web server accessing a common, compressed RAID. As demand increases from clinicians for Web access, more powerful servers may be needed to maintain performance. One way to accommodate this growing need is by increasing the server's processing power and RAM. Another way is simply to increase the number of servers available for the task. Beware, some vendors' software architectures will allow for these levels of upgrades, but others may run only as a single server, resulting in a much more expensive and limited upgrade path.

## CLIENT

The client device is quite variable (i.e., it can be any form of computer with access to the Internet—PDA, Tablet PC, or wireless laptop). In fact, possibly the only commonality necessary among clients is that they contain an HTTP browser to access the server. Further, the browser can be from a variety of manufacturers (Netscape Navigator and Microsoft Explorer are the most common). Because of the large file sizes and video requirements for display of medical images and because the most common client in the hospital setting is an IBM-compatible PC, minimum configurations for PC-based client machines accessing radiology images should be defined by the hospital and supported by the software provider. Again, there is a wide variety among Web PACS vendors regarding minimum configurations. Hospitals need to be certain that any potential Web PACS software provides

functionality to as many existing PCs as possible within their current infrastructure.

This is one of the most challenging aspects of ubiquitous Web distribution. Because any computer on the intranet can access the radiology Web server, it is often difficult to enforce a minimum client requirement. Therefore, performance is unpredictable. For a new institution, it is simple to recommend the purchase of a certain level of PC. But for most of us, there is an existing fleet of PCs throughout our institutions that would probably require upgrades to meet the core requirements. For a large institutional deployment of clients, it is best to evaluate users' needs individually to assess the intent of specific clinical access. Typically, some client computers are dedicated to specific clinicians while other systems are common and available to several users (ideally with individual logons and passwords).

## EXAMPLE: THE WEB EXPERIENCE AT MASSACHUSETTS GENERAL HOSPITAL

A Web-based primary interpretation client (Figure 19.7) and server solution were installed for image distribution (Figure 19.8) as a film alternative that would allow the hospital to cost-effectively distribute radiology images to all

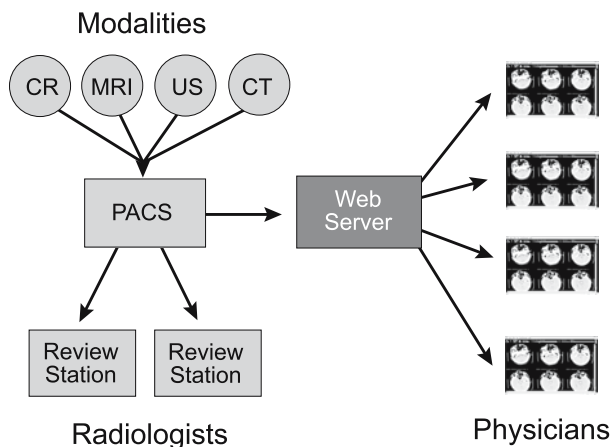


FIGURE 19.7


Web-based enterprise solution. CR indicates computed radiography; MRI, magnetic resonance imaging; US, ultrasound; CT, computed tomography.

- **PCT**—Private Communication Technology
- **SET**—Secure Electronic Transaction
- **S-HTTP**—Secure Hypertext Transfer Protocol
- **SSL**—Secure Sockets Layer
- **VPN**—Virtual Private Networks

**FIGURE 19.8**

Internet security options.

clinicians. The growing popularity of the Web guaranteed a degree of familiarity with the client software (a browser), which was easy to use and install with low support costs from the hospital information services department. It was felt that a Web server solution would layer onto any basic security system, such as firewalls, token-delivered coded access numbers, passwords and usernames, secure sockets layers (SSLs), and virtual private networks (VPNs) (Figure 19.9); and that these technologies, as they evolved, would enable the radiology department to ensure security and patient confidentiality regardless of the referring physician's access method and location. Further, integration with the institution's EMR has provided easy access to the enterprise for Health Insurance Portability and Accountability Act (HIPAA) tracking, central help-desk support, and integration with other departmental result-reporting applications. Incorporated into the Web-based solution is a JPEG2000 wavelet image compressor/decompressor that rapidly distributes compressed images throughout the enterprise with secondary access to lossless image data upon request. (No one, to date, has requested any.) The system is in its tenth year of deployment and continues to provide real-time image access to more than 5000 users with tremendous reliability.



**FIGURE 19.9**

Sample client Web image viewer for computed tomography with selected features from text depicted.

## CONCLUSION

Converting from film-based radiology to filmless radiology is challenging, to say the least. Many of these challenges exist outside the department of radiology. Replacing film means providing clinicians with another distribution mechanism. While there is great momentum within hospitals to keep the status quo, Web distribution offers many advantages over film. The successful deployment of Web servers and client viewers throughout the hospital enterprise will prove critical for the conversion to a truly digital radiology environment.

**IV**

**CLINICAL ISSUES**

DAVID S. HIRSCHORN

# PACS WORKSTATION SOFTWARE

---

---

STEVEN C. HORII

## BACKGROUND AND DEFINITIONS

Interpretation of radiographic images, for nearly the entire history of the specialty, has been based on the viewing of photographic images on a transparent or reflective substrate. Early images were photographic emulsions on glass plates, and recent ones are on transparent plastic polymer sheets. Fundamentally, however, the process is the same whether the images are on glass or plastic; x-radiation exposes areas of the film in inverse proportion to its absorption by the body interposed between the x-ray source and the film. This exposure results in a latent image in the silver halide that is converted to varying amounts of metallic silver through the development process. The advent of intensifying screens meant that much of the film exposure was due to the light emitted as the screens fluoresced in response to x-rays. The result was a tremendous effective amplification of the number of photons used to

expose the film compared to a nonscreen exposure. A consequence, though, was that the spatial resolution of the radiograph was considerably reduced. This change was accepted primarily because of the concomitant decrease in exposure necessary to make the radiographs.

In recent years, images of anatomy that are based on processes other than transmission radiography have also been incorporated into radiology. Most notably, these include nuclear medicine images that use radionuclides attached to molecules to target specific organs or functions, ultrasound that uses reflected sound waves and magnetic resonance imaging that uses radio frequency (RF) signals generated by the body in a high-strength magnetic field in response to pulses of external RF energy. Computed tomography (CT) uses a conventional x-ray source but reconstructs cross-sectional image planes through the body from the x-ray projected attenuation information. A popular digital replacement for conventional radiography is computed radiography (CR). Reusable phosphor plates replace conventional film-screen cassettes. The plates store the radiographic image as a latent image, and as the plates are scanned with a laser, the latent image fluoresces. This fluorescent output is digitized as the plate is scanned. The most current sensors are arrays of phototransistors that are either direct x-ray detectors or light-sensitive transistors coupled to a scintillator layer. This has been referred to as “direct capture” radiography as no film or phosphor plate intermediary is used.

These techniques require the computer processing of information to form the resulting images. The images, then, are in digital form. They can be turned into film images by converting the digital values into modulated light that exposes film, or they can be turned into softcopy images displayed on a display device, typically a cathode-ray tube (CRT), but increasingly a liquid-crystal flat-panel display. For the purposes of this chapter, a workstation is a device that displays images in electronic form instead of film. This display will be referred to as softcopy as opposed to the hardcopy that film represents.

## A BRIEF HISTORY

The concept of a workstation extends back at least as far as the 1950s in the history of computer science. Early workstations were terminals connected to mainframe computers. This was a change from “batch” computing, in which a user submitted a “job” to system operators and sometime later received the printed output. With workstations, users could interact with the computer

directly. Early workstations also had a place in military applications, primarily owing to Cold War concerns about nuclear attacks on the United States. The Cape Cod air defense system (in the early 1950s) in the United States was one of the first applications of user interaction with a graphical display. This system employed multiple CRT-based display stations with which the user interacted using a light pen. Many of the individuals involved in these early efforts went on to careers that involved furthering the notion of distributed, easy-to-use computing power. Work at the Xerox Palo Alto Research Center (PARC), in particular, led to many of the concepts that define the modern workstation. These early systems influenced development of the Apple Macintosh, X-Windows on UNIX machines, and, more recently, Microsoft Windows by introducing the fundamental concepts of the graphical user interface. Among these concepts are the idea of the pointing device (e.g., mouse or trackball) that initiates an action on a file or program by pointing to it and requiring a manual action (such as clicking a mouse button).

Expansion beyond these early mainframe systems required that some of the power of the centralized computer be distributed to smaller computers. A most important step in the history of computer science was the introduction of the microprocessor. This advancement in integrated circuits put the majority of the circuitry of a room-sized mainframe computer onto a series of silicon chips about 6 to 10mm on a side. Cost of these microprocessors was so low that true distributed computing became possible. Incorporation into personal computers and workstations took place fairly rapidly.

The history of workstations for medical imaging is less clear than for workstations used in more general computer science applications. To some extent, this depends on how one defines a workstation. Taking the definition to be one of a computer graphics-intensive display device intended to support interactive display and manipulation of images by physicians or other healthcare professionals, the first “workstations” were likely those used as “independent consoles” for CT scanners, nuclear medicine systems, and radiation therapy planning systems. No doubt many readers of this chapter can find the occasional example of the very early use of a computer graphics display device for medical imaging, but this chapter will concentrate more on references from the radiological literature.

Very early references on medical imaging workstations are from the mid- to late 1970s. Larsen described the application of an image-processing workstation to CT scans of the head in 1976. Of interest is that Larsen noted that one of the important ideas of the use of the image processing was that



it would be interactive so that the user would get rapid visual feedback from the processing performed. Lemke et al., describe what may be the first medical picture archiving and communications system (PACS) in the literature. The system was designed to perform image processing on head CT images for enhancement and automated extraction of features such as ventricles. It included workstations for image processing and evaluation, a computer network for communicating the images, and a filing system for storing them. Of importance is that this paper referred not only to the processing and display of images, but to the management of them and the integration of information sources for the medical record such as text and voice annotation of the images.

By the time of the first International Workshop and Conference on PACS for Medical Applications in 1982, the concept of a PACS had clearly spread through much of the radiology research community. The author and colleagues described a system at that meeting (much like the one outlined by Lemke). They constructed the workstation prototype from a minicomputer, graphics display subsystem, and magnetic tape drive. Dwyer and coworkers outlined many of the requirements for a radiology PACS, including the data on retrieval rates for radiological examinations that formed the basis of many designs for archives. Even from the beginning, many users and some vendors described the importance of user interfaces, integration with other information systems, and analysis of what it is that radiologists do when they read films.

## THE ROLES OF THE WORKSTATION

The workstation has one of the more difficult jobs within a PACS. While all components of a PACS are needed for the system to operate, the workstation has the unique function in serving as the human interface for the display role. As such, it has to support a user interface that is readily learned and used, yet be flexible enough to meet the varied needs of different health-care professionals. The film-based environment that is so familiar to radiologists has evolved to the point where the “user interface” is well known and the roles played by support personnel are transparent unless there is a problem (see Figure 20.1). Workstations also have to be able faithfully to display the information they are sent or, if they cannot, to provide some indication to the user that some change from original data may have occurred. The speed of workstation operations has to be high enough so as not to result in user frustration but not so high as to drive costs for the hardware to impractical levels. This author also believes that workstations have to be



**FIGURE 20.1**

This photograph is of the “workstation” for film-based reading. This film multiviewer shows up to two panels of four  $14 \times 17$ -inch ( $35 \times 43$  cm) films at one time. Switches allow any combination of the eight light boxes behind the films to be turned on or off. The footswitch on the floor causes the mechanism to move through the panels to display a particular set. The device on the right of the work surface is not a telephone but the digital dictation system. Radiologists are so used to this environment that it is hard for them to enumerate all of the tasks performed in addition to reading the films.

able to provide a clear advantage over film for the users, whether in speed, functions performed, or ability to do tasks not possible with film. If this is not the case, much of the drive to switch from hardcopy to softcopy will be absent.

It is important to examine in what roles a workstation will serve. The one best known to radiologists is the diagnostic role. Many would say that this is also the most demanding, as the workstation is being used for primary reading of examinations. Anything that detracts from the accuracy of the radiologist is likely to be a cause to reject the workstation. However, other

roles also parallel the use of conventional light boxes and films. There is the review role that the radiologist also does on examinations that have already been interpreted. This may be done for groups, as in a conference, or for individuals, as when consulting specialists seek a radiologist's opinion. There is a review function that the referring physician performs, usually on examinations that have already been interpreted by the radiologist. This is often done for learning and for confirming treatment decisions. A review use is also required by other healthcare professionals, including physical therapists, nurse practitioners, respiratory therapists, and others. These roles are important not only because they help determine the level of performance and functions that will be needed, but also because they will shape the user interface design. Certain workstation applications, such as those in an operating room, require special hardware configurations, designed to meet special electrical safety requirements and with a user interface tailored to hands-free (or sterile operable) controls.

The distinction between "diagnostic" and "review" workstations, the two major subclasses of user requirements, has been used to define hardware requirements and some software functions and was typically based on equipment location ("diagnostic" in radiology reading rooms; "review" elsewhere). However, with the increasing use of images for guiding therapy or planning and performing surgery, the diagnostic role now extends beyond radiology into other specialties. Rather than divide workstations into categories based on their location in a radiology department or outside of it, this author believes the classification should be based on the workstation application.

Different clinical specialists will likely have requirements for the workstation that vary from those of radiologists. The author conducted a study comparing medical intensive care unit (MICU) physicians with radiologists using a workstation for portable chest radiographs done on CR plates. He found that some functions were used more frequently by the radiologists than by their MICU colleagues. In particular, the radiologists used the grayscale adjustment ("window" and "level") more frequently than did the MICU physicians. However, the MICU physicians used the other workstation image manipulation functions (zoom, invert, and full-resolution display) more frequently than did the radiologists. In the task evaluated (reading portable chest radiographs), the radiologists were doing primary interpretation of the images. Such studies tend to be limited to specific image categories or reading conditions and should not be used to generalize a list of workstation functions. However, such studies may be useful in determining which functions are most readily selected under different workstation use scenarios.

## WORKSTATION HARDWARE

Aspects of workstation hardware are also discussed in other chapters of this book. However, a short overview of displays is provided here since the software interaction with the display devices is often close. Other sections of this chapter discuss aspects of the workstation hardware that influence, or are needed by, the workstation software.

While being displayed, the image is stored in the same sort of memory used for computer information: solid-state random-access memory (RAM). Since the type of CRT used in medical image displays does not store the image, the memory is used to “refresh” the display, typically at a rate of 60 or more times per second. The type of liquid crystal display most commonly used for medical imaging is the active matrix liquid crystal display (AMLCD, or simply LCD). These are the same as used in almost all laptop computers. Since each display element is a set of transistors that act as memory, the display is refreshed only if the image on the screen changes. Since the refresh process for CRTs and the update for LCDs need to be fast, separate memory called a frame buffer is part of the display system. In current workstations, this memory is on a separate circuit card that often has additional processing circuits on it to speed up the operations usually done on images (changing grayscale, displaying graphics elements, supporting zoom operations). Most workstations also have some form of storage for the images not currently in the frame buffer. This is usually magnetic disk, but large amounts of RAM also sometimes serve as intermediate storage. Images are transferred to this RAM area from magnetic disk, other storage devices, or the network, and can then be moved to the frame buffer much more rapidly than if they were transferred directly from the other storage.

Displays, whether CRT or LCD, are characterized by whether they are color or monochrome and by how large an array of pixels they can display. Typical resolutions are described by the number of pixels that can be displayed simultaneously. Note that the frame buffer memory may, in some circumstances, store a larger image pixel matrix than can be completely displayed on the connected display device. For reference, most personal computer displays are color,  $1024 \times 768$  pixels. This is referred to as a “three-quarter megapixel [MP]” display (the total is actually 786,432 pixels). Other display resolutions are often available on personal computers; increasing to  $1280 \times 960$  pixels results in a “1.2 megapixel” display (actual: 1,228,800 pixels). Medical displays are more typically monochrome with display matrix sizes ranging from  $1280 \times 1024$  (1.3 MP) to  $2560 \times 2048$  (5 MP). An intermediate display matrix of  $2048 \times 1520$  pixels (3 MP) is becoming more

popular for many clinical applications, including diagnostic work. Recent availability of a 9-MP display has resulted in some prototype workstations that use the display to show 2 CR images simultaneously.

Many PACS workstation vendors are moving toward a “standard” hardware configuration of 3 LCD displays per workstation: a single color display and a pair of 3-MP grayscale panels. A pair of 5-MP displays can be substituted for tasks that commonly rely on high resolution (some pediatric and adult musculoskeletal work, digital mammography, and some chest radiography). The color flat panel is typically used for display of integrated dictation system controls, radiology information system (RIS) or electronic medical record displays, and worklists. The color panel can also be used for the display of color-containing images, such as ultrasound, nuclear medicine, and three-dimensional (3-D) rendered studies, when needed.

Most workstation CRTs are being replaced with LCD flat panels for sales of new systems and retrofitted to older systems, particularly as CRTs reach their end of life or PACS are updated.

## **WORKSTATION SOFTWARE: SOFTWARE FOR OPERATIONS AND FUNCTIONS**

The software running on the workstation determines what functions it supports and how the user interacts with it. Before one examines workstation software design concepts, an examination of what users need to accomplish at workstations and how they interact with the software and hardware is a useful beginning. Workstation functions and the user interface are the two major areas of software design that will be considered in this section.

A basic requirement for workstations is speed. Some types of film alternators (multiviewers) can move films so fast that they go by in a blur (the transparent band types). Even the slower film alternators can move 8 sheets of film (amounting to some 96 to 120 images for many of the cross-sectional imaging methods) in some 3 to 5 seconds. This is a very high performance to match in a workstation. For the example of CR films, moving 8 of them in 3 to 5 seconds amounts to a data rate of some 10 to 16 megabytes per second. In general, an image display time of less than 2 seconds from the time the image is selected for display is a desirable target speed. Most users will not tolerate display times longer than this for routine displays. If they understand that an ad hoc request is taking longer because of the time to fetch the study from the archive, they are more willing to tolerate a longer

delay. It is desirable to provide either a progress indicator for, or a notice about, how long such a retrieve will take. This avoids the problem of the user being unaware about whether or not the system is working. With a multitasking workstation, a retrieve can be done in the background while other processes (for example, reading another study) can take place in the foreground. A feature that was implemented on the workstation developed at the University of Pennsylvania took advantage of common user actions. Kept on local disk would be the current examination plus up to 4 prior studies. However, icons for up to 4 studies older than the first 4 were also displayed. The software automatically displayed the current examination plus the most recent previous study, and users could select older examinations by clicking on the icons. Once a user selected the third prior study, the software would assume that even earlier studies might be desired and would initiate a request from the archive for the older studies. In this way, if the user got to the fifth older examination, it would either have been loaded from the archive or the retrieve would be in process. Since most users did not request studies beyond the second or third, this avoided filling local storage with examinations that were unlikely to be reviewed.

Workstations also have to provide for a navigation method for finding examinations to be read, comparison studies (older examinations of the same type), and correlation examinations (current or prior examinations of a different type). This navigation process is related closely to the user interface and the mental model assumed. The user interface issues are discussed subsequently. Table 20.1 lists commonly encountered navigation functions. On most PACS workstations, the study navigation process is handled through worklists. These are often developed specifically for a medical imaging department and may vary by section of the department and by who is reading. Typically, worklists include such basics as unread studies, emergency studies, and resident-interpreted studies for review. Since worklists will vary greatly by facility, it is more important to understand the concept of the worklist than to understand the specific types. The basic idea behind a worklist is to provide the radiologist with a list of work to be done. Ideally, the worklists are also managed by the PACS (or the RIS, or the PACS and RIS in conjunction) so that, for example, when a radiologist displays a study from a worklist of unread studies, it will be flagged so that another radiologist will see that it is being read and will progress to a different examination.

A method for adjusting grayscale is evident from most workstation designs and studies of those designs. Adjustment of window width and level has been consistently shown to be the image manipulation tool used most frequently. Other image manipulation tools typically include some form of

**TABLE 20.1**  
Navigation Functions

Worklist	Display a list of examinations to be interpreted (should be reachable by a single keystroke equivalent). The user should be able to sort the worklist chronologically, by patient name, or by examination type with sub-sorting by date or patient name. Searching for a particular patient should also be supported using any of the identifiers as search keys (e.g., patient name, medical record number, order/request number, etc.)
List all	Display a list of all the current day's examinations (or other time interval set by the user)—the daily list would be typical (also reachable by a single keystroke equivalent). It should be possible to sort this list as for the worklist.
Folder display	For each patient, show the analog of the master jacket
Image icons	In exam folders (within the patient folder), show icons of images in studies.
Next patient (exam)	Moves to the next patient or examination on the worklist.
Previous patient (exam)	Moves to the prior patient or examination on the worklist.
Consult	Interrupts current display, brings up "List all" or "Worklist", and allows selection and display of another patient or examination for consultation purposes.
Consult remote	For consultation between workstations. Invokes the display of the same study on two workstations and two cursors, one controlled by each workstation so that the two users may point to different areas on the image.
Resume last	Returns to exam interrupted by the Consult or Consult remote function (should return to images, format, and grayscale settings as prior to invoking either Consult mode).
Mark as read	Marks the examination as "read". This should alter the appropriate flag in the image database so that others looking at the worklist or patient list will see the marked examinations displayed as "read". This function should also: a) be restricted to use by those permitted to read examinations, and b) prevent more than one person marking the examination as "read".

TABLE 20.1

*Continued*

---

---

Compare	Brings up the patient folder so that a prior exam of the same type or other examination can be displayed for comparison and correlation purposes. Note that some examination types should have a site-configurable automatic display of prior examinations (e.g., chest radiographs will almost always be displayed with the prior examination). Display formats should be both site- and user-configurable.
Stack view/Tile view	For cross-sectional and time-sequenced images, allows rapid selection of the stack view (rapid sequential image display in a single image frame) or tile view (images arranged in sequence but spatially separated much as photographed on conventional film). Ideally, stack view should be implementable with a selection on any image in a tile view display; ending the stack view would then return that frame to its original image. Stack view should also be applicable between and within (different series; e.g., non-contrast/contrast, or pulse sequences in MR) examinations for comparison purposes.
Cine loop view	For time-sequenced images, allows variable-speed display of the images to show dynamic phenomena.
Rearrange/Save	Allows “drag and drop” rearrangement of the images displayed. While “hanging protocols” should be automated as much as possible, rearrangement is important for comparison purposes. The undo function should return the display to the original configuration. A “Save” function should store the information about the rearranged images so that the next time the same study is called up, it is displayed as rearranged.

---

---

Note: Also see Table 20.3. There are functions, such as “worklist”, that cross boundaries between image navigation and management.



magnification (either an overall magnify or the analog of a magnifying lens, or both), roaming (either moving around in the magnified image or moving the magnifying window), video inversion (“black bone”), and rotation and flipping of the image, all typical of those provided. Some workstations also perform fast image processing to support edge enhancement, grayscale histogram equalization, and image arithmetic (adding or subtracting images). Specialty workstations may add additional image manipulation functions. Table 20.2 summarizes these image manipulation functions.

Measuring features on images and calculating pixel statistics are 2 commonly available functions. These are deemed a necessity for CT workstations because of the heavy use of these operations on dedicated CT display consoles. Such measurements are diagnostically important and, in film-based operation, are often filmed as overlays on the image.

The rapid expansion of multidetector computed tomography (MDCT) has resulted in studies with large numbers of images. In many instances, the number of images per study has increased by an order of magnitude. Such large studies are proving a challenge to PACS and workstation developers. In addition to the much increased data volume, there is the problem of how to present studies of 500 or more images for interpretation. One result of this is an increased interest on the part of radiologists in having 3-D reconstruction tools built into the workstations usually used for cross-sectional imaging. At present, to create and view 3-D images, a separate graphics-intensive workstation is typically used. However, this means the radiologist has to move from one workstation to another; an undesirable situation.

A method for displaying the radiology reports of the examinations being viewed is essential. In subspecialty radiology practices, it is very useful to know how a correlative study was read when it is being viewed as part of the interpretation of an examination of a different type. In more generalist practices, it is still important to know how a colleague read a prior examination. Display of reports usually requires that some form of interface to the RIS or hospital information system (HIS) be operational.

To a large extent, a successful conversion from film-based reading to the use of workstations for reading will also involve the replacement with automated equivalents of many of the functions now done manually. Selection of prior studies to display, arrangement of images when displayed, and clearing of films from the alternator when the readings are done are all examples of image management functions that need to be supported in the workstation environment. If this is not done, the transition to workstations will also mean the migration of support tasks to the radiologist. While this might seem attractive at first because of potential personnel reductions, the radiologist's time is very expensive and his or her productivity is very dependent

**TABLE 20.2**  
Image Manipulation Functions

<b>Function</b>	<b>Description</b>
Window width	Adjusts range of pixel bits sent to the display system. Should always be accessible.
Window level (center)	Sets center of window width in the bits of the pixel. Should always be accessible.
Grayscale reset	Resets window width and level settings to initially displayed values.
Zoom/magnify	Invokes an operation that maps the displayed image into a larger area (zoom is typically by pixel replication; magnify uses pixel interpolation).
Place ROI	Places a preshaped (circle, square) or hand-drawn region of interest on the image. The controls should allow for the user to move the ROI (region of interest) around the image and change its size.
ROI magnify	Magnifies within a movable ROI. This should allow for variable degrees of magnification (may be in discrete steps) and should be movable in real time using the manual input control.
Grayscale invert	Reverses the meaning of the highest and lowest pixel values in terms of grayscale; for example, white bones would be displayed as black.
Window width and level presets	Single keystroke equivalent selectable preset window width/window level combinations, for example, lung, mediastinum, and bone for CT.
Image processing	Image-processing functions are not commonly used while reading cases; however, access to such functions as edge enhancement, histogram equalization, and application of user-defined filters could be provided through a button that brings up a menu of such functions available.
Undo	At least 1 level of being able to undo the function just completed. This should be accessible with a single keystroke equivalent. Note: An undo of some image-processing functions may be difficult to implement unless a copy of the "original" image is maintained.
Pixel statistics	Reports the number, mean, and standard deviation of pixels in an ROI.
Measure	Provides linear, area, and angle measurements on an image.
Annotate	Allows a user to add text annotation to an image. This function can be used with the display of an ROI or graphics arrows. Users should be able to toggle the annotation on and off. Such annotation should be saved with the image.

TABLE 20.3

## Image Management Functions

Delete exam (local)	Provides for deletion of an examination from the local storage (with appropriate warnings).
Autodelete	Provides for automated deletion of examinations from local storage (if applicable) based on user-set criteria.
Mark for teaching file	Flags the examination being read as one useful for teaching purposes.
Mark for nondeletion	Flags the examination being read as one not to be autodeleted from workstation local storage (if applicable). Note that a site-configurable expiration time for this “protection” is useful.
Redirect	Allows sending of an examination to a particular workstation. Even in a centralized storage PACS, this function would allow sending the examination as viewed, annotated, and measured. DICOM has a way of supporting this through the “presentation state storage” service class.
Scrapbook	Saves selected images (or pointers to images) in a file. Retrieving this file later would (optionally) first display the scrapbook images while providing access to the full study, if requested. DICOM “key images” are an example of what would be put into the “scrapbook”.
Print	Provides hardcopy printing of selected images (e.g., the Scrapbook) or a whole study. Print options (film, paper, 35 mm slide) depend on devices in the system. Print function permission should be configured by type of user.
Local storage statistics	If the system uses local storage for images, this function should display the percentage of capacity currently used.

on these support functions being accomplished. Simply moving these operations from clerks and other personnel to the radiologist will have a negative impact on productivity. Table 20.3 lists a set of these image management functions. It is the creation and management of worklists that encompass most of these functions.

**TABLE 20.3***Continued*

---

---

In progress notice	For time-consuming operations (such as retrieving an archived case) this function should display a progress indicator to give the user some idea when the operation will be completed.
Hanging protocol	Allows for site- and user-customizable arrangements of images, series, etc. as defaults for initial displays. For example, some radiologists prefer to read chest cases with the current films in the center two of four light boxes and the prior study on the outer two light boxes; others prefer the pair of frontal images on two adjacent light boxes and the pair of lateral images on the other two. It should be possible using electronic hanging protocols to support both of these possibilities and others. The arrangement used should be based on user login.
Associate report	If a radiology report is available for an examination on the daily list or for prior studies that will be displayed, a method for displaying that report should be implemented. An entry on the daily list should show report status (none, preliminary, reviewed, finalized or signed, referring physician reviewed). The latter status is one that is being added to the others at many institutions.

---

---

Note: Also see Table 20.1. There are functions, such as “worklist”, that cross boundaries between image management and navigation.

## **THE USER INTERFACE: HISTORICAL BACKGROUND AND ANALYSIS**

The “first approximation” approach to designing workstations has been to emulate what is done with film on light boxes or alternators. Although this approach has flaws, as will be discussed, it nonetheless provides a good starting point. One of the most useful tools in performing the analysis needed for this approach to workstation design is a detailed time-motion study. A very thorough such examination was done by Rogers et al, which should be read by anyone involved in workstation design. The authors found that the

reading process could be divided into 3 stages. The first stage is a general information-gathering period during which the patient's history, radiological examination history, and other data are reviewed and a rapid sort through the films of the current study is undertaken. The second stage consists of a rapid review of the image set directed at detecting candidate pathological states. The third stage focuses attention on those candidate abnormalities with the dismissal of some as nonpathological and others as significant. This final stage is what most would think of as the actual "reading" or diagnostic phase. Some further analysis of that first stage, gathering information, was done by Levin et al and Haynor and Saarinen. Levin and colleagues did a workflow analysis of all the activities required to bring together the information needed to support the interpretation of nuclear medicine ventilation/perfusion lung scans. This was done as an example to determine which steps could be done within a PACS, which were handled by other information systems, and which were manual operations. What emerged from this study is that there are extensive interactions with multiple information sources to support examination interpretation. Haynor and Saarinen looked more at the problem of finding any prior examinations of the same type that the patient has had and any correlative (cross-modality) examinations that the radiologist might want during reading. They found that inpatients had a higher frequency of prior studies compared with outpatients when considering plain films. Also, the inpatients had old studies that were more recent than those of the outpatients (an average of 15 days as opposed to 53 days). For correlative studies, with plain films, ultrasound, and neuroradiologic magnetic resonance (MR), the correlative study was a prior CT examination. For neuroradiologic CT, it was a prior MR or myelogram. What this study suggests is that relatively simple rules can be constructed for the automated retrieval of these needed cases from a PACS archive.

Note that the above studies are primarily for activities taking place in support of the reading task and secondarily for the actual reading process. Nonetheless, without such support processes, radiologists would be reluctant to interpret examinations. For a determination of the requirements of diagnostic workstations, there are several possible approaches. One involves interviewing radiologists in addition to watching what they do. Another is based on using the actions identified in time-motion studies or building functional models of the radiologist's working tasks and combining these with human-computer interaction principles. A third approach involves the building of models or prototypes, having them evaluated by radiologists, refining them based on such evaluation, and repeating this cycle until requested changes reach a minimum or the users express a high degree of satisfaction with the result.

The survey approach was taken by the author, by Hohman et al., in the University of California, Los Angeles (UCLA), Department of Radiological Sciences, and by Honeyman et al., among others. The work by the author was done when there were only prototype workstations in his department and most radiologists did not use them but did use the independent consoles of the CT machines to view images. The work by Hohman and colleagues was done using the developmental PACS in place at UCLA. Honeyman and colleagues took the approach of comparing features of 4 workstations installed in their facility and also surveying attendees of the workstation exhibits at the Radiological Society of North America InfoRAD area. Of interest is that despite the 10 years that elapsed between the Horii and Hohman papers, many of the responses were similar. From both studies, it is evident that controlling the grayscale of the display (window width and level adjustment) was the most frequently used function. For the Honeyman survey, the importance of window/level, magnification, and flip/rotate functions was thought to be high enough that users were told to assume that the workstation had these features. One major difference seen between the Horii and Hohman studies is the change in acceptance of the mouse and menu for control of the user interface. In the Horii study, the mouse ranked least desirable and exactly the opposite was true in the Hohman survey. This difference, the author believes, is readily explained by the rapid growth of personal computers with graphical user interfaces for which the mouse and menu is the predominant control mechanism. The Honeyman study also asked about "higher level" workstation functions—those that involve organization of images or their management. The most important function in that study was the ability to view multiple studies on the same patient simultaneously.

Either a time-motion study coupled with human-computer interaction principles, or functional requirements modeling based on more informal observation combined with ergonomic design is a popular method of arriving at workstation requirements. A very detailed such analysis was done by Meyer-Ebrecht and proposed to use a well-defined set of image-processing primitives to implement workstation functions. The previously cited paper by Rogers et al., analyzed the time-motion study they did and determined a basic set of functions that workstations supported by a PACS should perform. Included were access to patient history (text and prior images), window/level presets, controllable automatic progression through series, temporary image markers to return quickly to a marked image, display of multiple images, magnification, and measurement (pixel value and dimensional). With the possible exception of patient history (other than what is provided on the radiology requisition), most of these functions have been implemented. Of

interest is the “controllable automatic progression through series,” which from the description in the paper is familiar now as the “stack view” mode in which a user can move through a series of images very rapidly. For cross-sectional imaging, this display mode allows very rapid review of studies and provides a 3-D sense of anatomic structures.

Most of the authors who use the model and human factors methods arrive at similar conclusions; that a user interface developed from a simple or intuitive mental model is easier to learn and use with a minimum of errors. Also, there is the indication that the radiology reading task is a complex one, from the point of view of both the radiologist (the “internal” processes) and the infrastructure to support the radiologist (the “external” processes).

The idea of developing a prototype user interface and refining it by having different users work with it has also been well used. In effect, this same cycle happens with products installed by vendors but with a longer cycle time. In the prototyping method, trial user interfaces are constructed and refined before they are committed to “production” hardware. Among the researchers who used this method, the work by Belanger and van der Voorde is important in that a great deal of time was spent with radiologists and nonradiologists in evaluating or developing the user interface. Both Belanger at Ottawa Civic Hospital and van der Voorde at the Hospital of the University of Pennsylvania developed workstations that would be used extensively by referring physicians. Rather than assume that functional requirements were the same as for radiologists, they sought the advice of the physician population that would be using the workstation. The result in both cases was a high degree of acceptance by the nonradiologist users.

## SYSTEM DESIGN CONSIDERATIONS

The overall architecture of a PACS will determine the hardware design of the workstations connected to it. The typical components of a workstation are those of a personal computer: the computer “box” (contains the processor, memory, power supply, and slots for other circuit cards), the display/graphics processor card, the display device, and input devices (typically a keyboard and mouse). The design of the PACS will largely determine both the hardware specifications and software functions for the workstation. There are basically 2 major PACS designs, the centralized storage version and the distributed storage version. The design chosen for the Department of Defense Medical Diagnostic Imaging Support (MDIS) System is the archetype of centralized systems. All images are stored in a very-high-speed

magnetic disk array that is supported by a database management computer and a long-term optical disk archive. The high-speed, wideband nature of the storage system means that workstations do not have to store images locally and that the workstation hardware does not need to include a large-capacity magnetic disk of its own. The MDIS design does include local magnetic disk, but for storage of the workstation software and utilities, not images. Images are stored locally in the display memory only. The centralized design must have speed and bandwidth necessary to support fast retrieval of images from multiple workstations. In the MDIS system, for example, any workstation can retrieve images from the central storage system in 1 second or less per image. This performance also requires that the network connecting the central storage device and the workstations be very fast, or else the high speed of the storage is of little advantage.

The alternative design is a distributed system in which there may be one or more short- and long-term storage devices, but the images in current use are stored at the workstation. In this case, the workstation needs to have enough local magnetic disk to store as many images as needed for a typical reading session. This also must include any prior or correlative examinations retrieved from long-term storage. In some senses, this design trades off workstation local storage against network performance. Low-cost local area networks (LANs) can be used to implement the image transmission network, subject to constraints of the management system. Distributed systems also need to have good image management so that images are moved to where they will be needed (workstation local disk) to avoid long transmission delays resulting from network performance. Use of a distributed design with a moderate speed (about 10 megabits per second) LAN is possible with such good image management and also benefits from predictable loads and behaviors. If radiologists and other users tend to view images from the same locations and if the number of ad hoc requests (those not predicted from RIS or HIS information) is low, such designs can work quite well. These considerations also do not include the impact of data compression, particularly lossy compression, which can reduce data per image by a factor of 10 or more. Other chapters in this book discuss the compression issue in more detail. With recent availability of much faster LAN hardware (100 and 1000 megabits per second), many PACS designs are shifting to a centralized storage model as the network is less of a bottleneck. Countering this is the growth of examinations with large numbers of images; moving such studies across even fast LANs can be a problem when there are a large number of such studies to be moved.

The two models are not mutually exclusive; there are designs that are hybrids of both. In particular, there are PACS made up of interconnected smaller systems that may themselves be centralized-storage designs. Such



designs tend to be those that have “grown” rather than installed as a single system. In some senses, this is a good design, if somewhat engineering-intensive, as each subsystem can be optimized for its particular tasks.

Chapters 13 and 16 of this volume detail the designs and storage systems of PACS. The design alternatives are discussed here as well because they have an impact on workstation software. A current model of storage is very different from the “traditional” design in that it does not rely on a large-capacity robotic media library. Instead, the proposal is to keep all studies “online” for the required legal storage duration. In this model, storage is purchased as needed. This takes advantage of steadily declining prices of magnetic disk systems. The typical increase in magnetic disk storage capacity quoted by the industry is that it doubles about every 18 months. The impact of this sort of design on PACS workstations is the very rapid availability of even old studies.

In both the central and distributed storage models of PACS, the workstation software supports the various functions the user needs. The software to do all this requires many steps and is usually a large program as a result. An emerging model for PACS and workstations takes a very different approach. Instead of having all the functions implemented in a program that resides on the workstation computer, a fast server could perform many of these tasks; the workstation would then need to manage only requests from the user and responses from the server. As a result, the program to do this would be much smaller. The familiar example of this sort of design is the World Wide Web and the use of browsers. The browser program itself is simply an interpreter of user actions and the programming needed to convert user requests into the language of the Web and convert the server responses back into a user-friendly form. The Web uses a special language, hypertext markup language, or HTML, as the way of encoding what is requested and what is sent back for display. If special functions are needed (for example, software to display animation) at the user’s computer, the Web protocol allows these functions to be loaded into the user’s computer at the time they are needed. What this means for the user is that the browser software is a relatively small program compared to software that would be needed to carry out all the functions available on the Web. A system design that moves many of the computing (and storage) functions to the network and servers and that, as a result, relies on a smaller program on user machines is a “client-server” model with a “thin client.” The term “thin” refers to the fact that the software at the user computer does not have programming built in for many of its functions; it relies on the server for them (as opposed to “thick client” designs, in which many of the functions are coded into the workstation software).

There are PACS designs that take advantage of the Web concept. The user runs a browser on the workstation and the PACS server supplies the images and many of the database (e.g., worklist management) functions the user needs. Figure 20.2 shows two examples of commercially available Web-based PACS image viewers (GE Centricity Web, General Electric Medical Systems, Milwaukee, WI and Stentor iSite, Stentor, Inc., Brisbane, CA). The great advantage of these designs is that the workstation moves from being a specially configured computer to a commercial off-the-shelf one. However, for the high-quality displays needed for medical imaging, the display portion of these workstation computers is still often different from




FIGURE 20.2A

General Electric Centricity (General Electric Medical Systems, Inc., Milwaukee, WI) WWW-based viewer. The left side of the screen shows small icons of the images. The currently-displayed images correspond to highlighted icons. The very top of the image shows the Microsoft (Microsoft, Inc., Redmond, WA) Internet Explorer control bar. Various function controls are represented as buttons around the image display area.



FIGURE 20.2B

Stentor iSite (Stentor, Inc., Brisbane, CA) WWW-based viewer showing ultrasound images. Note that the very top of the image shows the same Internet Explorer control bar as for the GE Centricity application. What is different is the application running in the window. The bar above the images is used to show a timeline of the examination history for this patient. In this case, the patient has only this examination in the “folder”. The row of images below the timeline provide an overview of the entire study and, if there are more images than can be displayed on the screen, the row can be scrolled back and forth to view all the images. A double-click on any image expands it to a larger size. Note that there are fewer buttons for functions than shown in Figure 20.2A. This application implements most of the same functions with mouse clicks and position-sensitive menus.

consumer hardware. Another major advantage of these Web-based designs is that much of the software that had to be specially developed for PACS is already available for Web applications. An example is for security; the encryption of confidential medical information can be handled in a manner that is used on the Web for activities like credit card transactions.

## IMPLEMENTATION

There is little doubt that the most heavily discussed aspect of workstations for medical imaging is the user interface. In part, this is because there is wide opinion about how such an interface should work. The “classic” model of emulating the operation of a film alternator is now countered by the concept of going beyond what can be done with film. Manufacturers have tended to use the user interface design to differentiate one product from another because the hardware is often the same or similar across multiple manufacturers.

One difficulty with user interface design is that objective study of the resulting interface is difficult and time consuming. Users will find the obvious problems with a design very quickly; it is the more subtle difficulties that may not have immediately evident impact, and only thorough examination will reveal them. At best, they may cause problems such as longer reading time or complaints of fatigue. At worst, they will degrade diagnostic accuracy.

The reader is urged to become familiar with the early work on workstation user interfaces done by groups at the University of North Carolina (UNC) and the University of Arizona. The author believes that this work is still relevant to modern PACS workstation software. These groups have applied well-validated methods of industrial psychology to the evaluation of user interfaces for medical imaging workstations. The results are sometimes counterintuitive. Beard and colleagues at UNC Chapel Hill, for example, found that a very fast, single-screen workstation could be as effective for reading CT studies as multiple films on light boxes. This is despite the widely held opinion that multiple-monitor workstations are needed for CT reading. More recently, Reiner and colleagues have shown that the reading time decreases significantly when moving from a single-monitor to a 2-monitor workstation; but the decrease in reading time from 2- to 4-monitor workstations was not significant.

Krupinski and Lund at the University of Arizona conducted a very important study of visual gaze fixation during the image-reading task. They found that not only did it take longer to fixate on an abnormality on workstation displays compared to film, but overall viewing time was longer on the workstation. Of considerable importance is that they also found that while for film all eye fixations were confined to the image area, for the workstation display, about 20 percent of the fixations were in the area of the menu that displayed the image-processing functions. In other words, the radiologists set their gaze on a nonimage area for about a fifth of all fixations. If they were to spend the same amount of overall time looking at the image,

those additional gaze fixations would add to the viewing time per image. A more serious potential problem occurs if the gaze fixations on the control menus decrease the number of fixations on the image, possibly decreasing detection of abnormalities. This should give some pause to workstation designers who have continued to add controls and functions to on-screen menus or buttons.

From such research, several key points have emerged. A basic principle of human computer interface (HCI) design is that the user interface should follow a logical and simple mental model. Since the task of reading from a workstation is likely to be different from the film-reading one, the mental model of the workstation interface should also be simple to learn. If radiologists are confronted with an unfamiliar film alternator (or multi-viewer), they are likely to make a few mistakes until they learn how to advance or back up the panels and how to turn on and off the separate illuminators. The number of errors they make until they become facile with the system is a measure of how simple the interface is to learn. An example of a common user interface control is the familiar computer mouse. On a flat desk, moving the mouse away from the user generally moves the cursor on the screen in an upward direction, with mouse movement toward the user being down, and right and left movement corresponding to those directions. Switching these directions is very difficult for almost all users to get used to; the controls run counter to the mental model. This sort of control inversion can have disastrous consequences in other settings. Aircraft that use a joystick respond to pulling back on the joystick by climbing and to pushing forward on the joystick by diving. If the airplane becomes inverted and the pilot disoriented, the reversal of the control responses may not be realized; pulling back on the joystick (a typical response in an emergency) will then put the airplane into a dive instead of a climb.

This author (along with those who have done user interface analysis; Paul Chang and David Beard are two others who have also proposed this idea) has often noted a good test of a user interface: if a person can figure out how to use it without having to refer to a manual, it is probably based on a good mental model. One of the significant advances in personal computers was popularized by Apple Computer (Apple Computer Corporation, Cupertino, Calif) in its Macintosh computer line. Developers were told to adhere to a particular set of interface guidelines. This meant that most application programs operated in a similar fashion. If software developers adhered to these principles, it was possible for users to accomplish most basic operations in any program without having to read volumes of documentation. Unlike other operating systems of the time (e.g., DOS and Unix [Bell Telephone Laboratories, Holmdel, NJ]), the user did not have to remember com-

mands or the syntax of many of them. The Macintosh environment was, and is, highly graphical, and, though not the first graphical user interface (GUI), it became the first widely popular one, with Microsoft Windows (Microsoft Corporation, Redmond, Wash) rapidly following.

From the mental model of “How do I make this thing do what I want?” comes the development of logical controls to perform those tasks. It is this process, turning operations done in a conventional system into controls and their effects in the computer-based system, that constitutes development of a user interface. The trend in medical imaging workstations has been to adopt the GUI-based user interface. Figure 20.3 shows a workstation display designed for viewing ultrasound (Siemens KinetDx, Siemens Medical Systems, Malvern, PA). As noted from the work by Krupinski and Lund, however, that may not be the best possible choice. There have been, mostly in historical workstations, significant departures from the GUI paradigm. Examples include the General Electric Medical Systems Independent Console (G.E. Medical Systems, Milwaukee, Wis) for which multiple keys were traded off against a GUI. The plasma panel of that console did allow for flexible “soft” buttons to be used as well, vastly increasing the number of things that could be done without making the number of function keys impossibly large. The idea used for this console, that commonly used functions would be represented on dedicated keys and those that were needed for specific operations would be on the plasma display, has also been translated into more modern workstation designs. In addition, the control panels of many types of imaging equipment have a similar strategy so as not to overwhelm the user with buttons. Many current PACS workstations have “soft” keys displayed in a fixed portion of the menu on the monitor(s) to perform the common functions. Operation-specific functions are selectable from menu portions that change with the operation being done. The very commonly used functions, such as window width and level adjustment, are usually always available under mouse control. The work by Krupinski and Lund for display of radiographic images on workstations and that of Beard for CT images would argue for a minimal set of controls that do not require users to look at them to operate them. The author refers to such a control set as “low visual effort controls.” A low visual effort control set was not available from medical imaging workstation manufacturers at the time this chapter was written, though many current workstations allow “mapping” of commonly used functions to keyboard keys. If the function and position of these keys is memorized, many functions can be performed without having to find a control button on the screen. Another common feature of current user interfaces is the context- or position-sensitive menu. Clicking on the image with one of the mouse buttons will bring up a menu with image-



FIGURE 20.3

Siemens KinetDx (Siemens Medical Solutions, Malvern, PA) ultrasound PACS display. Unlike the displays of Figure 20.2, this is not a Web-based application. Note that the functions are clustered as “soft” buttons across the top of the images. The rectangular area in the upper left of the control bar is actually a shortcut to the worklist. Clicking on it results in a pull-down list sorted according to what was selected for the worklist (in this case, called “study list” on the control bar). Clicking on the “study list” button results in the display shown in Figure 20.5B.

manipulation specific functions. Figure 20.4 shows an example of “drop down” menus (GE Centricity Web, General Electric Medical Systems, Milwaukee, WI). Clicking on the border surrounding the image will bring up a different menu, for example, functions that affect the display of the whole study or a linkage of one image to another. This avoids having to display a large number of controls or buttons on the screen in addition to the images.

The navigation method most commonly used is based on the model of the paper worklist (usually posted on the side of the film alternator) and film




FIGURE 20.4

General Electric Centricity (General Electric Medical Systems, Inc., Milwaukee, WI) WWW-based viewer showing drop down menus. This is the same application shown in Figure 20.2A. In this instance, the right button on the computer mouse was clicked while the cursor was on the image. The menu near the center of the image then opened. By moving the cursor down through the menu, specific functions or actions can be invoked. In this case, the “Annotation” item was highlighted and this opened a second menu (the middle one) in which yet another choice (“create”) was highlighted. A third menu (the farthest on the right) shows the choices available under the “create” choice. In this example, “box ROI” (box region of interest) has been chosen. It is important to note that this is done with a “click and drag” followed by a single click and not with multiple clicks. Clicking on the “box ROI” choice would result in displaying a rectangular region of interest (for pixel value measurements, etc.) on the image.



jacket. This concept was put forward very early in the development of PACS by Feingold, Seshadri, and Arenson, who referred to such functions as the “folder manager.” In a hierarchical display, the electronic patient “jacket” is shown with the studies it contains. The user then selects the study (and sometimes the series) to be displayed. On most workstations, the worklist

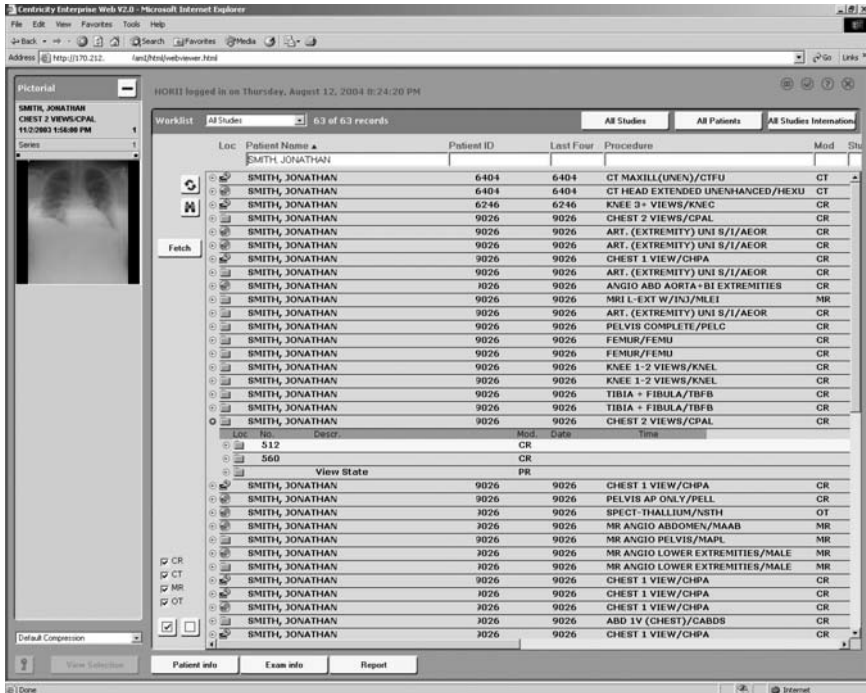


FIGURE 20.5A

General Electric Centricity (General Electric Medical Systems, Inc., Milwaukee, WI) WWW-based viewer worklist. In this example, “all examinations” was chosen as the worklist, but for a particular patient. Other choices would also be available by clicking on the small arrow next to the “worklist” box. The empty boxes above the list of examinations found would allow for searching and sorting the worklist based on criteria typed in those boxes. Notice that there are four “modality” check boxes in the lower left of the worklist screen. These are labeled “CR, CT, MR, OT” for the different image types (“OT” stands for “other”). Checking or unchecking these is a way to filter the worklist. The worklist shows the patient medical record number (the “last four” is used as some practices use the last four digits of the MRN for an identifier), the procedure, and the modality.

Patient Name	Patient ID	Accession #	Study Type	Modality	Date	# of Images	Location
SIMPSON, H	00000001	1000001	UABC	US	08/12/2004, 2:17 PM	41	Online
SIMPSON, H	00000001	1000002	UABC	US	08/12/2004, 1:16 PM	58	Online
SIMPSON, H	00000001	1000003	URPC	US	08/12/2004, 2:51 PM	67	Online
SIMPSON, H	00000001	1100001	UPLA	US	08/12/2004, 8:00 AM	35	Online
SIMPSON, H	00000001	1100002	URPL	US	08/12/2004, 2:40 PM	33	Online
SIMPSON, H	00000001	1100003	UABL,URG	US	08/12/2004, 1:49 PM	59	Online
SIMPSON, H	00000001	1010001	UPLA,URG	US	08/12/2004, 9:14 AM	69	Online
SIMPSON, H	00000001	1010002	Ultrasound Study	US	08/12/2004, 1:40 AM	31	Online
7 SIMPSON, A	00000002	2110001	UPFT	US	08/12/2004, 10:58 AM	37	Online
SIMPSON, H	00000001	1110001	Ultrasound Study	US	08/12/2004, 2:18 AM	26	Online
SIMPSON, H	00000001	1110002	UPLA	US	08/12/2004, 9:32 AM	44	Online
SIMPSON, H	00000001	1110003	UTWB	US	08/12/2004, 10:40 AM	15	Online
SIMPSON, H	00000001	2000001	UTWA	US	08/12/2004, 10:29 AM	24	Online
SIMPSON, H	00000001	2000002	UTWA	US	08/12/2004, 8:23 AM	31	Online
SIMPSON, H	00000001	2000003	UPST	US	08/12/2004, 9:43 AM	53	Online
SIMPSON, H	00000001	2100001		US	08/12/2004, 9:22 AM	9	Online
SIMPSON, H	00000001	2100002	UPST	US	08/12/2004, 9:19 AM	59	Online
SIMPSON, H	00000001	2100003	UPST	US	08/12/2004, 2:41 PM	73	Online
SIMPSON, H	00000001	2010001	UPLA	US	08/12/2004, 1:06 PM	42	Online
SIMPSON, H	00000001	2010002	UABC	US	08/12/2004, 1:48 PM	47	Online
SIMPSON, H	00000001	2010003	USPC	US	08/12/2004, 1:27 PM	13	Online
SIMPSON, H	00000001	2110001	UGEA	US	08/12/2004, 1:57 PM	62	Online
SIMPSON, H	00000001	2110002	UPLA	US	08/12/2004, 9:59 AM	35	Online
SIMPSON, H	00000001	2110003	UABC	US	08/12/2004, 10:13 AM	38	Online
SIMPSON, H	00000001	1200001	UEVC,ROU	US	08/12/2004, 11:20 AM	24	Online
SIMPSON, H	00000001	1200002	USPC	US	08/12/2004, 2:48 PM	12	Online
SIMPSON, H	00000001	1200003	UPLA	US	08/12/2004, 2:02 PM	42	Online
SIMPSON, H	00000001	1210001	URBP	US	08/12/2004, 11:15 AM	6	Online
SIMPSON, H	00000001	1210002	URBP	US	08/12/2004, 9:43 PM	8	Online
SIMPSON, H	00000001	1210003	UPST	US	08/12/2004, 10:06 AM	65	Online

FIGURE 20.5B

Siemens KinetDx (Siemens Medical Solutions, Malvern, PA) worklist. This is the worklist screen from a non-WWW ultrasound PACS. In this example, the worklist chosen is “Today’s” (meaning, “today’s examinations”). This is selected with the box near the upper right of the screen. In the upper center of the screen, there is a “quick search” box. This allows a search of the database by patient name or medical record number. The “study type” column displays the examination code used by the departmental RIS. There is a “modality” column as this PACS can display images from other imaging equipment. Note the scroll bar at the bottom of the worklist; there are other database information items that are not shown in this view but scrolling allows users to access them. Double-clicking on any patient in the list results in a display of the images.

shows a list of patients and selecting a patient then displays the “jacket.” Figure 20.5 shows examples of worklists implemented on different PACS workstations (GE Centricity Web, General Electric Medical Systems, Milwaukee, WI; Siemens KinetDx, Siemens Medical Systems, Malvern, PA; Stentor iSite Radiology; Stentor, Inc., Brisbane, CA). The worklist is a very powerful tool; with appropriate workstation software and an interface of the PACS to a radiology information system, worklists are used to help manage

Class	Patient Name	MRN	Exam Date	Patient Location
+	TEST, A	3955	07/29/2004 08:28:00	
+	TEST, B	3967	07/29/2004 08:16:00	
+	TEST, C	1099	07/29/2004 07:43:00	
+	TEST, D	7499	07/29/2004 07:33:00	
+	TEST, E	4117	07/29/2004 21:32:00	
+	TEST, F	3842	07/29/2004 20:06:00	
+	TEST, G	72378	07/29/2004 18:44:00	
+	TEST, H	59571	07/29/2004 18:14:00	
+	TEST, I	3919	07/29/2004 16:24:00	
+	TEST, J	3955	07/29/2004 16:11:00	
+	TEST, K	3676	07/29/2004 16:09:00	
+	TEST, L	3908	07/29/2004 15:49:00	
+	TEST, M	5270	07/29/2004 15:47:00	
+	TEST, N	3794	07/29/2004 15:43:00	
+	TEST, O	1481	07/29/2004 15:33:00	
+	TEST, P	1104	07/29/2004 15:31:00	
+	TEST, Q	8931	07/29/2004 15:03:00	
+	TEST, R	28145055	07/29/2004 14:50:00	
+	TEST, S	3050	07/29/2004 14:40:00	
+	TEST, T	2805	07/29/2004 14:36:00	
+	TEST, U	3986	07/29/2004 14:00:00	
+	TEST, V	4366	07/29/2004 13:45:00	
+	TEST, W	3802	07/29/2004 13:41:00	
+	TEST, X	8248	07/29/2004 13:17:00	
+	TEST, Y	1599	07/29/2004 13:10:00	
+	TEST, Z	4297	07/29/2004 13:03:00	
+	TEST, AA	4732	07/29/2004 11:31:00	
+	TEST, AB	2378	07/29/2004 11:20:00	
+	TEST, AC	4317	07/29/2004 11:06:00	
+	TEST, AD	9481	07/29/2004 11:03:00	
+	TEST, AE	0193	07/29/2004 11:00:00	
+	TEST, AF	4724	07/29/2004 10:52:00	
+	TEST, AG	7400	07/29/2004 10:41:00	
+	TEST, AH	6226	07/29/2004 10:32:00	
+	TEST, AI	9863	07/29/2004 10:31:00	
+	TEST, AJ	6801	07/29/2004 10:07:00	

FIGURE 20.5C

Stentor (Stentor, Inc., Brisbane, CA) iSite worklist. As for the other examples, the worklist selection was set to “all exams” rather than specific choices. Sorting the list based on entries in the various columns is possible with typed entries (or, if there is a small triangle as for “exam date”) by clicking on the column title. Double clicking on any patient name will bring up the images for that patient. Though all three of these worklist displays look different, the functions they support are the same.

workflow. That is, examinations to be read can be directed to particular workstations or shown on the worklist of a particular radiologist when he or she logs on. Worklists are also used to show referring physicians which examinations have been read and have reports available or may be custom-tailored to display a list of only their patients who have imaging studies. A “lockout” typically prevents 2 radiologists from inadvertently trying simultaneously to read the same examination while still allowing “viewing” access. There is potentially more that can be done with worklists. Through using information from the RIS, for example, information such as the time the

patient has been waiting, “stat” status, or the patient’s appointment time with the referring physician could be used to help the radiologist prioritize reading order. Paying more attention to the design of the navigation process is one area in which workstations can be made to improve on film-based methods. Beard developed a rapid prototyping and evaluation method in part for examining the human-computer interactions involved in the navigation process.

A significant part of image navigation is some form of automated arrangement of the images in an examination when there are multiple images. This has come to be called the “hanging protocol” after the manner in which films are arranged on light boxes or alternator panels for reading. For sectional imaging, such as CT, MR, and nuclear medicine positron-emission tomography and single photon emission computed tomography, the ability to sort images by series, by time of acquisition, by anatomic location, and by imaging protocol (such as pulse sequence for MR) is of great importance. In the conventional film-based operation, this sorting is often done by the technologist as the examination is filmed. Automation of such sorting facilitates interpreting these examination types at workstations. For plain radiographic images, the hanging protocols are also of some importance. In the scenario of a trauma patient who has multiple radiographs done, sorting them simply by examination time may be wholly inadequate. For trauma patients in particular, the radiologist may request additional views done after the first images are obtained. Displayed strictly in chronological order, these additional images will appear either at the end of the examination or the beginning, but not grouped with the earlier views of the anatomic area in question. Other criteria, such as anatomic area being radiographed, also have to be considered. In addition, the sorting of sectional images and the arrangement of radiographic ones is also governed to some extent by the individual desires of the radiologist. All these requirements dictate that such sorting and arrangement be both automated and customizable.

Controls and displays in general could benefit the user if they were customizable. The number of images displayed simultaneously for sectional imaging, the window width and level preset values, the initial image display parameters (very useful for MR imaging), and even the location and arrangement of on-screen controls could all be based on who logs in, so that personalized setups are automatically set. If this sounds extravagant, readers should recall that many such functions were done in film-based operation and were done in a fashion “transparent” to the radiologist by technologists, file room clerks, and radiology residents who quickly learned the preferences of the radiologist with whom they were working.

An important aspect of user interfaces is error recovery. In the best of circumstances, users may still make mistakes. How a user interface allows the correction of such errors is important; if the error-correction process is deemed too taxing by the user, it may be ignored and errors may propagate. Aside from preventing erroneous conditions that are detectable, there is the matter of reversing a function or operation that was inadvertently selected. There are 2 levels of such reversal; the “Undo what I just did” version that can be thought of as allowing a user to back up a step, and the “Take me back to where I started” choice that resets starting conditions.

In the case of the undo function, some things, such as selecting a preset image arrangement choice, can be reversed fairly simply. Others, particularly operations that manipulate pixel data, may be irreversible. Some forms of edge sharpening, for example, cannot readily be reversed. In some instances, warnings to the user prior to executing such an operation may be warranted. Erasing images is a typical example (though a well-designed user interface will provide a safety option for even this action).

The reset function, analogous to a panic button, is useful when a user has gotten so far off the desired functional direction that only resetting to the starting conditions will be helpful. Such a reset, however, should ordinarily not be the equivalent of shutting down and restarting the workstation; a well-designed reset should just reestablish starting conditions and do so quickly.

A manner of quantifying the user effort required to invoke undo and reset functions is based on a model proposed in the MDIS Specification. This model relies on the number of keystrokes or their equivalent (a mouse button click) used to implement workstation functions and has also been used by Beard in developing workstation designs. With the increasing use of computer mice or trackballs for workstation controls, another interesting metric for user effort is to determine the total distance the mouse or trackball was moved to accomplish various tasks. Larger distances may translate into more user fatigue or repetitive motion injuries. For users contemplating the purchase of PACS workstations, determining the number of keystroke equivalents needed to perform common functions is a good method of comparing the effort involved in using different workstations. For example, if a workstation required 6 keystroke equivalents to perform the window width and level functions, this would be, by almost anyone’s measure, far too cumbersome. Typically, such a commonly used function would be either always available or 1 to 2 keystroke equivalents away at a maximum.

User interfaces are a complex subject, and many believe a matter of personal choice. Nonetheless, principles including simplicity of the mental model

behind the interface and minimum actions to perform common functions are applicable across many users and levels of experience. To provide some guidance in developing the workstation section of a PACS request for proposal (RFP), Table 20.4 lists some questions that should be asked of PACS users and vendors. Note that this list is not exhaustive and may have different content based on local experience and requirements. The use of a workstation has human factors concerns that extend beyond the workstation proper. The next section examines the environment in which the workstation is used and stresses its importance on both user comfort and performance.

## WORKSTATION ENVIRONMENTS

### THE WORKSTATION ENVIRONMENT DEFINED

Since interpreting of radiological images has evolved using film illuminated by light boxes, the reading rooms in which such interpretation is done are also suited to such work. The square footage devoted to film viewers, their spacing, power requirements, heat dissipation, and ambient lighting are all based on devices with high luminance. The traffic-flow patterns of people to put films on and remove films from these devices, the radiologists who read the cases, and the referring physicians who come to review the images with the radiologist are also part of the reading room design. In many cases, reading rooms are ad hoc affairs, fit in wherever space is left over or can be taken from other departmental functions.

The problem for workstations is that they do not operate like film-viewing devices. Their luminance is much lower. Though personnel do not have to access them to put images up, they still have to be usable for review of studies with referring physicians or with residents who are being trained. Power requirements and heat dissipation may be lower, but the quality of the electrical power often has to be higher. The voltage spikes and sags, and the electrical noise that light boxes can tolerate might cause damage to a workstation's electronics. There will be more electrical cables and network wiring to be considered. Work surfaces have to accommodate control panels and other input devices and still have sufficient room for the inevitable manual tasks required, such as reading charts or taking notes.

The workstation environment, then, consists of all the elements of the reading room plus the heating, ventilation, and air conditioning and communications and electrical power that are supplied from outside. Aside from the physical layout, the workstation environment is also dynamic and needs

**TABLE 20.4**  
Basic Workstation Questions

<b>Question</b>	<b>Comments</b>
Who is going to use this workstation? Please answer with a list of all users and about what percentage of the time they will use the workstation.	Establishes the user population to consider.
What types of examinations will be interpreted or reviewed at this workstation?	A list of examinations normally read/reviewed. This helps establish display system requirements and number of monitors needed.
What comparison and correlation examinations are needed?	This list should be exhaustive—failing to include, for example, color flow ultrasound may mean that the workstation will have no color display capability.
With what frequency will (a) comparison and (b) correlation examinations be used?	May need data from existing work patterns.
How many examinations per day will be read or reviewed at this workstation?	To determine workload.
What is the projected growth rate in this number per day?	To allow for growth—failure to consider this may result in obsolete workstations very quickly.
About how many hours per day do you spend reading the examinations you are answering these questions about?	This helps establish a reading rate—how many examinations per unit time need to be read.
Do you have a relatively “set” way of arranging the films of an examination?	This helps determine whether automated hanging protocols can be used.
If so, who does this arranging now?	This will determine whether this task is done by the viewer or by someone else.
At present, does the way films are arranged vary from user to user in your setting?	To determine whether hanging protocols need to be customized on a per-user basis.

**TABLE 20.4**  
*Continued*

<b>Question</b>	<b>Comments</b>
Do you work from a list of examinations to be read/reviewed?	Establishes need for a worklist.
If so, who or what system creates this list?	Manual or automatic worklist at present—determines if an interface to an external system may be needed for this. For example, the worklist is generated from the RIS.
Once an examination is read or reviewed, do you need to notify someone or some system of that?	Determines how the user notes an examination as “read.” This may be automatic—see the follow-up question, next.
What system, or whom, do you notify, or is this triggered from your dictation?	Most RIS software includes a function that changes the status of an examination when a report on it is dictated, when that report is transcribed, and when it is finalized. This question is to determine whether a workstation would need to communicate with the RIS about this so that the examination will be marked as “read” or removed from the worklist. If the user reports this to a person, that person’s tasks need to be elaborated.
How do you get comparison and correlation examinations?	To determine who does this—this function likely needs to be automated. This question is to help make sure that this function is not left out.
What image manipulation (window width and level, bright light, grayscale alteration, etc.) do you, or does someone else, do presently? Please make the answer to this question a list of things done.	Establishes image manipulation functions used with film. Tell the user that window width and level changes or presets used by technologists when filming an examination are examples of image manipulation done by someone else for the user.

*(Continued)*



**TABLE 20.4**  
Basic Workstation Questions (*Continued*)

<b>Question</b>	<b>Comments</b>
What measurements do you make, or are made for you, on the examinations you read/review? Please make the answer to this question a list of measurements done.	Example: pixel statistics (ROI) for CT.
If an error is made in the process of performing the examination (for example, the wrong patient name is “flashed” on the film or the right/left marker is incorrectly placed), who corrects such errors and about how long does that process take?	To establish how quality control/quality assurance is done in some PACS; the QC-QA steps may take longer than for conventional film.
Are most errors with examinations caught before you see the films to read?	Helps determine control points for image movement—that is, the QC-QA workstation the technologist uses may be used to correct errors so that they are rarely sent to the reading workstation.
Do you now read/review, and will you be reading/reviewing, examinations with others? If so, how many?	Determines the need for group viewing and teaching functions.
Have you used electronic workstations for reading or reviewing examinations before?	Helps establish the experience of the user.

to account for the movement of personnel and access to the workstations and the people using them.

If nothing else, lighting requirements are very different from film-based to workstation reading. The low luminance of workstations compared to light boxes means that the whole of the ambient lighting has to be lower for workstations. The adverse effects of high ambient lighting have been long known in radiology and were described both before and after the introduction of softcopy displays. These studies all demonstrated that diagnostic accuracy was compromised as a result of high ambient light levels. The curved surfaces of most CRT monitor screens means that glare lighting

sources have to be controlled carefully. This concern also applies to LCD displays, albeit somewhat less so because their surfaces are flat and often covered with an antiglare coating. The lower overall ambient light may require that task lighting be provided for reading or writing activities. This author has described some guidelines for reading room design for workstations. These include indirect lighting (bounced and diffused from the ceiling or wall), avoiding placement of workstations opposite each other, and the use of partial room partitions or barriers between workstations to control both sound and light spill. If workstations and film viewing have to coexist in a single reading room, then partitions between light boxes or alternators and workstations are a must, and placing alternators and workstations opposite each other is absolutely contraindicated.

A well-established quality assurance process for workstation displays should help equalize ambient lighting requirements across reading areas. The DICOM standard (Part 14) describes a process that supports a mathematically defined function that maps the digital input to display devices into perceptually linear output. What this means is that even across different display types (CRT, LCD, film), if calibrated to this standard, equal changes in digital pixel value will result in the same perception of changes in gray level of the displayed pixel. Many of the vendors of medical imaging displays now support hardware and software to make calibration to this standard at least semiautomated. The American Association of Physicists in Medicine (AAPM) formed Task Group 18 with the specific charge of assessing displays used for medical imaging. Their work is a superset, and makes use, of the DICOM Part 14 standard. Readers are referred to the report as it is a comprehensive view of items important to evaluation and quality assurance of medical imaging displays.

Architects familiar with design of radiology departments should be consulted as they have a very good understanding of traffic flow based on room functions. Many architectural offices themselves now use workstations instead of conventional drafting tables, so they have some familiarity with the problems that workstations pose. This author would caution, however, that design principles used for office environments in which workstations are used may not be applied successfully to the medical imaging environment. Many office-based workstation tasks involve very-high-contrast subject matter (e.g., text displays), and higher levels of ambient light are not only tolerated but needed. The ideas about imaging workstation furniture, particularly seating and work surfaces, do have some counterpart in office environments and so can be used as guidelines. An architect with a particular interest in healthcare facilities, Rostenberg has written a book on design of radiology departments that is more up-to-date as of the time this is being

written than other books on such design. The author of this chapter, and Reiner and colleagues, have also written about the importance and impact of reading room design on radiologist performance.


There are anecdotal, but numerous, reports of increased fatigue when reading at workstations compared to film on alternators. The reasons for this are not fully known, but a review of the work of Rogers et al., shows that radiologists spend time performing actions that take them away from the alternator, or at least serve as a break in the reading process. Examples cited by Rogers include reviewing the patient history including the requisition and other paperwork, sorting through films and hanging them, and removing a film from the view box for "hot light" viewing or repositioning. With workstations, particularly those that do support an RIS interface, these tasks are done at the workstation. This author believes that at least some component of the fatigue is due to the lack of the short breaks afforded by the tasks associated with the reading process for film. The increased time spent looking at nonimage areas to perform operations, as noted by Krupinski, may also contribute to this feeling of fatigue since the amount of eye movement involved in the reading task is increased.

Better attention to both the ergonomics of the workstation and the workstation environment should help increase acceptance of workstations for primary diagnosis and decrease user physical discomfort while doing so.

## THE FUTURE OF WORKSTATIONS

One problem with the development of user interfaces and specification of workstation functions is that the model used is film-based operation. While this is a good base on which to build, it tends to obscure the things that can be done with digital imaging workstations that cannot be done with film. Some techniques, such as window width and level adjustment, have become a routine function on workstations that is not possible with film without photographing the image at multiple window width and level settings (which is typically still done in film-based CT reading). For ultrasound, nuclear medicine, and cardiac angiography, cine or dynamic displays can be done with film or videotape, but they are much easier to do on workstations. The "stack view" mode, in which cross-sectional images are viewed rapidly as though moving through a stack, is also readily accomplished on most workstations.

The author has experimented with a technique from the aerial reconnaissance field, sometimes called the "window shade" display. This would potentially allow a display of 2 images to be compared on 1 monitor. A



**FIGURE 20.6**

These are two images of the same screen and show a display method that the author has experimented with for comparison of two images. There are two images placed “on top of” each other. By sliding a bar (vertically in this example, but also configurable horizontally or as a region of interest box) more or less of the image “underneath” is revealed. No attempt was made to align the images (as is apparent from the differing position of the ribs). This allows comparison without having to move one’s gaze from one monitor to another. It might also be useful for showing a processed, different grayscale adjusted, or annotated image together with the original.

movable bar or region of interest (ROI) would show the image “underneath” (see Figure 20.6). Though formal studies of this technique have not been done in radiology, those radiologists who tried the prototype expressed interest in it. One advantage is that the viewer does not have to turn his or her head to view the comparison image on another monitor.

A method employing the registration and subtraction of 2 temporally spaced images (temporal subtraction) has been demonstrated by Kano and colleagues from the University of Chicago. In this technique, the 2 images to be subtracted are first registered so that differences due to patient position are minimized. The subtraction image then shows areas of change as black. Such areas can be used to direct the viewer’s attention during inspection of the image to be read.

Also pioneered by the University of Chicago has been the use of computer-aided detection (CAD). Software for automated detection of breast masses is now in increasing clinical use; there are commercial vendors of such system software. At present, such software is not intended to replace the radiologist but rather to reduce the false negative rate in circumstances in which the consequences of the false negative are severe. The programs can display what they detect as a graphic overlay on the image. The size or

shape of the graphical elements used can indicate the computed likelihood of a real lesion based on the decision criteria the program uses. The radiologist then evaluates these areas on the image and decides whether the detected structure represents pathology.

Aside from functions and methods employed to enhance usability or diagnostic accuracy, workstations must also address the problem of image availability in diverse locations. Haskin noted some years ago that most hospitals have a very large number of light boxes, which implied that there are a large number of locations in which images are viewed. Contemporary healthcare systems involve very large networks of facilities of varying sizes, from physician offices to large medical centers. Since the patient population is served by combinations of these facilities, some manner of distributing images and other information is needed to support the level of service that such networks are advocating.

The Web-based workstation described above is rapidly proving to be a useful way to provide very wide distribution of images to clinical users. Experience with the Internet has allowed intranets to be developed. Essentially, these are local (or even wide) area networks that are either private or that have secure connections between them and the Internet. Chapter 14 of this book discusses networks to a greater extent.

For historical interest, the use of personal computers with a browser to access medical information has been described by several authors. Movement of the Internet Web search service software to intranets means that the capability to search even very large databases will be available for intranets.

Workstations also permit on a desktop what used to require a main-frame computer: reconstruction of two-dimensional (2-D) slice data into 3-D images that may be manipulated in real-time. This is, in itself, a large topic deserving of much greater explanation than can be given here. Computer graphics methods add yet another dimension to abilities beyond what film can accomplish. The reference by Vannier provides an overview of the subject. A current trend is to support these complex graphics-intensive operations at the same workstation for which conventional projection radiographic as well as 2-D cross-sectional imaging is interpreted. The increasing power of the personal computers and graphics display cards used in workstations is making it possible to incorporate into the workstation software that previously needed special high-performance graphics hardware. The client-server approach is also being used, with the very high performance of a centralized graphics processor being made available with a thin client application on the workstation (e.g., TeraRecon, Inc, San Mateo, Calif).

It is important to know about another development in workstation software. Though a small effort in comparison to commercial vendors' software, there are a number of software packages available for free that perform many of the functions of commercial software. Some of these are also "open source" developments, and those who wrote the software are making the source code (the programming of the software, not just the "runtime" part) so that others can add to the effort. The operating system Linux is an example of an open source software project that has been very successful. The topic of open source software development is a very extensive one and is well beyond the scope of this chapter, but the future importance for medical imaging is unmistakable.

The author has collected a number of references specifically to free software for viewing DICOM images. He has personally used a program called IrfanView for converting various image formats, ezDICOM as a PC-based viewer that has many workstation features, and Osirix, which is a very powerful Macintosh-based image viewer. The reader is encouraged to do Internet searches for "DICOM viewers" as a way to find the latest such offerings.

It is likely that displays will continue to evolve. Plasma and LCD displays are now prevalent in consumer televisions. Some of these are physically large enough and have the ability to display pixel matrices sufficient for viewing medical images from a distance and so may find use in conference and operating rooms. Other flat-panel and novel display technologies, notably the organic light emitting diode (OLED) and micromirror (also known as the digital light processor, DLP [DLP, Texas Instruments Inc, Plano, Tex]) are likely to find application in medical imaging. The micromirror display is already widely used in the digital projectors and some large-screen projection television sets. The micromirror is a set of microscopic mirrors and drive electronics fabricated on a silicon wafer using the same techniques used to make other integrated circuit chips. Unlike flat panels, this display uses the mirrors to modulate a beam of light reflected from their surface. The OLED potentially allows for very thin, "flexible" displays. These could be applied to a wide variety of surfaces or built into monitor cases. Unlike LCDs, the OLED is a light-emitting device, so it does not need backlights. Potential production costs are low compared to LCDs. New display technology can mean new applications for workstations as they can be used in locations and for purposes not possible with other workstation displays.

## CONCLUSION

Workstations are replacing or will replace conventional film for the diagnostic and review tasks for medical imaging. For this to happen on a large scale, the barriers to acceptance (lack of adaptability, automation, and integration) will have to be overcome. This, in turn, requires that the users of workstations expend more effort in understanding their work in detail and with a knowledge of the human factors involved. Manufacturers need to use this understanding and knowledge in the development of the next generation of workstations. Department administration and the architects who plan departments should learn, preferably by firsthand observation, that the environments into which workstations will be placed are critical to their use and acceptance.

Those who are impatient for workstations to replace film-based reading should recall that radiologists have a hundred-year experience with film. There are those who pessimistically note that the “paperless office” and “bookless library” have not come to pass despite the availability of the technology to make these possible should also note that where electronic methods are superior to paper and books and where the applications developed are simpler to use, computer-based technology has prevailed. The examples of electronic mail and computerized literature searches readily come to mind. This author would agree, though, that more widespread acceptance of workstations is dependent on their further evaluation and refinement and that it is the efforts of those who have pioneered electronic imaging in medicine that are enabling this progress.

## REFERENCES

- Alter AJ, Kargas GA, Kargas SA, Cameron JR, McDermott JC. The influence of ambient and viewbox light upon visual detection of low-contrast targets in a radiograph. *Invest Radiol.* 1982;17:402-406.
- American College of Radiology. *ACR Standard for Teleradiology.* Reston, VA: American College of Radiology; 1994.
- Arenson RL, Chakraborty DP, Seshadri SB, Kundel HL. The digital imaging workstation. *Radiology.* 1990;176:303-315.
- Asher RW, Martin H. Cathode ray devices. In: Luxenberg HR, Kuehn RL, eds. *Display Systems Engineering.* New York: McGraw-Hill; 1968:237-276.
- Augarten S. *Bit by Bit: An Illustrated History of Computers.* New York: Ticknor and Fields; 1984:263-281.

- Baxter B, Ravindra H, Norman RA. Changes in lesion detectability caused by light adaptation in retinal photoreceptors. *Invest Radiol.* 1982;17:394.
- Beard D, Pizer S, Rogers D, et al. A prototype single-screen PACs console development using human-computer interaction techniques. *Proc SPIE.* 1987;767:646-653.
- Beard DV. Designing a radiology workstation: a focus on navigation during the interpretation task. *J Digit Imaging.* 1990;3(3):152-163.
- Beard DV, Walker JGF, Bell I, Cromartie R. Evolved design of a radiology workstation using time-motion analysis and the keystroke model. *SPIE Phys Med Imaging.* 1989;1091:121-131.
- Beard DV, Hemminger BM, Perry JR, et al. Interpretation of CT studies: single-screen workstation versus film alternator. *Radiology.* 1993;187(2):565-569.
- Beard DV, Pisano ED, Denelsbeck KM, Johnston RE. Eye movement during computed tomography interpretation: eyetracker results and image display-time implications. *J Digit Imaging.* 1994;7(4):189-192.
- Bell CG. Toward a history of (personal) workstations. In: Goldberg, A, ed. *A History of Personal Workstations.* New York: ACM Press; 1988:4-36.
- Chang PJ, Hoffman E. Multimodality workstation featuring multiband cine mode and real-time distributed interactive consultation. RSNA; 1993:InfoRAD Exhibit 9507WS.
- DeJesus EX. The searchable kingdom. *Byte.* 1996;22(6):92NA1-92NA12.
- Derbyshire K. Beyond AMLCDs: Field emission displays? *Electronics Design.* October 1994:56-66.
- Duerinckx AJ, ed. *Picture Archiving and Communications Systems for Medical Applications, Parts I and II.* Bellingham, WA: SPIE; 1982. *Proc. SPIE*; vol 318:
- Dwyer SJ III, Templeton AW, Anderson WH, et al. Salient characteristics of a distributed diagnostic imaging management system for a radiology department. *SPIE Phys Med Imaging.* 1982;318:194-204.
- el-Saden SM, Hademenos GJ, Zhu W. Assessment of intraaxial and extraaxial brain lesions with digitized computed tomographic images versus film: ROC analysis. *Acad Radiol.* 1997;4(2):90-95.
- Engelbart D. The augmented knowledge workshop. In: Goldberg A, ed. *A History of Personal Workstations.* New York: ACM Press; 1988:187-232.
- Feingold E, Seshadri SB, Arenson RL. Folder management on a multimodality PACS display station. *SPIE Phys Med Imaging.* 1991;1446:211-216.
- Feingold ER, Grevera GJ, Mezrich RS, et al. Web-based radiology applications for clinicians and radiologists. *SPIE Phys Med Imaging.* 1997;3035:60-71.
- Fujino T. *Simulation and Computer Aided Surgery.* Chichester, UK: John Wiley and Sons; 1994.
- Greenes RA. Toward more effective radiologic workstation consultation: Design of a desktop workstation to aid in the selection and interpretation of diagnostic



- procedures. In: *Proceedings of the Eighth Conference of Computer Applications in Radiology*. American College of Radiology; 1984:554–561.
- Greenes RA, Bauman RA. The era of healthcare reform and the information super-highway. *Radiol Clin North Am*. 1996;34(3):463–468.
- Goldberg A, ed. *A History of Personal Workstations*. New York: ACM Press; 1988.
- Gray HT, Sune CT, Jones GW. Silicon field emitter arrays for cathodoluminescent flat-panel displays. *J Soc Info Display*. 1993;1(2):143–146.
- Hafner K, Lyon M. *Where Wizards Stay Up Late: The Origins of the Internet*. New York: Simon and Schuster; 1996.
- Haskin ME, Haskin PH, Laffey PA, et al. Data versus information: which should we exchange? *SPIE Phys Med Imaging*. 1985;536:37–42.
- Hohman SA, Johnson SJ, Valentino DJ. Radiologists' requirements for primary diagnosis workstations: Preliminary results of task-based design surveys. *SPIE Phys Med Imaging*. 1994;2165:2–7.
- Honeyman JC, Arenson RL, Frost MM, et al. Functional requirements for diagnostic workstations. *SPIE Phys Med Imaging*. 1993;1899:103–109.
- Horii SC. Electronic imaging workstations: ergonomic issues and the user interface. *Radio-Graphics*. 1992;12:773–787.
- Horii SC, Horii HN, Mun SK, Benson HR, Zeman RK. Environmental designs for reading from imaging workstations: ergonomic and architectural features. *J Digit Imaging*. 1989;2(3):156–162.
- Horii SC, Kishore S, Stevens JF, et al. Overlapped image display method: a technique for comparing medical images on a workstation. *SPIE Phys Med Imaging*. 1994;2164:456–466.
- Horii SC, Feingold ER, Kundel HL, et al. PACS workstation usage differences between radiologists and MICU physicians. *SPIE Phys Med Imaging*. 1996;2711:266–271.
- Huo Z, Giger ML, Vyborny CJ, Wolverton DE, Doi K, Schmidt RA. Computerized classification of malignant and benign masses on digitized mammograms. *Acad Radiol*. 1998;5(3):155–168.
- Jost RG, et al. An electronic multiviewer. In: *Proceedings of the Eighth Conference of Computer Applications in Radiology*. American College of Radiology; 1984:304–311.
- Kano A, Doi K, MacMahon H, Hassell DD, Giger ML. Digital image subtraction of temporally sequential chest images for detection of interval change. *Med Phys*. 1994;21(3):453–461.
- Krupinski EA, Lund PJ. Comparison of film vs. monitor viewing of CR films using eye-position recording. *Proc SCAR '96*. 1996:269–274.
- Kundel HL. How much spatial resolution is enough? A meta-analysis of observer performance studies comparing plain films and digital hard copy. *SPIE Phys Med Imaging*. 1993;1899:86–89.

- Lampson BW. Personal distributed computing: the Alto and Ethernet software. In: Goldberg A, ed. *A History of Personal Workstations*. New York: ACM Press; 1988:293–335.
- Larsen GN. *Interactive Image Processing for Computerized Tomography* [PhD thesis]. Columbia: Department of Electronics and Electrical Engineering, University of Missouri at Columbia; 1976.
- Lawrence A. Acoustic design. In: Ruck NC, ed. *Building Design and Human Performance*. New York: Van Nostrand Reinhold; 1989:117.
- Lemke HU, Stiehl HS, Scharnweber N, Jackel D. Applications of picture processing, image analysis, and computer graphics techniques to cranial CT scans. In: *Proceedings of the Sixth Conference on Computer Applications in Radiology and Computer/Aided Analysis of Radiological Images*. IEEE Computer Society Press; 1979:341–354. Reprinted in: *J Dig Imaging* 2003;16(1):11–28.
- Leung KT, Ho KBT, Chao W, et al. Image navigation for PACS workstations. *SPIE Phys Med Imaging*. 1995;2435:43–49.
- Li X, Valentino DJ, So GJ, et al. World Wide Web telemedicine system. *SPIE Phys Med Imaging*. 1995;2711:427–439.
- Licklider JCR. Some reflections on early history. In: Goldberg A, ed. *A History of Personal Workstations*. New York: ACM Press; 1988:117–130.
- Maguire GQ Jr, Zeleznik MP, Horii SC, et al. Image processing requirements in hospitals and an integrated systems approach. *SPIE Phys Med Imaging*. 1982;318:206–213.
- Mann S. Cyborg seeks community. *Technol Rev*. 1999;102(3):36–42.
- Mascarini C, Ratib O, Trayser G, Ligier Y, Appel RD. In-house access to PACS images and related data through World Wide Web. *SPIE Phys Med Imaging*. 1995;2711:531–537.
- MDIS Technical Development Team. *Medical Diagnostic Imaging Support System: Acquisition Document*. Huntsville, AL: US Army Engineer Division; 1991.
- Motamedi ME, Wu MC, Pister KSJ. Micro-opto-electro-mechanical devices and on-chip optical processing. *Opt Eng*. 1997;36(5):1282–1297.
- Perry J, Prior F. Purchasing a PACS: an RFP toolkit. In: Siegel EL, Kolodner RM, eds. *Filmless Radiology*. New York: Springer-Verlag; 1999:33–84.
- Ravindra H, Norman RA, Baxter B. The effect of extraneous light on lesion detectability: a demonstration. *Invest Radiol*. 1983;18:105.
- Robb RA. Volume visualization and virtual reality in medicine and biology. In: Lemke HU, Vannier MW, Inamura K, Farman AG, eds. *Proceedings of CAR '98: Computer Assisted Radiology and Surgery*. Amsterdam: Elsevier; 1998:131–142.
- Rogers DC, Johnston RE. Effect of ambient light on electronically displayed medical images as measured by luminance-discrimination thresholds. *J Opt Soc Am*. 1984;A4:976.

- Ross DT. A personal view of the personal work station. In: Goldberg A, ed. *A History of Personal Workstations*. New York: ACM Press; 1988:54–111.
- Rossman K. Image quality. *Radiol Clin North Am*. 1969;7(3):419–433.
- Rostenberg B. *The Architecture of Imaging*. Chicago: American Hospital Publishing; 1995.
- Sampsell JB. An overview of the digital micromirror device and its application to projection displays. *SID International Symposium Digest of Technical Papers*. 1993;24:1012.
- Say DL, Hedler RA, Maninger LL, et al. Monochrome and color image-display devices. In: Benson KB, ed. *Television Engineering Handbook*. New York: McGraw-Hill; 1986:12.1–12.53.
- Sezan MI, Yip KL, Daly SJ. An investigation of the effects of uniform perceptual quantization in the context of digital radiography. *SPIE Phys Med Imaging*. 1987;767:622–630.
- Scott WW Jr, Bluemke DA, Mysko WK, et al. Interpretation of Emergency Department radiographs by radiologists and emergency medicine physicians: teleradiology workstation versus radiograph readings. *Radiology*. 1995;195:223–229.
- Sherr S. *Electronic Displays*. 2nd ed. New York: John Wiley and Sons; 1993:201–340.
- Shyankay J. Diamond films used in flat panel displays. *R&D Magazine*. April 1995: 44.
- Slasky BS, Gur D, Good WF, et al. Receiver operating characteristic analysis of chest image interpretation with conventional, laser printed, and high-resolution workstation images. *Radiology*. 1990;174:775–780.
- Spindt CA, Brodie I, Hamphrey L, Westerberg ER. Physical properties of thin-film field emission cathodes with molybdenum cones. *J Appl Phys*. 1976;47(12): 5248–5251.
- Taira RK, Wong C, Johnson D, et al. Design of a graphical user interface for an intelligent multimedia information system for radiology research. *SPIE Phys Med Imaging*. 1995;2435:11–23.
- Taylor JH. Vision. In: Parker JF, West VR, eds. *Bioastronautics Data Book*. 2nd ed. Washington, DC: NASA; 1973:611–665.
- ter Haar Romeny BM, Raymakers J, van Waes PFGM, et al. The Dutch PACS project: Philosophy, design of a digital reading room and first observations in the Utrecht University Hospital in the Netherlands. *SPIE Phys Med Imaging*. 1987;767:787–792.
- Thacker CP. Personal distributed computing: the Alto and Ethernet hardware. In: Goldberg A, ed. *A History of Personal Workstations*. New York: ACM Press; 1988:267–289.
- van der Voorde F, Arenson R, Kundel H, et al. Development of a physician-friendly digital image display console. *SPIE Phys Med Imaging*. 1986;626:541–548.

- Vannier MW, Marsh JL. Three-dimensional imaging, surgical planning, and image-guided therapy. *Radiol Clin North Am.* 1996;34(3):545-563.
- Vining DJ. Virtual endoscopy: is it reality? *Radiology.* 1996;200:30-31.
- Xu XW, Doi K, Kobayashi T, MacMahon H, Giger ML. Development of an improved CAD scheme for automated detection of lung nodules in digital chest images. *Med Phys.* 1997;24(9):395-403.

# BREAST IMAGING, COMPUTER-AIDED DETECTION, AND COMPUTER- ASSISTED CLASSIFICATION

---

---

SHALOM S. BUCHBINDER

**M**ammography, the primary screening and diagnostic tool for the detection and diagnosis of breast cancer, in many ways represents the extremes on the timeline in the shift of radiology toward full digital imaging. Mammography was the first modality to be digitally converted and to apply

computer-associated artificial intelligence to diagnosis. However, it was also the last major modality to have a commercially marketed digital system. Furthermore, in many departments mammography is the last division converted to digital format.

The delay in the development and deployment of full field digital mammography (FFDM) primarily relates to the exacting spatial resolution required in order to detect and correctly analyze microcalcifications. Many malignant microcalcifications may be as small as 0.3 mm in size, 1 order of magnitude smaller than most radiographically identified abnormalities. Controversy still exists regarding the optimum, cost-effective acquired pixel size for the appropriate visualization and analysis of microcalcifications. For example, a 50 or 100  $\mu\text{m}$  pixel pitch might enable the identification of small clusters of microcalcifications. However, a 25  $\mu\text{m}$  pixel pitch may enable a more rigorous analysis of the morphology of the individual microcalcifications, thus improving specificity by allowing better discrimination between benign and malignant microcalcifications. However, with each halving of acquired pixel pitch, there is a quadrupling of the number of pixels, which substantially increases the cost of storage and time required for processing and network transmission of the data.

A second major challenge faced by digital mammography is the need for superb contrast over a wide dynamic range. The breast contains benign, dense fibroglandular tissue and lucent fatty tissue as well as masses and/or microcalcifications that may be of benign or malignant origin. Uniformly dense mammograms and even mammograms with a combination of different densities can make discriminating subtle areas of malignancy most challenging. Improving visualization requires improved contrast resolution, which entails optimizing the bit depth of each pixel since contrast resolution is a function of the bit depth. Improved contrast resolution must also be optimized from a cost-benefit perspective similar to that for spatial resolution. The requirements for a wide dynamic range, with optimal spatial and contrast resolution, have been met through extensive research and confirmed with the approval process's mandatory equivalency testing, allowing for the ultimate introduction of digital mammography into the marketplace. Because of these technological advances, new applications are being developed for FFDM. For example, tomosynthesis shows promise as a technique to improve visualization of breast cancer.

In addition to the acquisition of the digital mammogram, downstream requirements for processing (as well as computer-aided detection [CAD] and potentially computer-assisted classification [CAC]) had to be developed. The rigorous requirements related to the display for softcopy interpretation, which is the preferred manner for digital interpretation, and production of

high-resolution hard copies (as needed) required extensive evaluation and approval. The data storage requirements also had to be met in a cost-effective manner. FFDM storage requirements, by virtue of the optimum contrast and spatial resolution needed, are particularly large and hence costly. Facilities with robust picture archiving and communication systems (PACS) usually have a cost-effective archive. However, smaller radiology practices or freestanding mammography boutiques may need to employ a specialized mammography PACS or utilize offsite digital archiving, which adds significantly to the cost of FFDM.

Due to these requirements, the overall cost for complete FFDM systems is substantially higher than for conventional mammography. This cost has been partially mitigated by increased patient throughput available with FFDM, which is in the range of 150% to 200% of standard film-screen examinations. In addition, increased reimbursement for each case is allowed by most health insurers, which is frequently 150% of the standard rate.

Now that FFDM has become a reality, the major benefits of the digital image—CAD, CAC, and PACS—can be exploited. The digital image may allow improved accuracy with the utilization of the CAD and CAC software. Integration into a PACS allows optimized distribution and storage. However, as these benefits are fundamentally related to the digital image, similar benefits may potentially be obtained from computed radiography (CR) technology once its ability to provide diagnostic accuracy is proven to approximate conventional film-screen images and FFDM. Computed radiography mammography, when introduced, potentially has the benefit of reduced costs, which in the competitive mammography marketplace may exert downward pressure on the cost of FFDM.

## **COMPUTER-AIDED DETECTION AND COMPUTER-ASSISTED CLASSIFICATION**

Early and accurate detection of breast cancer is important for optimizing therapy and affording the patient a better prognosis. Mammography, in conjunction with physical examination, is the most effective tool available today for mass screening for the early detection of breast cancer. It can potentially reduce mortality by as much as 40% to 63%. Various studies have shown that a significant number of biopsy-proven malignancies are not detected by mammography. Mammography has particular difficulty in detecting breast cancers in women with dense breast tissue that may obscure the lesion or make it more difficult to identify even if visible. Subtle areas of architectural distortion can also be challenging. Other false negatives may relate to visible

cancers being misinterpreted or overlooked because of inexperience, distraction, or fatigue caused by the large number of examinations presented to the limited number of mammographers or by suboptimal film quality. Therefore, methods that can overcome these difficulties in interpreting mammograms are of major importance. It has been shown that breast cancer detection rates can be substantially improved by using second readers. However, in many instances it is not feasible to perform a second read by human mammographers. Potentially, even after a second reader's interpretation, significant numbers of missed cancers might still exist and may be detected by the addition of computer software.

For this reason CAD appears to be one of the most important or, as some suggest, the most important collateral benefit from the availability of digital mammography, since it allows cost-effective and efficient acquisition of the digital image to computer software that can then serve as the second reviewer (or even the third reviewer after the double reading by the human mammographers) to optimize the detection of abnormalities. CAD can overcome much of the difficulty and expense inherent in a reading by a second mammographer and may conceptually be even more accurate than some mammographers in detecting abnormalities.

These detection systems are popularly known by the acronym CAD, which has been used to represent varied abbreviations including computer-aided detection, computer-aided diagnosis, computer-assisted detection, and computer-assisted diagnosis. In this chapter, CAD will be used to represent computer software that detects abnormalities. In addition, the data can also be analyzed by specialized software that can help classify the likelihood of malignancy of any identified abnormalities. These systems will be called CAC in this chapter, representing computer-assisted classification.

This idea of using computers in facilitating the mammography examination is not a new one. In August of 1967, Winsberg, Elkin, et al., from the Albert Einstein College of Medicine of Yeshiva University, authored an article in *Radiology* titled "Detection of Radiographic Abnormalities in Mammograms by Means of Optical Scanning and Computer Analysis," in which they noted, "Because of the problems inherent in the routine viewing of large numbers of examinations of presumably asymptomatic patients we have proposed the automation of reading of the radiographs by means of optical scanning and computer interpretation." It has taken over 35 years to even approximate this idealized vision of mammography.

After the digital image is acquired, software algorithms are employed to meticulously review the data and search for areas that appear to be different from the surrounding tissue. These algorithms can identify microcalcifications and masses. However, the relatively rare but visually challenging



areas of architectural distortion are not as well identified by CAD, requiring more research. Integration of data from both breasts, which would allow identification of potential asymmetry, is also presently under development. Integration of prior mammograms may also be useful in identifying developing densities or other interval changes.

The sensitivity of detection of abnormalities, the prime role of CAD, is related to a great extent to the specificity of the system. For example, using just the number of microcalcifications as a discretionary feature, a system set for a minimum threshold of 5 microcalcifications per area may discount early malignancies that manifest themselves with only 4 microcalcifications. However a threshold level set at 3 microcalcifications would identify those malignant clusters as well as any other clusters of 4 microcalcifications. Thus, to maximize true positive identifications (reduce false negatives) it is necessary to allow for more false positives, although too many false positives may be disturbing and potentially discredit the CAD system for some users. Potentially, practices whose referral base has little tolerance for a low sensitivity and will readily accept lower specificity can have their needs met by increasing the false positives allowed. A “dial up or down” feature to allow for practices or mammographers to set their own sensitivity/specificity level would be a desirable feature.

The utilization of a CAD system requires a digital image. This digital image can be derived from 3 different mechanisms: (1) manual digitization of a conventional film-screen mammogram, which is labor intensive, (2) a FFDM machine, or (3) CR (potentially). Conceptually, the end result is a digital image representing the underlying breast tissue. However, there are inherent problems associated with digitization of conventional film-screen mammograms. This is because of the lack of reproducibility of the digitized image. The use of calibration strips can allow for an approximation, but each time an analog image is digitized, the results are slightly different, and this could affect the results of the CAD analysis. A slightly different input due to fluctuations in the digitization process could yield a very different output, which could create medicolegal difficulties. Thus, it is simpler to deal with images that are inherently digital instead of trying to digitize analog ones. Inherently digital images can also be sent more efficiently through a CAD system because they require no digitization step (although no definitive studies have yet been published on this). For this reason, the CAD-marked images should not have to be archived, as long as the version of the CAD system used is documented and can be reproduced as needed in the future in conjunction with the original digital mammography data.

Because there are 2 different primary signs of breast cancer, microcalcifications and masses, it is necessary to have 2 different detection algorithms.

The microcalcification algorithm identifies and evaluates pixels that may represent calcium. The individual microcalcifications can be outlined, their size quantified, and the features of each analyzed. The cluster as a whole is also analyzed. The minimum number of adjacent calcifications within a prescribed area may be calculated and compared to the software's minimum threshold level deemed appropriate for analysis. Mathematical models are then used to evaluate the combined data and determine which cluster of microcalcifications should be presented to the mammographer for further evaluation. With respect to the evaluation of masses, the software can search for an area with different characteristics from the background. It can refine its analysis with evaluation of the mass's density and its border characteristics, such as spiculation.

CAD has been successful in detecting approximately 80% to 90% of all malignancies. However, the systems at present appear to be more sensitive in detecting clusters of microcalcifications than masses, with microcalcifications detected in the upper 90% range and masses detected in the 80% to 90% range. The less frequently occurring architectural distortion is relatively poorly identified. As can be expected, CAD will benefit some radiologists more than others. Experience appears to be inversely proportional to the amount of improvement seen with CAD.

Computer-aided detection is primarily a sensitivity-enhancing tool, with potentially deleterious effects not only to specificity, but also to recall rates and positive predictable value. Recall rates appear to have a modest deterioration that appears because of the sensitivity improvement. Similarly, there have been only minimal effects on positive predictive value (PPV). This is because the additional false positive biopsies recommended are offset by the additional true positives. When the two are combined, the PPV approaches the original PPV, but with more cancers detected. Therefore, it is better if abnormalities detected by either the mammographer or CAD can be directed to a CAC system. Such a classification system can then evaluate and exclude many questionable areas, thereby making the correct diagnosis as accurate, sensitive, and specific as possible while maintaining reasonable recall rates. Similar to CAD, these classification systems also need to have different algorithms to classify microcalcifications and masses, but they do not use the same algorithms as those used for detection. Once detected, each microcalcification is characterized according to several properties, including brightness, area, length, and shape. After these quantitative parameters are obtained, the microcalcification cluster can be classified into each of 3 major subdivisions: first is a group that characterizes variables in morphology of the individual calcifications; second is a group that reflects the heterogeneity of the morphologic features related to the individual calcifications; third

is a group that characterizes variables in the distribution of the calcifications within a cluster. The spatial orientation appears to be most helpful in distinguishing benign from malignant clusters. Artificial intelligence then can analyze all the information and compare it to known statistics representing frequency of malignancy. The system then can give a final assessment as to the likelihood of an abnormality being malignant (see Figure 21.1). This analysis cannot be performed by the mammographer because many of the underlying features analyzed are too small (malignant microcalcifications are in the 0.3 mm range) and the calculations too complex, making the analysis beyond human capability. Algorithms for evaluating masses include analysis of features that characterize shape, definition of the margins, and various aspects of spiculation (see Figure 21.2). The information derived can be used to approximate the likelihood of malignancy (high, intermediate, or low) of any abnormalities identified.

Prototype classification schemes that give gross estimates of malignancy are currently being evaluated and show promise in initially being used


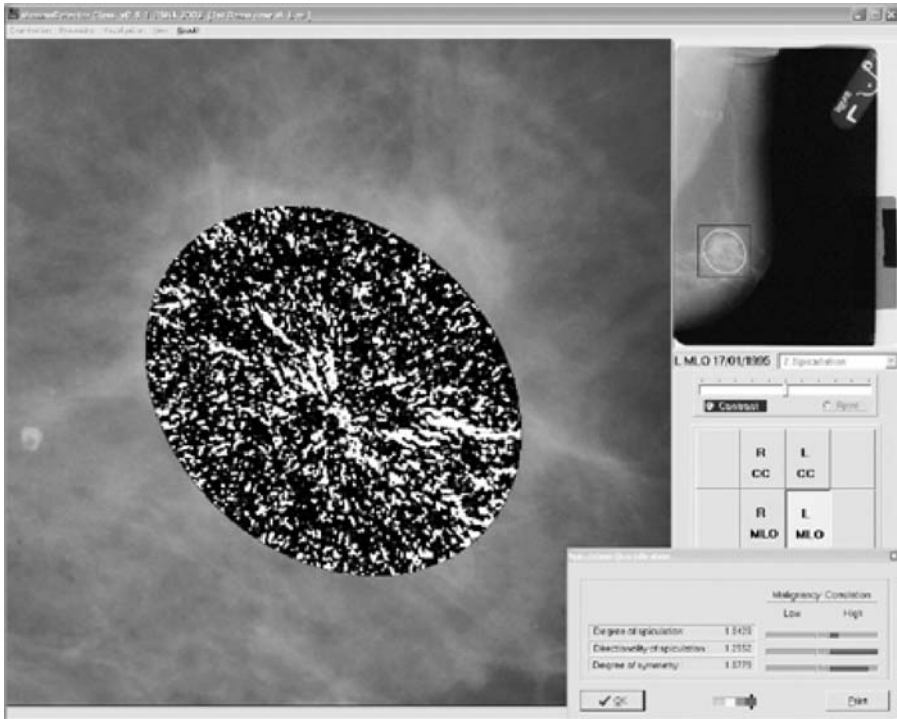


FIGURE 21.1

Computer analysis of a cluster of microcalcifications. (Courtesy of Professor Isaac Leichter, Siemens Computer Aided Diagnosis Ltd.)



**FIGURE 21.2**

Computer analysis of a breast mass. (Courtesy of Professor Isaac Leichter, Siemens Computer Aided Diagnosis Ltd.)

by physicians to potentially (and at their discretion) modify Breast Imaging Reporting and Data System (BIRADS) assessments. Systems of this nature may be particularly helpful in evaluating abnormalities that have a preliminary diagnosis of probably benign. With more refinements, the CAC systems may ultimately be able to give a probability range of an abnormality's potential to be malignant, for example, 20% to 25% chance. The referring physician and patient can then incorporate these data to more accurately assess risk-benefit relationships of the various options for management of the identified abnormality. In addition, combination systems with CAD and CAC, in tandem, are also being evaluated and appear to be on the horizon to aid the mammographer's interpretation.

Ultimately, robust algorithms combining CAD and CAC systems may be employed to prescreen the large number of screening mammograms from the patients recommended to have this examination. Only selected cases that

fall outside predetermined threshold levels for probability of disease will then need to be evaluated by the mammographer for final assessment. Similar computer methods have been applied medically for cervical cancer screening and also in nonmedical commercial endeavors. Despite the significant research that has already been undertaken, particularly over the past decade, additional research is required to fulfill the promise envisioned by Winsberg, Elkin, et al., in 1967 to make mammography as sensitive, specific, readily available, and cost effective as possible.

## PACS AND BREAST IMAGING

The integration of breast imaging into PACS is associated with several unique challenges with respect to softcopy image interpretation, communication for consultation, and archival storage. When an organization commits to FFDM, it must decide how the images will be interpreted. The preferred method of interpretation is via softcopy display, although some practices may still prefer printed film displayed on a view box. Softcopy reading requires high-resolution monitors to enable the physician to visualize the minute details for accurate diagnosis. Printed film interpretation of images derived from FFDM requires a specialized high-resolution printing capability.

If the decision is to use the digital image for only softcopy interpretation, no hardcopy production should be necessary unless specifically required for distribution outside the institution. Using the PAC system's existing archive appears intuitively to be the most appropriate method. However, during a transition phase (which in mammography usually takes up to 3 years), previous images are needed for comparison in order to optimize diagnosis. The comparison with the previous examinations enables appropriate diagnoses of any new masses or microcalcifications or allows for the identification of any developing densities. Therefore, in the interim, adjacent to softcopy monitors used for FFDM acquired images, a view box for the conventional historical image review is necessary. An alternative is the digitization of the historic films with subsequent archiving into the PACS system and visualization as necessary on a softcopy monitor for direct comparison to the FFDM image. This process, in addition to being costly and tedious, has the inherent limitation of potential loss of information related to the digitization process. Therefore, most institutions have chosen to have a dual-method display: softcopy for the current FFDM images and films on a view box for the historical image comparison.

With respect to distribution of images, 3 separate situations must be considered. Many facilities run screening programs. Approximately 5% to

15% of patients who are seen at a screening facility will require immediate additional imaging, usually performed within 2 weeks at a diagnostic facility. This will require the screening images to be available in order to appropriately work up the abnormality. Other patients who have had FFDM performed will need to have images distributed to referring physicians, such as breast surgeons or other consultants (and potentially to the operating room during a surgical procedure). There are also images that the patient may wish to distribute for a second opinion to physicians not affiliated with the original practice. With respect to patients who have had screening images, the site where the screening images were obtained may either be FFDM or conventional film-screen, and the site where the patient will undergo the diagnostic study may likewise have either of the 2 types of imaging equipment. As with the routine comparison with historic films, 1 set of images may be on softcopy display and the second set of images may be on a film view box.

Distribution to referring physicians and/or the operating room within the institution for gross visualization (not primary diagnosis) may be performed over the PACS enterprise distribution system (typically a Web browser-based interface) even though the display systems on most of the clients is not suitable for primary diagnosis. For most exams, this would appear to be adequate particularly for abnormalities such as large masses or clusters. For off-site distribution, facilities must use high-resolution copiers, which need to be specifically approved by the FDA for use with mammography images obtained with FFDM. These images can be shipped as any other images and/or given to the patients for personal distribution. Direct electronic image distribution to off-site locations would be desirable provided that necessary issues, such as network security, bandwidth, and Health Insurance Portability and Accountability Act (HIPAA) regulations, are worked out.

Mammography images, by virtue of their high spatial resolution and need for contrast over a large dynamic range, require relatively large amounts of storage. Currently, mammography images are stored with lossless compression. No studies have been performed yet with respect to storing mammography in a lossy form. This burden on a PACS system is substantial and clearly has to be calculated into the size of storage requirements when an organization with a PACS is contemplating FFDM integration or when an FFDM system is bought for a freestanding organization without a PACS already in place.

It is unclear whether it is really necessary to store the outputs of CAD and CAC. So long as the version of the algorithms used at the time of diagnosis is clearly documented, the results should be completely reproducible.

This is analogous to the use of window and leveling in computed tomography (CT) exams. Even though the radiologist often views CT images under 3 different window/level settings, there is no need to store a separate copy of the images for each window/level setting used. Rather, applying the standard settings to the original pixel data will consistently yield the same results. So, too, knowing which versions of CAD and CAC were used should permit presentation of the marking data at any future time exactly as they were presented to the mammographer at the time of initial interpretation.

In summary, mammography has had a long history with the digital revolution, but it is only now with the advent of FFDM that the benefits of CAD, CAC, and PACS can be realized.

## REFERENCES

- Australian Health Technology Assessment Commission. Review of automated and semi-automated cervical screening devices—a summary. Canberra, Australia: Commonwealth Department of Health and Family Services; 1998.
- Baker JA, Rosen EL, Lo JY, et al. Computer-aided detection (CAD) in screening mammography: sensitivity of commercial CAD systems for detecting architectural distortion. *AJR Am J Roentgenol*. October 2003;181:1083–1088.
- Bamberger P, Novak B, Fields, S. Optimizing parameters for computer aided diagnosis of microcalcifications in mammography. *Acad Radiol*. 2000;7:406–412.
- Beam CA, Sullivan DC, Layde PM. Effect of human variability on independent double reading in screening mammography. *Acad Radiol*. 1996;3:891–897.
- Bird RE, Wallace TW, Yankaskas BC. Analysis of cancers missed at screening mammography. *Radiology*. 1992;184:613–617.
- Birdwell RL, Ikeda DM, O'Shaughnessy KF, Sickles EA. Mammographic characteristics of 115 missed cancers later detected with screening mammography and the potential utility of computer-aided detection. *Radiology*. 2001; 219:192–202.
- Brem RF, Baum J, Lechner M, et al. Improvement in sensitivity of screening mammography with computer-aided detection: a multiinstitutional trial. *AJR Am J Roentgenol*. September 2003;181(3):687–693.
- Buchbinder S, Novak B. Quantitative characterization of mass lesions on digitized mammograms for computer assisted diagnosis. *Invest Radiol*. 2000;35:366–372.
- Buchbinder S, Leichter I, Bamberger P, et al. Analysis of clustered microcalcification using a single numeric classifier extracted from mammographic digital images. *Acad Radiol*. November 1998;5:779–784.
- Buchbinder S, Leichter I, Lederman R, et al. Can the size of microcalcifications predict malignancy of clusters in mammography? *Acad Radiol*. 2002;9:18–25.

- Buchbinder S, Leichter I, Lederman R, et al. Computer-aided classification of BI-RADS category 3 breast lesions. *Radiology*. 2004;230:820–823.
- Burhenne LW, Wood S, D’Orsi C, et al. Potential contribution of computer-aided detection to the sensitivity of screening mammography. *Radiology*. 2000;215:554–562.
- Cady B, Michaelson JS. The life sparing potential of mammographic screening. *Cancer*. 2001;91:1699–1703.
- Ciatto S, Del Turco MR, Morrone D, et al. Independent double reading of screening mammograms. *J Med Screen*. 1995;2:99–101.
- Destounis SV, Dinitto P, Young-Logan W, et al. Can computer-aided detection with double reading of screening mammograms help decrease the false-negative rate? Initial experience. *RSNA*. 2004;578:584–585.
- Duffy SW, Tabar L, Chen HH, et al. The impact of organized mammography service screening on breast carcinoma mortality in seven Swedish counties. *Cancer*. 2002;95:458–469.
- Fields S, Lederman R, Buchbinder S, et al. Improved mammographic accuracy with CAD assisted ranking of lesions. In: Proceedings of the 6th International Workshop on Digital Mammography, June 2002, Bremen, Germany.
- Fields S, Leichter I, Lederman R, et al. Improving the Performance of Mammographic Assessment by the use of an Advanced Two-tiered CAD/CAC System. In: Proceedings of the 7th International Workshop on Digital Mammography, June 2004, Charlotte, NC, USA.
- Freer T, Ulissey M. Screening mammography with computer-aided detection: prospective study of 12,860 patients in a community breast center. *Radiology*. 2001;220:781–786.
- Helvie MA, Hadjiiski L, Makariou E, et al. Sensitivity of Noncommercial Computer-aided Detection System for Mammographic Breast Cancer Detection: Pilot Clinical Trial *RSNA*. 2004;208:214.
- Improving the performance of mammographic assessment by the use of an advanced two-tiered CAD/CAC system. *Digital Mammography*. 2002.
- Kopans D. *Breast Imaging*. 2nd ed. Philadelphia: Lippincott-Raven; 1998.
- Leichter I, Fields S, Nirel R, et al. Improved mammographic interpretations of masses using computer-aided diagnosis: *Eur Radiol*. 2000;10:377–383.
- Leichter I, Lederman R, Buchbinder S. Mammographic lesions: what diagnostic role does the individual mc’s shape play compared to the cluster geometry? *AJR Am J Roentgenol*. 2004;182:705–712.
- Leichter I, Lederman R, Novak B, et al. The use of interactive software program for quantitative characterization of microcalcifications of digitized film-screen mammograms. *Invest Radiol*. June 1999;34:394–400.
- Lederman R, Leichter I, Buchbinder S, et al. Stratification of mammographic computerized analysis by BIRADS categories, *Eur Radiol*. 2003;13:347–353.



- Malich A, Azhari T, Bohm T, et al. Reproducibility: an important factor determining the quality of computer-aided detection (CAD) systems. *Eur Radiol.* 2000;36:170–174.
- Malich A, Marx C, Facius M, et al. Tumor detection rate of a new commercially available computer-aided detection system. *Eur Radiol.* 2001;11:2454–2459.
- Nelson MT, Lechner MC. Comparison of three commercially available FDA approved computer-aided detection (CAD) systems [abstract]. *Radiology.* 2002; 225(P):600.
- Oestmann JW. Digital mammography. In: *Radiological Diagnosis of Breast Disease.* Friedrich M, Sickles EA, eds. Berlin: Springer Verlag; 1997:65–73.
- Petrick N, Sahiner B, Chan HP, et al. Breast cancer detection: evaluation of a mass detection algorithm for computer-aided diagnosis—experience in 263 patients. *Radiology.* 2002;224:217–224.
- Tabar L, Vitak B, Chen HC, et al. Beyond randomized controlled trials; organized mammographic screening substantially reduces breast carcinoma mortality. *Cancer.* 2001;91:1724–1731.
- Thurfjell EL, Lernevall KA, Taube AA. Benefit of independent double reading in a population-based mammography screening program. *Radiology.* 1994;191: 241–244.
- Winsberg F, Elkin M, Macy J, et al. Detection of radiographic abnormalities in mammograms by means of optical scanning and computer analysis. *Radiology.* 1967;89:211–215.
- Zheng B, Hardesty L, Poller WR, et al. Mammography with computer-aided detection: reproducibility assessment—initial experience. *Radiology.* 2003;228:58–62.

# THREE- DIMENSIONAL IMAGING IN RADIOLOGY

---

---

GORDON J. HARRIS

**T**he purpose of this chapter is to explore the growing need for three-dimensional (3-D) imaging services, as well as the challenges and opportunities associated with implementing 3-D imaging in conjunction with a picture archive and communication system (PACS) environment. The advantages of 3-D imaging services are discussed, as well as the workflow involved in integration of 3-D imaging with PACS and radiology information systems (RIS), the staff and equipment needed for providing 3-D services, and issues with proper coding of claims and obtaining reimbursement from insurance.

Radiological images were initially reviewed as transmission x-ray films, and, during the past 30 years, as cross-sectional images with the advent of computed tomography (CT) and magnetic resonance (MR) imaging. During

the first century of radiology (1896 to 1996), all these images were read primarily in the same paradigm, that is, viewed on films printed with a fixed preselected window and level setting and displayed before a lighting source such as a light box. Thus, the image review process for CT and MR imaging for the radiologist was similar in its workflow to that for x-ray films. Primarily during the past decade, PACS softcopy reading has replaced film at many institutions. While these images can be interactively adjusted for window/level, zoom, and so forth, PACS softcopy reading, as with film reading, generally is limited to viewing the transmission or cross-sectional images in their fixed orientations as provided by the original scan. Recent improvements in computing power, networking speeds, and image standards for digital image communications (Digital Imaging and Communications in Medicine [DICOM]) have allowed medical images to be transferred among modalities, image archives, and vendor computer-imaging platforms relatively seamlessly and quickly. These advances, combined with a dramatic improvement in the speed, usability, and affordability of 3-D workstation hardware and software, make it feasible to apply computer image-processing techniques for 3-D imaging and computer-aided diagnosis (CAD) within the routine clinical workflow of a PACS environment. This may involve image post-processing on specialized 3-D imaging workstations now available from a variety of vendors, or 3-D imaging techniques may be integrated and available directly through the PACS workstations themselves.


Three-dimensional medical imaging has revolutionized both radiological diagnosis and surgical planning. Three-dimensional reconstructions provide lifelike views that can quickly summarize the relationship among anatomic structures to facilitate accurate diagnosis or to identify the optimal treatment plan before the patient is in the operating room (OR). The benefits include decreased exploratory time in the OR, less damage to healthy tissues, and a lower risk of complications for the patient, all of which contribute to reduced surgical morbidity. Increased diagnostic sensitivity across all specialties and the likelihood of shorter OR time per procedure contribute to reduce costs. Thus, the quality of care is greatly enhanced, often with a decrease in cost compared to alternative procedures such as catheter angiography. For example, prior to the establishment of the Massachusetts General Hospital (MGH) 3D Imaging Service in February 1999, the standard protocol at MGH and other institutions for assessing laparoscopic living renal donor candidates involved a catheter angiogram to assess the renal vasculature (number, location, and relationships among renal arteries and veins) plus an intravenous pyelogram (IVP) to assess the configuration of the ureters and renal collecting system. These procedures involved exposing the patient (in this case a healthy renal donor candidate) to a significant health risk

involving anesthesia with associated risk of complications, catheterization with associated risk of vessel perforation, as well as a high radiation dose, and a high cost (several thousand dollars for the 2 procedures). Through 3-D imaging post-processing of high-resolution CT scans, the current protocol involves a CT angiography scan (CTA) to assess the renal vasculature in place of the catheter angiogram, plus a delayed-phase CT urogram (CTU), which replaced the IVP. Both the CTA and CTU are acquired in a single outpatient CT contrast study (see Figure 22.1). Thus, a single, contrast-enhanced CTA/CTU outpatient procedure, costing a few hundred dollars with minimal risk and exposure to the patient, has replaced 2 expensive and invasive procedures costing several thousand dollars and exposing an otherwise healthy patient to significant risk. Similarly, diagnostic 3-D CTA imaging has also replaced catheter angiography for assessment of cerebral aneurysms in our institution as well as many other hospitals that perform diagnostic 3-D CTA (see Figure 22.2). Ultimately, 3-D imaging provides these patients with improved care at significantly lower risk and cost and with improved diagnostic clarity than the older diagnostic protocols.



**FIGURE 22.1**

Left: CTA (CT angiography) of the renal arteries. Right: CTU (CT urography) of the ureters and collecting system (kidneys, bladder, etc.). These 2 images were created from the same CT contrast scan session with 1 contrast bolus at different time delays. These 3-D images depict in high-resolution anatomic detail the information required in assessing renal donors for transplants and have replaced more expensive and invasive procedures (catheter angiography plus intravenous pyelography [IVP]).



**FIGURE 22.2**

CTA showing a right middle cerebral artery aneurysm. The right orbit and skull base have been digitally segmented out from the image to allow a view of the aneurysm and feeding artery. CTA can replace catheter angiography for assessing cerebral aneurysms and provide clear views of vascular geometry to allow critical treatment planning decisions.

Multidetector CT (MDCT) has opened new possibilities and demand for 3-D imaging. Rapid acquisition using thin slices and improved rendering algorithms facilitate exquisite 3-D images and reformats. The pace of progress is being hastened by rapid developments in MDCT technology with availability of 4-, 8-, 16-, 32, and 64-slice scanners and more recently volume CT. From a practical standpoint, developments in MDCT scanning allow multiple thin slices to be acquired with increased z-axis coverage resulting in an extraordinary increase in the quantity of acquired data. In parallel with the developments in MDCT, MR imaging has also advanced such that MR studies often involve larger numbers of images, with thinner slices, scanned with multiple series of varying scan sequences and scan orientations, all of which contribute to increased image review workload to the

radiologist. Review of large numbers of images from modern MDCT and MR scanners, often hundreds or even thousands of slices per study, poses significant challenges to radiologists' efficiency and may simply be impractical for most radiology departments. Having 3-D capability available for diagnosis and surgical planning allows the radiologists and referring physicians to first get a summary view of the entire anatomy in a few concise anatomically clear 3-D views and then refer back to the original two-dimensional (2-D) data for comparison and confirmation. Thus, 3-D reconstruction is more frequently becoming a valuable technique to summarize in a concise and clear way the overwhelming number of slices produced by modern CT and MR scanners.

## **CLINICAL THREE-DIMENSIONAL IMAGING SERVICES**

### **THREE-DIMENSIONAL IMAGING TECHNIQUES AND CLINICAL APPLICATIONS**

The variety of tools and techniques comprised by 3-D imaging may involve a single image-processing operation or a combination of image segmentation, filtering, visualization, and feature-extraction algorithms to extract the clinical information of interest from the scan. There are many resources available for review of image-processing techniques. Computed tomography and MR collect a series of 2-D slices through the region of the body under study. These 2-D slices can be "stacked" and combined into a single 3-D volume of data. However, this requires relatively thin slice acquisition for high-quality 3-D imaging. The thinner the slices and the closer to isotropic voxel dimensions, the better the resolution of the reformatted images. For example, most 3-D protocols involve 0.5- to 3.0-mm thick slice data, while thicker slices, such as 5 mm, produce blurred reconstructions and reformats with a stepping artifact orthogonal to the original acquisition plane.

To create post-processed images from the original CT or MR source images, it is necessary to load the slice data into a computer that has 3-D image post-processing software. This may involve loading the scan onto a separate 3-D imaging workstation, running a 3-D software program integrated on the PACS workstation, or loading a client-server-based 3-D software program remotely through a thin-client or Web-based applet. There are several vendors providing various solutions such as these to enable advanced 3-D image post-processing of CT or MR source 2-D datasets.

Once the data are imported into a 3-D imaging system (workstation, PACS plug-in, client-server, etc.), the 3-D volumetric dataset can be resampled in different planes or visualized through various volumetric viewing techniques. The simplest type of 3-D operation is multiplanar reconstruction (MPR), in which the 3 primary orthogonal views are produced. If these views are sampled with a slice thickness greater than the original slice thickness (for example, with volume-rendered orthogonal slabs), this becomes multiplanar volume reconstruction (MPVR). The reconstructed planes can also be rotated and images created in any oblique plane or produced in a curved reformat to create a reconstructed view that follows a curved structure such as a blood vessel, the spinal cord, or the inferior alveolar nerve canal.

The next level of volumetric visualization involves maximum-intensity projection (MIP), average-intensity projection (AveIP), or minimum-intensity projection (MinIP), which display the maximum, average, or minimum value through the volume along each ray drawn from a particular viewpoint. These techniques are good for instances where the interpreting physician does not want the image blocked at the surface, for example, in visualizing vascular stenosis or craniosynostoses. However, MIP views do not provide good depth cues since the brightest voxel can be at any distance from the viewpoint. Therefore, MIP images are often created in a rotational batch series to provide parallax depth cues through rotational orientation changes. Volume and surface rendering provide an anatomical model that allows selected tissues to be either displayed or stripped away to reveal the structures of interest. Surface rendering involves segmenting the image into a binary mask from which a 3-D surface is created. The segmentation can involve various techniques to select the surface to be rendered, such as setting a threshold and/or manually carving portions out from the desired anatomy of interest to be viewed. Because the surface is created based on a binary mask, no voxel intensity information is appreciable in the surface rendering since all surface points are displayed uniformly, with differences resulting only from lighting and depth cues. Volume rendering involves application of opacity maps based on the image histogram values, as well as different grayscale and/or color lookup maps, such that a voxel may be displayed with a variable opacity as well as variable intensity/color depending on the voxel values that are displayed in a continuous range. Surface rendering and volume rendering are useful in viewing bony anatomy, organ surfaces such as brain gyral/sulcal patterns, as well as for visualizing blood vessels for planning treatment approaches for aneurysms. Volume rendering also allows variable viewing of different tissues such that multiple components of the

image may be represented in different brightness, color, or relative transparency, such as bone versus tendons or blood vessels.

The specific type of image reconstruction depends on the anatomy, the purpose (diagnosis or surgical planning), and the customer needs (radiologist or surgeon). In general, diagnostic radiologists often prefer MIP views rather than volume- or surface-rendered views because image data are retained and can be retrieved by window/level adjustments and rotation, while the surface- and volume-rendered views make some objects transparent in a user-dependent manner. Maximum-intensity projection views have the shortcoming, however, that objects within a high-signal object may be obstructed from view, since only the brightest pixel along a ray is displayed. For example, a contrast-enhanced lesion with a necrotic nonenhancing core will have only the surface enhancement displayed since the necrotic core will not be visible through the surrounding enhancement in the MIP view. Average-intensity projection views may allow a more transparent view in such cases and can be used as a simulation of transmission x-ray imaging views. Surgeons, however, often prefer surface-rendered views because they are more anatomically realistic for surgical planning purposes. However, these views may create a threshold artifact that may give a false impression of abnormality, which is why radiologists often feel more comfortable with MIP or AveIP views rather than surface- or volume-rendered views when looking at vascular structures, for example.

### THREE-DIMENSIONAL IMAGING IN A PACS ENVIRONMENT: MASSACHUSETTS GENERAL HOSPITAL EXAMPLE

Currently, the 3D Imaging Service at MGH processes an average of 100 exams per day, which represents approximately 10 percent of the hospital's CT scan volume, plus about 20 percent of the MR and ultrasound (US) volume. The 3D Imaging Service processes exams for vascular, orthopedic, chest, breast MR, gastrointestinal/genitourinary, emergency, and pediatric exams. However, about half of the 3-D exams are neuroradiology studies of the brain, head, and neck. Many neuroradiology exams involve studies of the carotid and/or cerebral vasculature with CT or MR angiography (CTA; MRA). In addition, a variety of neurology, neurosurgery, and oral-maxillofacial exams are reconstructed with 3-D imaging, including brain surface renderings for neurosurgical planning, stroke and aneurysm imaging, brain tumor volumetric measurements, as well as imaging of facial malfor-



mations, craniosynostoses, complex facial fractures/trauma, and surgical planning for dental implants, extractions, and nerve damage evaluations. The 3-D Imaging Service also processes vascular studies of the body, including scans of the aorta, renal, hepatic, and/or pancreatic vasculature and peripheral runoff CTA and MRA exams. High-quality, consistent, reliable 3-D reconstructions are essential to diagnosis and treatment planning of vascular imaging. For example, aortic aneurysm, liver resection, and living liver and renal donor assessments are greatly aided by presurgical 3-D-imaging vascular evaluation and treatment planning on computer rather than exploration in the OR. Three-dimensional imaging is also effective in assessing pancreatic cancer and associated vascular involvement and for virtual endoscopy, such as virtual colonography for assessing colon polyps.

At MGH, we are able to manage this high volume of clinical 3-D image processing as a direct result of the successful implementation of a PACS system with a high-speed image network (1000 base T or gigabit). When the 3-D Imaging Service was established in February 1999, the MGH radiology department was already nearly entirely filmless. All CT and MR images were available both on PACS for radiologists and on a hospital-wide enterprise image server, via a Web-browser plug-in applet, for all referring physicians. This plug-in applet was also integrated with the electronic medical record (EMR) such that all images are available online and are a mouse click away from the patient medical history, clinical test results, and clinical reports.

The successful implementation of the 3-D Imaging Service required that the clinical workflow, protocols, billing, and staffing issues were established in a way that would enable routing clinical 3-D imaging with rapid turnaround, and consistent, high-quality, clinically relevant post-processed images. The 3-D Imaging Service has worked closely on an ongoing basis with hospital billing and compliance, coding, and PACS to establish workflow, and with the radiologists, referring clinicians, and CT and MR scan departments to establish clinical 3-D and scan protocols optimized for clinically useful 3-D imaging.

Within the PACS environment at MGH, the routine 3-D image routing begins with the original 2-D scan being completed at the CT or MR scanner. The exam descriptor is then appended with a 3-D modifier (“/3D”) either by the scheduling department or by the scan technologist to indicate that the study requires 3-D post-processing. The scan technologist sends the scan to a PACS relay, where the scan header information is checked against the patient information in the RIS and corrected for errors if necessary. Once the scan information is validated at the relay, the original 2-D scan is transmitted to a PACS gateway, where it is available for review from PACS terminals throughout the radiology department. The original scan is also

transmitted to the enterprise image server, where the original 2-D scan is made available for review on PCs throughout the hospital via a Web browser-based applet, as well as integrated with the EMR (see Figure 22.3A).

The 3-D exam modifier (“/3D”) informs the PACS routing table to automatically send the scan to the 3-D lab and also sends a message to the 3-D lab pager with the study identifier to indicate that the scan is pending post-processing. If the “/3D” modifier is inadvertently left off the exam descriptor, the 3-D technologists can locate the study through a program that searches the RIS system for exam descriptors that involve 3-D imaging as standard protocol or by manual searching throughout the day of the RIS exam lists. Once a scan has been identified for 3-D post-processing by one of these methods, the 3-D workflow next involves downloading the original

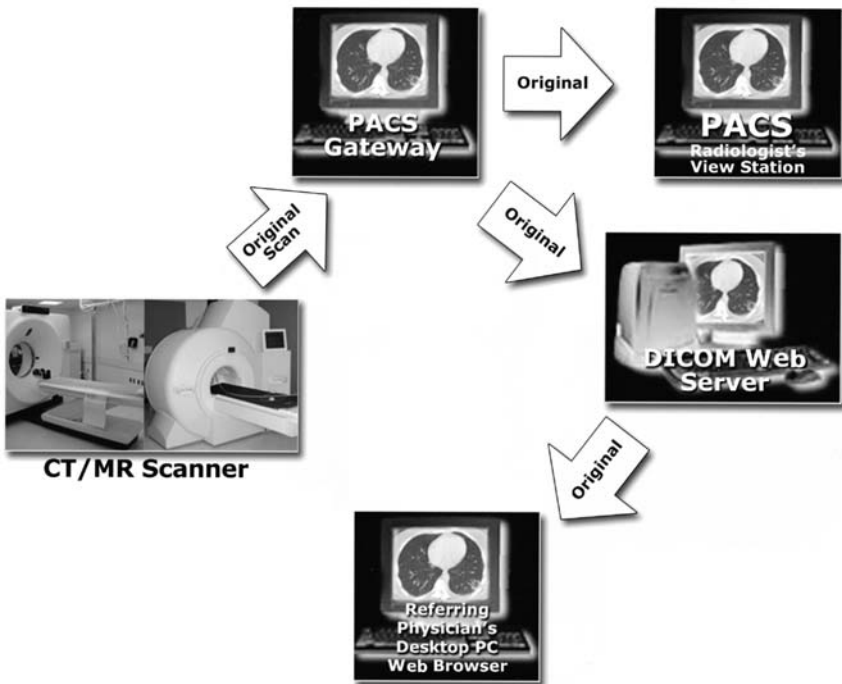


FIGURE 22.3A

Workflow of standard CT or MR 2-D imaging in a PACS environment. The original scan is transmitted from a CT or MR scanner to a PACS gateway, where it may be forwarded to an enterprise-wide image web server. The images are viewable from the radiologist's PACS review station, and also from the physician's desktop PC via web-browser DICOM viewer.

2-D scan from the PACS archive and loading the 3-D volume image data into 3-D imaging software. Once the 3-D images have been post-processed, a summary series of 3-D images is produced, and this series is sent back to the PACS gateway, where the original 2-D study resides. The 3-D technologist completes the 3-D exam code appropriate for the study in the RIS. This 3-D series has a new unique identifier (UID), with the same study number (“accession number”), such that the PACS routing recognizes that this is a new “series” within the same “study.” The PACS appends the new series to the original study, and also forwards the new series to the enter-

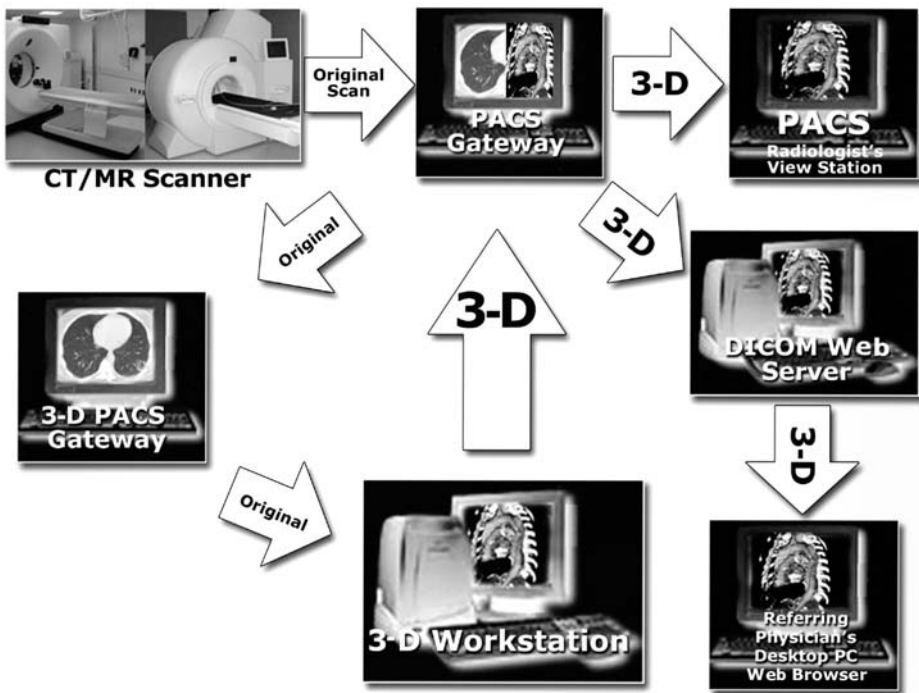


FIGURE 22.3B

Three-dimensional imaging workflow diagram involves the same pathways for the original 2-D CT or MR scan. The original 2-D scan is also routed to a 3-D imaging workstation or software where the 2-D study is converted into a 3-D dataset and a series of 3-D images are created and sent back to the PACS gateway and forwarded to the enterprise-wide image Web server. The 3-D images become a new series attached to the original study and are viewable with the original 2-D images at the radiologist’s PACS review station and at the physician’s desktop PC via Web browser DICOM viewer.

prise image server so that the new 3-D series is available both on the PACS and at the physicians' desktop computer together with the original 2-D scan (see Figure 22.3B). The 3-D processing may take 20 minutes to an hour, and either it is available at the time of the preliminary clinical read, or the 3-D images are reviewed afterward and comments are added to the report as an addendum associated with the primary exam report. It is necessary to report that the 3-D images were reviewed and indicate any additional findings in order to bill for the 3-D exam code when appropriate.

## THREE-DIMENSIONAL IMAGING SERVICE— STAFFING ISSUES

There are 2 major hurdles to developing a consistent, reliable, and cost-effective 3-D imaging service at a facility: (1) determining who will perform the 3-D processing (i.e., appropriate staffing) and (2) capturing sufficient revenue to cover the costs of the service, (i.e., reimbursement). One consideration in determining the cost-effectiveness of the service would be the cost of *not* having a dedicated 3-D imaging service. For example, if a radiologist were to spend a few hours per day performing 3-D processing, he or she would be diverted from reading additional clinical cases, which would have generated added revenue, and the facility might need to hire additional radiologists or moonlighters to read exams. Likewise, if CT or MR technicians were to perform 3-D imaging between scans, it would slow down the scanner throughput, also decreasing revenue. A primary objective in establishing a 3-D imaging service should be to allow the technologists and radiologists to focus on their primary revenue-generation modes for the practice and hospital, namely scanning the patient and acquiring the images, and interpreting the exams, respectively. Meanwhile, a dedicated staff for 3-D imaging services can provide an efficient and consistently processed set of 3-D views. By having a dedicated 3-D technologist processing exams, it is also possible to create greater uniformity and consistency across exams, as opposed to having 3-D images produced by a variety of radiologists or technologists who are not trained to produce a clinically determined set of 3-D imaging protocols. However, smaller facilities may not have the resources or clinical volume to support a 3-D technologist, and thus, 3-D may be performed on a more ad hoc basis by a radiologist or technologist, may be outsourced to an outside 3-D imaging service, or might not be available or performed at all.

Many complex 3-D exams, CTA, for example, require 45 to 60 minutes or more to process all the associated views, while other, less-complicated

3-D exams may take only 15 to 30 minutes each to process. If the clinical volume at a facility is sufficient, it would be cost effective to have one or more dedicated 3-D technologists process these exams (or alternatively, consider outsourcing the 3-D image processing to a centralized 3-D service). A dedicated 3-D technologist should be hired as an experienced, trained, and certified CT or MR technologist who can work closely with the radiologists and referring physicians to create 3-D protocols that generate the most useful views for each type of exam. The dedicated 3-D technologist is akin to a computer graphics medical artist who can take the cross-sectional scans and craft them into a beautiful set of concise, anatomically detailed 3-D views. To accomplish this, the 3-D technologist must understand the cross-sectional anatomy and pathology in depth, as well as scan artifacts, and must learn to operate complex computer software. This is a very complex set of skills and experience, and thus the training period for a dedicated 3-D technologist ranges from 3 to 6 months at our institution.

Many hospitals, including MGH, have established a 3-D imaging service that has become an integral part of the clinical workflow for radiologists and referring clinicians. The 3-D Imaging Service at MGH has been in operation since February 1999, with a staff that has grown to 6 full-time 3-D technologists and 2 image-analysis specialists, a billing coordinator/administrator, a technical director, a systems administrator, and a director. The 3-D Imaging Service at MGH is supported entirely by revenue generated from these services.

## WHEN AND HOW TO CHARGE FOR THREE-DIMENSIONAL IMAGING

### BILLING COMPLIANCE

The second hurdle in justifying a 3-D imaging service is to determine how to capture revenue to which the facility is entitled in order to support the 3-D imaging staff and equipment cost and to demonstrate cost-effectiveness. Carrier reimbursement from Medicare and other health insurers is the main source of revenue to support 3-D imaging, *but it is critical that services only be billed to Medicare and other payers when it is appropriate to do so*. In all circumstances, the 3-D reconstructions must be performed on software that has the required FDA clearances when billing Medicare.

At MGH, Medicare coverage guidelines are used as the benchmark for all carriers. Because payer mix varies widely among hospitals, it is important

to confirm the similarity between commercial/private carriers in a specific geographical region with Medicare. This discussion will focus on Medicare reimbursement as the revenue framework for 3-D imaging.

Medicare reimbursement is based on the medical necessity of the service and procedural code definitions. For example, medical necessity of CTA or MRA can be supported by the need to image the vasculature. This should be reflected in the radiology report by documenting the patient's history and/or indications for the exam as specified by the referral or patient's medical record. Appropriate codes must be used to reflect the services provided.

With the exception of CTA, MRA, and CT colonography, in which the 3-D post-processing is considered part of the primary exam and not billed separately (see following sections for more detail on billing codes), 3-D imaging is a separately reportable and billable service with CT, MR, or US procedural codes (3-D add-on CPT code 76375).<sup>\*</sup> However, for reimbursement, Medicare requires that *both* the initial exam *and* the 3-D imaging are medically necessary.

Medical necessity can be supported by the radiologist's documentation of the patient history or indications for the exam. In addition, the medical necessity of 3-D imaging must be supported in 1 of 3 ways. The first scenario is a direct order from the requesting physician for 3-D imaging in addition to the primary exam request (e.g., a surgeon requests a CT or MR exam with 3-D reconstruction for surgical planning purposes). The second category is the presence of positive findings by the radiologist during an ordered exam requiring additional views for clarification or confirmation of diagnosis; both the positive findings and the reason for the additional view(s) must appear in the radiology report to support the medical necessity of the 3-D imaging. The third approach involves carefully monitored radiology departmental policies, procedures, and protocols. The radiologist writes a protocol policy for an exam to be associated with 3-D imaging that presents medical necessity, reasonableness, and standard of care as supported in the medical literature. The policy is then reviewed for medical necessity, coding accuracy, and billing compliance by the chief of radiology, billing and coding offices (professional and technical), compliance offices (professional and intermediary), and radiology administration. When the policy is approved, a notice is sent to all referring physicians to notify them of the medical necessity and rationale for performing 3-D in association with a particular exam. The referring physicians are also informed that a 3-D charge will be billed to the beneficiary in

---

\* CPT<sup>®</sup> is a registered trademark of the American Medical Association.

addition to the charge for the primary requested exam. These policies are restricted to specific diagnostic situations in which omitting 3-D imaging would lower diagnostic confidence below practice standards and compromise patient care. The notice also needs to provide the requesting physician with *directions* on who to contact with questions or concerns regarding the medical necessity of the 3-D exam. Once the new policy notice has been communicated to requesting physicians, the policy is made available on request in the department of radiology and placed in the radiology policies and protocol manual. To monitor ongoing billing compliance, this type of policy is subject to rigorous interdepartmental review on a yearly basis.

## REIMBURSEMENT FOR THREE-DIMENSIONAL IMAGING

For 3-D reconstruction (other than CTA, MRA, and CT colonography as discussed below), the image post-processing is billed as an add-on to the primary CT, MR, or US exam under CPT code 76375, “coronal, sagittal, multiplanar, oblique, 3-dimensional and/or holographic reconstruction of computerized tomography, magnetic resonance imaging, or other tomographic modality.” This CPT code was originally created for 2-D reformatting using “Coronal, Sagittal, Multiplanar and/or oblique Reconstruction” of CT scans, generally performed on the scanner console at the time of the scan. Then, in 1988, 3-D was added to the code description, and in 1998, the description was expanded to include other tomographic modalities such as MR imaging and 3-D US.

Insurance reimbursement is an important factor in developing a self-supporting clinical 3-D imaging service. For Medicare outpatients, depending on whether the technical costs (3-D staff, 3-D workstations, etc.) are funded by the radiology practice or through the hospital, the technical component of the service would either be billed under the Medicare Hospital Outpatient Prospective Payment System (HOPPS) ambulatory payment classification (APC) for hospital billing, or under the Medicare Physician Fee Schedule (MPFS) under the relative value unit (RVU) system for nonhospital billing. The technical component for hospital-based billing under HOPPS for CPT 76375 is approximately \$92 (2004 national average) and for free-standing clinics or where the 3-D staff and equipment is funded by the professional organization, billing under the MPFS is about \$160. (At MGH, the 3-D Imaging Service was established and funded through the physicians’ organization, with services billed under MPFS, not hospital billing.) The professional component (“-26” billing) is minimal (~\$9).

There is a separate reimbursement code (G0288) for aortic aneurysm stent planning and surveillance. This code was originally instituted in July 2001 specific to a particular centralized service vendor's analysis system. In 2003, the code was revised to be more general. This code has a technical reimbursement only for hospital billing under HOPPS (no professional component), with national average of \$267 in 2005, and a technical reimbursement only for non-hospital billing under MPFS (TC only, no professional component), of \$425 national average reimbursement.

### REIMBURSEMENT CHANGES: COMPUTED TOMOGRAPHY ANGIOGRAPHY, MAGNETIC RESONANCE ANGIOGRAPHY, AND COMPUTED TOMOGRAPHY COLONOGRAPHY

The Medicare payment for 3-D images that are created as part of a CTA, MRA, or CT colonography, which do not allow separate billing and payment for the add-on code (CPT 76375), has been the source of financial concern. A detailed description of the problems associated with CTA and MRA reimbursements appears in a recent article. To summarize here, prior to 2001, appropriate billing for CTA included 2 separate CPT codes: CT for the specific anatomy involved, plus the 3-D reconstruction add-on code, CPT 76375. Thus, 3-D reconstructions for CTA, as well as for MRA, could be appropriately billed separately, as with other reconstructions, as an add-on to the primary exam. However, effective January 2001, the American Medical Association's (AMA) CPT editorial panel created new codes specifically for CTA indicating that 3-D reconstructions must be performed or these exams are not considered CTA. The code descriptor for CTA includes the phrase "without contrast material(s), followed by contrast material(s) and further sections, including image post-processing." Thus, the add-on CPT code for 3-D reconstruction may not be billed with CTA or MRA as of 2001.

Unfortunately, when the Centers for Medicare and Medicaid Services (CMS) established payment rates for the new CTA codes, it did not recognize the additional facility costs of performing CTA as compared with ordinary CT (i.e., the additional 3-D imaging reconstruction in a CTA). Therefore, CMS assigned CTA the same payment rate as the lesser CT service, effectively eliminating all reimbursement for 3-D post-processing of CTA and MRA; of course, the facility had to bear the costs of the additional equipment and staff time to create the 3-D post-processed images that distinguished CTA from CT.



This change in coding for CTA and MRA has had a significant impact on the cost justification for 3-D imaging services. CTA and MRA account for half of all 3-D exams processed at MGH. Prior to the coding change in 2001, the added \$160 reimbursement for 3-D reconstructions under CPT 76375 supported the costs of 3-D workstations, technologist time, and extra images for the 3-D Imaging Service, but since 2001, the add-on 3-D code can no longer be billed separately for these exams. The 3-D Imaging Service at MGH negotiated with the hospital finance department in 2001 to develop a reimbursement-sharing model that allows for the 3-D Imaging Service to obtain a portion of the total hospital payments for the CTA and MRA services to support the 3-D processing costs of providing this service.

During the past 2 years, the 3-D Imaging Service at MGH, together with the American College of Radiology, has worked to successfully petition CMS to allow for a differential payment for CTA by creating a new CTA reimbursement APC (0662) separate from the CT without and with contrast APC 0333, effective 2003. However, to set reimbursements, CMS looks at hospital charges, adjusted by the hospital-specific cost-to-charge ratio, to determine an “imputed cost.” It is these costs that are the basis for each APC payment rate.

Unfortunately, CMS uses Medicare claims data to determine the payment rate for the 2 separate APC groups, and these data, based on 2001 and 2002 Medicare claims, were not a valid reflection of differential hospital resource use. An analysis of the 2001 and 2002 claims data found that less than half of all hospitals that submitted CTA claims charged more for CTA than CT—clearly indicating that claims were not reflective of the additional resources needed. Since a CTA involves all the costs of CT *plus* the additional 3-D post-processing, CTA should always be a higher-cost procedure relative to CT, and thus, the claims data was not possibly accurate if CTA was being charged less than CT at most hospitals. This may have partly resulted from erroneous and conflicting coding advice in 2001 from the AMA on whether the 3-D charge (CPT code 76375) could be billed with CTA. Thus, despite the separation of CT and CTA into 2 APC groups to allow for differential payment, the net differential was only \$5 in 2003, and only \$24 in the 2004 proposed rule.

Considering that Medicare rates are based on hospital charges, adjusted by a cost-to-charge ratio, it is imperative that all hospitals consider the charges for the new CTA codes carefully. Advocates for CTA urge hospital radiology administrators to consult with their chargemaster managers and revise the CTA charges at their hospital so that CTA charges reflect the CT plus the 3-D reconstruction costs. If hospitals that are charging erroneously (charging less for CTA than for CT) will make this adjustment to reflect the

added costs involved with CTA versus CT, then the claims data will be more valid, and, over the next few years, the differential for CTA reimbursements will rise accordingly. Otherwise, it will be more challenging for new 3-D labs to justify their added costs of staff and equipment to the hospital, and for hospitals to justify the lost revenue of shifting from more invasive procedures such as catheter angiography to less expensive and less invasive 3-D imaging procedures such as CTA and MRA, which translates into better patient care.

Similarly, in July 2004, new codes were created for screening (0066T) and diagnostic (0067T) CT colonography, which include the 3-D post-processing component. Previously, diagnostic CT colonography was properly coded and billed as a CT abdomen and pelvis scan plus 3-D reconstruction (76375) add-on code. Currently, there is no reimbursement value associated with these codes, and for the most part, there are no payments from Medicare or other insurers. We encourage hospitals performing CT colonography to raise this issue with their local Medicare Part B carriers to try to create reimbursement policies for these exams.

## CONCLUSION

Three-dimensional imaging in radiology has made major inroads in the past decade in routine clinical practice. These images both are visually compelling and have clinical impact. However, there is a cost and a learning curve involved in producing optimal clinical 3-D imaging. The benefits of implementing 3-D imaging services can include improved confidence in diagnosis and treatment/surgical planning, reduced cost and invasiveness of procedures, presurgical planning rather than exploratory surgery, and ultimately less risk to the patient and improved patient care.

## REFERENCES

- American Medical Association. *CPT Assistant*. April 1996;6(4):11.
- American Medical Association. *CPT Assistant*. July 2001;11(7):3.
- American Medical Association. CPT 2001 Coding Symposium—Radiology, Commonly Asked Questions, 2001.
- Berland LL, Smith JK. Multidetector-array CT: once again, technology creates new opportunities. *Radiology*. 1998;209:327–239.
- Bogetti JD, Herts BR, Sands MJ, Carroll JF, Vogt DP, Henderson JM. Accuracy and utility of 3-dimensional computed tomography in evaluating donors for adult living related liver transplants. *Liver Transpl*. 2001;7:687–692.

- Center for Medicare and Medicaid Services. CMS Program Manuals—Medicare Carrier (Pub 14), Part 3—Claims Process. October 25, 2002;Transmittal 1776:1.
- Center for Medicare and Medicaid Services. CMS Program Manuals—Medicare Carrier (Pub 14), Part 3—Claims Process: 15021 Ordering of Diagnostic Tests. October 25, 2002;Transmittal 1776:1–5.
- Cline HE, Lorensen WE, Ludke S, Crawford CR, Teeter BC. Two algorithms for the three-dimensional reconstruction of tomograms. *Med Phys*. 1988;15:320–327.
- Damadian R. Tumor detection by nuclear magnetic resonance. *Science*. 1971;171:1151–1153.
- Drebin RA, Carpenter L, Hanrahan P. Volume rendering. *Comput Graph*. 1988;22:65–74.
- Fenlon HM, Nunes DP, Schroy PC, Barish MA, Clarke PD, Ferrucci JT. A comparison of virtual and conventional colonoscopy for the detection of colorectal polyps. *N Engl J Med*. 1999;341:1496–1503.
- Fishman EK, Horton KM. Imaging pancreatic cancer: the role of multidetector CT with three-dimensional CT angiography. *Pancreatol*. 2001;1:610–624.
- Glasser OWC. Roentgen and the discovery of the Roentgen rays. *AJR Am J Roentgenol*. 1995;165:1033–1040.
- Gonzalez RC, Woods RE. *Digital Image Processing*. Boston: Addison-Wesley; 2002.
- Harris GJ, Kassing P, Linton P, Levy-Zuckerman P. Economics of the third dimension. *Decis Imaging Econ*. November 2003:32–38.
- Herfarth C, Lamade W, Fischer L, et al. The effect of virtual reality and training on liver operation planning. *Swiss Surg*. 2002;8:67–73.
- Hounsfield GN. Computerized transverse axial scanning (tomography). 1. Description of system. *Br J Radiol*. 1973;46:1016–1022.
- Hu H, He HD, Foley WD, Fox SH. Four multidetector-row helical CT: image quality and volume coverage speed. *Radiology*. 2000;215:55–62.
- Kaufman A, Cohen D, Yagel R. Volume graphics. *Computer*. 1993;26:51–64.
- Kawamoto S, Montgomery RA, Lawler LP, Horton KM, Fishman EK. Multidetector CT angiography for preoperative evaluation of living laparoscopic kidney donors. *AJR Am J Roentgenol*. 2003;180:1633–1638.
- Klingenbeck-Regn K, Schaller S, Flohr T, Ohnesorge B, Kopp AF, Baum U. Sub-second multi-slice computed tomography: basics and applications. *Eur J Radiol*. 1999;31:110–124.
- Lauterbur PC. Cancer detection by nuclear magnetic resonance zeugmatographic imaging. *Cancer*. 1986;57:1899–1904.
- Levoy M. Volume rendering. *IEEE Comp Graph Appl*. 1990;10:33–40.
- Lorenson W, Cline H. A high resolution 3-D surface construction algorithm. *Comp Graph*. 1982;21:163–169.
- Mansfield P, Morris PG, Ordidge R, Coupland RE, Bishop HM, Blamey RW. Carcinoma of the breast imaged by nuclear magnetic resonance (NMR). *Br J Radiol*. 1979;52:242–243.

- Robb RA. *Biomedical Imaging, Visualization, and Analysis*. New York: Wiley-Liss; 1999.
- Rubin GD. 3-D imaging with MDCT. *Eur J Radiol*. 2003;45(suppl 1):S37–S41.
- Rubin GD, Beaulieu CF, Argiro V, et al. Perspective volume rendering of CT and MR images: applications for endoscopic imaging. *Radiology*. 1996;199:321–330.
- Russ JC. *The Image Processing Handbook*. Boca Raton, FL: CRC Press, 1992:445.
- Sonka M, Hlavac V, Boyle R. *Image Processing, Analysis and Machine Vision*. Pacific Grove, CA: Brooks/Cole Publishing; 1999:57.
- Suri JS, Setarehdan SK, Singh S. *Advanced Algorithmic Approaches to Medical Image Segmentation : State-of-the-Art Applications in Cardiology, Neurology, Mammography and Pathology*. New York: Springer-Verlag; 2002:660.
- ter Haar Romeny BM, Zuiderveld KJ, Van Waes PF, et al. Advances in three-dimensional diagnostic radiology. *J Anat*. 1998;193:363–371.
- Tillich M, Hill, Desbrun M, Schroder P, et al. Prediction of aortoiliac stent-graft length: comparison of measurement methods. *Radiology*. 2001;220:475–483.
- Wigmore SJ, Redhead DN, Yan XJ, et al. Virtual hepatic resection using three-dimensional reconstruction of helical computed tomography angioportograms. *Ann Surg*. 2001;233:221–226.
- Wood Z, Desbrun M, Schroder P, et al. Semiregular mesh extraction from volumes. *Proc Vis 2000*. 2000:275–282.

# VOICE RECOGNITION

---

---

MICHAEL J. MARDINI • AMIT MEHTA

Perhaps no new technology introduced into radiology has evoked more debate and raw emotion than the use of speech recognition technology for radiology reporting. As adoption of the technology has increased, the controversy has increased, with arguments both for and against its overall value to a radiology service. Among those that use the technology, approximately 20% strongly oppose it, 20% strongly favor it, and the remaining 60% fall somewhere in between. Regardless of who is right, there is no arguing the fact that successfully adopting speech technology has a positive effect on many of the areas in which a radiology service is measured. The reduced report turnaround time, decreased costs associated with producing a report, and overall decrease in confusion about report availability and distribution all support the adoption of this technology. One of the largest issues has been the question of the effect on radiologists' productivity and workload. Does it slow the radiologist down to unacceptable production levels, as some claim, or is it just a matter of changing some work habits to allow the user to maintain and ultimately improve productivity, as others claim? Real results from sites using this technology appear to support the latter.

The underlying cause for successful adoption of this technology seems to be the belief that it is no longer enough just to interpret images and dictate findings. Now that we have the technology to do so, it is necessary to take the extra step to communicate *all* findings, not just critical ones, to clinicians and patients *efficiently*. With increased competition for business and the availability of images through Web distribution, it simply comes down to being clinically relevant to patient care. If the report is unavailable but the images are, the clinician down the street may feel that he can view and interpret them to his satisfaction and disregard the radiologist's report entirely. At what point do he start billing for the service?

## HISTORY

The use of computers to recognize human speech has a long history within medicine. In the late 1980s, radiologists and other medical subspecialists started to use expensive dedicated hardware systems that employed specialized vocabulary to recognize reports dictated in a discrete speaking style. After several years of use, it was deemed that voice recognition technology was not mature enough to handle the high-volume demand for the transcription requirements of most radiology practices. The late 1980s and early 1990s saw limited usage and general nonacceptance of voice recognition technology for commercial applications. However, despite the lack of widespread use, development and distribution continued. By 1994, speech recognition systems in American English running on computers with increased processing power had progressed, and speech recognition engines had vastly improved. These developments led to accuracy rates that became acceptable for commercial applications, especially medicine.

It became apparent once development began that no specialty was better suited to voice recognition applications than radiology for the following reasons:

- ▶ A strong need to reduce report turnaround time for improved service provided operational motivation.
- ▶ Financial pressures and a shortage of qualified medical transcriptionists presented a strong financial motivation.
- ▶ A vocabulary that was limited and predictable provided the ability to develop systems that were accurate enough for high-volume use.
- ▶ A defined number of stationary users provided an environment for a supportable application.

- ▮ High-volume, daily use created an opportunity for users to become proficient relatively quickly.

Currently, there are more than 900 radiology installations of various voice recognition systems throughout the United States and continued interest throughout the world. The technology has encountered both acceptance and rejection based on several factors. As the technology becomes more available as an option on picture archiving and communication system (PACS) and radiology information system (RIS) workstations, the barriers to adoption continue to recede.

## ACCEPTANCE AND CONTINUED INTEREST

Speech recognition software packages offer many advantages that promote acceptance by radiologists. First, continuous speech recognition systems require limited amounts of learning and adaptation compared to other transcription systems and methods. The systems are designed to conform to people's most natural way of communicating and essentially do not require the user to alter this method short of speaking more clearly. Second, the ability of the software package to integrate almost seamlessly into existing radiology and hospital information workflows makes the transition easier. The biggest barriers are often resistance to change and fear of technology. As these are overcome, the benefits become more apparent and user acceptance increases. This, in turn, encourages developers to continue development.

Several factors drive interest in speech recognition. First, continued development coupled with increasing processing power lead to improved accuracy rates and the easier use of natural speech. Second, a shortage of medical transcriptionists is occurring in most medical markets. This forces healthcare institutions and practice groups to seek alternative strategies to foster growth and maintain services. Third, there is an immediate return on investment, with decreased operating costs and improved services. Fourth, when thoroughly analyzed, this technique does not require physicians to drastically change their practice in terms of transcription, a feature setting it far apart from competing technologies. Fifth, pressure to be clinically relevant by providing results quickly has increased dramatically as images become easily available through Web distribution. Finally, integration on the desktop with PACS and RIS systems to provide seamless workflow and improved reporting capabilities such as the creation of multimedia and structured reports will further attract attention.

## RESISTANCE TO ADOPTION

Despite the bright future for this technology, several factors delay its widespread implementation. Physicians and people, in general, have an inherent resistance to change. As the next generation of medical practitioners who have been trained on computer systems begin to practice, we will see a change in the dissemination of new technologies. Users who resist the technology claim that they are forced to do more work by editing their own reports. However, an analysis of the current dictation-to-transcriptionist method uncovers the same editing requirements, so why the complaints? The fact is that radiologists typically browse through a report returned from transcription 24 hours later. They may uncover glaring and obvious typographical errors but have a slim chance of correcting errors involving content because they may not remember what they said. For example, a user will most likely miss a mistake of a transcriptionist typing “left” instead of “right” unless “left” was misspelled. In contrast, a report dictated and edited immediately using voice recognition usually receives a more careful review of the resulting text not simply because the software makes different types of mistakes but also because the content is still at the forefront of the radiologist’s mind. This is a very important difference and often overlooked when evaluating this technology. Furthermore, when the time spent reviewing the transcribed reports is added to the time spent initially dictating them, it is unclear whether the total is really any different than that spent with voice recognition systems. Over the next few years, as the technology continues to improve, workflow options increase through integration with RIS and PACS systems, and continued market pressures for improved service and decreased costs drive need, this technology will become ubiquitous and a routine part of a radiologist’s tools.

## CURRENT OFFERINGS

There are a number of sources for this technology. The most common are companies that have built stand-alone applications that include the components and workflow options needed for radiology reporting. The most prevalent of these in the market are Dictaphone’s PowerScribe and Agfa’s Talk Technology TalkStation. These 2 products comprise over 90% of the systems in use for radiology today. Both of these companies are undergoing some changes that have hindered their ability to continue developing and supporting the product offerings and have created an opportunity for some of the newer players to enter the market, such as Lanier/Medquist’s SpeechQ



(Mount Laurel, NJ), Provox's VoxReports (Roanoke, VA), and Commissure's RadWhere (New York), which combines structured reporting with voice recognition.

Solutions may also be obtained through various RIS and PACS vendors. Most have simply integrated with one or both of the above products, while others have opted to integrate a speech engine such as IBM ViaVoice (Wizard Software, Pittsburgh, PA), Philips SpeechMagic (Vienna) or Dragon NaturallySpeaking (Scansoft, Burlington, MA) into their own reporting workflow application. Regardless of how the application is purchased, there are 6 key components to any viable solution to bear in mind:

1. The core speech engine
2. The language model
3. The application and workflow
4. The interface and integration components
5. Implementation services
6. Ongoing support and customization

## CORE SPEECH ENGINE

The speech engine is the starting point for any application driven by speech recognition. Think of the speech engine as a keyboard replacement. There are different types of speech engines. Some recognize only commands and can work on a number of platforms, from PDAs to a PC. Many are built for telephonic applications such as airline scheduling and customer service answering systems. The continuous speech recognition engine is the one used for dictation solutions in medicine. This type of speech recognition allows users to speak in their natural style and will also allow a properly designed application to be controlled by voice command. This is among the most advanced speech technology and among the most processing-power intensive.

Speech engines use a number of proven algorithms and models to recognize speech. Words are spoken into a microphone. Some microphones transmit the analog sound signal to the sound card of the computer, which then digitizes the signal for processing. Others digitize the analog sound signal immediately within the microphone itself and then send the digital signal to the computer, typically via the universal serial bus (USB) port. The latter method seems more robust, as it eliminates the distortions caused to the analog signal from surrounding interference (common in a hospital

setting) as it travels down the wire to reach the sound card. Digital signals are essentially unaffected by such interference. The speech engine “listens” to these digital signals and compares them to acoustic models stored in the software. The engine tries to find the best word match for the acoustic signal to recognize it and convert it to text. Speech engines can typically handle a total vocabulary of more than 60,000 words. However, there are very few applications in which using more than 15,000 is necessary. Shakespeare’s total works comprised fewer than 13,000 words. The average college graduate uses a vocabulary of fewer than 7,000 words in everyday conversation. More specifically, an analysis of more than 4 million radiology reports collected from more than 15 different hospitals produced fewer than 25,000 unique words, including proper nouns.

Nonetheless, building acoustic and language models is a time-consuming process that costs many person-hours to successfully complete. In fact, speech recognition technology has been under development since 1966. It is the longest-running continuously funded research and development (R & D) effort in the history of IBM. Philips has made a significant R & D effort, too, and there are literally dozens of other development efforts continuing. Perhaps the biggest indicator that the time for speech recognition has come is that Microsoft has built its own speech development labs and is likely outspending all other engine developers combined to advance this technology for the human-machine interface.

## LANGUAGE MODEL

The language model is a layer on top of the speech engine that allows for the recognition of specialty-specific terms and reporting structures. Application providers build these language models using programming tools that are included with the speech engine tool kits they have integrated into their applications.

There are actually 3 layers to the language model. The first typically contains a base vocabulary of 30,000 to 60,000 words that come from the manufacturer. These are from everyday English or New York Times English. They include many medical terms as well. This layer also contains the acoustic models associated with those words. The acoustic models are what allow an engine to differentiate one spoken sound from another. Commercial off-the-shelf packages typically provide only this base language model.

The second layer is used to build specialty models that allow application providers to customize and weight the vocabulary for certain topics or

subject matters. This layer is in addition to the base model. Hence, in most cases, radiology-specific applications still contain many nonradiology words that are also active in the vocabulary. However, the words in the customization layer will be more heavily weighted than those in the base model when the engine is working to recognize text, resulting in improved accuracy. The better the effort to build this layer is, the better the overall accuracy of the product will be. It is common to use in excess of 2 million reports to build a language model for medical reporting.

The final layer is the user-specific layer. This layer gets built and optimized as the user dictates and corrects words. If a user adds a word or dictates in a specific style, this layer will learn the style, resulting in improved accuracy. In most cases, 2 to 4 weeks of use will optimize this layer for optimum accuracy. There is a method of speeding this process with a utility inherent in some products that will allow a user to feed up to 3 megabytes (MB) of reports to learn new words and reporting style from previously dictated documents. A properly built language model is very important to a successful implementation of speech recognition.

## APPLICATION COMPONENT

The application is the most important component of a solution. A functional application for radiology reporting is comprised of many things. These include a database, an interface engine, a desktop workflow application for the radiologist, administrative tools, voice file management tools, and distribution and integration capabilities. Figure 23.1 represents how data flows in a properly designed reporting application.

The database includes all the data elements necessary for receiving order data and allowing for the creation and distribution of a report and associated information. It should mirror many of the tables in the RIS database. By integrating workstations into a network, a user can dictate at any PC and not be restricted to a particular workstation. The network link also allows integration with the RIS through the voice recognition server and ultimately into the hospital information system (HIS) and PACS.

## FEATURES NEEDED FOR SPEECH REPORTING

Besides the ability to support speech recognition and communicate within the existing information technology (IT) infrastructure, there are a few basic features a speech reporting solution should have, as follows.

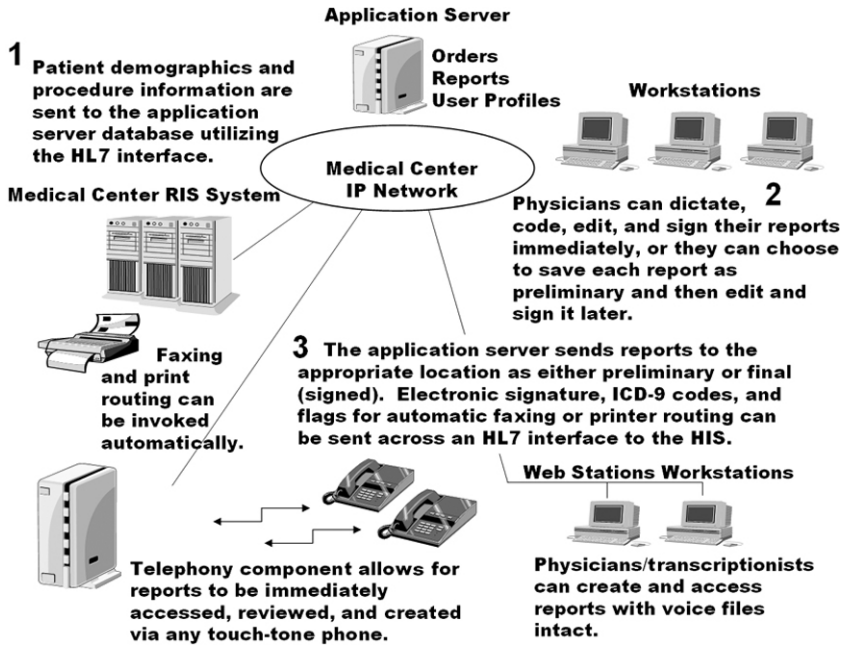


FIGURE 23.1

Diagram shows how data flows in a properly designed reporting application. HL7 indicates Health Level Seven; IP, Internet protocol; RIS, radiology information system.

## SYSTEM SECURITY AND LOG-ON

The voice recognition software package must ensure that security is maintained by requiring each radiologist to enter a user identification code and password to sign on and begin dictation. This is also a utility in the speaker identification and verification function to ensure high accuracy rates. In addition, most systems require the radiologist to enter the same password or another predefined password in order to sign off on the dictated report and allow its transfer to the RIS and HIS. This second password feature ensures that the user is the designated physician identified by the system as dictating the report. Further security measures include the inability to log on to multiple workstations within a facility, which ensures that user workstations are not left unattended. The system will also not allow an order to be open more than once at any one time.

## STANDARD REPORTS

A speech recognition system should allow members of the radiology department to create predefined reports for individual radiologists or the institution. These predefined reports may be categorized for normal studies or commonly performed studies. For example, there is typically a predefined report for the normal chest radiograph, which would describe the normal cardiomedial silhouette and clear lungs. Standard reports should be easily called up either by voice command or mouse click. Some systems may also bring up standard reports based on the type of exam being interpreted.

## TEMPLATES AND MACROS

As an extension of the standard report function, many systems allow the creation of standard reports with customizable templates and fields. The template capability gives the radiologist the flexibility to create a form with blank areas that get filled in during dictation. For items such as procedures, the radiologist is able to dictate the necessary components to fill in the blanks. This feature has great value when describing different dosages or instruments. Most systems contain a feature-filled macro/template creation/editing utility. Using this function during report generation greatly improves the efficiency of the radiologist and decreases the time required to dictate reports.

## CUSTOMIZABLE FIELDS

Most packages permit custom definitions by the institution of multiple fields associated with a report. These fields may include ICD-9, current procedural terminology (CPT), Breast Imaging Reporting and Data System (BIRADS), American College of Radiology (ACR)/National Electrical Manufacturers Association (NEMA) codes, or ACR pathology identifiers. The data entered into these fields may be shared with other information systems, such as the HIS-RIS integration interface. The integration of these fields into a speech recognition solution allows many collateral benefits. For example, an institution may generate a database that can be used for research and education purposes. By defining cases by ACR codes, trainees can retrospectively review cases with selected words. By employing ICD-9 codes, radiology billing services are greatly facilitated. The various uses of these fields are innumerable.

## DESKTOP INTEGRATION

Desktop integration with PACS or RIS workstations has become increasingly important as the proliferation of PACS has advanced. There are many ways for a speech reporting solution to sit on a radiologist's desktop. In a nonintegrated environment there are 2 workstations, each with a mouse and a keyboard. These systems do not communicate, and the user must select and sign cases individually within each application. In a semi-integrated environment, there are still 2 separate workstations. However, they communicate such that order selection and signing happen just once, typically on the PACS but sometimes on the RIS. In a fully integrated environment, the applications reside on a single workstation and communicate almost transparently on the workstation. From a usability perspective, this is the ideal scenario. However, PACS and speech recognition are 2 of the most demanding applications for a PC to run. The author has not yet seen these 2 applications run simultaneously on a workstation and consistently perform optimally over an extended period of time. There are also issues with sharing screen real estate, although this can be overcome by adding more inexpensive color LCD displays. Currently, the most reliable integration is semi-integration. This allows for the use of context sharing and a single workflow point while also providing enough PC horsepower for both applications to run optimally.

## RIS/HIS INTERFACE

Any system needs to provide links to the existing IT infrastructure, allowing seamless integration with the RIS, HIS, PACS, and the billing system. Most software packages incorporate back-end transparent interfaces that allow the speech recognition system to query the RIS for demographic data as well as report status. In addition, the system allows the upload of dictated reports to the RIS and ultimately to the HIS. To achieve this task, most software packages contain Health Level Seven (HL7)-compliant application programming interfaces that use standard formats and protocols. Depending on the level of HL7 interface, most systems allow the radiologist to create a worklist by modality, date, or wildcard categories. Here is an explanation of how HL7 messaging works and some sample messages related to a reporting system.

The HL7 protocol is a simple messaging system in which all data are transferred as ASCII data, not unlike a simple text file. Each transaction consists of a message unit (file) that consists of segments (lines or rows) that

consist of fields. The fields are formatted according to their HL7 data type. Each message, segment, and field is of variable length and uses ASCII characters for delimiters. An example portion of a message follows:

```
MSH|^~\&|Radiology|RDW|Radiology||19960412|ORU|||||
PID|||0123456||Smith^James^T||19610122|M
```

The first segment is a message header (MSH). It contains the delimiter characters to be used in the message (^~\&|), the message type (ORU), and other message control information. The PID segment contains patient demographics, such as name and date of birth (DOB). Notice the format of the DOB, "19610122." This is an HL7 data type called TS (time stamp). The TS is defined as YYMMDDHHMM[SS]. There are numerous HL7 data types, ranging from tightly restricted formats such as TS to very loose restrictions such as ST, which is any string data. Also notice that some fields contain no data. These are optional fields. For a more detailed description of HL7, refer to the HL7 Standard Specification (Version 2.1 or higher).

There are 2 types of communications that occur between a RIS and a speech reporting solution. The first is the sending of an order to the reporting system so that it is aware of what needs to be dictated. These are very similar to orders sent from an RIS to a PACS. The second is the sending of a result from the reporting system to the RIS. An example of an order message that an RIS might send is:

```
MSH|^~\&|||||ORU|0123456
PID|||012-345||Doe^John^P||19610122|M|||1 Fairway
Lane|||||012-34-
5678OBR|||9901|26345^Head,left
view|||199604181202|||||||1234^Welby
```

This example contains basic patient data, exam type, accession number, and the ordering physician information. An example of a result message sent to the RIS is:

```
MSH|^~\&|||||ORU|2345678
PID|||012-345
ORC|RE
OBR|||9901|201A2|||||||||||||199604201055|||F
|||||||D12345
```

OBX|1|FT|201A2&BODY^PA & Lat. Chest||FINDINGS:  
Comparison, 03/01/91.

OBX|2|FT|201A2&BODY^PA & Lat. Chest||The patient is  
status post right mastectomy. The lungs are clear.

OBX|3|FT|201A2&IMP^PA & Lat. Chest||No active  
cardiopulmonary disease. No interval change since  
03/01/91.

DG1|1||174.1^ MALIG NEOPLASM BREAST-  
CENTRAL^ICD9^^|||F

DG1|2||611.9^ BREAST DISORDER NOS^ICD9^^|||F

Many variables and options can be a part of the messaging between an RIS and a reporting system. The above examples are the most basic of messaging that occurs.

## IMPLEMENTATION

Once a particular solution has been identified, the next task is to prepare the site for implementation to ensure success. The institution must first decide what its objectives are and what a successful implementation should look like. This varies depending on differences among sites. Information technology infrastructure must be analyzed to ensure that the system can be deployed everywhere it is needed. A champion should be appointed from among the radiologists to ensure that his or her colleagues are capable and willing to adapt to using the system. Without exception, the single most important component of a successful implementation is a strong commitment from clinical leadership in the department. If the chairman and section heads are not fully committed to succeed, the chances of a successful implementation are small.

## RADIOLOGIST TRAINING

The training of radiologists to use voice recognition software and workstations must occur prior to its wide-scale use in a department. Radiologists must first be able to navigate the basic functions of the operating system; second, they must be versed in the use of computer input devices such as a keyboard, mouse, microphone—whether head mounted or a hand style—and possibly a bar code reader. Next the radiologist must become familiar



with the software interface of the voice recognition package. Many packages offer navigation by voice, which requires the user to remember the names of each function, such as “Accept and sign” or “Save as preliminary.” Other packages require mouse clicks to perform these same functions. Some allow a variety of methods and prompt the user with these commonly used commands to decrease the need to memorize commands immediately. Usually, the champion of the transition is a key radiologist who facilitates learning and eases the implementation. Such an individual can be instrumental in promoting acceptance and training colleagues in the use of the new system. The vendor will typically schedule enough training for all users. Follow-up training some weeks after go-live is highly recommended.

### TRAINING MATERIALS AND TOOLS

Providing users with a written review of the steps they have performed serves to reinforce computer-based training. Vendors should provide written material that is graphically intensive and easy to follow. It should be in a quick-reference format and easily accessed on the workstation.

### TECHNICAL SUPPORT

All devices stop working at times. Whether the affected component is a mouse, a monitor, or the entire workstation, problems invariably occur at one time or another to a PC. This is obvious to all users and inherent in all computers despite precautions taken to prevent downtime. However, a reporting system plays a mission-critical role in the life of a radiology department. Downtime can cost thousands of dollars in radiologist time and can have a severe negative effect on service. In most cases, vendor help-desk support is simply not enough to tend to the needs of a radiologist with a down workstation. Having in-department or hospital support staff that can respond to basic issues in real time is an expense but well worth the effort.

### OPERATIONS

From the initial thought of employing a voice recognition system to all stages during the deployment, close attention to operational planning is necessary. This attention must be generated both departmentally as well as individually.

The first impressions of voice recognition systems are that they are slow and often cause significant delays with large volumes. It is always difficult to foster change in a department, especially when the ultimate benefits of improved report turnaround times and departmental cost savings seem to occur at the expense of the radiologists' time. To successfully overcome these hurdles, there must be a general understanding, especially by those who champion such an effort, of the impact a system shift of this magnitude will have on the radiologists' daily workflow. Only with acknowledgment of these gains by the designated leadership will it be feasible to achieve milestones and communicate the required steps to the members of the department. Beyond these operational issues, specific issues must be addressed during the integration of voice recognition into a department.

### SPECIFIC DELIVERY

During the installation of a system, 4 operational issues should be considered prior to its full-time use in the practice. One first needs a firm grasp of the hardware and infrastructure requirements necessary to create the support systems that would allow seamless integration into the department. Second, individuals need to be identified who can assist radiologists with technical issues promptly and effectively. Third, timely operational checks with close follow-up are necessary to monitor that users are productive with the new system. By ensuring that individual users are effectively operating with the new system, enterprise benefits can be realized. The final operational requirement is a plan for the removal of the legacy dictation system so that penetration of the newer system is guaranteed.

### OPERATIONAL CHECKS

For continued success in voice recognition system implementation, the department must make a commitment to guarantee continued support and use. First and foremost, technical support must be quick, knowledgeable, and courteous. The time that lapses between the request for support and the response to a "distressed" user should be monitored because once fully implemented, a nonfunctioning voice recognition system is costly to both the productivity and the morale of the department. A workstation that crashes usually results in the complete cessation of workflow because of its absolute necessity, especially once 100% penetration has been achieved. Often, users will resort to using legacy dictation systems if they are avail-

able, but this further hampers the operational benefits of the system. Second, the support team must follow up on all outstanding events to prevent recurrence of common problems. There should be a well-documented contingency plan in the event of system failure and a means of communicating this to both radiologists and other users of the system.

## REMOVAL OF OLDER DICTATION SYSTEMS

During the phase-in of the voice recognition system, a point is reached at which a well-organized and timely removal of legacy dictation systems must occur. The removal of these alternate dictation systems ensures both primarily use of the voice recognition system as well as a commitment on the part of the department to the radiologists. It also demonstrates that a well-planned contingency structure is in place in the event of failure. Experience has demonstrated that every department has users who resist change and will continue to use legacy systems if they are available (and in some cases even if they are not readily available). However, the removal of older dictation systems once the voice recognition system is in place, coupled with encouragement from department heads and senior management, allows usage to steadily increase and goals to be reached.

## COST SAVINGS

Several areas of cost savings are associated with the implementation of a voice recognition system. The obvious and most apparent in the radiology practice is the transcription cost. With the direct dictation and transcription of the report, the number of personnel who handle the transcription process can be decreased. The cost savings are realized not only in terms of salaries and benefits but also in the host of other costs associated with personnel.

The indirect cost savings are realized through the use of computer systems. Although this may appear as an added cost due to the purchase of hardware, the actual savings come from the multiple uses of a desktop PC. With the integration of the medical record to include radiology, pathology, and other image-based specialties, the practice of radiology is undergoing drastic changes in the availability of information to the radiologist. Housing the voice recognition system on a conventional desktop PC allows other information agents, including references, paging systems, and HIS-RIS applications, to be coupled to this system. Adding these services to a unifying system saves both time and physical space.

## REPORT TURNAROUND

There is no arguing that report turnaround times decrease dramatically when using a speech recognition system. In most cases, times drop by more than 90%. If coupled with proper distribution methods, reports can be delivered to ordering physicians before the patient leaves the department or arrives home. Considering the level of service this represents and the fact that images are now available via the Web from most PACS, it becomes a necessity to implement this type of service to remain current and competitive.

## CONCLUSION

There is no doubt that radiology must take the path that leads to improved service and more efficient delivery of results. The demands of more timely and improved patient care are greater than ever. The documentation process must be improved, and no technology has demonstrated itself to be more of an aid toward attaining this goal than speech recognition. Perhaps speech in combination with some clinical content and structured reporting will be the pinnacle of functional reporting systems. The technology has improved dramatically over the years. Vendors have learned and advanced their applications to be true workflow improvement tools. Now is the time to consider implementing speech recognition as a means to improve service and remain clinically relevant to patient care.

# ORDER ENTRY IN RADIOLOGY

---

---

DANIEL I. ROSENTHAL

**E**lectronic order entry for diagnostic imaging offers many potential advantages for both the radiology department and its referring community. Despite this, it is not a widely available service, especially for outpatient practices. This lack of penetration may be due to the cost and complexity of initial deployment. However, it may also represent failure to comprehend the full range of possible benefits or to understand the barriers to acceptance by clinicians and radiology personnel.

Implementation of an electronic system is a major investment in time and effort. It is also an exercise in change management in regards to both the referring practices and the radiology department. Electronic order entry may be perceived either as positive or negative depending on whether the users acknowledge the problems that the system is designed to address. For example, it is very difficult to persuade referring offices to comply with a system whose primary benefit is to enhance radiology billing. For this reason, in our large (more than 100 radiologists, 5 locations) academic medical center, we believe that it is very important that the order-entry system be designed with the needs of both user groups in mind. Only then will the initial effort of learning a new procedure appear worthwhile.

At present, there are no widely accepted, radiology-specific electronic order-entry systems on the market. We have created our own electronic order-entry system that functions very well. It is described here to demonstrate what can be accomplished when such a system is tailored for radiology instead of being just a part of a generic order-entry system. Our system was designed empirically and has undergone constant growth and modification as new capabilities were recognized. It is hoped that systems such as this will soon be commercially available.

## PRINCIPLES

Our outpatient Web-based order-entry system was designed with several basic principles in mind:

1. The minimum essential information needed to perform the examination must be captured through the use of mandatory fields. This includes patient, doctor, and examination identification, as well as enough clinical information to assign at least one ICD-9 code (required for billing).
2. All information is electronically transferred to the radiology information system (RIS).
3. Total time for use of the system must not exceed that for a manual system. Keystrokes are kept to a minimum. Since use of the system requires access to a computer, it may be perceived as less convenient than a paper system for ordering, especially for inexperienced users. However, improved scheduling more than compensates for this. Therefore, we have found it important for initial user acceptance to focus on those examinations that require advance appointments (most procedures other than plain radiography).
4. Not every examination can be included. We do not think that electronic order entry is appropriate for urgent or “stat” cases because of the risk that they may not be noted rapidly enough. Interventional procedures require more detailed communication than the system is designed to handle.
5. Some adjustments by “expert users” will continue to be required. It appears unreasonable to burden the users with the full range of possible procedure codes. The final assignment is done within the radiology department either by schedulers or by technologists.

## DESIGN

In the mind of the user, the process of ordering a diagnostic imaging test probably includes the process of scheduling. Different individuals may be responsible for these 2 components. The physician or nurse may write an order in a patient record. The task of scheduling the appointment is often left for a secretarial assistant. We think that it is important that the electronic system provide for both functions, as the greatest advantages for the referring practices are in scheduling convenience.

Our system requires a password-protected log-on for all users, including both physicians and office assistants. At the time of initial registration, each user must indicate for whom he or she is authorized to order exams. Physicians generally order only for themselves. Office staff may sometimes order examinations for a few physicians. This information is stored in the order-entry database and can be edited if required. For each subsequent use, the user sees only the authorized names on a pull-down menu.

Patient identification is done by means of the medical record number. The system verifies the validity of the number and returns the patient name for additional verification.

After identification of the doctor and the patient, the main ordering screen appears (Figure 24.1). This is a listing of modalities (computed tomography [CT], magnetic resonance imaging [MRI], etc.). A click on the modality icon reveals a pull-down menu of examinations (chest, abdomen, head, etc.). When the proper examination is selected, the specific ordering screen appears.

All ordering screens have the same basic structure, intended to mirror clinical practice. A checklist of indications is offered, permitting selection from 3 different categories (Figure 24.2):

1. A sign or symptom
2. A known diagnosis
3. An abnormal previous test

Selection of at least 1 of these boxes is mandatory for scheduling to proceed. Multiple selections may be made, and a free-text box (optional) is provided for additional information. In addition, a listing of common special considerations is offered. These are intended to minimize the need for free text and include common requests such as the use of contrast, conditions to rule out, requests for “wet readings,” or additional physicians to be sent reports. Selection of 1 or more of the special considerations is optional.

Massachusetts General Hospital  
Department of Radiology

Patient: \_\_\_\_\_ MRN: 2219795 Ordering Phys: Rosenthal, Daniel

Exam	Ordering Phys	Schedule Location/Time	Special Considerations	History
<input type="button" value="Start Again"/> <input type="button" value="Print"/>				
<b>CT SCAN</b> SELECT ONE HEAD CT ABD/PELVIC CT CHEST CT EXTREMITY CT SPINE CT FACE OR SINUS CT NECK CT DIAGNOSTIC CT COLONOGRAPHY SCREENING CT COLONOGRAPHY PET CT			<b>MRI</b> SELECT ONE <b>Bone Densitometry</b> SELECT ONE <b>Spine, Bone &amp; Procedures</b> SELECT ONE	<b>Ultrasound</b> SELECT ONE <b>Mammography</b> SELEC <b>Nuclear Cardiology Exams</b> SELECT ONE
<b>FOR CONSULTATION PLEASE</b> <b>CALL:</b> MUSCULOSKELETAL 726-7717 CHEST 724-4254 ABDOMINAL 726-8396 NEURORADIOLOGY 726-8320 NUCLEAR 726-8350 PEDIATRICS 724-4207				

FIGURE 24.1

The opening screen: a listing of imaging modalities. Clicking on the desired modality produces a pull-down listing of examination types.

The lists of common indications and special considerations for each examination were developed using a multistep process. Billing submissions from the years prior to order entry were reviewed to identify commonly used terms. These lists were validated against published listings of medical appropriateness. The lists were then verified by subspecialist radiologists. Finally, prior to creation of the electronic system, the lists of indications were piloted in paper form. Indications were added as needed and removed if not used. Modification of these lists is ongoing.

Once the necessary information has been provided, the user has the option to save the order or proceed to schedule an appointment. Physicians may prefer to have their office staff do the latter.

Clicking on the Schedule icon brings up a screen showing a list of sites at which the procedure is available. A calendar is displayed for the selected site. Days with available appointments are identified in color. Clicking on



the hyperlink shows a listing of available appointments. A limited number of appointments are reserved for same-day use. These times are hidden until the day for which they are offered.

Selection of an appointment time secures the time, removes it from the list of available times, and completes the process. A reminder slip is printed for the patient with instructions for examination preparation. If 1 of our satellite facilities has been selected, driving directions are included.

This system functions independently of the RIS. To avoid conflicts, appointments made available on the order-entry system are blocked in the

Requisition Form - Microsoft Internet Explorer

Patient: \_\_\_\_\_ MRN: 2219795 Ordering Physician: Rosenthal, Daniel

Save/Complete Cancel

### HEAD CT

Policy on CONTRAST: Scans are performed according to departmental protocols selected by the radiologist unless otherwise specified.

Special Considerations, Check If Appropriate

**Contrast**

Contrast MUST NOT BE USED because (Required): \_\_\_\_\_

3D  BUN/Creat (If Known) \_\_\_\_\_

Dissection Protocol (schedule chest, abdomen and pelvis in same time slot)  EVT Protocol (schedule abdomen and pelvis in same time slot)

Pregnant  Radiation planning

Reformats (sagittal/coronal)  Stereotactic

Send additional report copies to: \_\_\_\_\_  Research Study (7 digit # begins w/9) \_\_\_\_\_

At least one box MUST be selected from either of the following groups

**SIGNS / SYMPTOMS**

<input type="checkbox"/> Ataxia	<input type="checkbox"/> Concussion mild or moderate acute, no neurological deficit
<input type="checkbox"/> Convulsions	<input type="checkbox"/> Coordination changes, new or progressive
<input type="checkbox"/> Dementia	<input type="checkbox"/> Dizziness
<input type="checkbox"/> Headache chronic with progressive worsening	<input type="checkbox"/> Headache migraine or chronic
<input type="checkbox"/> Headache sudden onset or Thunderclap	<input type="checkbox"/> Hearing changes
<input type="checkbox"/> Hyperprolactinemia	<input type="checkbox"/> Mental Status change (after trauma)
<input type="checkbox"/> Pain in face	<input type="checkbox"/> Speech changes
<input type="checkbox"/> Swelling, mass or lump	<input type="checkbox"/> Syncope/fainting
<input type="checkbox"/> TIA with transient neurological disturbance	<input type="checkbox"/> Vision changes
<input type="checkbox"/> Weakness-right/left/both	

**KNOWN DIAGNOSES (NOT Rule/out!)**

<input type="checkbox"/> Aneurysm	<input type="checkbox"/> Intracranial hemorrhage
<input type="checkbox"/> Neoplasm CNS primary (specify) _____	<input type="checkbox"/> Neoplasm non-CNS primary (specify) _____
<input type="checkbox"/> Neoplasm-Primary Unknown	<input type="checkbox"/> Sub-dural hemorrhage

Information for radiologist (Only 140 Characters Allowed):

Save/Complete Cancel

FIGURE 24.2

A typical order screen, showing the categories of information: special considerations, signs and symptoms, known diagnoses, and additional information. Selection of at least 1 sign, symptom, or known diagnosis is mandatory for scheduling to proceed.

RIS scheduling template. To prevent times from being unused, blocked times are released for general scheduling 24 hours ahead, and thus the majority of our order entry appointments are available at least one day into the future. A small number of appointments is reserved for same-day availability.

Order-entry data are stored in a Structured Query Language (SQL) database. A Visual Basic script allows the transfer of information from Radiology Order Entry (ROE) to the RIS. Fields are transferred automatically by pressing a single function key, similar to a cut-and-paste operation, thereby drastically reducing the errors invariably associated with data entry. Technical details such as assignment of proper examination codes and system resources are done manually by radiology department personnel.

## ADVANTAGES

### ACCURACY

The order-entry system has resulted in important improvements in the accuracy of the information provided to the radiology department. Although identification of the patient was rarely a problem with the paper system, doctor identification was much more problematic. This occurred because in our large system many physicians have similar names and because their written signatures were often hard to read. Electronic orders eliminate this problem.

Handwritten history presented problems for similar reasons and because the clerical staff who transcribed it might be unfamiliar with the medical vocabulary. Information absolutely necessary for performance of the test (side, level) might be overlooked. The mandatory fields of the order-entry system have made this a problem of the past.

Specific requests for the radiology department (such as “wet readings”) might be overlooked or their significance not comprehended. In addition, there was a further problem due to the fact that the written forms might not arrive in advance of the scheduled appointment.

### CONVENIENCE

Our system offers significant conveniences for the referring office and the patient. The alternative method for obtaining an appointment requires a telephone call to the radiology department. During peak hours, there may be a delay until a service representative is available. Information communicated by telephone may be incomplete or poorly understood. The telephone

communication is difficult to structure, resulting in information that might not be helpful for interpretation or not suitable for billing. Delays in the process often require the clinician's office to schedule the appointment after the patient has left, leading to yet another communication problem, and an increased chance of a no-show.

Because the electronic system is fast and easy to use, patient appointments tend to be arranged while the patient is in the physician's office. This facilitates an interactive process that allows the patient to participate in selection of the appointment, thus minimizing the need to reschedule and the risk of no-show.

Our referring offices prefer not to access multiple systems. Therefore, we provide access to reports through the order-entry system. An icon on the main page allows the referring office to view and print reports on their patients as well as the status of cases not yet performed or reported. Reports not previously viewed by them are searchable separately from those already seen, helping them to keep track of patient flow.

Although we have not as yet done this, in principle, some appointments (such as screening mammography) could be scheduled directly by the patient.

There are also significant advantages for the radiology department. The most obvious advantage is that less time is needed to schedule an appointment by ROE than is needed to schedule an appointment by telephone. Therefore, fewer service representatives can schedule more procedures. In addition, the clerical function of entering data into the RIS takes place without delaying the referring office or patient.

Perhaps less obvious is the fact that the order-entry system becomes a useful conduit for information concerning radiology services. We use banners to announce the availability of new sites and new procedures [such as positron emission tomography-computed tomography (PET-CT)] and to provide tips for earliest appointments. For example, banners have been used to indicate that the earliest appointments for a particular type of exam are available at a particular satellite center. In addition, the choice of defaults helps to guide the selection of appointments to those sites with greatest availability. Finally, because it keeps track of cancellations and rescheduled procedures, the order-entry system provides important insight into the ways in which different referring practices use (and misuse) our scheduling service. This knowledge has proven valuable in our education and marketing operations.

## REQUIRED INFORMATION

The era of managed care has brought about specific information requirements that are often extraneous to the traditional care process. For example,

to submit a bill for diagnostic imaging services, an ICD-9 code must be provided. The ICD-9 code classifies the patient abnormality (a sign, symptom, or pathological state). In our experience, it has been impossible to educate the referring offices to the need for specific information that might be used for ICD-9 coding. Prior to the order-entry system, large numbers of requests continued to arrive with only the name of the condition that was being sought (e.g., “rule out tumor”) but lacking any information concerning patient complaint.

In addition, certain insurance plans require precertifications for certain procedures and provide the precertification based on clinical information. The order-entry system ensures that the required information is collected and can be provided in electronic form should the insurance plan be willing to accept it.

## HEALTH INSURANCE PORTABILITY AND ACCOUNTABILITY ACT OF 1996

We believe that the electronic order-entry system represents a clear improvement in compliance with the Health Insurance Portability and Accountability Act of 1996 (HIPAA). Unlike our prior telephone-based scheduling system, the identification of the user is verified through a secure identifier (ID). Report-viewing privileges are limited to those doctors for whom scheduling is performed. Although our system permits adjustments to the list of physicians for whom scheduling and report viewing are performed, logs are maintained of each interaction with the system, and inappropriate use can be identified.

## SPECIAL CAPABILITIES

Several additional special capabilities of electronic order entry should be noted. Decision support (advice about the appropriate use of imaging studies) can be provided at the time the examination is requested. Our system provides a “utility score” for the requested examination based on the history provided, similar to the American College of Radiology appropriateness criteria. This function is targeted to the physicians who use the system, as clerical staff are not at liberty to change orders and are unlikely to serve as conduits for this type of advice. It is too early to determine what effect, if any, this information may have on clinical practice.

In addition, our system provides information concerning potential duplication. If a patient has had a similar examination within the past month, a warning is displayed on the ordering screen. Similar examinations are not necessarily of the same type. For example, an alert concerning possible duplication would appear if a CT scan of the head is requested when an MRI of the head had recently been done.

## PROBLEMS

As the use of the electronic system has grown, it has been necessary to allocate a steadily increasing fraction of the potential times to it. Because we lack a direct interface to the RIS, this requires continual balancing.

Radiologists have occasionally complained that the histories provided by the checklists sound generic. It has also been suggested that the checklists are too easy to use and that perhaps unqualified personnel are taking the “path of least resistance,” resulting in histories that may not be accurate. It is difficult to verify these concerns or to determine whether similar issues existed prior to order entry; however, ongoing surveillance of the quality of information is required.

An unanticipated source of resistance has been from within the radiology department. Modality managers may regard the order-entry system as lessening their control over their environment. By demystifying the scheduling process, the role of the expert is diminished. Further, the ability to do favors for colleagues by providing access to limited appointment times is a source of pride and power, especially in an environment of scarcity.

## THE FUTURE

The development of large healthcare networks and the increasing growth in imaging volume have undermined the type of personal referrals that once formed the basis of radiology practice. Radiology staff, and even the radiologists, may not know the physicians who send patients. Indeed, in some larger practices, radiology office staff may not even know their own radiologists. Multiple physicians may have the same or similar names, especially in large academic centers. Knowledge of medical terminology may be limited, leading to misspelled or even incomprehensible requests. Such trends will probably continue, making development of sophisticated knowledge-based, error-resistant systems ever more important.

It has become increasingly clear that tight integration of information systems is essential to capture their full benefits. For example, use of a “broker” to coordinate Digital Imaging and Communications in Medicine (DICOM) images stored on picture archiving and communication systems (PACS) with patient information in the RIS introduces complexity that may result in use of outdated information. Newer system designs are tending to incorporate the broker function directly into the RIS. Ideally, the order-entry system should communicate directly with the RIS. In the near future, our RIS will likely be able to handle such interaction, thereby simplifying the whole process even further.

Implementation of a computerized order-entry system in a neonatal intensive care unit has already been shown to shorten the radiology response time (interval between order and image display) by almost 25%. Integration of technology will bring many further gains. Doubtless the future evolution of continuous speech recognition systems will include order entry, making these systems even faster and more convenient.

As progress is made in coordination of the components, there will be improved understanding of the possible failure points, leading to development of system tools that understand and monitor information flow.

Perhaps, as reality comes closer to concept, radiology order-entry systems may be integrated with electronic patient records, appropriateness criteria, and practice guidelines. Such computer-based practice guidelines might also include additional didactic material, such as images, videos, sounds, simulations, and links to bibliographic databases. Given patient data, these future systems might even select the most appropriate imaging study automatically.

## REFERENCES

- Clark KW, Melson DL, Moore SM, et al. Tools for managing image flow in the modality to clinical-image-review chain. *J Digit Imaging*. September 2003;16(3):310–317.
- Cordero L, Kuehn L, Kumar RR, Mekhjian HS. Impact of computerized physician order entry on clinical practice in a newborn intensive care unit. *J Perinatol*. February 2004;24(2):88–93.
- Langer S. Radiology speech recognition: workflow, integration, and productivity issues. *Curr Probl Diagn Radiol*. May–June 2002;31(3):95–104.
- Mulvaney J. The case for RIS/PACS integration. *Radiol Manage*. May–Jun 2002;24(3):24–29.

- St. Anthony's Guide to Code Linkage, CPT and ICD-9 CM Code Connection*. St. Anthony's Publishing; January, June 1997.
- Tjahjono D, Kahn CE Jr. Promoting the online use of radiology appropriateness criteria. *Radiographics*. November–December 1999;19(6):1673–1681.

# DIGITAL TEACHING FILES AND EDUCATION

---

---

KHAN M. SIDDIQUI • BARTON F. BRANSTETTER IV

**T**he learning process in radiology requires systematic study of a large knowledge base of medical images. The Residency Review Committee (RRC) for diagnostic radiology of the Accreditation Council for Graduate Medical Education (ACGME) mandates resident education and competency requirements. These requirements include core competencies stressing the importance of computer-aided applications in medicine and the use of online resources for educational and patient care purposes. Patient care requirements include the use of information technology (IT) to support both patient care decisions and patient education. Practice-based learning and improvement requirements demand that residents use IT to access online medical information and support their own education. As part of competency in professionalism, residents are expected to show ongoing professional development. Residency programs are expected to teach residents computer applications in radiology, practice management, and quality improvement.



Residency programs are also required to develop a system to assess and document residents' competence in these areas. ACGME requirements also state that "[t]eaching files (electronic or film) must be available for use by residents and these files should contain a minimum of 1000 cases that are *actively maintained* and *continually enhanced* with new cases" (emphasis added). This underscores the importance of addressing teaching file requirements during the transition to filmless imaging.

Traditionally, case-based hardcopy teaching files have been created from radiology studies of real patients. Such teaching files consist of radiology images accompanied by relevant clinical data and a short write-up of the pathological condition. Libraries of representative teaching files are highly desirable in radiology teaching programs, but the organization and maintenance of collections can become problematic.

With the advent of picture archive and communication systems (PACS), the acquisition, viewing, and organization of medical images for teaching and research should be simplified, because such data are digital from the start. Moreover, in the electronic environment, powerful and versatile tools have been developed and can be customized to maintain, organize, and search for images. However, no PACS vendor now provides a solution for the creation of electronic teaching files (ETF) as a commercial option. Hence, the potential for PACS as a teaching and research tool has not been maximized. Many teaching hospitals and universities have designed in-house solutions, employing a variety of methods to create their own ETF libraries.

## MEDICAL IMAGE RESOURCE CENTER

The Radiological Society of North America (RSNA) has defined standards for a medical image resource center (MIRC). A set of extensible markup language (XML) schemas for information exchange, such as MIRC query, MIRC query result, MIRC site index, and MIRC document, have been developed. By complying with these schemas, individual teaching file servers around the world can be linked together through the World Wide Web to form a global MIRC community. Thus MIRC has the potential to become the worldwide standard for teaching and research data exchange, in the same way that the Digital Imaging and Communications in Medicine (DICOM) standards have been accepted and incorporated into PACS.

MIRC was originally conceived as a large, central database of images and related clinical information that would be maintained by the RSNA for open access by the medical imaging community. It soon became evident that a more effective strategy would be to leverage existing online electronic

resources in addition to creating new ones. A decision was made to create a virtual community of medical image libraries that would benefit from a distributed index mechanism and also support a distributed search mechanism—in other words, to create a kind of radiologic Napster. The final result was MIRC, an open community of medical imaging libraries and teaching files from around the world, sharing information through a common query format. The system has evolved to provide support for a community of co-operating libraries that can be managed at multiple levels, ranging from an individual working at a laptop or personal computer (PC) to a level at which a user can access global content as if it were in a single library. The virtual libraries provide all kinds of digital information, including teaching files, clinical and technical documents, electronic presentations, and imaging datasets for research (Figure 25.1).

The objective of the MIRC project is to support the production, storage, indexing, and distribution of medical imaging resources, such as teaching files, scientific and technical documentation, research images and datasets, and clinical trials data.

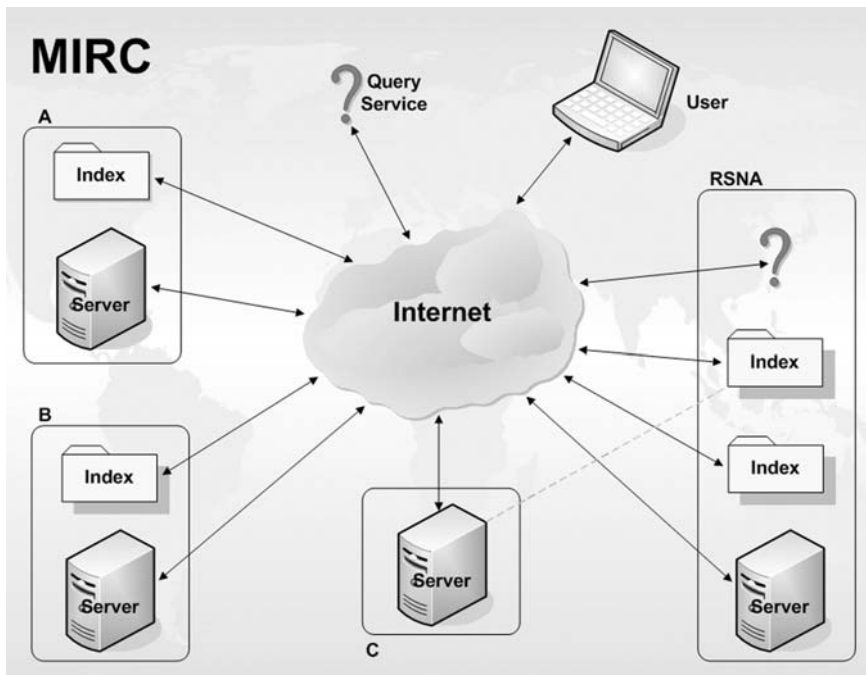


FIGURE 25.1

Medical image resource center (MIRC) community diagram.

Authors may use a MIRC-defined format to construct teaching files and other documents in a common structure that allows libraries to index the documents in medically meaningful ways (Figure 25.2). The RSNA has also developed a MIRC authoring tool, known as MIRCat, as well as a Web-based MIRC authoring service.

An indexing mechanism provides users great flexibility in searching the MIRC community. Users can perform free-text searches on the contents of documents, as well as structured searches on patient criteria (e.g., sex, age) or image criteria (e.g., modality, anatomical region, storage format, compression), all through a standard Web browser (Figure 25.3). Any MIRC site can function as either an access point for users (called a query service) or an indexed information library (called a storage service) or both. A query service provides a point of access to the entire MIRC community. It provides a query form to the user, distributes the search criteria to all selected storage services, collates the responses, and presents them to the user. A storage service responds to the query received from the query service, searches its index for documents meeting the search criteria, and returns abstracts and locations of the matching documents to the query service.

The screenshot displays the 'Khan's MIRC Site' interface. At the top left is the 'MIRC Medical Imaging Resource Center' logo. The top right corner shows the version 'Version 121h' and links for 'Help', 'MIRC Home', 'MIRC News', 'MIRC Forum', and 'Documentation'. Below the logo is a 'MIRC Query' section. On the left, there is a 'Submit Query' button and three checkboxes: 'Display as unknowns', 'Case navigator', and 'Randomize results'. The main area features a list of MIRC sites: 'Khan's MIRC Site', 'The RSNA MIRC Site', 'Casimage Teaching File', 'The Indiana University MIRC Site', 'Mallinckrodt MIRC Storage Service', 'The MedPix MIRC Storage Service', and 'MyPACS.net: Teaching File'. Below the list are tabs for 'Basic', 'Document', 'Content', 'Clinical', 'Image', and 'Patient'. A 'Free Text Query:' section contains input fields for 'Title:', 'Author:', 'Abstract:', and 'Keywords:'. To the right of the search fields are buttons for 'RadLex', 'Author Service', 'Submit Service', 'Admin Service', and 'Login'. At the bottom right, there is a 'Powered by TOMCAT' logo.

FIGURE 25.2

MIRC teaching file (mirctf) case display format.




FIGURE 25.3

MIRC query page used to search through the MIRC community.

As of this writing, the international MIRC community includes 10 sites, from Switzerland to Singapore, accessible through the RSNA MIRC site, plus others running the RSNA MIRC software privately. The RSNA MIRC software now supports teaching files, research datasets, scientific and technical documents, and clinical trial image acquisition and distribution.

The current software release supports public, restricted, and private teaching files, all on the same server, allowing authors to use MIRC as a personal archive as well as a publication mechanism to group, department, and global audiences. The software now includes a server-based editor, facilitating online creation of teaching file cases and other documents without any special software on the author's system.

The current software release also includes a DICOM storage service that automatically creates clinical trial documents in response to the receipt of images from local modalities or participating clinical trial sites. It allows distribution of DICOM images to other participating clinical trial sites in a wide range of topologies. The clinical trial service also supports research dataset acquisition, with automatically created documents that can be edited through the author service. The service also allows capture of commentary

on images and studies, making them more easily searchable. More information about MIRC is available at <http://mirc.rsna.org>.

## OTHER TEACHING FILE RESOURCES

A wide range of online teaching file resources is available to radiologists in training or in clinical practice who want to increase their breadth of knowledge or who wish to use these resources as electronic decision support tools during daily practice or to anonymize and store interesting cases for departmental or personal teaching file collections. A simple search for the keywords “radiology teaching file” on <http://www.Google.com> at the time of preparation for this chapter yielded 81,100 hits. Some of the more popular teaching file resources are listed in Table 25.1.

**TABLE 25.1**  
Selected Digital Teaching File Resources

Digital Teaching File	Web Address
MedPix	<a href="http://rad.usuhs.mil/index.html">http://rad.usuhs.mil/index.html</a>
MyPACS	<a href="http://www.mypacs.net">http://www.mypacs.net</a>
CasImage	<a href="http://www.casimage.com">http://www.casimage.com</a>
EuroRad	<a href="http://www.euro-rad.org">http://www.euro-rad.org</a>
CTisUS	<a href="http://www.ctisus.com">http://www.ctisus.com</a>
BrighamRad	<a href="http://brighamrad.harvard.edu/">http://brighamrad.harvard.edu/</a>
Auntminnie—Indiana University	<a href="http://www.auntminnie.com">http://www.auntminnie.com</a>
Mallinckrodt	<a href="http://gamma.wustl.edu/home.html">http://gamma.wustl.edu/home.html</a>
American College of Radiology Learning File	<a href="http://www.learningfile.com">http://www.learningfile.com</a>
Thomas Jefferson Neuroradiology	<a href="http://popeye.rad.tju.edu:8080/mirc/query">http://popeye.rad.tju.edu:8080/mirc/query</a>
University of California at San Francisco	<a href="http://128.218.59.127/">http://128.218.59.127/</a>
GlobalRad	<a href="http://www.globalrad.com">http://www.globalrad.com</a>

*Source:* Data from Branstetter BF, Siddiqui KM, Lionetti DM, Chang PJ. Defining a digital teaching file workflow: specifications for software development. *J Digit Imaging.* 2003;16(suppl 1):37–40.

## **SPECIFICATIONS FOR DEVELOPMENT OF TEACHING FILES**

Despite advances toward an all-digital environment in academic radiology departments, the teaching file—perhaps the most distinguishing and tradition-honored feature of medical imaging pedagogy—has been slow to adapt to the full range of informatics possibilities. One complicating element has been the requirements imposed by compliance with the Health Insurance Portability and Accountability Act (HIPAA). Because most traditional teaching files are not sufficiently anonymized, they are noncompliant with HIPAA requirements. To get around this difficulty, many academic institutions have developed in-house digital teaching files (DTFs); Web-based and standardized formats have been developed for sharing cases between institutions. However, most PACS vendors have yet to incorporate DTFs directly into their PACS offerings. Without an easy teaching file solution integrated within the PACS system to allow users to create teaching file cases easily and effectively, user compliance and integration of teaching files into all-digital department academic activities will continue to pose difficulties.

## **INTEGRATING THE HEALTHCARE ENTERPRISE TEACHING FILE INFORMATION EXPORT PROFILE**

The Integrating Healthcare Enterprise (IHE) initiative is a project designed to advance the state of data integration in health care. Sponsored by the RSNA and the Healthcare Information and Management Systems Society (HIMSS), it brings together medical professionals and the healthcare information and imaging systems industry to agree on, document, and demonstrate standards-based methods of sharing information in support of optimal patient care. A proposal for a teaching file information export profile is under consideration by the IHE that, once released, will address uniform methods for capturing information for teaching file cases across various PACS vendors.

Although the IHE profile will address workflow implications for creating a case at the PACS workstation, specific data captured for cases may differ for each user or institution. General format and data elements for a DTF are detailed in Table 25.2. Included in the following sections are proposed specifications that should be addressed by an ideal DTF software package; the focus is on the importance of workflow integration.

TABLE 25.2

## General Format for Digital Teaching Files

Basic data elements for a DTF (mandatory fields)

- a. Patient name\*
  - b. Medical record number\*
  - c. Accession numbers\*
  - d. Category
    - ▶ Users should be permitted to develop their own tree structures for categorizing diseases, although using some standard terminology will make integration with other DTF solutions easier. Note that a patient may appear more than once, either because a disease fits 2 categories or because a patient has 2 diseases.
    - ▶ It is important to have a different field that defines or differentiates an adult from a pediatric patient. Simply defining a category for pediatric radiology may not be sufficient, because this would not divide the topics into anatomic- or system-based categories.
    - ▶ Some teaching file vendors fail to define a hematology/oncology category, which is important because some multisystem malignancies are not covered in other standard categories.
  - e. Diagnosis
  - f. Modalities
    - ▶ Included in “modality” is the encoding of detailed image attributes. For magnetic resonance, in particular, a list of pulse sequences should be generated. Ideally, this list would be generated automatically as part of the PACS integration. The list of modalities and possible attributes must continually expand to accommodate new technologies.
2. Secondary data elements (useful, but not mandatory)
    - a. Keywords: Keywords are used to search or filter.
    - b. User flags
      - ▶ An arbitrary number of flags should be available to users, and these should be user definable. Examples of user flags include “board review,” “follow up,” or “potential AFIP [Armed Forces Institute of Pathology] case.”
    - c. Publication flag: To prevent duplicate publication within the medical literature.
    - d. American College of Radiology code
    - e. Patient gender
    - f. Date of birth\*
    - g. Date of exam\*
  3. Key images: Access to the full examination should be maintained. Key images may differ when a patient appears more than once in the database.

*Source:* Data from Branstetter BF, Siddiqui KM, Lionetti DM, Chang PJ. Defining a digital teaching file workflow: specifications for software development. *J Digit Imaging.* 2003;16(suppl 1):37–40.

\* Patient identifiers may be hidden or encoded, depending on the specific use.

## TYPES OF TEACHING CASES

At an academic institution, there are many situations in which a physician needs a group of radiology studies to be stored together. Digital teaching files are widely discussed but are only one application of the more generic “case collection.” Case conferences, teaching conferences, tumor boards, morbidity and mortality conferences, clinical research protocols, and board reviews are among the many uses for case collections. A complete DTF software package would focus on the needs of an institutional teaching file but be flexible enough to accommodate other uses. The critical element that is lacking in current DTF solutions is integration. Digital teaching file software requires an efficient workflow so that cases can be rapidly added to lists during readouts, with a minimum of interruption. Familiar authoring tools for annotation of images can be applied later when these will have less workflow impact. None of the current DTF solutions specifically addresses the usual academic reading room workflow, and most are not integrated with the PACS.

## WORKFLOW INTEGRATION

Flagging a case for the teaching file should be an instinctive task, completely integrated into the workflow of the daily reading session. Flagged cases should be assigned an owner, who should be reminded at a convenient time that teaching file cases are available for processing. A 1-line description (e.g., “otosclerosis vs. Paget’s”) would be useful to direct the owner to the appropriate images while reviewing a case. The DTF database should be available on the same machine that displays the PACS images so that interesting comparisons and counterexamples can be used to reinforce cases during readout and so that users can benefit from the DTF as a decision support tool.

## SHARING CASES WITH COLLEAGUES

In an academic setting, the identified PACS user is often the resident in training, but it is the attending radiologist who wants the cases stored. Similarly, a radiologist may find a case that would be of greater interest to a colleague than it is to him- or herself. It is critical that users be permitted to alert others to an interesting case at the time the case is chosen. This is another aspect of workflow integration.



## PRIVATE FILES

Users may wish to keep individual cases private but make the rest of the teaching file available to trainees. Such private cases might include board review cases, test cases, or cases pending publication. Private cases should be interspersed with public cases within the category tree rather than relegated to a distinct section of the database.

## PUBLIC FILES

Authorized users within an institution should have full access to “public” files for teaching purposes. Presentation of cases may be simple (a few key images) or complex (unknown cases for case-based learning). Anonymizing cases is critical for HIPAA compliance.

## DIVISIONAL FILES

The divisional file is the digital counterpart of the conventional film-based teaching file. Cases in the divisional file have been contributed by members of that division for the purpose of teaching residents and fellows. Ideally, there is an individual gatekeeper from each division who verifies the quality and legitimacy of each case and prevents excessive duplication of content.

## PUBLISHED FILES

Sharing of files across institutions provides greater opportunities for resident education, particularly in rare diseases. Sharing files for research collaboration is also desirable. Cases that are available outside an institution are considered “published cases.” This should not be confused with formal publication in the medical literature.

## **MEDICAL IMAGE RESOURCE CENTER COMPATIBILITY**

Although many different teaching file indexing schemas are available, it is crucial that vendors and institutions adhere to a standard format for cross-compatibility and easy case querying. The RSNA MIRC project described

previously is currently the only schema with the potential to be a worldwide standard for teaching file indexing and development. Any vendor DTF solution should have the capacity to comply with the MIRC document schema so that individual institutions can make anonymized public cases available for review to the entire radiology community when they choose to do so.

## **ELECTRONIC MEDICAL RECORD INTEGRATION**

For a DTF system to provide a robust and crucial decision support role in day-to-day radiology practice, it is important that the system be able to communicate with multiple databases and information systems throughout the hospital. This will enable the DTF system to notify users when the status for a specific case is changed or some needed information is available to update the case.

### **HOSPITAL INFORMATION SYSTEM INTEGRATION**

Using the IHE profiles, DTF software should be able to query the hospital information system (HIS) to obtain pathology results, laboratory results, and related data.

### **NOTIFICATION**

Cases that have pending results (e.g., biopsies performed under image guidance) should be flagged. The system should poll for pending results on a regular basis and notify the owner of the case when results are received. Similarly, when a follow-up radiological examination is expected on an interesting case, the owner of the case should be notified when the follow-up images become available. The DTF system should be able to do this while keeping the cases anonymized, without compromising the patient's identity.

### **SUPPLEMENTAL IMAGES**

Images from nonradiological modalities (endoscopy, pathology, or clinical images), as well as digitized images from outside sources (e.g., noninstitutional examinations), should be accepted by the DTF system. These supplemental digitized radiographic images may be supplied in a variety of

formats (e.g., JPEG, GIF, TIFF, PNG, MPEG, AVI) but are best converted to a DICOM format with appropriate headers for ease of manipulation (e.g., window and level changes, etc.).

## PERSISTENCE

Some cases, such as those for teaching conferences or planning conferences, are needed for only a short time. Others, such as teaching files, require permanent archiving. Institutions that archive old examinations must either exempt DTF cases from removal or make copies of the DTF cases in a separate section of the PACS.

## ANONYMIZATION

### HEALTH INSURANCE PORTABILITY AND ACCOUNTABILITY ACT

The Health Insurance Portability and Accountability Act (HIPAA) was designed to protect patient privacy by restricting access to identifiable patient information or protected health information (PHI) for treatment, planning, or operations only. These regulations explicitly define 18 demographic elements that might be used to uniquely identify a patient. Teaching files fall outside treatment, planning, or operations and are intended to be viewed by large numbers of trainees; therefore, they should not contain PHI. Most film-based teaching files were developed before the HIPAA regulations were formulated and contain identifiable data that may render them noncompliant. Patients should be assigned identification numbers unique to the DTF system. Because there is often a need to obtain additional data from the radiology information system (RIS) or HIS at a later time, the system must be able to decode the unique identifier for the purpose of querying the electronic medical record (EMR). Institutions and departments should specifically address these operational issues and their countermeasures in their local HIPAA policy. A complete description of patient care-related tasks and events should be in place to cover daily user interactions with the PHI for creating and updating teaching files. Some institutions request waivers from patients to allow the use of their PHI in teaching tasks. This strategy may satisfy HIPAA compliance locally, but disclosure of the PHI outside the covered entity may still be a HIPAA violation. In a similar

process, a patient could conceivably sign away rights to privacy and PHI security and allow his or her studies to be shared over the Internet.

## EXEMPTIONS

Some case collections, such as those for treatment planning conferences, may be exempt from anonymization so that clinicians can appropriately care for patients. Such collections should have limited persistence in the DTF database or be de-identified after a certain time period.

## DEPTH

Some teaching files are thoroughly researched and include discussions of clinical presentation, lengthy differentials, and literature references. But sometimes the user needs only a pointer to find the case for future reference. The clinical information present in the different types of teaching file documents can be categorized as follows:

### LIST

The list is a quick pointer to the case, with fields for diagnosis and keywords. No images are needed, and no annotation is performed.

### KEY IMAGE

In this case, a few key images are selected. There may be annotation on those images, but no lengthy prose is associated. These are useful for quick examples to demonstrate a point during teaching rounds.

### FULL DIGITAL TEACHING FILE

Complete teaching cases with images, history, findings, differential diagnosis, final diagnosis, discussion, and references are a staple of teaching institutions. The DTF software should support links to related teaching file cases or to Internet references.

## QUALITY ASSURANCE

### AUTHORIZATION

Full DTF cases should undergo review by a subspecialist before inclusion in the formal institutional teaching file. A division might allow any member to sponsor a teaching file or might assign a single individual to act as a gatekeeper. Robust review and notification tools are necessary to ensure that teaching file cases are reviewed in a timely manner upon submission.

### PEER REVIEW

More than 1 expert may vouch for the quality of a single case. Secure tracking of cross-institutional peer review would allow users to assess the quality and accuracy of a published case.

### USER REVIEW

Accuracy of material is not the only measure of the quality of a teaching case. Users should be able to rate the teaching quality of the case presentation.

The validity of the averaged results may be questionable, but this gives users further guidance when sifting through the myriad of cases that may eventually be available over the Internet.

## AUTHORING TOOLS

### INTEGRATION

Authoring tools need not be fully integrated into the PACS. When the PACS allows a case to be flagged for later processing, that processing can be performed at another site (back in the office or in the residents' lounge) by passing the information for the flagged case to the enterprise-wide, Web-based PACS viewer. This keeps high-end PACS resources free for readouts.

### THIN AND THICK CLIENTS

Web-based thin clients have the advantage of universal availability and would be useful if users are expected to author cases from outside the institution.

Thin-client applications do not require a great deal of processing power and can be installed on almost any PC. A Web browser–based client eliminates the need to install the DTF application on various computers, because it can be accessed from anywhere using a standard Web browser. However, thick-client applications require much more processing power and therefore are available only at specialized workstations or beefed-up PCs. They may provide more robust and advanced image-processing tools and may be the preferred solution within an institution, although thin-client applications can provide equally robust tools.

## IMAGE ANNOTATION

Original image data, without demographic or image attribute text, should be stored in the DTF. Annotations should be stored in a separate layer (or layers), so that they may be hidden during self-tests or teaching conferences. Vector encoding of annotations is generally more compact than raster encoding. Possible annotations should include at least key image ordering (the order in which images should be presented), arrows, circles, text, and measurements.

## USER ROLES

### TRAINEE

The turnover among trainees is high, and they are expected to create cases under the supervision of an attending imaging specialist. The attending can then take ownership of the case when the trainee leaves. Although this creates a small additional workload for the sponsoring attending at the time the case is created, it prevents interesting cases from being lost when the trainee leaves for another institution.

### ATTENDING

Attending or staff radiologists have the highest level of access. This level of access may not be needed for individuals outside the radiology department, unless the DTF system also provides case storage solutions for other departments, such as cardiology, dermatology, or pathology. Departmental policy

should govern the cases that are left behind when an attending leaves the institution.

## RESEARCHERS

Researchers may make anonymous patient lists by utilizing a teaching file system as a research file system for a research project or clinical trials, but these cases should have a limited persistence that is directly related to the expected length of the research project. Appropriate institutional research board (IRB) consent should be secured before creating any list for research purposes.

## POST-PROCESSING

### CASE COLLECTIONS

Users should be able to conveniently group cases for teaching conferences or lecture preparation. These groupings can cross the boundaries of the category tree structure. Boolean searches incorporating keyword, diagnosis, American College of Radiology code, and/or user-defined fields are needed.

### PRESENTATION TOOLS

The DTF software should provide tools to efficiently create case-based exams, assign teaching file creation (e.g., case of the week), and present cases in a classroom or conference setting. If these tools are Web based, they will be available throughout the institution and should obviate the need to hand carry media or hardware. Users, particularly medical students and residents, should be able to request full DTF cases from a given category to be presented as unknowns for the purpose of self-testing.

### PUBLISH TO PRESENTATION SOFTWARE

Integration with established presentation software would allow a user to take, for example, all the key images from a case collection and send them into Microsoft PowerPoint (Microsoft Corporation, Redmond, Wash.) in a compatible image format, 1 image per slide, preserving annotation.

## BURN A CD

The DTF solution should have the ability to quickly store case collections on CDs, either as key images only or as entire studies. A simple case viewer should also be burned on each CD for portability.

Maintaining an active teaching file is a requirement for the accreditation of radiology residency programs. Academic institutions should consider the importance of integrated teaching file software when choosing a PACS vendor. In the future, institutions should demand that vendors provide conformance with the IHE teaching file information export profile. The proposed software specifications for digital teaching files may be helpful to programmers and PACS vendors in understanding the functionality needed for DTFs in academic practice. The importance of workflow analysis and integration cannot be overemphasized.

## IMAGING INFORMATICS EDUCATION

If one searches for the meaning of medical informatics, a number of divergent and confusing definitions can be found. One definition states that medical informatics is an emerging discipline focusing on the study, invention, and implementation of structures and algorithms to improve communication, understanding, and management of medical information. The end objective is the coalescing of data, knowledge, and the tools necessary to apply such data in the medical decision-making process. Another definition describes medical informatics as an evolving scientific discipline that deals with the collection, storage, retrieval, communication, and optimal use of health-related data, information, and knowledge. In the end, the common denominators in most definitions are information management and technology, which are at the core of informatics. For the imaging community, information management relates to medical images and extends to the manner in which imaging data are acquired, stored, distributed, interpreted, and communicated. Hence, the field of imaging informatics could be defined as being concerned with “management of information from the point at which a physician decides to order an imaging study to the point at which an informed interpretation of the results of that study are delivered back to the physician.” An imaging informatician is an individual who bridges the knowledge and communication gaps among radiologists, physicists, and IT personnel. Although such a skill is not listed in most job descriptions for radiologists, it is becoming more and more important as imaging departments transition from film-based to filmless imaging, adopt and integrate



PACS with the HIS and RIS, and add new modalities that rely heavily on IT.

## EVOLUTION OF IMAGING INFORMATICS AS A RADIOLOGY SUBSPECIALTY

Advantages that occur when 2 different institutions merge to form a new entity can result in much more rapid change over a given period of time than can be achieved by a single entity. When different scientific disciplines borrow from one another and combine to form a new specialty, this can also result in dramatic advances in a relatively short period of time.

From its origins in mathematics and military ballistics, digital technology has evolved to become an integral part of endeavors as diverse as communication, commerce, entertainment, and medicine. The field of computer science continues to evolve into a wide variety of subspecialties. It is becoming increasingly evident that just as amalgamation between the field of medicine and the study of the physics of the x-ray resulted in the field of radiology, the union between radiology and IT is resulting in a new specialty: imaging informatics.

This emerging subspecialty has practical applications that transcend the traditional modality/anatomic-based radiology subspecialties. As medical imaging undergoes a paradigm shift from traditional film-based to filmless operation, many of the traditional concepts of medical imaging become obsolete. The fundamental processes of image acquisition, display, storage, distribution, interpretation, and reporting undergo fundamental change with the implementation of digital imaging, PACS, advanced image processing, decision support tools, speech recognition, and structured reporting. With accompanying changes in image acquisition, larger and more complex imaging datasets are being created that bring new challenges for radiologists. One obvious example is the widespread adoption of multi-detector computed tomography (CT), which has increased the number of individual CT images by 3- to 6-fold for routine studies. The radiologist who continues to review and interpret these imaging datasets in the conventional tile mode format will quickly become overwhelmed. Such challenges transcend all modalities, anatomic regions, and imaging subspecialties. The angiographer must cope with complex CT and magnetic resonance (MR) angiography, the breast imager with digital mammography and computer-assisted diagnosis (CAD) software, the oncologic imaging specialist with positron emission tomography (PET)/CT fusion datasets, and

the chest specialist with dual-energy subtraction and tomosynthesis. The field of imaging informatics can address many of these challenges that transcend the various imaging subspecialties.

## LATENT POWER OF IMAGING INFORMATICS

“Information is power,” according to an old adage, but in the case of the radiology community, where do we go to acquire the information we will need to maintain power? With these new computer-based applications, many traditional concepts of image acquisition, archiving, interpretation, and reporting will disappear. If we are to maintain our power and role as medical imaging experts, we need to learn how technology and information management define the digital imaging practice. To accomplish this task, in-depth training is required and should begin during residency. Ironically, few if any radiology residency programs offer formal imaging informatics training, resulting in a distinct knowledge void among new radiology residency graduates. This is particularly detrimental in the private practice radiology community, which often relies on the most recently trained radiologists as champions of new technologies.

While some in the private practice radiologist community are lamenting their relative lack of expertise in imaging informatics, academic radiology is not far behind. Academic radiology has long been the bastion of radiology “purists,” who define and set the standard of practice for the radiology community at large. When new modalities or applications are adopted, it is representatives from these academic programs who define how these new applications will become incorporated into everyday practice. This includes such mundane aspects as the establishment of exam protocols and indications and extends into broader areas, such as radiologist education and certification. Traditionally, these academic challenges have been assigned to the various modality and anatomic-based subspecialists within the academic ranks. When a new modality, such as magnetic resonance imaging (MRI), is introduced, a new breed of subspecialists is born from within the existing academic ranks. By the very nature of being an early adopter of this new technology and performing the pioneering research in its use, these academicians define the new subspecialty and serve as benefactors to the radiologist community at large.

The problem with imaging informatics is that research and education have not kept up with either the speed of development or the rapid imple-

mentation of the technology. Modality vendors market and sell multidetector CT scanners, for example, to customers who have almost no understanding of the fundamental change in the application and the challenges that will inevitably result. To date, little if any research has been performed to identify best-practice guidelines for this technology as it relates to optimal imaging plane, slice thickness, compression algorithm, or default display—essential elements in interpretation quality and effective workflow. The same vendors that sell these scanners also offer the same customers elaborate advanced image-processing workstations, but in many cases this only adds a layer of confusing options to the review/interpretation dilemma. Who is the responsible party for image processing and reconstruction? Is this the role of the technologist or radiologist? How can decision support software be incorporated into the process to facilitate improved diagnosis and workflow? Appropriate research has not been performed to answer these basic questions, and the result is that many imaging practitioners have expensive, data-intensive technologies but little understanding of their optimal use.

The same challenges within the imaging departments also extend throughout the filmless/paperless medical enterprise. With the implementation of the EMR and telemedicine, new challenges exist in terms of patient confidentiality, data security, and best-practice standards. These clinical issues go beyond the expertise and training of IT staff and must largely rest in the hands of medical experts. What other medical group is better equipped to address these challenges than radiologists? The key, of course, is to ensure that the radiologist community is well educated in all aspects of imaging informatics and takes a proactive role in related research.

## MAKING THE CASE FOR INFORMATICS RESEARCH

The good news, however, is that nature abhors a vacuum, which creates unlimited opportunity for the willing (and naive). Residents and fellows in training have the unique opportunity to quickly become experts in this burgeoning discipline. A number of important research opportunities exist with great relevance and potential significance to the radiology community at large.

To date, the relative lack of awareness of imaging informatics within the radiology community has extended to radiology residents, whose role is to become the thought leaders of the profession in the near future. It is

therefore critical that radiologists in training understand the clinical, economic, and political importance of this emerging subspecialty. To accomplish this, imaging informatics education and research must be prioritized within academic radiology programs. Imaging informatics research has the dual role of legitimizing the subspecialty as well as solidifying its place within the radiology domain. An intimate synergy exists between technology and research within radiology. Whoever controls the research controls the technology. And whoever controls the technology can dictate its clinical application. In the past, other medical specialties have effectively seized control of imaging technologies by co-opting research applications. One example can be seen in cardiac angiography and echocardiography, which were once the domain of the radiologist and are now firmly under the control of the cardiologist. It is therefore imperative that all radiologists, especially those in training, understand the long-term implications of what is at stake and the tremendous potential for those who embrace the challenges that lie ahead.

The best method to inform and create understanding of imaging informatics is to educate as many participants in the field as possible. Imaging informatics must become an integral part of resident education, and institutions and radiology departments must modify radiology residency curricula to incorporate informatics education. A suitable curriculum would describe the expected training issues through a 4-year radiology residency. Knowledge and experience should be imparted through didactic lectures, informal review of topics during reading sessions, review of journal articles, outside conferences, and self-study programs. Many informatics topics overlap directly with physics instruction and other imaging subspecialties, where modifications to existing curricula should be considered. Included here is a sample 4-year imaging informatics curriculum currently in place in the radiology department at Geisinger Medical Center (Danville, Penn.) (Table 25.3). This curriculum adheres to the ACGME guidelines for dividing the course of study to cover patient care, medical knowledge, practice-based learning and improvement, interpersonal and communication skills, professionalism, and systems-based practice.

This outline is designed to provide guidance for development of informatics curricula as they relate to resident education. More detailed training is needed when a professional chooses imaging informatics as a career. The curriculum for such training should incorporate more detailed understanding of digital imaging technologies as well as computer sciences. Multiple advanced fellowship programs are available to provide such training, and a list of current imaging informatics training programs is available on the Web site for the Society for Computer Applications in Radiology.

TABLE 25.3

## Outline of Informatics Curriculum for Radiology Resident Education

**First Year of Training**

The resident entering the first year of residency should be encouraged to develop at least moderate general computer skills. Unlike clinical skills, the basic use of PACS, dictation, and information systems software can be performed without in-depth knowledge of structure and function. Because the use of these systems is essential for clinical tasks, the first-year curriculum should focus on practical use issues.

*Patient Care*

Basic use of the PACS workstation

Patient worklist functionality

Knowledge and use of RIS and HIS user interface

*Medical Knowledge*

Instructions on basic physics of MR, CT, general radiography, ultrasound, and nuclear medicine should be given and coordinated with the physics curriculum

*Practice-Based Learning and Improvement*

Use of the Internet as a decision support tool to search for knowledge from existing online education resources

*Interpersonal and Communication Skills*

Electronic communication skills, including personal and group e-mail

Basic instructions in use of Microsoft Word

Preparing and using PowerPoint for presentations

Clinical use of digital dictation, speech recognition, and structured reporting

Concise, meaningful, and accurate reporting of radiology findings

*Professionalism*

Patient privacy issues as they relate to electronic data

Ethical use of downloads

*System-Based Practice*

Use of RIS for accessing patient data

Use of HIS for accessing clinical information, messaging, patient orders, and progress notes

TABLE 25.3

*Continued***Second Year of Training**

Advanced and more sophisticated use of clinical PACS, dictation systems, and information systems should be stressed during the second year. Emphasis should be placed on increasing speed and efficiency in the use of these systems without distraction from image viewing. Basic informatics instructions should begin on multidetector CT issues, three-dimensional (3-D) workstation functions, and computed radiography/direct radiography and should be coordinated with the physics curriculum.

*Patient Care*

Clinical applications of 3-D/advanced image-processing workstation

Use of alternate PACS/human interface devices

*Medical Knowledge*

Principles of multidetector CT

Basic PACS architecture, central versus distributed files, thin-client, browser-based systems

Principles of network, bandwidth, and wireless connectivity

Principles of computed radiography

*Practice-Based Learning and Improvement*

Informatics as a subspecialty

Decision support: preparing and using personal and system teaching files

*Interpersonal and Communication Skills*

Reporting skills with emphasis on context, conciseness, and readability

Use of macros for radiology reporting

*Professionalism*

Knowledge of computer viruses and other potential issues related to public computer use

*System-Based Practice*

DICOM basics

Integration of PACS, speech recognition, RIS, and HIS

---

---

*(Continued)*

**TABLE 25.3**

Outline of Informatics Curriculum for Radiology Resident Education  
(Continued)

**Third Year of Training**

Instruction on informatics topics should be continued at the intermediate level. It is expected that knowledge and understanding of the underlying principles of informatics will enhance clinical use of workstations.

*Patient Care*

Principles of digital mammography

Decision support: use of computer-assisted diagnosis in mammography

*Medical Knowledge*

PACS storage issues: RAID, pre-fetch algorithms, future storage applications, cost comparisons

Principles of direct radiography

*Practice-Based Learning and Improvement*

Advantages and disadvantages of PACS, speech recognition, and structured reporting

Economic issues in PACS and informatics

*Interpersonal and Communication Skills*

Principles of speech recognition and structured reporting: advantages and disadvantages

*Professionalism*

HIPAA regulations as they pertain to radiology practice

Audit principles in PACS, RIS, and HIS

*System-Based Practice*

Basics of the IHE profiles

Failover strategies for image interpretation and reporting

Principles of teleradiology

**TABLE 25.3***Continued***Fourth Year of Training**

Advanced informatics topics should be taught in continuum from the previous year. Many topics in the fourth year can be customized for those residents with special interest in a specific clinical area, research topic, or informatics as a subspecialty.

*Patient Care*

Workstation design and ergonomics

Image display principles

Decision support: computer-assisted detection for chest CT

*Medical Knowledge*

Image compression principles, in particular as they relate to clinical interpretation

Advanced direct radiography applications: dual-energy subtraction, temporal subtraction, tomosynthesis

Disaster recovery of images and other data

*Practice-Based Learning and Improvement*

Research opportunities in informatics

Future directions of image storage and interpretation

*Interpersonal and Communication Skills*

Future efficiencies in radiology reporting

*Professionalism*

Biometric security: log-in and password issues

*System-Based Practice*

Evaluation of PACS and other software for purchase

Use of requests for proposal

Principles of HL7 as these relate to imaging issues

Operation issues for healthcare IT as they relate to imaging



## CONCLUSION

Imaging practitioners today are faced with the important task of making critical decisions concerning technology implementation and integration, radiologist workflow, imaging economics, and standards for quality and security. Demands for imaging services are at an all-time high, and the supply of qualified radiologists cannot keep up. The common denominator to addressing these needs is technology, which provides radiologists with an opportunity to improve quality, timeliness, and security of services. Unfortunately, technology is a double-edged sword. It must be used judiciously to succeed in maximizing potential benefit, but if used inappropriately might actually be deleterious.

Although radiology has always been technology driven, many of the new digital applications in medicine go beyond the experience and expertise of practicing radiologists. If radiology is to maintain a technology leadership role within the medical enterprise, additional computer-based training and expertise are required. New technologies will play an important role in future radiology practice and will require dedicated research to enhance integration and application of information technologies within the clinical imaging practice. This provides the foundation for imaging informatics as a new and uniting radiology subspecialty.

## REFERENCES

- Bartholmai BJ, Erickson BJ, Hartman TE, et al. The electronic imaging technology specialist: the role of a new radiology subspecialty for the 21st century. *J Digit Imaging*. 2002;15(suppl 1):184–188.
- Bernardini A, Alonzi M, Campioni P, Vecchioli A, Marano P. IHE: Integrating the Healthcare Enterprise, towards complete integration of healthcare information systems. *Rays*. 2003;28:83–93.
- Bramble JM, Insana ME, Dwyer SJ 3rd. Information retrieval for teaching files: a preliminary study. *J Digit Imaging*. 1990;3:164–169.
- Branstetter BF, Siddiqui KM, Lionetti DM, Chang PJ. Defining a digital teaching file workflow: specifications for software development. *J Digit Imaging*. 2003; 16:37–40.
- Carr CD, Moore SM. IHE: a model for driving adoption of standards. *Comput Med Imaging Graph*. 2003;27:137–146.
- Channin DS. Integrating the healthcare enterprise: a primer. Part 6: the fellowship of IHE: year 4 additions and extensions. *Radiographics*. 2002;22:1555–1560.

- Chronaki CE, Zabulis X, Orphanoudakis SC. I2Cnet medical image annotation service. *Med Inform (Lond)*. 1997;22:337–347.
- Flanders AE, Carrino JA. Understanding DICOM and IHE. *Semin Roentgenol*. 2003;38:270–281.
- Goldberg DJ, DeMarco JK, Parikh T. Internet-based interactive teaching file for neuroradiology. *AJR Am J Roentgenol*. 2000;175:1371–1373.
- Halsted MJ, Perry LA, Cripe TP, et al. Improving patient care: the use of a digital teaching file to enhance clinicians' access to the intellectual capital of inter-departmental conferences. *AJR Am J Roentgenol*. 2004;182:307–309.
- Health Informatics in New Zealand. Available at: <http://www.hinz.org.nz>. Accessed July 15, 2004.
- Henderson B, Camorlinga S, DeGagne JC. A cost-effective Web-based teaching file system. *J Digit Imaging*. 2004;17:87–91.
- IHE (Integrating the Healthcare Enterprise). Available at: <http://www.rсна.org/ihe/>. Accessed July 15, 2004.
- IHE integration profiles: guidelines for buyers. *Radiol Manage*. 2001;23:37–40.
- Khorasani R, Lester JM, Davis SD, et al. Web-based digital radiology teaching file: facilitating case input at time of interpretation. *AJR Am J Roentgenol*. 1998;170:1165–1167.
- Lim CC, Yang GL, Nowinski WL, Hui F. Medical Image Resource Center—making electronic teaching files from PACS. *J Digit Imaging*. 2003;16:331–336.
- Macura KJ, Macura RT, Morstad BD. Digital case library: a resource for teaching, learning, and diagnosis support in radiology. *Radiographics*. 1995;15:155–164.
- Markivee CR. Automated filing of “teaching” radiographs under the ACR Index of Roentgen Diagnosis. *AJR Am J Roentgenol*. 1990;154:659.
- Nagy P, Leblanc DA, Bartholmai B, Branstetter BF, Siddiqui KM. The Web Warrior's Workshop. 2004. Available at: <http://www.scarnet.net/2004presentations.html/>. Accessed July 15, 2004.
- Piraino D, Recht M, Richmond B. Implementation of an electronic teaching file using Web technology. *J Digit Imaging*. 1997;10:190–192.
- Radiological Society of North America. MIRC committee publications. 2004. Available at: <http://www.rсна.org/mirc/index.html>. Accessed July 15, 2004.
- Radiological Society of North America (RSNA). MIRC documentation index. 2004. Available at: <http://mirc.rсна.org/mircstorage/documents/documentation/MIRCdocumentationindex.xml>. Accessed July 15, 2004.
- Radiological Society of North America. MIRC Web site. 2004. Available at: <http://mirc.rсна.org/>. Accessed July 15, 2004.
- Reiner B. Introduction: researching the impact of information technologies on radiology technologists: “taking it to the streets.” *J Digit Imaging*. 2002;15:118–120.
- Rode D. Understanding HIPAA transactions and code sets. *J AHIMA*. 2001;72:42–49.

- Rodriguez TA. HIPAA privacy regulations clarified: let calm prevail. *J Med Pract Manage.* 2003;19:61–66.
- Rosset A, Muller H, Martins M, Dfouni N, Vallee JP, Ratib O. Casimage project: a digital teaching files authoring environment. *J Thorac Imaging.* 2004;19:103–108.
- Rosset A, Ratib O, Geissbuhler A, Vallee JP. Integration of a multimedia teaching and reference database in a PACS environment. *Radiographics.* 2002;22:1567–1577.
- Seshadri SB, Arenson RL. The impact of PACS on research and education. *Int J Biomed Comput.* 1992;30:263–266.
- Seshadri SB, Arenson R, Khalsa S, Brikman I, Voorde F. Prototype Medical Image Management System (MIMS) at the University of Pennsylvania: software design considerations. *J Digit Imaging.* 2003;16:96–102.
- Siddiqui KM. A practical guide to planning and implementing privacy and security systems. In: Dwyer SJ III, Reiner BI, Siegel EL, eds. *Security Issues in the Digital Medical Environment.* 2nd ed. Great Falls, VA: Society for Computer Applications in Radiology; 2004:49–58.
- Siegel E, Reiner B. Electronic teaching files: seven-year experience using a commercial picture archiving and communication system. *J Digit Imaging.* 2001;14:125–127.
- Siegel EL, Reiner BI, Siddiqui KM. Clinical impact of security on radiology. In: Dwyer SJ III, Reiner BI, Siegel EL, eds. *Security Issues in the Digital Medical Environment.* 2nd ed. Great Falls, VA: Society for Computer Applications in Radiology; 2004:1–8.
- Siegel EL, Diaconis JN, Pomerantz S, Allman R, Briscoe B. Making filmless radiology work. *J Digit Imaging.* 1995;8:151–155.
- Siegel EL, Perry J, Langlotz C, Reiner BI, Siddiqui KM. The RSNA's MIRC and RadLex initiatives: 2004 update. 2004. Paper presented at: 21st Meeting of the Society for Computer Applications in Radiology; May 2004; Vancouver, British Columbia, Canada.
- Sinha S, Sinha U, Kangarloo H, Huang HK. A PACS-based interactive teaching module for radiologic sciences. *AJR Am J Roentgenol.* 1992;159:199–205.
- Society for Computer Applications in Radiology. Radiology Informatics Fellowship List. 2004. Available at: <http://www.scarnet.net/residents/Fellowshiplist.html>. Accessed July 15, 2004.
- Vanderbilt University Medical Center. What is Biomedical Informatics? Available at: <http://www.mc.vanderbilt.edu/dbmi/informatics.html>. Accessed July 15, 2004.
- Weiss DL, Siddiqui KM, Scopelliti J. Radiology residency informatics curriculum: incorporating the six general competencies of the ACGME. 2004. Paper presented at: Radiological Society of North America 90th Scientific Assembly and Annual Meeting; November 28–December 3, 2004; Chicago, Ill.

# TELERADIOLOGY

---

---

GILES BOLAND • JONATHAN T. SCHLAKMAN •  
JAMES H. THRALL

**T**he conversion from analog to digital methods in the medical imaging world and the emergence of widely available mechanisms to quickly and affordably transmit digital data over large distances have fueled the rapid growth of teleradiology. In modern radiology departments in both the academic and private sectors it has become commonplace to select the location of a fully functioning picture and archiving communications system (PACS) workstation based on manpower and workflow considerations rather than proximity to the site of image acquisition. In addition, affordable scaled-down teleradiology solutions have allowed referring physicians and radiologists to access imaging from the convenience of their homes and offices.

This chapter summarizes some of the key issues in image acquisition, transmission, and interpretation. These summaries are followed by a discussion of the clinical practice of teleradiology, sharing knowledge gained from practical experience that can be useful in helping a teleradiology program meet its full potential. Diagrams of a typical teleradiology system are shown in Figures 26.1 and 26.2.

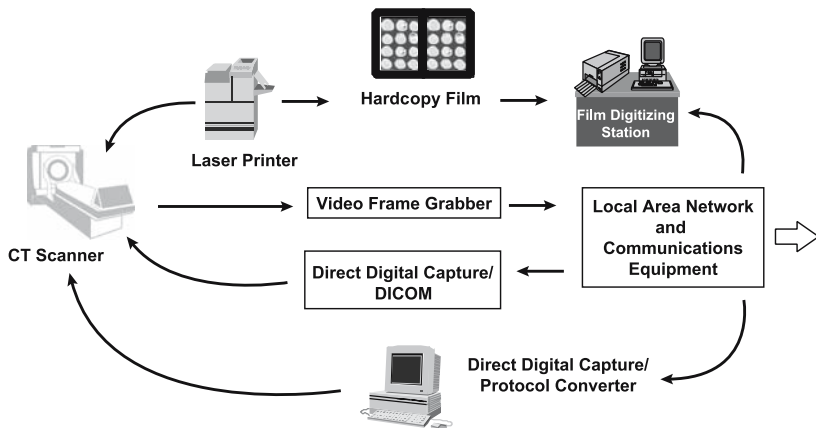


FIGURE 26.1

Schematic diagram of different approaches to image acquisition in teleradiology.

## IMAGE ACQUISITION AND IMAGE DIGITIZATION

Although any type of medical image may be transmitted by teleradiology, all images must be in a digital form before transmission can occur. Conventional hardcopy images from any modality can be digitized by special high-

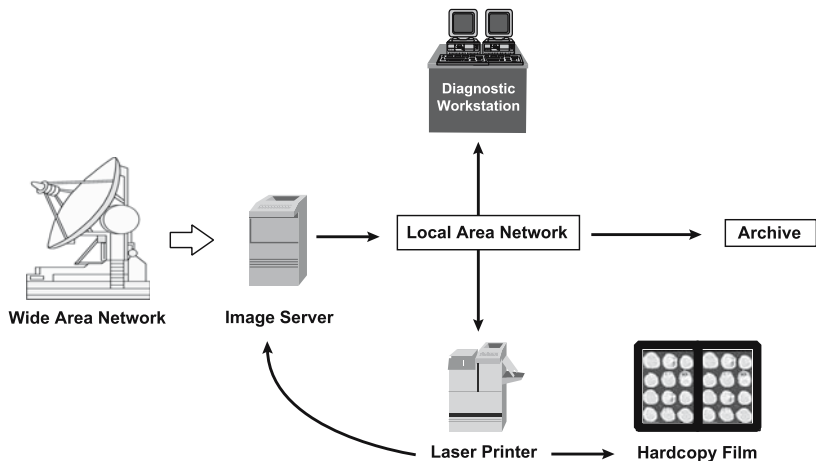


FIGURE 26.2

Schematic layout of a teleradiology interpretation site.

resolution laser or charge-coupled device (CCD) scanners (Figure 26.1). Simplistically, film scanners for teleradiology function similarly to fax machines by scanning analog data and converting it into digital form. Charge-coupled device scanners using the same technology as video cameras have tiny photocells that acquire data from a transilluminated film. Laser scanners offer better signal-to-noise ratios than CCD scanners, leading to superior contrast resolution. They are, however, more expensive. Additional studies are required to determine whether the superior contrast resolution offered by laser scanners is diagnostically significant.

Another alternative to image digitization of hardcopy films is the use of a video camera to view a film on a light box. The analog video signal can then be converted into a digital format. This approach was very common in the early days of teleradiology but has been largely abandoned due to inadequate image quality.

Many images are inherently digital: computed tomography, magnetic resonance, ultrasound, nuclear medicine, computed radiography, digital radiography, and digital fluoroscopy. All can be directly linked to a teleradiology system if they are in a standard format (Figure 26.1). Fortunately, more and more imaging devices are complying with the ACR-NEMA (American College of Radiology and National Electrical Manufacturers Association) DICOM 3 standard (Digital Imaging and Communications in Medicine—version 3). The DICOM 3 standard is important to teleradiology because a direct digital connection can be made from the image source to the teleradiology server and then from the teleradiology-receiving computer to a diagnostic workstation. Furthermore, DICOM 3 offers no loss of the full 12-bit dataset (2056 grayscales) generated from digitally acquired images, exhibits no image degradation, and has full capability to adjust image window and level settings. The DICOM 3 standard is a great step forward, although implementation of the standard is still not uniform. Some manufacturers may be DICOM compliant, but the standard allows enough flexibility that implementations can vary from one manufacturer to another. However, these incompatibilities are usually resolved with the cooperation of the respective vendors.

In practice, many digitally acquired images cannot be directly linked to teleradiology systems because they are acquired on older equipment that is not DICOM compliant. Different manufacturers of imaging equipment historically used proprietary file formats and communications protocols, which prevent direct interfacing to communications networks. These digitally acquired data need to be converted into the standard DICOM 3 format before image transmission can occur (Figure 26.1).

Various solutions are available to convert proprietary digital data into an acceptable form for transmission over a teleradiology network. One of the most commonly used methods is to simply take the hardcopy rendition of a digital modality and digitize the image with a laser or CCD digitizer. Although this is not scientifically elegant, it overcomes a number of problems. However, it is less desirable because of the image degradation that invariably occurs during an analog-to-digital conversion. Another alternative is to use a video frame grabber wherein the video signal output that is sent to an imaging console is converted to digital form (Figure 26.1). Such devices are commonly used to connect ultrasound machines to PACS, even today. It is also possible to use a protocol converter, which is a special computing device that converts proprietary image data to the DICOM 3 compliant format.

## IMAGE COMPRESSION

File sizes for typical digitized medical images are large (Table 26.1). Transmission of this volume of data requires significant bandwidth (the capacity of a communication medium to carry data). Therefore, these file sizes may be too large for teleradiology to be effective, both practically and economically. To reduce the amount of digital data to be transmitted, the digital data can be compressed prior to transmission.

Table 26.2 illustrates the potential benefits of compression. Using the example of a typical radiographic exam with 4 images, the table shows that

**TABLE 26.1**  
File Sizes for Digital Medical Images

<b>Modality (MB)</b>	<b>Image Matrix</b>	<b>Study</b>	<b>Images/File Size</b>
Mammography	2294 × 1914 × 12	4	32
Plain radiographs	2048 × 2048 × 12	4	25
Fluoroscopy	1024 × 1024 × 8	18	19
Computed tomography	512 × 512 × 12	30	12
Magnetic resonance imaging	256 × 256 × 12	100	10
Ultrasound	256 × 256 × 8	24	1.6
Nuclear medicine	128 × 128 × 8	24	0.4

**TABLE 26.2**  
Effect on Transmission Times (minutes) of Compression for  
Digitized Radiographs (25 MB)

Compression Ratio	Bandwidth				
	14.4 Kb/s <sup>†</sup>	28.8 Kb/s	56 Kb/s	1.45 Mb/s (T1)	155 Mb/s (ATM)
1:1	231.5	115.7	58.8	2.2	0.02
2:1	115.7	57.8	29.4	1.1	0.01
10:1	3.1	11.5	5.9	0.22	0.002
20:1	11.5	5.7	3.0	0.11	0.001
30:1	7.7	3.8	2.0	0.07	0.0006
60:1	3.8	1.9	1.0	0.04	0.0003

<sup>†</sup> Kilobits per second.

it takes 231.5 min to transmit at 14.4Kb/s, the bandwidth of most older standard phone lines and modems. This is completely impractical, and teleradiology systems use both compression and access to higher bandwidths (Table 26.3) to reduce transmission times to manageable levels.

Compression can be “lossless” (reversible), with compression ratios typically in the range of 3:1, and the original dataset can be fully regenerated, or “lossy” (irreversible), where much higher compression ratios are possible. Compression ratios of at least 10:1 are generally required before data compression can have a significant economic effect. Although lossy compression requires some loss from the original dataset, several studies have shown that compression ratios of 20:1 or higher can be achieved without sacrificing diagnostic image content.

The Joint Photographics Experts Group (JPEG) standard is the most commonly employed technique, but a number of alternatives exist, including wavelets, which have proven to be a practical and powerful compression tool. Both JPEG and wavelets can be used in either a lossy or lossless mode. JPEG compression is the only technique currently supported by the DICOM 3 standard. JPEG’s principal advantages are that it is inexpensive, widely acceptable to most computing platforms, and implemented in both hardware and software. It does, however, suffer from “block artifacts” (artificial edges created between pixel blocks to which the human eye is



**TABLE 26.3**  
Effect of Bandwidth on the Transmission Times of Medical Images

Modality	Bandwidth				
	14.4 Kb/s*	28.8 Kb/s	56 Kb/s	1.544 Mb/s (T1)	155 Mb/s (ATM)
Plain radiographs	231.5 <sup>†</sup>	115.7	58.8	2.2	0.02
Fluoroscopy	176	88	44	1.6	0.02
Computed tomography	111	55.5	27.8	1.0	0.04
Magnetic resonance imaging	94	47.5	23.5	0.9	0.01
Nuclear medicine	6.4	3.2	1.6	0.03	0.0003

\* Kilobits per second.

<sup>†</sup> All transmission times in minutes.

sensitive), particularly at higher compression ratios. Wavelet compression may be advantageous at higher compression ratios, which may be required for implementation of a teleradiology system that is both practical and economical. The DICOM working group on compression (Group 4) is evaluating alternative compression techniques such as wavelets for inclusion in the standard.

A major advantage of wavelet compression over JPEG compression is that it permits substantially higher compression ratios while maintaining image quality. This has practical implications for high-volume teleradiology, particularly from international sites, where the cost of data transmission becomes a significant factor in the overall cost of the teleradiology system. Several studies have now confirmed that compression ratios of up to 20:1 are diagnostically acceptable. Film transmission times can therefore be significantly reduced, thereby permitting images to reach the referral site within a fraction of the time at significantly reduced costs (Table 26.2). Real-time teleradiology therefore becomes a reality, enabling remote consultations while the patient is still in the doctor's office or emergency room (ER) at the remote site.

Wavelet compression algorithms consist of 3 basic functions: (1) transformation, (2) quantization, and (3) lossless coding. Once discrete wavelet transformation has occurred, the image is processed using quantization, which is the lossy step in the compression algorithm. Care must be taken at

this stage not to lose image quality because of the lossy step in the process. Finally, lossless coding is performed, which removes redundant information from the dataset. Once compressed, the image is transferred to the remote site, where it requires decompression before it can be interpreted. The lossless step that was performed as the last step in image compression is reversed. Quantization cannot be reversed, as it is a lossy function. Finally, the image is reconstructed using an inverse transformation process, which is then available for image interpretation.

Although image degradation occurs at higher compression ratios using wavelet compression, these artifacts produce less image degradation than JPEG artifacts at similar compression ratios. Practical image quality is maintained up to compression ratios of 20:1 and even higher.

## IMAGE TRANSMISSION

Several forms of transmission media exist with different bandwidths (Table 26.3), including: conventional dial-up telephone lines (9.6 to 28.8 kilobits per second [Kb/s]), switched digital service (56Kb/s), frame relay (up to T1), integrated services digital network, or ISDN (128Kb/s), T1 lines (1.54Mb/s), digital signal level 3, or DS3 (44.736Mb/s), and asynchronous transfer mode, or ATM (typically 155Mb/s). Although ATM can be extremely fast, it is expensive. As broadband has penetrated the home market, most homes in the developed world can access the Internet via either digital subscriber line (DSL) over a telephone wire or over a television cable, with speeds varying from 750Kb/s up to 10Mb/s. However, actual transmission speeds depend on the total route the data need to travel; the longer the distance, the longer it will typically take. It also varies with the amount of competing traffic on the network segments involved. Dedicated T1 and ISDN lines are variably deployed and may be expensive for small-volume teleradiology. Ultimately, the choice of line depends on the customer's needs, the volume of studies to be transmitted, the types of studies, turnaround times, and expected peak activity. If all the images are transmitted during one period of the day, then higher bandwidth may be required. As indicated in Table 26.3, if sufficient bandwidth is available, transmission times become quite short even for uncompressed large image files.

Transmission of digitized data requires communication equipment. The nature of the equipment depends on the communication medium being used. This may be a modem for conventional telephone lines, a terminal adapter for an ISDN line, a channel service unit (CSU) for a T1 line, a DSL modem, or a cable modem.

## IMAGE INTERPRETATION

Once images are received from the wide area network (WAN) at the interpretation site, they can be sent directly to an interpretation workstation or to an image server that permits distribution within the institution (Figure 26.2). Archiving or storing the images for long periods may not be required, as is necessary for PACS. Provided images are online for several days to handle delays in processing (i.e., weekends and holidays) or to respond to clinical questions pertinent to the case; longer-term storage usually is not necessary because the images are typically archived in the department of origin.

The type of monitor used to read the studies depends on whether full primary readings are being performed or the system is being used for on-call emergencies. The American College of Radiology (ACR) recommends a minimum resolution of  $2000 \times 2000$  pixels  $\times$  12 bits for image acquisition for conventional radiographs. However, this does not mean that the monitor must be capable of displaying all the pixels of the image at once. In fact, the de facto standard in radiology departments for the display of 5-megapixel (MP) radiographs is a 3-MP display. Even on 5-MP displays, 5-MP radiographs are almost never displayed at their full native resolution. They are almost always minified to make room for other elements of the screen such as buttons and menus, and when a landscape-oriented portable chest x-ray is displayed on a portrait-oriented display (a common occurrence), the image is shrunk even more to make it fit. The requirements for teleradiology are certainly not more stringent than those used in the radiology department. However, even 3-MP monitors are expensive and may be impractical for general home use when on call. For lower-resolution images (CT, MRI, ultrasound, and nuclear medicine),  $512 \times 512$  pixels  $\times$  8-bit resolution is adequate, for which most monitors used for home personal computers suffice.

## ACCURACY OF TELERADIOLOGY

Improvements in teleradiology systems and the imperatives of contemporary medical practice have led to the widespread use of teleradiology. Two of the major concerns by early adopters have been whether radiologists would accept softcopy review of images on computer workstation screens versus traditional hard copy, and whether the accuracy of interpretation was sufficient to justify deployment of clinical teleradiology. Both questions have

now been answered to variable degrees. First, radiologists have readily adapted to computer viewing to the point that many departments are using it not only for teleradiology cases but also for routine work. At Massachusetts General Hospital (MGH), substantially all modalities, including plain radiography, fluoroscopy, angiography, CT, MRI, ultrasound, and nuclear medicine, are routinely acquired, viewed, interpreted, and managed in a digital electronic environment. In most cases, hardcopy films are no longer printed, and there is a growing consensus that in many applications the ability to optimize image contrast and intensity greatly facilitates the radiologist's practice.

With respect to accuracy of interpretation of softcopy versus hardcopy, the issues are more complex. For the relatively low-resolution cross-sectional modalities, including CT, MRI, ultrasound, and nuclear medicine, the accuracy of softcopy interpretation has, for all practical purposes, not been an issue. It has not been the subject of intense clinical research investigation. Radiologists have simply switched with a collective subjective impression of equal or superior interpretive accuracy. The relatively low resolution of these studies is easily encompassed by the resolution of commonly available monitors. The ability to use flexible viewing formats, such as "stack mode," which allows the radiologist to move quickly back and forth through a series of tomographic sections, has supported the change in practice from hardcopy to softcopy viewing of these modalities. In the teleradiology application, softcopy viewing of these modalities is no different remotely than locally, once acquisition of the original digital datasets has been accomplished.

With respect to conventional radiographic images, the question of accuracy is less clear. A significant body of literature compares accuracy of interpretation of original hardcopy radiographs with digitized radiographs viewed in softcopy mode. As pointed out by Larson et al., among others, there is an inevitable loss of spatial and contrast resolution in the digitization process, with some subjective decrease in image quality. However, in their study of the sensitivity for detecting nodules, pneumothoraces, interstitial lung disease, and fractures, there were no statistically significant differences between the original hardcopy analog radiographs and the digitized images viewed on a  $1280 \times 1024$ -pixel matrix monitor equipped with an 8-bit/pixel grayscale display. An important point in Larson et al.'s study is that the overall accuracy for detection of conditions such as nodules and pneumothoraces is less than optimal in the first place, averaging 60% and 77%, respectively. These numbers are in keeping with the literature and indicate that the real issue is finding more accurate ways of detecting subtle abnormalities regardless of viewing method. Computer-assisted diagnostic

techniques will probably be required to boost detection of subtle nodules and pneumothoraces from historic levels.

Another problem in assessing the accuracy of softcopy interpretation of radiographs is the low resolution and low quality of much of the equipment, including digitizers and workstations used for these analyses. For example, papers by Ackerman and Scott demonstrated statistically lower detection rates for digitized radiographs versus analog radiographs for pneumonia and fractures. However, the digitizer used in their studies had a spot size of 210  $\mu\text{m}$  compared with the 100- $\mu\text{m}$  spot size used in Larson et al.'s work. Also, the monitor resolution was either 1280  $\times$  1024 or 1200  $\times$  1600 pixels. Each of these is lower than the minimum standard recommended by the American College of Radiology of 2.5 line pairs per millimeter (lp/mm) at a 10-bit depth. This line-pair resolution requires a equates to a roughly 2000  $\times$  2500-resolution matrix for a 14  $\times$  17-in radiograph.

Optimistically, some studies with high-resolution monitors have demonstrated substantial equivalency between hardcopy and softcopy interpretation. Goldberg's prospective study of 685 cases of double-read softcopy and hardcopy produced overall agreement in 97% of cases without a statistical advantage for either approach. Ten cases were judged to be false negative by softcopy interpretation, 3 cases were judged to be hardcopy interpretation errors, and 1 case was unresolved as observer variation. In Rajavi's series of 239 pediatric cases, no significant difference was found between hardcopy interpretation and use of high-resolution 2000  $\times$  2000-pixel softcopy images on a CRT monitor. The clearly agreed-upon exception to softcopy viewing is mammography, which requires higher spatial resolution than is available on current monitors.

The major problem in sorting through the literature on the accuracy of softcopy interpretation of radiographs is the lack of standardization. Digitizers of widely varying spot size and workstations with different spatial resolutions, luminosities, and functional characteristics have been used, and there is no widespread use of standard sets of reference images, making true comparisons between observers difficult. Work on defining the necessary parameters for optimal softcopy viewing continues. In the meantime, the practice of softcopy interpretation of radiographs is rapidly becoming widespread as departments of radiology become increasingly digital. There is also no question that the quality of radiographs obtained by computed radiography or direct digital capture are superior to radiographs obtained in original analog hardcopy format and secondarily digitized by a CCD or laser digitizer. The optimal configuration for teleradiology of all modalities is direct digital capture with transmission of the full digital dataset.

## APPLICATIONS OF TELERADIOLOGY

Table 26.4 summarizes current clinical, research, and educational uses of teleradiology. Although exact numbers are not available, several thousand teleradiology systems are deployed in the United States. Many of these, if not the majority, are “lower end” or entry-level systems used by radiologists to provide emergency or on-call coverage for their practices. The home workstation is often a personal computer, so that CT, MRI, ultrasound, and nuclear medicine images up to  $512 \times 512$  pixels or somewhat more can be

---

---

**TABLE 26.4**  
Applications of Teleradiology

1. On-call coverage—Emergency/24-hour
    - Hospital to home
    - Inter- and intrainstitutional
  2. Primary Interpretations
    - Freestanding imaging centers
    - Rural hospitals and clinics
    - Imaging centers within regional delivery systems
    - Nursing homes, other special care facilities
  3. “Reverse” teleradiology
  4. Second opinions/consultations
  5. Access to subspecialty expertise at academic medical centers
    - Domestic and international consultations
  6. Image processing
  7. Utilization management
  8. Quality assurance
  9. Overreads by second radiologist
  10. Research
  11. Image data collection and management
    - Image analysis
  12. Teaching files
  13. Care presentations
  14. Online journals
- 
-

displayed at full resolution. These systems often allow image magnification so that plain films may be viewed at nominal full resolution according to the American College of Radiology standards, as noted. Full resolution is equivalent to roughly  $2000 \times 2500$  pixels.

On-call applications are proving extremely valuable in the clinical practice of radiology. Radiologists are able to provide more rapid consultations than would be possible if physical travel to the hospital were required. Radiologists can cover multiple institutions simultaneously, and subspecialists within a group can provide on-call coverage more flexibly. Many teleradiology companies have formed in recent years that offer emergency radiology coverage during overnight hours and weekends. These services can help radiologists face the challenge of meeting ever-increasing on-call demands in the setting of a relative shortage of radiologists. As round-the-clock service is impractical for solo practitioners or small groups, teleradiology offers the opportunity for contemporaneous interpretation without having to be on call an undue or impractical percentage of the time.

One of the best documented teleradiology programs in the United States providing services to rural hospitals and clinics is at the University of Iowa. Franken et al., have studied patterns of use and have tried to determine what value is added by radiologists' consultation versus having studies reviewed by nonradiologists. They have confirmed that higher accuracy is achieved by radiologists. For example, in one of their teleradiology series, radiologists demonstrated a 92% versus 86% sensitivity compared with family practitioners in interpreting studies of the chest and extremities.

Another growing application of teleradiology is coverage for free-standing imaging centers, outpatient clinics, nursing homes, and smaller hospitals. In these applications, common themes are improved coverage, improved access to subspecialist radiologists, and lower cost of service provided. Using the experience with the MGH as an example, teleradiology is used for all the purposes cited. Prior to teleradiology, if a clinician in a neighborhood health center needed an immediate interpretation, films were put into a taxicab and taken to the MGH main campus. This cumbersome and expensive solution was never satisfactory and is being replaced by teleradiology, with essentially real-time, online interpretative services becoming available throughout the MGH service region. Routine interpretations from the affiliated imaging centers had been available only through batch reading of cases brought twice a day to the MGH main campus or one of the larger satellite centers. Now, every case can be read immediately by a subspecialist radiologist. The MGH experience highlights the observation that teleradiology will be used extensively by regional integrated delivery systems to improve radiology coverage.

The rapid growth in freestanding imaging centers, often dedicated to cross-sectional imaging, has provided another impetus for teleradiology. Hundreds of these centers across the country have been established by non-radiologist entrepreneurs, who then approach the radiology community for professional interpretation support. Teleradiology is an efficient way for groups of radiologists to work with these imaging centers to expand their practices and benefit from the growth in demand for imaging studies without unduly disrupting the logistics of their practices, especially in the case of hospital-based radiology groups.

A number of institutions, including MGH as well as the US military, are exploring the concept of "reverse teleradiology" by sending cases from larger centers to smaller centers. In some situations, small departments require a radiologist to do procedures or to meet requirements for direct supervision but do not generate enough work to keep them busy or justify their cost. By sending cases from busier departments, the presence of the on-site radiologist can be cost justified, with the added benefit of improving the quality of care offered at the smaller location. At MGH, this model provides coverage for a series of high-tech imaging centers that the department has established in the community surrounding Boston. This allows the department to place a staff radiologist at each location without loss of productivity.

Teleradiology also provides access to subspecialty expertise at academic medical centers or to second opinions and consultations among larger radiology groups. This practice is growing both domestically and internationally and is being highly influenced by the Internet. Patients, physicians, third-party payers, and government agencies all need second opinions from time to time. Web sites on the Internet are being established to solicit and accommodate these opinions. In other cases, academic practices, larger radiology groups, and commercial enterprises dedicated to teleradiology are making their services available for second opinions and consultations.

The increasing importance of image post-processing suggests another application of teleradiology that has been used in the early days of CT and may become common. Smaller facilities without on-site access to three-dimensional rendering software or personnel with image-processing expertise can readily access both via teleradiology. Multisite institutions send image sets to central processing.

The possibility for radiology groups to help each other in their quality assurance programs is another potential use of teleradiology. In this model, a statistical sampling of cases from 1 group can be sent to another group via teleradiology for a quality assurance overread. This model is also useful for helping radiology groups initiate services in a new modality. For example, if



a group that has not provided MRI services begins doing so, as they come up to speed they can use teleradiology to send cases to a more experienced group for an overreading function. This was a common practice in the early days of MRI, although the cases were often sent as hard copy.

Research and educational applications of teleradiology are less well publicized than direct clinical applications. The radiology department at MGH has used teleradiology to manage clinical trials in which imaging is an important part of the data collection. This application is likely to grow dramatically with the increased use of imaging-based surrogate endpoints in drug trials and the establishment of the ACRIN (American College of Radiology Imaging Network) by the American College of Radiology for the purpose of organizing clinical trials aimed at assessing the efficacy of imaging technologies.

Educational applications of teleradiology are burgeoning over the Internet. Many academic departments have Web sites with extensive teaching files. These files are often augmented by case presentations. Some sites offer radiologists the opportunity for online continuing medical education credits. The American Board of Radiology is exploring the possibility of administering its certifying examinations electronically. Major journals have extensive Web sites providing a range of services for their subscribers and even for nonsubscribers, and it is likely that the applications of remote education will become among the most prevalent uses of teleradiology.

## MEETING THE GOALS OF TELERADIOLOGY

Ultimately, the purpose of off-site image interpretation is to provide better and faster delivery of patient care. When organizing a teleradiology program, it is important to take the necessary steps to ensure that this objective is met and that referring physicians, imagers, and patients are highly satisfied.

### SPEED

The ability to quickly move images from place to place lends itself to the provision of fast and responsive imaging interpretation. The implementation of a teleradiology program should be associated with improved delivery of services by those who utilize radiology most frequently.

A prerequisite for fast report turnaround times is the ability of the radiologist to access images quickly. This is crucial for teleradiologists

providing emergency coverage. An appropriate compression mechanism and degree should be chosen to maximize speed of transmission without diminishing practical image quality. The range of size of the typical exam files that will be sent should be taken into account. This can be dependent on the type of scanners utilized by a particular practice, as new, multislice CT scanners can generate very large datasets. The teleradiology system should be tested in advance at the site in which it will be stationed using the workstation and bandwidth connection that will be utilized. Testing should be done during the same time of day that the system will be used, as this can have a significant effect on speed at both the sending and receiving sites when certain Internet connectivity options are employed.

For the truly emergent patients it is very helpful to have a system that allows the radiologist to see the images as they are being downloaded locally onto a workstation. Some systems allow lower-resolution images to be visualized almost immediately while the case is in the process of downloading. This allows the radiologist to quickly visualize a crucial finding such as a ruptured aortic aneurysm in seconds and begin to take the necessary steps to ensure that prompt care is provided to the patient. It avoids the frustration of staring at a blank computer screen while a case is being transmitted after having been told by a technologist or trauma surgeon that a patient may have a critical finding.

Also important in emergency teleradiology is the ability to choose the parts of a case to download preferentially. In the age of multislice CT scanning, it is not uncommon to have a study that contains 750 or more images. If transmission speed is at all an issue given a particular bandwidth connection, it can be critically important to be able to download standard images in a single plane without waiting for the entire study to arrive. In this way a radiologist can begin to interpret an exam and spot most of the important findings during the time that images in multiple additional planes and specialized algorithms and delayed images are being downloaded in the background. Furthermore, a single computer error is less likely to stop the entire lengthy process and leave an interpreter without any images to view.

Speed can also be enhanced using a system that can be configured to display the most recently completed examinations at the top of a patient list. This allows a radiologist to begin to access a study before a technologist even notifies him or her of the case. In some settings this can save a significant amount of time. In the emergency setting, technologists should be encouraged to communicate with the radiologist as soon as there is a new study to be read and not wait to batch studies. Batching studies can be particularly troubling when using a teleradiology system that allows just 1 exam to be transmitted at a time. These steps can avoid the uncomfortable phone call

from a technologist informing the radiologist that there are 3 patients' studies to be read and there is an ER physician demanding results on the first patient, who returned from the CT scanner over an hour ago.

The ability to download multiple cases in parallel is an important component of a teleradiology system for workflow and patient care issues. At the worst end of the spectrum are systems that allow for the download of only 1 series of 1 study at a time. This means that only after reviewing the axial images of a cervical spine CT can a radiologist request the sagittal or coronal reconstructions and wait for them to arrive. Only at the completion of this process can the next trauma patient's exam be requested. An old exam likewise cannot be retrieved or reviewed simultaneously with a current exam on these systems. In practice, this will decrease the number of cases that will be compared to old exams and could impact negatively on patient care. In sharp contrast, other teleradiology systems are available that can download multiple exams or components of exams at the same time. On some of these systems, old exams and old reports are easy to locate and can be synchronized by table position with a current study for convenient comparison.

In choosing a mechanism to secure patient information and the hospital's network, it is important to balance security concerns with issues of speed and reliability. Cumbersome access mechanisms with frequent automatic terminations may provide security at the cost of valuable minutes in report turnaround times and more frequent downtimes. It is important to work closely with the information technology department to ensure adequate security, while at the same time allowing for convenient, reliable, and fast access by the teleradiologist.

## MANPOWER

Manpower is perhaps the most critical component for the consistent delivery of timely care. This can be relatively easy to calculate for elective imaging needs, but quite difficult for emergency coverage. By its very nature, emergency radiology comes in bursts. Maintaining the manpower requirements to provide immediate attention to potentially critical patients during these bursts can be very inefficient for the vast majority of time when the ER pace is not as frenzied. This can be true for small departments where even a single radiologist is not needed for most of the overnight hours and for larger departments where a single radiologist may have difficulty keeping up during the busiest spurts, but 2 radiologists would be unnecessary for the vast majority of the time.

Efficiency and responsiveness could be optimized if radiologists interpreting elective cases could be quickly recruited for emergency work when

needed. Though this is typical for many practices during the daytime, it can be highly impractical for most departments to achieve 24 hours per day. Over the last several years, some radiology practices have made this happen by using teleradiology to its full advantage. These practices have made the night shift attractive by establishing an office in an exotic or otherwise desirable part of the world where it is daytime during the overnight hours at the local hospital. Experts in emergency radiology modalities voluntarily relocate or rotate to this office for blocks of time while maintaining an affiliation with the home practice.

Newly formed teleradiology companies offer radiology departments the ability, in effect, to hire the portion of a dedicated night radiology workforce that they need. Ideally, this workforce should be composed of experts in ER radiology working with a schedule or in a location that is conducive to their being rested and alert. A potential disadvantage of this type of arrangement is the perception amongst colleagues of having “outsourced” the undesirable portion of the radiologists’ responsibilities. Additionally, studies done overnight are interpreted by radiologists unfamiliar to the referring doctors. These disadvantages can be overcome, however, if there is adequate communication between the night radiologists and the referring doctors and if higher-quality reports can be generated with faster turnaround times.

## RELATIONSHIPS WITH COLLEAGUES

The knowledge that radiologists are working from home or from a beach-front office in Hawaii does not necessarily impart an endearing first impression to nonradiologists that spend long hours each day in hospital wards. To maximize the success of a teleradiology practice, it is imperative to provide these doctors with benefits that they would not have expected from their radiology departments prior to the implementation of teleradiology. In an elective setting this can be done by maximizing the percentage of studies that are read by subspecialist radiologists. In an emergency setting, increased availability of a radiologist during off-hours will be appreciated. In many departments, referring physicians have grown accustomed to the limited availability of an on-call radiologist. Emergency room physicians quickly learn to appreciate having a radiologist colleague that is freely available for consultation, protocol questions, report clarifications, and interpretations of timelier follow-up examinations. Faster report turnaround times that can be facilitated with the use of voice recognition technology or on-site transcription will be welcomed in any clinical setting. Outpatient doctors that begin to receive same-day reports from subspecialists on patients sent to an

imaging center and ER physicians that start to get head CT results before their patients return from the scanner are more likely to write complimentary letters to a radiology chairman than to be bothered by where a radiologist may be sitting while the images are being interpreted.

Teleradiology should not result in weaker relationships between radiologists and colleagues in other fields. If referring physicians do not see a teleradiologist's face, effort should be made to ensure that they recognize their teleradiologist's voice and personality. In addition to the usual phone calls for unexpected or emergent findings, teleradiologists should consider other occasions to call the ordering doctors such as for interesting cases or for follow-up. Relationships build confidence, which is critical for any radiology department. This can also help improve the clinical histories that are made available for the teleradiologist and increase the amount of times that meaningful follow-up is received.

## PATIENT CARE

The benefits of teleradiology that result in better service will by their nature impact patient care. Ultimately, it is the patients who benefit most by having their imaging interpreted by a subspecialist. In many instances, teleradiology has put an end to common practices that used to be quite difficult for many patients. For instance, a modern teleradiology program should guarantee that ER patients no longer have to lie for hours in cervical spine collars waiting for their spine imaging to be reviewed by a radiologist in the morning. In the not-too-distant past this was a common occurrence in many community hospitals because it was difficult to quickly move large radiological examinations from place to place. In the most emergent settings, a well-run teleradiology practice will ensure that a radiologist is immediately available to turn his or her full attention to patients who have sustained significant trauma or present for other critical indications.

Night call has often been considered by many to be the least desirable part of a radiologist's practice. Many departments demand that all full-time members share this responsibility equally. This can mean that subspecialized radiologists who do not routinely interpret important emergency imaging modalities such as CT in the daytime are forced to provide coverage for these modalities at night. With no backup readily available, vital decisions on ER and trauma patients must be made by these radiologists all night long. Transmitting images to ER radiology experts who are in favorable time zones or are working dedicated shifts can greatly improve nighttime coverage. This can also prevent having an important patient examination sent to be read by

a radiologist who was just woken from sleep after having worked the previous 18 hours.

The capabilities of the teleradiology system can also have a surprisingly important impact on the quality of patient care. When making purchasing decisions it can be easy to reason that one will be able to make do with certain inconveniences while off-site. While this may be true for very low-volume practices, this is certainly not the case for busy ones. There are systems currently in use that require a combination of greater than 10 mouse clicks and button pushes to change from a standard to a subdural window. Over time, patient care will likely be compromised if a radiologist begins thinking in the midst of a busy stream of pending studies that standard brain windows are probably adequate in order to avoid the time and loss of concentration needed to change windows. If a teleradiology system is not capable of calculating a region of interest (ROI) measurement, it is extremely difficult to expect that a radiologist will telephone a technologist to have this done as frequently as may be appropriate. Similarly, features that are available but are clumsy or inconvenient may not be utilized as often as needed. A good system should give the radiologist the ability to easily window/level, change between frequently used preset values, and electronically measure, cross-reference, magnify, and synchronize images for easy comparison.

Redundancy in a teleradiology system is crucial to minimize downtime and delays in patient care. Alternate options should be considered for each component of a teleradiology system, including computers, monitors, power sources, phone lines, fax machines, and Internet connectivity at both the sending and receiving sites. If downtime does occur, a backup plan should be in place to get the images as quickly as possible to a radiologist who can interpret them before there is any negative outcome for the patient or dissatisfaction from the referring physician.

In the current market, there are affordable teleradiology systems and adequate bandwidth options that allow off-site radiologists to rapidly and reliably access high-quality, convenient-to-manipulate images. These readily available technologies will facilitate the realization of the many potential benefits of bringing radiological examinations to radiologists that are alert and expert in the type of modality or body system being imaged.

## JOB SATISFACTION

Working off-site can add many administrative steps to a radiology operation. Hiring adequate support staff to assist the radiologist is an essential component of a busy teleradiology practice. Receipt of information on cases to be

read, transcription, faxing reports to appropriate destinations, and logging patient information are functions that can be shared by support personnel and facilitated by available technologies. Minimizing the amount of time that a radiologist's attention is diverted from radiology to administration can maximize both quality of patient care and job approval. Teleradiology has great potential to increase the level of career satisfaction for many radiologists. Subspecialists can more easily dedicate themselves to their subspecialties. Consultation on difficult or interesting cases is possible even for radiologists who primarily work alone or geographically far from subspecialists. Radiologists wishing to work off-hours or to work from home can more easily do so. Radiologists wanting to travel can live in almost any location around the world. These many advantages may entice some of the more talented medical students into exploring radiology as a career choice, ultimately strengthening the future of the field.

## LEGAL AND SOCIOECONOMIC ISSUES

Table 26.5 lists legal and socioeconomic issues affecting telemedicine. Most of these are still unresolved to variable degrees.

### LICENSURE AND CREDENTIALLING

A conflict between, on the one hand, the ability of teleradiology and other telemedicine services to transcend geographic barriers technologically and,

---

---

**TABLE 26.5**

Legal and Socioeconomic Issues

1. Medical licensure and credentialing
  2. Malpractice insurance coverage
  3. Jurisdictional control over malpractice suits
  4. Confidentiality of medical records
  5. Physician-patient relationship
  6. Technical and clinical practice standards
  7. Reimbursement by third parties
  8. Turf issues among radiologists
- 
-

on the other, the legal responsibilities of medical licensing authorities within individual states has increased. Both the American Medical Association and the American College of Radiology have adopted policies that recommend full licensure in states where teleradiology or other forms of telemedicine are practiced. The ACR policy says, “[S]tates and their Medical Board should require a full and unrestricted medical license in the state in which the examination originates, with no differentiation by specialty, for physicians who wish to regularly practice telemedicine.” This means that a physician must be licensed in both the state of his or her residence and in the other state where the image originated. Many states have either passed legislation or implemented rulings by state medical boards to this effect. These laws and rulings are summarized in Chapter 8. Requirements for full licensure are pending in several additional states, and a number of states are considering a limited license or registration program, with the proviso that the provider physician be licensed in another state. The Federation of State Medical Boards favors special-purpose licensure providing telemedicine services.

In addition to licensure, most hospitals require teleradiology providers to apply for staff credentials. The new Joint Commission on Accreditation of Healthcare Organizations (JCAHO) guidelines have a provision allowing hospitals to use telemedicine services of physicians that are credentialed at a distant site that is JCAHO certified even if they are not credentialed at the originating site. It is not yet clear how many hospitals will feel comfortable enough using the services of physicians that they have not credentialed to take advantage of this provision.

## MALPRACTICE INSURANCE COVERAGE

Radiologists providing teleradiology services must have malpractice insurance that will extend coverage to lawsuits brought in all of the states in which they are providing coverage. As noted by Berlin, it is an established legal principle that the state in which the injury occurs and that has the greatest connection to the injury possesses jurisdiction over the lawsuit. Unfortunately, many malpractice insurers have regional limitations on their coverage. These have been set up due to the difficulties in defending a lawsuit in a remote location where the insurance company does not have relationships with malpractice attorneys that will defend the case and where the defending physician may be seen as an outsider. In addition, the current malpractice climate has turned many regions into very difficult malpractice environments that carriers are reluctant to become involved in. For these reasons, a radiologist practicing in many different states will often have to turn to excess



and surplus-lines carriers, which can be very expensive. Furthermore, radiologists practicing outside of the United States may face additional difficulty finding malpractice carriers that have confidence that an overseas radiologist will return to the state of origin to properly defend a possible lawsuit. As the field of teleradiology continues to advance and malpractice risk levels become clarified, it is likely that insurers will gain enough comfort to respond to the unique coverage needs of teleradiology.

## PATIENT CONFIDENTIALITY

Ensuring patient confidentiality is of utmost importance, and formal procedures for doing so have been mandated with the passing of HIPAA (Health Insurance Portability and Accountability Act) legislation by congress. This legislation provides guidelines for the protection of patient privacy that hospitals, imaging centers, and doctors must adhere to. For teleradiology, the image-compression process encrypts the image in a way that prevents casual theft, and Internet-based systems can be protected using virtual private networks or secure sockets layer technology. With a secure mechanism for image transmission in place, a teleradiology office must take the usual steps for patient privacy that other medical offices would take.

## STANDARDS

As noted in the telemedicine report to Congress from the US Department of Commerce, there are no specialty-generated technical standards, protocols, or clinical guidelines for telemedicine, with the exception of the ACR, which has developed practice guidelines for teleradiology. The American Medical Association, the American Telemedicine Association, and a number of other professional societies are currently developing practice standards for telemedicine. The Department of Commerce report also notes that the Food and Drug Administration (FDA) is the lead federal agency with responsibility for ensuring the safety and effectiveness of telemedicine devices marketed in the United States.

The ACR has deemed that several criteria in the teleradiology process should meet minimum standards. (The ACR guidelines are included as an appendix to this chapter.) It would be prudent for anyone providing teleradiology services to review the ACR standards. Highlights from the guidelines include:

- ▶ For small matrix studies (e.g., CT, MRI, ultrasound, nuclear medicine, digital fluorography, and digital angiography),  $512 \times 512 \times 8$ -bit acquisition datasets should be used; for large matrix images (e.g., conventional radiographs), datasets should be a minimum of 2.5 lp/mm at a minimum 10-bit depth.
- ▶ Minimum display resolution is not specified, although the display device should be capable of displaying all acquired data. This does not mean that the entire image needs to be able to be displayed at full native resolution, but rather that the software allows the user to zoom and pan the image sufficiently to view any part of an image at its native resolution. Thus a 3-MP display used for radiography in regular PACS workstations can certainly be used for teleradiology. The same would apply for 1–2-MP displays that are used for small matrix studies (nonradiography).
- ▶ Image annotation (patient demographics and examination information) should be included.
- ▶ Image display stations should provide functionality for window width and level settings, inverting contrast, image rotation, measurement function, and magnification.
- ▶ Display brightness of at least 50 foot-lamberts should be achieved.
- ▶ Compression ratios should be displayed with the image (no specification for what type or level of compression).
- ▶ Reasonable measures must be taken to protect patient confidentiality.

## REIMBURSEMENT

There has been fairly uniform recognition of teleradiology for reimbursement by private insurance companies and government programs including Medicare and Medicaid. The exception to this is for interpretations rendered outside the United States, for which Medicare will not pay. Some private insurers have followed Medicare on this issue as well. The ACR task force on international teleradiology put out a report in April 2004 recognizing that there is no inherent technological difference between teleradiology interpretations generated in the United States and those generated outside the United States. This task force felt that payment for international teleradiology reports is appropriate provided that the following criteria are met:

1. The person interpreting the examination and submitting the report to the referring physician are one and the same.
2. The person rendering the report is licensed in the state and credentialed as a member of the medical staff at the institution performing the examination and receiving the report.
3. The person performing the interpretation and rendering the report is available for consultation.
4. The report meets the guidelines for diagnostic reports as promulgated by the ACR.
5. The ACR Technical Standard for Teleradiology is met.

Despite the ACR recommendations, it remains to be seen whether Medicare will alter its policy. In the meantime, the ACR report also indicated that it is unethical for a radiologist who did not personally interpret the images in a study to sign his or her name to a report in a manner that would indicate that this radiologist interpreted the study.

## TURF ISSUES

The radiology community is divided on the issue of teleradiology. Although the general concept has been widely embraced throughout different practice settings—from the largest academic institutions to solo practitioners in private practice—the practice of teleradiology continues to create controversy. Some people in the radiology community have expressed concern that larger groups, including academic medical centers, will use their financial resources, prestige, and clout to invade the turf of smaller groups, who regard their geographic service area as inviolable. In defending this view, the point is made that it is important for radiologists to provide more than just interpretations and to take part in the medical life of the respective hospital or healthcare community. Counterarguments include the view that teleradiology provides patients and their attending physicians the opportunity to access more highly subspecialized experts than may be available in the local community, and that it is not appropriate for individual radiologists to stand between the patient and such expert opinions.

It is far from clear how the issues of mistrust and suspicion on the part of community radiologists versus the desire of larger subspecialist groups to offer their expertise more broadly will be resolved. However, one opportunity open to community radiologists is to come together among themselves to form larger networks in which individuals can begin to subspecialize and,

as discussed above, where smaller groups can provide coverage to one another for a variety of purposes. The one thing abundantly obvious in the Internet age is that patients will increasingly expect electronic access when they see it as beneficial to their interests and needs. Doctors, including radiologists, will be ill-served to stand in the way of that new and increasing imperative.

## CONCLUSION

Teleradiology has transitioned from its infancy and is no longer a technological curiosity. It is already part of daily practice. It is here, it works, and it will have an enduring role in shaping the future of radiology. It is likely that, in less than a decade, the term “teleradiology” will be obsolete for the simple reason that remote interpretations and consultations will be such an integral part of radiology practice that referring to them will neither require nor occasion a special term.

## ACKNOWLEDGMENT

Portions of this chapter were published as Thrall JH, Boland G. Telemedicine in practice. *Semin Nucl Med.* 1998;28:142–157 (with permission of Elsevier).

## REFERENCES

- Aberle DR, Gleeson F, Sayre JW, et al. The effect of irreversible image compression on diagnostic accuracy in thoracic imaging. *Invest Radiol.* 1993;28:398–403.
- Ackerman S, Gitlin JW, Gayler RW, Flagle CD, Bryan RN. Receiver operating characteristic analysis of fracture and pneumonia detection: Comparison laser-digitized workstation images and conventional analog radiographs. *Radiology.* 1993;186:263–268.
- ACR Bulletin.* 1997;53:13.
- American College of Radiology. ACR standard for teleradiology. In *ACR Practice Guidelines and Technical Standards.* Reston, VA: ACR; 2004:709–718.
- Andrus WS, Bird KT. Teleradiology: evolution through bias to reality. *Chest.* 1972;62:655–657.
- Baer L, Cukor P, Jenike MA, Leahy L, O’Laughlen J, Coyle JT. Pilot studies of telemedicine for patients with obsessive-compulsive disorder. *Am J Psychiatry.* 1995;152:1383–1385.

- Baer L, Jacobs DG, Cukor P, O'Laughlen J, Coyle JT, Magruder KM. Automated telephone screening survey for depression. *JAMA*. 1995;273:1943-1944.
- Berlin L. Malpractice issues in radiology: telereadiology. *AJR Am J Roentgenol*. 1998;170:1417-1422.
- Bird KT. Cardiopulmonary frontiers: quality health care via interactive television. *Chest*. 1972;61:204-207.
- Black-Schaffer S, Flotte TJ. Current issues in telepathology. *Telemed J*. 1995;1:95-106.
- Calcagni DE, Clyburn CA, Tomkins G, et al. Operation Joint Endeavor in Bosnia: telemedicine systems and case reports. *Telemed J*. 1996;2:211-224.
- Cawthon MA, Goeringer F, Telepak RJ, et al. Preliminary assessment of computed tomography and satellite teleradiology from Operation Desert Storm. *Invest Radiol*. 1991;26:854-857.
- Dakins D. Market targets 1997. *Telemed Telehealth Networks*. 1997;3:25-29.
- Elam EA, Rehm K, Hillman BJ, Maloney K, Fajardo LL, McNeill K. Efficacy of digital radiography for detection of pneumothorax: comparison with conventional chest radiography. *AJR Am J Roentgenol*. 1992;158:509-514.
- Fosberg DA. Quality assurance in teleradiology. *Telemed J*. 1995;1:38-44.
- Franken EA Jr, Bergus GR, Koch TJ, Berbaum KS, Smith WL. Added value of radiologist consultation to family practitioners in the outpatient setting. *Radiology*. 1995;197:759-762.
- Franken EA, Harkens KL, Berbaum KS. Teleradiology consultation for a rural hospital: patterns of use. *Acad Radiol*. 1997;14:492-496.
- Fuchs M. Provider attitudes toward STARPAHC, a telemedicine project on the Papago Reservation. *Med Care*. 1979;17:59-68.
- Gagner M, Begin E, Hurteau R, Pomp A. Robotic interactive laparoscopic cholecystectomy. *Lancet*. 1994;343:596-597.
- Goldberg MA. Teleradiology and telemedicine. *Radiol Clin North Am*. 1996;43:647-665.
- Goldberg MA, Rosenthal DI, Chew FS, Blickman JG, Miller SW, Mueller PR. New high resolution teleradiology system: prospective study of diagnostic accuracy in 685 transmitted clinical cases. *Radiology*. 1993;186:429-434.
- Goldberg MA, Pivovarov M, Mayo-Smith WW, et al. Application of wavelet compression to digitized radiographs. *AJR Am J Roentgenol*. 1994;163:463-468.
- Goldberg MA, et al. Effect of 3-D wavelet compression on the detection of focal hepatic lesions at CT. *Radiology*. 1997;202:159-165.
- Gomez E, Poropatich R, Karinch MA, Zajtchuk J. Tertiary telemedicine support during global military humanitarian missions. *Telemed J*. 1996;2:201-210.
- Gorden JW, Knapp CF, Sanders JH. Ophthalmologic electronic imaging and data transfer. *J Ky Med Assoc*. 1991;89:115-117.
- Grigsby J, Kaehny MM, Sandberg EJ, Schlenker RE, Shaughnessy PW. Effects and effectiveness of telemedicine. *Health Care Financ Rev*. 1995;17:115-131.

- Ho BKT, et al. A mathematical model to quantify JPEG block artifacts. *SPIE Phys Med Imaging*. 1993;1987:169-274.
- Hubble JP, Pahwa R, Michalek DK, Thomas C, Koller WC. Interactive video conferencing: a means of providing interim care to Parkinson's disease patients. *Mov Disord*. 1993;8:380-382.
- Jutra A. Teleröntgen diagnosis by means of videotape recording. *AJR Am J Roentgenol*. 1959;82:1099-1102.
- Kvedar JC, Edwards RA, Menn ER, et al. The substitution of digital images for dermatologic physical examination. *Arch Dermatol*. 1997;133:161-167.
- Larson A, Lynch DA, Zeligman B, et al. Accuracy of diagnosis of subtle chest disease and subtle fractures with a teleradiology system. *AJR Am J Roentgenol*. 1998;170:19-22.
- Mullick FG, Fontelo P, Pemble C. Telemedicine and telepathology at the Armed Forces Institute of Pathology: history and current mission. *Telemed J*. 1996;2:187-194.
- Murphy RLH, Bird KT. Telediagnosis: a new community health resource. *Am J Public Health*. 1974;64:113-119.
- Murphy RL Jr, Fitzpatrick TB, Haynes HA, Bird KT, Sheridan TB. Accuracy of dermatologic diagnosis by television. *Arch Dermatol*. 1972;105:833-835.
- Perednia DA, Allen A. Telemedicine technology and clinical applications. *JAMA*. 1995;273:483-488.
- Rajavi M, Sayre JW, Taira RK, et al. Receiver-operating-characteristic study of chest radiographs in children: digital hard-copy film versus 2K32K soft copy images. *AJR Am J Roentgenol*. 1992;158:443-448.
- Rayman BR. Telemedicine: military applications. *Aviat Space Environ Med*. 1992;63:135-137.
- Sanders JH, Tedesco FJ. Telemedicine: bringing medical care to isolated communities. *J Med Assoc Ga*. 1993;82:237-241.
- Satava RM. Robotics, telepresence and virtual reality: a critical analysis of the future of surgery. *Minim Invasive Ther*. 1992;1:357-363.
- Satava RM. Virtual reality surgery simulator: the first steps. *Surg Endosc*. 1993;7:203-205.
- Satava RM, Jones SB. Virtual reality and telemedicine: exploring advanced concepts. *Telemed J*. 1996;2:195-200.
- Scott WW Jr, Rosenbaum JE, Ackerman SJ, et al. Subtle orthopedic fractures: teleradiology workstation versus film interpretation. *Radiology*. 1993;187:811-815.
- Scott WW Jr, Bluemke DA, Mysko WK, et al. Interpretation of emergency department radiographs by radiologists and emergency physicians: teleradiology workstation versus radiographic readings. *Radiology*. 1995;195:223-229.

- Slasky BS, Gur D, Good WF, et al. Receiver operating characteristic analysis of chest image interpretation with conventional, laser printed and high-resolution workstation images. *Radiology*. 1990;174:775-780.
- US Department of Commerce. *Telemedicine Report to Congress*. January 31, 1997.
- Van Moore A. *Report of the ACR Task Force on International Teleradiology*. April 30, 2004.
- Wilson AJ, Hodge JC. Digitized radiographs in skeletal trauma: a performance comparison between a digital workstation and the original film images. *Radiology*. 1995;196:565-568.
- Wittson CL, Affleck DC, Johnson V. Two-way television group therapy. *Ment Hosp*. 1961;12:22-23.
- Zajtchuk JT, Zajtchuk, R. Strategy for medical readiness: transition to the digital age. *Telemed J*. 1996;2:179-186.

## APPENDIX

### ACR STANDARD FOR TELERADIOLOGY

#### ACKNOWLEDGMENT

Reprinted with permission of the American College of Radiology. No other representation of this article is authorized without express, written permission from the American College of Radiology.

The American College of Radiology, with more than 30,000 members, is the principal organization of radiologists, radiation oncologists, and clinical medical physicists in the United States. The College is a nonprofit professional society whose primary purposes are to advance the science of radiology, improve radiologic services to the patient, study the socioeconomic aspects of the practice of radiology, and encourage continuing education for radiologists, radiation oncologists, medical physicists, and persons practicing in allied professional fields. The American College of Radiology will periodically define new practice guidelines and technical standards for radiologic practice to help advance the science of radiology and to improve the quality of service to patients throughout the United States. Existing practice guidelines and technical standards will be reviewed for revision or renewal, as appropriate, on their fifth anniversary or sooner, if indicated. Each practice guideline and technical standard, representing a policy statement by the College, has undergone a thorough consensus process in which it has been subjected to extensive review, requiring the approval of the Commission on Quality and Safety as well as the ACR Board of Chancellors, the ACR Council Steering Committee, and the ACR Council. The practice

guidelines and technical standards recognize that the safe and effective use of diagnostic and therapeutic radiology requires specific training, skills, and techniques, as described in each document. Reproduction or modification of the published practice guideline and technical standard by those entities not providing these services is not authorized. 1994 (Res. 21) Revised 1995 (Res. 26)

Revised 1998 (Res. 35)

Revised 2002 (Res. 11) Effective 1/1/03

## ACR TECHNICAL STANDARD FOR TELERADIOLOGY PREAMBLE

These guidelines are an educational tool designed to assist practitioners in providing appropriate radiologic care for patients. They are not inflexible rules or requirements of practice and are not intended, nor should they be used, to establish a legal standard of care. For these reasons and those set forth below, the American College of Radiology cautions against the use of these guidelines in litigation in which the clinical decisions of a practitioner are called into question. The ultimate judgment regarding the propriety of any specific procedure or course of action must be made by the physician or medical physicist in light of all the circumstances presented. Thus, an approach that differs from the guidelines, standing alone, does not necessarily imply that the approach was below the standard of care. To the contrary, a conscientious practitioner may responsibly adopt a course of action different from that set forth in the guidelines when, in the reasonable judgment of the practitioner, such course of action is indicated by the condition of the patient, limitations on available resources or advances in knowledge or technology subsequent to publication of the guidelines. However, a practitioner who employs an approach substantially different from these guidelines is advised to document in the patient record information sufficient to explain the approach taken. The practice of medicine involves not only the science, but also the art of dealing with the prevention, diagnosis, alleviation and treatment of disease. The variety and complexity of human conditions make it impossible to always reach the most appropriate diagnosis or to predict with certainty a particular response to treatment. It should be recognized; therefore, that adherence to these guidelines will not assure an accurate diagnosis or a successful outcome. All that should be expected is that the practitioner will follow a reasonable course of action based on current knowledge, available resources, and the needs of the patient to deliver effective and safe medical care. The sole purpose of these guidelines is to assist practitioners in achieving this objective.



## I. INTRODUCTION AND DEFINITION

Teleradiology is the electronic transmission of radiologic images from one location to another for the purposes of interpretation and/or consultation. Teleradiology may allow more timely interpretation of radiologic images and give greater access to secondary consultations and to improved continuing education. Users in different locations may simultaneously view images. Appropriately utilized, teleradiology may improve access to radiologic interpretations and thus significantly improve patient care. Teleradiology is not appropriate if the available teleradiology system does not provide images of sufficient quality to perform the indicated task. When a teleradiology system is used to render the official interpretation,<sup>1</sup> there should not be a clinically significant loss of data from image acquisition through transmission to final image display. For transmission of images for display use only, the image quality should be sufficient to satisfy the needs of the clinical circumstance.

This standard defines goals, qualifications of personnel, equipment guidelines, licensing, credentialing, liability, communication, quality control, and quality improvement for teleradiology. While not all-inclusive, the standard should serve as a model for all physicians and healthcare workers who utilize teleradiology. A glossary of commonly used terminology (not reprinted here) and a reference list are included.

## II. GOALS

Teleradiology is an evolving technology. New goals will continue to emerge. The current goals of teleradiology include:

- A.** Providing consultative and interpretative radiologic services.
- B.** Making radiologic consultations available in medical facilities without on-site radiologic support.
- C.** Providing timely availability of radiologic images and image interpretation in emergent and nonemergent clinical care areas.

---

<sup>1</sup> The ACR Medical Legal Committee defines official interpretation as that written report (and any supplements or amendments thereto) that attach to the patient's permanent record. In healthcare facilities with a privilege delineation system, such a written report is prepared only by a qualified physician who has been granted specific delineated clinical privileges for that purpose by the facility's governing body upon the recommendation of the medical staff.

- D. Facilitating radiologic interpretations in on-call situations.
- E. Providing subspecialty radiologic support as needed.
- F. Enhancing educational opportunities for practicing radiologists.
- G. Promoting efficiency and quality improvement.
- H. Providing interpreted images to referring providers.
- I. Supporting telemedicine.
- J. Providing supervision of off-site imaging studies.

### III. QUALIFICATIONS OF PERSONNEL

The radiologic examination at the transmitting site must be performed by qualified personnel trained in the examination to be performed. In all cases this means a licensed and/or registered radiologic technologist, radiation therapist, nuclear medicine technologist, or sonographer. This technologist must be under the supervision of a qualified licensed physician.

It is desirable to have a Qualified Medical Physicist and/or image management specialist on site or as consultants.

#### A. PHYSICIAN

The official interpretation of images must be done by a physician who has:

1. An understanding of the basic technology of teleradiology, its strengths and weaknesses (as well as limitations), and who is knowledgeable in the use of the teleradiology equipment.
2. Demonstrated qualifications as delineated in the appropriate American College of Radiology (ACR) guidelines or standards for the particular diagnostic modality being transmitted through teleradiology.

#### B. RADIOLOGIC TECHNOLOGIST, RADIATION THERAPIST, NUCLEAR MEDICINE TECHNOLOGIST, OR SONOGRAPHER

The technologist, therapist, or sonographer should be:

1. Certified by the appropriate registry and/or possess unrestricted state licensure.
2. Trained to properly operate and supervise the teleradiology system.

### C. QUALIFIED MEDICAL PHYSICIST

A Qualified Medical Physicist is an individual who is competent to practice independently in one or more of the subfields in medical physics. The ACR considers that certification and continuing education in the appropriate subfield(s) demonstrate that an individual is competent to practice one or more of the subfields in medical physics and to be a Qualified Medical Physicist. The ACR recommends that the individual be certified in the appropriate subfield(s) by the American Board of Radiology (ABR).

The appropriate subfields of medical physics for this standard are Therapeutic Radiological Physics, Diagnostic Radiological Physics, Medical Nuclear Physics, and Radiological Physics.

The continuing education of a Qualified Medical Physicist should be in accordance with the ACR Practice Guideline for Continuing Medical Education (CME).

### D. IMAGE MANAGEMENT SPECIALIST

1. The image management specialist is an individual who is qualified to assess and provide problem-solving input, initiate repair, and coordinate system-wide maintenance programs to assure sustainable high-image quality and system function. This individual would also be directly involved with any system expansion programs.

2. This specialist should be available in a timely manner in case of malfunction to facilitate return to optimal system functionality.

## IV. EQUIPMENT SPECIFICATIONS

Specifications for equipment used in teleradiology will vary depending on the individual facility's needs, but in all cases it should provide image quality and availability appropriate to the clinical need.

Compliance with the ACR/NEMA (National Electrical Manufacturers Association) Digital Imaging and Communication in Medicine (DICOM) standard is strongly recommended for all new equipment acquisitions, and consideration of periodic upgrades incorporating the expanding features of that standard should be part of the continuing quality improvement program.

Equipment guidelines cover two basic categories of teleradiology when used for rendering the official interpretation: small matrix size (e.g., com-

puted tomography [CT], magnetic resonance imaging [MRI], ultrasound, nuclear medicine, digital fluorography, and digital angiography) and large matrix size (e.g., digital radiography and digitized radiographic films).

**Small matrix:** The data set should provide a minimum of  $512 \times 512$  matrix size at a minimum 8-bit pixel depth for processing or manipulation with no loss of matrix size or bit depth at display.

**Large matrix:** The data set should allow a minimum of 2.5 lp/mm spatial resolution at a minimum 10-bit pixel depth.

## A. ACQUISITION OR DIGITIZATION

Initial image acquisition should be performed in accordance with the appropriate ACR modality or examination guideline or standard.

**1. DIRECT IMAGE CAPTURE** The entire image data set produced by the digital modality in terms of both image matrix size and pixel bit depth, should be transferred to the teleradiology system. It is recommended that the DICOM standard be used.

## 2. SECONDARY IMAGE CAPTURE

- a. **Small matrix images:** Each individual image should be digitized to a matrix size as large as or larger than that of the original image by the imaging modality. The images should be digitized to a minimum of 8 bits pixel depth. Film digitization or video frame grab systems conforming to the above specifications are acceptable.
- b. **Large matrix images:** These images should be digitized to a matrix size corresponding to 2.5 lp/mm or greater, measured in the original detector plane. These images should be digitized to a minimum of 10 bits pixel depth.

**3. GENERAL REQUIREMENTS** At the time of acquisition (small or large matrix), the system must include annotation capabilities including patient name, identification number, date and time of examination, name of facility or institution of acquisition, type of examination, patient or anatomic part orientation (e.g., right, left, superior, inferior), and amount and method of data compression.

The capability to record a brief patient history is desirable.

## B. COMPRESSION

Data compression may be used to increase transmission speed and reduce storage requirements. Several methods, including both reversible and irreversible techniques, may be used, under the direction of a qualified physician, with no reduction in clinically significant diagnostic image quality. The types and ratios of compression used for different imaging studies transmitted and stored by the system should be selected and periodically reviewed by the responsible physician to ensure appropriate clinical image quality.

## C. TRANSMISSION

The type and specifications of the transmission devices used will be dictated by the environment of the studies to be transmitted. In all cases, for official interpretation, the digital data received at the receiving end of any transmission must have no loss of clinically significant information. The transmission system shall have adequate error-checking capability.

## D. DISPLAY CAPABILITIES

Display workstations used for official interpretation and employed for small-matrix and large-matrix systems should provide the following characteristics:

1. Luminance of the gray-scale monitors should be at least 50 foot-lamberts.
2. Lighting in the reading room that can be controlled to eliminate reflections in the monitor and to lower the ambient lighting level as much as is feasible.
3. Capability for selecting image sequence.
4. Capability of accurately associating the patient and study demographic characterizations with the study images.
5. Capability of window and level adjustment, if those data are available.
6. Capability of pan and zoom functions.
7. Capability of rotating or flipping the images provided correct labeling of patient orientation is preserved.

8. Capability of calculating and displaying accurate linear measurements and pixel value determinations in appropriate values for the modality (e.g., Hounsfield units for CT images), if those data are available.
9. Capability of displaying prior image compression ratio, processing, or cropping.
10. The following elements of display:
  - a. Matrix size.
  - b. Bit depth.
  - c. Total number of images acquired in the study.
  - d. Clinically relevant technical parameters.

When the display systems are not used for the official interpretation, they need not meet all the characteristics listed above.

#### E. ARCHIVING AND RETRIEVAL

If electronic archiving is to be employed, the guidelines listed below should be followed:

1. Teleradiology systems should provide storage capacity sufficient to comply with all facility, state, and federal regulations regarding medical record retention. Images stored at either site should meet the jurisdictional requirements of the transmitting site. Images interpreted off-site need not be stored at the receiving facility, provided they are stored at the transmitting site. However, if the images are retained at the receiving site, the retention period of that jurisdiction must be met as well. The policy on record retention must be in writing.

2. Each examination data file must have an accurate corresponding patient and examination database record that includes patient name, identification number, examination date, type of examination, and facility at which examination was performed. It is desirable that space be available for a brief clinical history.

3. Prior examinations should be retrievable from archives in a time frame appropriate to the clinical needs of the facility and medical staff.

4. Each facility should have policies and procedures for archiving and storage of digital image data equivalent to the policies for protection of hard-copy storage media to preserve imaging records.

## F. SECURITY

Teleradiology systems should provide network and software security protocols to protect the confidentiality of patients' identification and imaging data consistent with federal and state legal requirements. There should be measures to safeguard the data and to ensure data integrity against intentional or unintentional corruption of the data.

## G. RELIABILITY AND REDUNDANCY

Quality patient care may depend on timely availability of the image interpretation. Written policies and procedures should be in place to ensure continuity of teleradiology services at a level consistent with those for hard-copy imaging studies and medical records within a facility or institution. This should include internal redundancy systems, backup telecommunication links, and a disaster plan.

## V. LICENSING, CREDENTIALING, AND LIABILITY

Physicians who provide the official interpretation of images transmitted by teleradiology should maintain licensure as may be required for provision of radiologic service at both the transmitting and receiving sites. When providing the official interpretation of images from a hospital, the physician should be credentialed and obtain appropriate privileges at that institution. These physicians should consult with their professional liability carrier to ensure coverage in both the sending and receiving sites (state or jurisdiction).

The physician performing the official interpretations is responsible for the quality of the images being reviewed.<sup>2</sup>

Images stored at either site should meet the jurisdictional requirements of the transmitting site. Images interpreted off-site need not be stored at the receiving facility, provided they are stored at the transmitting site. However,

---

<sup>2</sup> The ACR Rules of Ethics state: "it is proper for a diagnostic radiologist to provide a consultative opinion on radiographs and other images regardless of their origin. A diagnostic radiologist should regularly interpret radiographs and other images only when the radiologist reasonably participates in the quality of medical imaging, utilization review, and matters of policy which affect the quality of patient care."

if images are retained at the receiving site, the retention period of that jurisdiction must be met as well. The policy on record retention should be in writing.

The physicians who are involved in practicing teleradiology will conduct their practice in a manner consistent with the bylaws, rules, and regulations for patient care at the transmitting site.

## VI. DOCUMENTATION

Communication is a critical component of teleradiology. Physicians interpreting teleradiology examinations should render reports in accordance with the ACR Practice Guideline for Communication: Diagnostic Radiology.

## VII. QUALITY CONTROL AND IMPROVEMENT, SAFETY, INFECTION CONTROL, AND PATIENT EDUCATION CONCERNS

Policies and procedures related to quality, patient education, infection control, and safety should be developed and implemented in accordance with the ACR Policy on Quality Control and Improvement, Safety, Infection Control, and Patient Education Concerns appearing elsewhere in the ACR Practice Guidelines and Technical Standards book.

Any facility using a teleradiology system must have documented policies and procedures for monitoring and evaluating the effective management, safety, and proper performance of acquisition, digitization, compression, transmission, archiving, and retrieval functions of the system. The quality control program should be designed to maximize the quality and accessibility of diagnostic information.

A test image, such as the SMPTE test pattern, should be captured, transmitted, archived, retrieved, and displayed at appropriate intervals, but at least monthly, to test the overall operation of the system under conditions that simulate the normal operation of the system. As a spatial resolution test, at least  $512 \times 512$  resolution should be confirmed for small-matrix official interpretation, and 2.5 lp/mm resolutions for large-matrix official interpretation.

As a test of the display, SMPTE pattern data files sized to occupy the full area used to display images on the monitor should be displayed. The overall SMPTE image appearance should be inspected to assure the absence of gross artifacts (e.g., blurring or bleeding of bright display areas into dark



areas or aliasing of spatial resolution patterns). Display monitors used for primary interpretation should be tested at least monthly. As a dynamic range test, both the 5% and the 95% areas should be seen as distinct from the respective adjacent 0% and 100% areas.

The use of teleradiology does not reduce the responsibilities for the management and supervision of radiologic medicine.

# INDEX

---

---

## A

Abnormality, communication, speed of, 114  
Accelerated graphics port (AGP), 178  
Access time, hard drive, 322  
Acronis, 311  
Active-active clustering, 308  
Administrator accounts, 315  
Agfa-Gevaert, computed radiography (CR) systems, 203–204  
*Alberts v. Devine*, 167  
American College of Radiology (ACR)  
  on digital image data management, 151–152  
  on image storage, 166–167, 190–191  
  NEMA 2.0 standard, 190–191  
  standard (Appendix), 550–560  
  standard, overview of, 150–151, 544–545  
American Medical Association (AMA),  
  teleradiology policy, 152–153  
Analog-to-digital converters (ADC), 196

Annotation capabilities, images, 150  
Antivirus software, 313  
Appendices, request for proposal (RFP), 63  
Appletalk protocol, 275  
Application entities (AE), 191  
Application program interface (API), 304  
Architecture. *See* PACS architecture  
Archiving. *See* Disaster recovery; Image archive; Storage  
Arrival process, queuing analysis of, 86–87  
Artifacts  
  computed radiography (CR), 209–210  
  double exposures, 209–210  
  lightbulb effect, 209  
Asynchronous transfer mode (ATM), 292–293  
  advantages of, 292, 299  
  cards, 275  
  compared to Ethernet, 292  
  limitations of, 293, 529

- AT advanced (ATA) interface, 323–324
- Authentication, 314
- Autorouting, image manager, 53
- Average-intensity projection (AveIP) images, 452
- B**
- Bandwidth, 284–287  
 and cabling, 285  
 defined, 526  
 networks/speeds, 285  
 transmission media differences, 529
- Big bang approach, implementation, 57
- Binary digits, and CPU operation, 176–177, 284
- Bits, 284
- Black bone technique, 206–207, 396
- Bottlenecks, 288–289
- Breakeven analysis, 135–137
- Breast cancer  
 computer-aided diagnosis (CAD), 423–424, 436–441  
*See also* Mammography, digital
- Bridging, networks, 278–280, 282
- Bright-lighting film, 202
- Broker, RIS/PACS interface, 22
- Browsers, Internet, 404
- Bus, computer function, 177–178
- Bytes, 284
- C**
- Cables  
 networks, 274  
 and transmission speed, 285
- Calibration, monitors, 355–358
- Capital expenditures, 137–138  
 types of, 43, 138
- Cash flow, 126–128  
 costs, 127–128  
 payback period, 128–129
- Cassetteless computed radiography (CR) systems, 204
- Centera, 264
- Centralized model, 259–261, 402–404
- Central processing unit (CPU), operation of, 176–177
- Charge coupled device (CCD) scanners, 196, 525
- C language, 179
- Clerical personnel  
 PACS, impact on job, 77–79  
 workflow analysis, 77–79
- Client devices, 379–380
- Clientless backup, 330
- Clinical archive, 252–253
- Clinical review workstation, request for proposal (RFP) description, 55–56
- Clinical review section, request for proposal (RFP), 49
- Clinicians  
 PACS impact on, 112–114  
 time and cost-savings, 142
- Clustering  
 active-active clustering, 308  
 clustered failover model, 307–308  
 PACS components, 255
- Compact disc (CD)  
 digital teaching files (DTFs)  
 information, 511  
 storage media, 336
- Competition  
 impact of PACS implementation, 143  
 market share analysis, 37–39
- Compression methods. *See* Image compression
- CompuServe, Inc. v. Patterson*, 165
- Computed radiography (CR), 200–211  
 contrast enhancement algorithms, 206–207  
 contrast resolution, 201–202  
 detective quantum efficiency (DQE), 200–201  
 development of systems, 202–203  
 compared to digital radiography, 19–20, 211, 215–217  
 dynamic range control processing algorithm, 208  
 image acquisition process, 197–199  
 image artifacts, 209–210  
 image compression, 240  
 image production flowchart, 199  
 image segmentation algorithms, 205–206  
 increase in use of, 197  
 multiscale image contrast amplification (MUSICA), 208

- noise, 202
- operation of, 18
- quality assurance-quality control (QA-QC), 220
- compared with screen-film, 199–200, 202
- separation of functions of, 18
- spatial frequency processing algorithms, 207–208
- spatial resolution, 201
- system features, 199–200
- technological advances, 203–205, 221
- uses of, 18
- Computed tomography (CT)
  - image compression, 241
  - multidetector computed tomography (MDCT), 108, 396, 450–451, 512
  - 3-D scans, 449–450
- Computer-aided diagnosis (CAD)
  - classification schemes, 439–440
  - as diagnosis confirmation, 423–424
  - effectiveness rate, 438
  - image requirements, 437
  - limitations of, 238
  - for mammography, 423–424, 436–441
  - as measure of compression, 238, 240
- Computer functions
  - application programs, 180–181
  - bus, 177–178
  - computer languages, 179–180
  - input, 174
  - memory, 175–176
  - operating systems, 178–179
  - output, 175
  - processing, 176–177
- Confidentiality
  - HIPAA privacy rule, 158, 506, 544
  - record-keeping, 166–167
- Content addressable storage (CAS), 264, 334–335
  - companies offering, 264
- Contracts
  - maintenance contract, 61–62
  - negotiation with vendors, 70–71
- Contrast enhancement algorithms, 206–207
- Contrast resolution. *See* Resolution
- Core servers. *See* Servers
- Costs, 43–44, 127–128, 137–138
  - capital expenditures, 43, 137–138
  - fixed, 135
  - irrelevant costs, 127–128
  - sunk costs, 127
  - variable, 135–136
- Cost-savings, 138–143
  - clinician time, 142
  - film-costs, 32, 36, 140
  - film library cost, 140–141
  - lost/misplaced film, 141
  - multiple sites, reduction of, 142
  - multisite image reading, 142–143
  - network costs, 300–301
  - patient outcomes, 143
  - per-exam basis example, 138–139
  - redos, 141
  - speech recognition system, 481
- Credentialing, 543
- D**
- Database
  - backups, 312
  - information stored, 250
- Database manager
  - functions of, 20
  - optimal system, components of, 52–53
  - request for proposal (RFP) description, 52–53
  - software used, 20
- Data link layer, networks, 275
- Data transfer rate, hard drive, 323
- DAX Archiving Solutions, 311
- Decision-making
  - decision support and order entry, 490
  - decision support tools, 107
  - decision model, vendor selection, 66–69
- Default display protocols (DDPs), 105
- Default shares, disabling, 316
- Dent v. West Virginia*, 158
- Design of system, 402–406
- Desktop integration, speech recognition system, 476
- Detective quantum efficiency (DQE), 200–201
- Diagnostic review section, request for proposal (RFP), 49

- Diagnostic workstations
  - operation of, 24–25, 390
  - request for proposal (RFP) description, 54–55
  - software options, 55
- DICOM
  - autorouting destinations, 53
  - digital mammography standard, 366–367
  - header, example of, 191–192
  - and IHE initiative, 84
  - image acquisition standards, 20, 50–51, 150, 189–194
  - image compression standard, 235–236
  - language structure, 191
  - lookup tables, 359
  - medical image resource center (MIRC), 499
  - system readiness and vendor, 51
  - worklist standards, 102, 193–194, 251–252
- DICOM servers
  - functions of, 252
  - as image servers, 305–307
- Digital, meaning of, 10
- Digital cameras, 196
- Digital driving level (DDL), 355–356
- Digital format
  - clinical effectiveness, 11
  - HIPAA requirement, 10, 53
  - needs in medical environment, 11–12
  - quality issues, 10
- Digital images
  - compression methods, 150, 221–222
  - construction of, 184–185
  - enhancement tools, 106–107, 186–187
  - future imaging modalities, 220–221
  - grayscale, 185–186
  - image analysis, 187
  - image manipulation tools, 150, 393, 396–397
  - image processing, 186–187
  - image understanding, 187
  - matrix categories, 149
  - volumetric navigation, 107, 109
- Digital Imaging and Communications in Medicine (DICOM) standards. *See* DICOM
- Digital linear tape (DLT), 311–312
- Digital radiography (DR), 211–217
  - compared to computed radiography (CR), 19–20, 211, 215–217
  - cost factors, 214
  - direct digital radiography (DDR), 19, 212
  - image compression, 240
  - image production process, 212–213
  - image quality, 215
  - indirect digital radiography, 213
  - operation of, 18–19
  - system features, 213–215
- Digital teaching files (DTFs), 500–511
  - authoring tools, 508–509
  - case collections, 510
  - CDs for case collections, 511
  - development specifications, 501
  - divisional files, 504
  - flagged cases, 503
  - full file, 507
  - general format for, 502
  - image annotations, 509
  - key images, 507
  - lists, 507
  - medical record integration, 505–506
  - MIRC compatibility, 504–505
  - patient privacy, 506–507
  - presentation tools, 510
  - private files, 504
  - public files, 504
  - published files, 504
  - quality assurance, 508
  - teaching cases, types of, 503
  - users of, 509–510
  - Web resources, 500
- Digital visual interface (DVI), 360, 360–361
- Digitization devices, 17
- Direct attached storage (DAS), 253–254, 328–329
  - set-up diagram, 329
- Direct digital radiography (DDR), 19, 212

- Direct image capture, 150  
benefits of, 195  
elements of, 16–17, 194–195  
viewing levels, 195
- Disaster recovery, 262–267  
database backups, 312  
disk-to-disk strategy, 262–264  
and exam priority, 267  
and mean time to failure, 324–325  
operating system imaging, 311  
as PAC archive, 252  
removable media, 264–265, 311–312  
request for proposal (RFP) description, 51–52  
solution requirements, 266–267  
uninterruptible power supply (UPS), 310–311  
*See also* Security
- Discount rate, 131, 133
- Discrete cosine function (DCT), 233
- Disk-to-disk, disaster recovery strategy, 262–264
- Display protocols. *See* Hanging protocols
- Displays. *See* Monitors
- Distributed file system (DFS), 334
- Distributed model, 257–259, 403–404
- Doctor-patient relationship, radiologist malpractice, 164
- Double exposures, computed radiography (CR), 209–210
- Downtime contingency plan, 43
- DVD storage, 336–337
- Dynamic range control processing algorithm, 208
- E**
- Edge enhancement, 107
- EDictation, 471
- Education. *See* Digital teaching files (DTFs); Imaging informatics
- Electronic data interconnect (EDI), and XML, 375–376
- E-mail, authorized-only delivery, 314
- Encryption  
new technology file system (NTFS), 316  
user authentication, 314  
*See also* Security
- Enterprise medical record (EMR),  
purpose of, 11
- Ergonomics  
guidelines, 362  
workstation environment, 417, 420–422
- Error-checking  
data transmission, 150  
RIS-HIS interface, 192–193, 218
- Error recovery, by user, 416
- Ethernet  
compared to ATM, 292  
cards, 275  
complex network, example, 295–298  
gigabit Ethernet, 299–300  
resource reservation protocol (RSVP), 299  
shared, 276–278  
speeds, 285–286
- “Everyone” group, avoiding, 315
- ExaGrid, 264
- Exams, RIS scheduling of, 13–14
- Expenses, types of, 43–44
- Extensible markup language (XML), 375–376
- F**
- FAT files, 328
- Fault tolerance, PACS architecture, 254–257
- Federal regulations. *See* Legal/regulatory issues
- Fiber-optic cable, 274, 285
- Fibre Channel, 323–324  
and storage area network (SAN), 330–331
- Filers. *See* Network attached storage (NAS)
- File systems, 327–328  
FAT file system, 328  
new technology file system (NTFS), 316, 328
- Film, cost-savings, 32, 36, 140
- Film digitizers, 17, 196–197  
charge coupled device (CCD) scanners, 196, 525  
laser scanners, 197, 525
- Filmless operation, 101
- Film library, cost-savings, 140–141

- Financial analysis, 125–143
  - breakeven analysis, 135–137
  - cash flow, 126–128
  - cost reduction, 32, 36, 138–143
  - costs, 43–44, 137–138
  - expenses, types of, 43–44
  - internal rate of return, 133–135
  - net present value (NPV), 129–133
  - objectives of, 126
  - operating expenses, 43–44
  - PACS trade-off, 137
  - See also* Impact analysis
- First approximation approach, 399
- 5-MP display, 353, 370, 391–392, 530
- Fixed costs, 135
- Fixed noise, computed radiography (CR), 202
- Fluoroscopy, image compression, 241–242
- Food and Drug Administration (FDA)
  - Medical Devices Amendment (1976), 154–155
  - teleradiology policy statement, 155–156
- Frame-grabbing image acquisition, 16
- Frame mode, 107
- Fuji, computed radiography (CR) systems, 203–204
- Future demand, impact analysis, 39–41
- Future trends
  - image acquisition, 220–222, 225
  - image compression, 242–245
  - order entry, 491–492
  - PACS networks, 299–301
  - workstations, 422–425
- G**
- Gamma correction table, 359
- Gateways, 276
- Gigabit Ethernet, 299–300
- Gigabits, 285
- Goals
  - request for proposal (RFP), 47
  - strategic, 29
- Graphical processing unit (GPU), 178
- Graphical user interface (GUI), 179
- Graphics cards, 358–361
  - medical-grade, 358–359
  - slot type, 358–359
  - video random access memory (VRAM), 359
- Grayscale, 185–186
  - adjustment, 393, 396
  - future enhancement, 221
  - medical-grade grayscale, 361
- Grid servers, 334
- Grid storage, 264, 334
  - set-up example, 335
- H**
- Hand v. Tavera*, 164
- Hanging protocols, 104–106
  - default display protocols (DDPs), 105
  - defined, 399
  - methods of use, 105, 415
- Hard drives, 321–323
  - access time, 322
  - data transfer rate, 323
  - highest capacity, 323
  - interfaces, 323–324
  - operation of, 321–323
  - RAID controller, 325–326
- Hardware
  - servers, companies offering, 309
  - workstations, 391–392
- Health Care Financing Administration (HCFA), 156–157
  - Internet Security Policy, 156–157
- Health Insurance Portability and Accountability Act (HIPAA), 157–158
  - digital format requirement, 10, 53
  - impact of, 145–146
  - and order-entry systems, 490
  - privacy rule, 158, 506
  - security rule, 158
  - transaction rule, 158
- Health Level Seven (HL7)
  - and IHE initiative, 84, 102
  - as RIS-HIS standard, 193, 476–477
  - server, functions of, 250–251
- Helium-neon (HeNe) gas laser scanner, 197
- Hierarchical storage management (HSM), 254–255, 264–265, 334, 339–340
- Hospital information system (HIS)
  - method

- scheduling, 14
  - speech recognition system interface, 476–478
  - Human visual system (HVS)-based compression, 244–245
  - Hypertext markup language (HTML), 404
  - Hypertext transport protocol (HTTP), 375
- I**
- Image acquisition, 16–20, 189–225
    - common problems, 218–219
    - and computed radiography (CR), 18–20, 200–211
    - DICOM standards, 20, 50–51, 150, 189–194
    - and digital radiography (DR), 18–20, 211–217
    - direct capture, 16–17, 150, 194–195
    - and film digitizers, 17, 196–197
    - frame-grabbing, 16
    - future trends, 220–222, 225
    - image digitization, 17
    - mammography, digital, 364–367
    - with multidetector scanner, 108
    - non-digital methods, 196
    - quality control procedures, 219–220
    - request for proposal (RFP) description, 50–51
    - RIS-HIS data transfer modes, 193–194
    - secondary capture, 150
  - Image archive
    - clinical archive, 252–253
    - comparison of approaches, 265
    - disaster recovery, 252, 262–267
    - operation of, 20–21
    - software used, 20
    - standards, 150, 152
    - See also* Disaster recovery; Storage
  - Image compression, 231–245
    - benefits of, 526–527
    - computed radiography (CR), 240
    - computed tomography (CT), 241
    - and CPU power, 232
    - DICOM standard, 235–236
    - digital radiography (DR), 240
    - fluoroscopy, 241–242
    - future view, 242–245
    - irreversible (lossy) compression, 150, 232–235, 252, 527
    - JPEG2000, 235–236, 240–241, 252, 381, 527
    - magnetic resonance imaging (MRI), 241
    - mammography, 239–240
    - maximization of compression ratio, 237–238
    - medical images, bits per pixel, 235
    - nuclear medicine, 242
    - quality factor, 232
    - reversible (lossless) compression, 150, 231–232, 527
    - trade-offs, 236–237
    - ultrasound, 242
    - wavelet transform, 222, 233–234, 528–529
  - Image displays. *See* Monitors
  - Image distribution, 374–382
    - cost factors, 378
    - hospital intranet, 376–377
    - image distribution servers, 253
    - Internet. *See* Web-based distribution
    - mammography, digital, 442
  - Image management
    - ACR standard, 151–152
    - functions of, 21
    - image management specialist, 149
    - software, 399
    - workflow flowchart, 23
  - Image manager
    - autorouting capabilities, 53
    - functions of, 53
    - purge by, 53
    - request for proposal (RFP) description, 53
  - Image manipulation tools, 150, 393, 396–397
  - Image retrieval, mammography, digital, 367–368
  - Images
    - annotation capabilities, 150
    - overload, control of, 108–110
    - transmission methods, 529
    - See also* Digital images



- Image segmentation algorithms, 205–206
- Image servers, DICOM, 305–307
- Imaging informatics, 511–519
  - challenges related to, 513–514
  - defined, 511
  - imaging informatician, role of, 511
  - as radiology subspecialty, 512–513
  - and resident education, 515–519
- Imaging modalities
  - request for proposal (RFP)
    - descriptions, 48–49
    - types of, 48–49
- Imaging plate (IP) computed radiography (CR), 197–198
- Imaging technology, benefits of, 10
- Impact analysis, 32–37
  - on film-handling activities, 36–37
  - on film-related expense, 32, 36
  - of future demand, 39–41
  - market share analysis, 37–39
  - See also* Financial analysis
- Implementation plan
  - big bang approach, 57
  - phased approach, 57–58
  - request for proposal (RFP) description, 57–58
  - of strategic plan, 31–35
- Indirect digital radiography, 213
- Information objects (IO), 191
- Input devices, for computers, 174
- Installation of system, request for proposal (RFP) description, 59–60
- Insurance coverage
  - Medicare coverage, 458–505, 545–546
  - out-of-state malpractice incidents, 166
  - 3-D imaging, 460–463
- Integrating the Healthcare Enterprise (IHE) initiative
  - on compatibility, 102
  - purpose of, 84
- Interfaces
  - hard drive, 323–324
  - RIS-HIS error-checking, 192–193, 218
  - RIS-PACS interface, 22–23
  - speech recognition system interface, 476–478
  - See also* User interface
- Internal rate of return, 133–135
  - tables, 134–135
- International Classification of Disease (ICD-9), order entry codes, 490
- International Organization for Standardization (ISO), network model, 270–272
- International Shoe v. Washington*, 165
- Internet, 374–376
  - browsers, 404
  - elements of, 374
  - protocols, 375
  - virtual private networks (VPNs), 300
- Internet Security Policy
  - authentication procedures, 157
  - Health Care Financing Administration (HCFA), 156–157
- Interpretation process, 106–110
  - clinical information, retrieval of, 106–108
  - image overload, control of, 108–110
  - volumetric imaging issues, 109–110
- Intranet
  - components of, 376
  - cost factors, 378
  - image distribution, 376–377
  - pros/cons of, 377
  - protocols, 377
- Introduction section, request for proposal (RFP), 47
- IPX protocol, 275–276, 278
- Iron Mountain, 263, 265
- Irrelevant costs, 127–128
- Irreversible (lossy) compression, 150, 232–235, 252, 527
  - transformation in, 233
- ISDN lines, 529
- Itemized pricing, vendors, 64
- J**
- Java, 180
- Joint Photographic Experts Group (JPEG), 233
  - compression standard, 235–236, 240–241, 252, 381, 528
- Jukebox, 20
- Just noticeable difference (JND), 354

**K**

- Key images, digital teaching files (DTFs), 507
- Kodak, computed radiography (CR) systems, 202–204
- Konica, computed radiography (CR), 203

**L**

- Landscape mode, 360
- Language model, speech recognition system, 472–473
- Languages, computer, 179–180
- Laser scanners, 197
- Layers
  - file system, 327–328
  - layer 3 switching, 300
  - networks, 272–275
- Legal/regulatory issues, 145–167
  - American College of Radiology (ACR) standard, 151–152, 544–545, 550–560
  - American Medical Association (AMA) policy, 152–153
  - confidentiality/security, 167, 544
  - credentialing, 543
  - Digital Imaging and Communications in Medicine (DICOM), 20, 84, 149
  - Food and Drug Administration (FDA) regulation, 154–156
  - Health Care Financing Administration (HCFA) policy, 156–157
  - Health Insurance Portability and Accountability Act (HIPAA), 157–158
  - Health Level Seven (HL7), 84, 102, 193
  - insurance coverage issues, 165–166, 545–546
  - malpractice issues, 163–165, 543–544
  - matrix categories, 149
  - National Electrical Manufacturers Association (NEMA), 149
  - record-keeping regulations, 166–167
  - standards/practice guidelines, 544–545
  - state initiatives, 163
  - state licensure, 158–163, 542–543
- Legatto, 265

- Licensure, state requirements, 158–163, 543
- Life cycle management, 254–255, 264–265
- Lightbulb effect, 209
- Lighting environment, for workstations, 420–421
- Light photon noise, computed radiography (CR), 202
- Linear media, 176
- Line spread function (LSF), computed radiography (CR), 201
- Linked stack mode, 107
- Liquid crystal displays (LCDs), 348, 391–392
  - digital visual interface (DVI), 360
  - luminescence, 356–357
  - medical-grade, 357–358
  - mount, 358
- Lists, digital teaching files (DTFs), 507
- Load-balancing model, 305–307
- Lockouts, 414
- Lookup tables, 206, 359
- Lossless compression, 150, 231–232, 527
- Lossy compression, 150, 232–235, 252, 527
  - transformation in, 233
- Lost film, PACS cost-saving, 141
- Luminance, monitors, 354–355, 361
- Lumisys, computed radiography (CR) systems, 203

**M**

- Macros, speech recognition system, 475
- Magnetic resonance imaging (MRI), image compression, 241
- Magneto-optical (MO) drives, 337
- Mainframe systems, 386–387
- Maintenance/support of system, request for proposal (RFP) description, 61–62
- Malpractice claims, 163–165
  - out-of-state actions, 165–166, 543–544
  - state laws, 165
  - successful, elements of, 163–164
- Mammography, false negatives, 435–436
- Mammography, digital, 363–371, 441–443
  - challenges related to, 434–435, 441–442

- Mammography, digital (*cont.*)  
  computer-aided diagnosis (CAD),  
    423–424, 436–441  
  DICOM standard, 366–367  
  image acquisition, 364–367  
  compression, 239–240  
  image display, 368–369  
  image distribution, 442  
  image retrieval, 367–368  
  levels of output, 365–366  
  multiple systems, 370–371  
  printing, 368–369  
  quality assurance mandate, 89  
  workstation, 370
- Management of system  
  and strategic plan, 42–43  
  *See also* Image management; Image  
  manager
- Market share analysis, 37–39
- Matrix categories  
  future advances, 221  
  image interpretation, 149
- Maximum-intensity projection (MIP)  
  images, 452–453
- McKinney v. Schlatter*, 164
- Mean time to failure, 324–325, 338
- Mean time to recovery, 326
- Medical Devices Amendment (1976),  
  154–155
- Medical Diagnostic Imaging Support  
  (MDIS), 402–403
- Medical-grade grayscale, 361
- Medical image resource center (MIRC),  
  496–500  
  community diagram, 497  
  DICOM images, 499  
  indexing and searching, 498  
  international community, 499  
  objectives of, 497  
  teaching files compatibility, 504–505
- Medicare coverage  
  teleradiology criteria, 545–546  
  3-D imaging, 458–463
- Megabits, 284
- Megabytes, 284, 320
- Memory  
  computer, types of, 175–176  
  random access memory (RAM), 391
- Mental model, user interface, 408–409
- Microsoft SQL Server, 20
- Minimum-intensity projection (MinIP)  
  images, 452
- Mirrored storage system, 100–101
- Modality worklist (MWL), 251–252
- Modulation transfer function (MTF),  
  computed radiography (CR),  
    201
- Monitors, 348–362  
  calibration, 355–358  
  development of, 348  
  digital visual interface (DVI),  
    360–361  
  5-MP display, 353, 370, 391–392,  
    530  
  graphics cards, 358–361  
  liquid crystal displays (LCDs), 348  
  lookup tables, 359  
  luminance, 354–355, 361  
  mammography, digital, 368–369  
  medical-grade grayscale, 361  
  mounting of, 358  
  number for PACS system, 349–350  
  operating room displays, 362  
  physical size, 350–351  
  pivoting, 360  
  pixel pitch, 350–352  
  radiologic imaging requirements,  
    351–354  
  resolution/pitch, 350–351, 353, 530  
  standards, 150
- Moore, Gordon, 320
- Moore's law, 230, 320
- Motherboard, 177
- Mottle, computed radiography (CR),  
  202
- Multidetector computed tomography  
  (MDCT), 396, 512  
  scanner, 108  
  3-D images, 450–451
- Multi-institutional PACS, 115–116  
  WAN virtual radiology department,  
    115–116
- Multiscale image contrast amplification  
  (MUSICA), 208
- Multiserver model, 304, 306
- MUMPS operating system, 178

- N**
- National Electrical Manufacturers Association (NEMA), 149
    - on image storage, 149, 190–191
  - NaturallySpeaking, 471
  - Navigation process, 393–396, 415
  - Near-line storage, 320–321
  - Near line-storage, 51–53
  - Near-line storage
    - automatic archive, 53
    - optical storage, 336–337
    - request for proposal (RFP) description, 51–52
    - tape storage, 337–338
  - Needle-phosphor technology, 205
  - Negligence, malpractice issue, 164–165
  - Net present value (NPV), 129–133
    - formula, 130
    - internal rate of return, 133–135
    - risks of PACS, 131–133
  - Network Appliance, 264
  - Network attached storage (NAS), 253–254, 263–264, 332–334
    - companies offering, 264
    - distributed file system (DFS), 334
    - set-up diagrams, 329, 333
  - Network infrastructure, request for proposal (RFP) description, 56
  - Networks
    - bandwidth, 284–287
    - bottlenecks, 288–289
    - bridging, 278–280, 282
    - building, aspects of, 287–289
    - cabling, 274
    - cards, types of, 275
    - data link layer, 275
    - future trends, 299–301
    - gateways, 276
    - ISO standard, 270–272
    - packets, 286–287, 292
    - physical layer, 274
    - postal system example, 273
    - protocols, 275–276
    - routing, 280–282
    - shared Ethernet example, 276–278
    - simple, components of, 271–275
    - switching, 283, 287, 289
    - upgrading, 289–290
    - See also* PACS networks
  - New technology file system (NTFS), 316, 328
  - Noise, computed radiography (CR), 202
  - Nuclear medicine, image compression, 242
- O**
- Offline storage, 321
  - Online storage, 320
  - Operating expenses
    - financial analysis, 43–44
    - types of, 43
  - Operating room displays, 362
  - Operating systems
    - operating systems imaging, 311
    - patches, 312–313
    - security, 312–316
    - types of, 178–179, 309–310
  - Operational fit, and PACS architecture, 256
  - Optical storage, 336–337
    - compact disc (CD) storage, 336
    - DVD storage, 336–337
    - magneto-optical (MO) drives, 337
    - storing/handling, 337
  - Oracle, 20
  - Order entry, 483–492
    - advantages of, 488–491
    - basic principles, 484
    - decision support, 490
    - design of system, 485–488
    - future trends, 491–492
    - and HIPAA compliance, 490
    - ICD-9 code, 490
    - limitations of, 491
  - Orex CR system, 203
  - Organic light emitting diode (OLED), 425
  - OS/2 operating system, 178–179
  - Output devices, for computers, 175
- P**
- Packets, networks, 286–287, 292
  - PACS
    - benefits of, 3–4
    - database manager, 20

- PACS (*cont.*)
- financial analysis for, 43–44
  - image acquisition, 16–20
  - image archive, 20–21
  - image distribution, 374–382
  - image management, 21
  - image retrieval tools, 106–107
  - image routing, 14
  - multi-institutional PACS, 115–116
  - request for proposal (RFP), 45–71
  - and RIS, 4, 13–15, 22–24
  - and speech recognition system, 4
  - strategic plan, 28–41
  - technical strategic plan, 41–43
  - Web-based system, 24–25
  - workflow engineering, 74–92
  - workflow flowchart, 23
- PACS architecture, 250–267
- centralized model, 259–261, 402–404
  - combined distributed/centralized model, 261–262
  - distributed model, 257–259, 403–404
  - fault tolerance, 255–257
  - and operational fit, 256
  - scalability, 256
  - and versatility, 256
- PACS components, 15–16, 250–255
- archive system, 252–253
  - clustering of, 255
  - database, 250
  - DICOM server, 252
  - Health Level Seven (HL7) server, 250–251
  - image distribution servers, 253
  - modality worklist (MWL), 251–252
  - monitors, 348–362
  - servers, 250, 304–310
  - storage, 253–254
  - workstations, 24–25, 253
- PACS networks, 291–301
- asynchronous transfer mode (ATM), 292–293
  - complex, components of, 295–298
  - disaster recovery, 310–312
  - future trends, 299–301
  - security measures, 312–316
  - simple, example of, 294–295
  - star typology, 297–298
  - compared to typical networks, 291
  - wide-area networks (WANs), 293–294
- Paperless operation, 101
- Passwords, 313–314
- construction of, 314
  - protection of screensaver, 315
  - speech recognition system, 474
- Patches, operating system, 312–313
- Patient information
- appointment scheduling, 487–488
  - error-checking, 192–193, 218
  - indications, lists of, 485–486
  - order entry system, 485–488
- Patient outcomes, improved, 143
- Payback analysis
- PAC pros/cons, 129
  - payback period, 128–129
- PCI Express slot, 359
- Penalty box, 218
- Phased approach, implementation, 57–58
- Photomultiplier tube (PMT), 198–199
- Photostimulable phosphor (PSP)-based image sensor, 197–199
- Physical layer, networks, 274
- Physical security, 314–315
- Physicians, PACS impact on, 112–114
- Picture archiving and communication systems (PACS). *See* entries under *PACS*
- Pivoting, 360
- Pixels, 184–185
- monitor pixel pitch, 350–352
  - subpixels, 355
- Portrait mode, 352, 360
- Powerquest, 311
- PowerScribe, 470
- Power supply, uninterruptible, 310–311
- Pre-fetch, 103–104
- process of, 14
  - strategies, 103–104
- Presentation tools, digital teaching files (DTFs), 510
- Printing, mammography, digital, 368–369
- Privacy
- digital teaching files (DTFs), 506–507
  - HIPAA rule, 158, 506
- Privacy Act (1974), 156–157

- Processing, computer function, 176–177
- Productivity, quality assurance, 89–91
- Projection methods
- computed radiography (CR), 18–20, 200–211
  - digital radiography (DR), 18–20, 211–217
  - film digitizers, 17, 196–197
- “Promotion of Quality Telemedicine, The”, 152
- Protocols
- HL7, 476–477
  - Internet, 375
  - intranet, 377
  - networks, 275–276
- Protocol stack, networks, 272
- Q**
- Quality assurance-quality control (QA-QC)
- digital teaching files (DTFs), 508
  - federal mandate, 89
  - image acquisition, 219–220
  - in productivity/workflow, 89–91
- Queuing analysis, 85–88
- of arrival process, 86–87
  - issues to address, 88
  - and queue characteristics, 87–88
  - of service process, 87
- R**
- Radiologic technologists (RTs)
- PACS, impact on job, 80–83
  - in paper-based environments, 101
  - time, cost-savings, 142
  - workflow of. *See* Workflow analysis
- Radiologists
- education. *See* Digital teaching files (DTFs); Imaging informatics
  - malpractice issues, 163–165, 543–544
- Radiology
- necessity of, 184
  - traditional digital infrastructure, 13
  - See also* Teleradiology
- Radiology information system (RIS). *See* RIS
- RAID. *See* Redundant array of inexpensive disks (RAID)
- Random access memory (RAM), 176, 391
- Readiness assessment, for PACS, 30–31, 41–42
- Reading room, PACS set-up, 116–117
- Record-keeping, legal requirements, 166–167
- Redos, cost-savings, 141
- Redundant array of inexpensive disks (RAID), 20–21, 176, 325–326
- hierarchical storage management (HSM), 254–255
  - online storage, 51
  - versus pre-fetch, 103–104
  - pros/cons of, 100
  - RAID 0, 325
  - RAID 1, 325
  - RAID 3, 325
  - RAID 4, 325–326
  - RAID 5, 326
  - RAID 6, 326
  - RAID 10, 325
- Regional Bell Operating Companies (RBOCs), leased networks, 293–294, 296–297
- Regius, 203
- Regulatory issues. *See* Legal/regulatory issues
- Removable media
- companies offering, 265, 311–312
  - disaster recovery strategy, 264–265
  - security measures, 315
- Reporting process, 110–112
- speech recognition system, 110–111, 475, 482
  - structured reporting, 111–112
- Request for proposal (RFP), 45–71
- appendices, 63
  - basic guidelines, 46–47
  - clinical operations overview, 47–49
  - clinical review section, 49
  - clinical review workstation, 55–56
  - contract negotiation, 70–71
  - database manager, 52–53
  - decision model, use of, 66–69
  - diagnostic review, 49
  - diagnostic review workstation, 54–55
  - DICOM Query/Retrieve functionality, 53

- Request for proposal (RFP) (*cont.*)
  - disaster recovery subsystem, 51–52
  - distribution of, 65
  - image acquisition, 50–51
  - image distribution via Web server, 55
  - image manager, 53
  - imaging modalities, workflow
    - description, 48
  - implementation plan, 57–58
  - installation of system, 59–60
  - introduction section, 47
  - near-line storage, 51–52
  - network infrastructure description, 56
  - objectives of, 46
  - response time, 60–61
  - RIS interface, 53–54
  - site-specific operations, 48
  - site visits, 67, 70
  - storage system, 51
  - strategic goals, 47
  - system architecture, 50
  - system reliability/uptime, 60–61
  - system throughput/performance, 57
  - technologist Q/A workstation, 56
  - telerradiology operations, 50
  - training program, 58–59
  - vendor evaluation/selection, 65–70
  - vendor information/selection, 62–63
  - vendor maintenance/support, 61–62
  - vendor response form, 63–64
  - warranty period, 61
- Reset function, 416
- Resolution
  - computed radiography (CR), 201–202
  - monitors, 350–351, 353, 530
- Resource reservation protocol (RSVP), 299
- Response time, request for proposal (RFP) description, 60–61
- Return-on-investment (ROI), 37–39
- Reverse contrast scale, 206–207
- Reverse telerradiology, 535
- Reversible (lossless) compression, 150, 231–232, 527
- Review workstations, 25, 390
- RIS, 13–15, 22–24
  - error-checking HIS interface, 192–193, 218
  - functions of, 13
  - and HL7 server, 250–251, 476–477
  - image routing, 14
  - operational features, 13–15
  - PACS interface with, 22–23
  - pre-fetching, 14
  - request for proposal (RFP) description, 53–54
  - RIS-HIS data transfer mode, 193–194
  - speech recognition system interface, 476–478
- Risks
  - net present value (NPV) calculation, 131–133
  - of PAC implementation, 130–133
- Robustness, PACS architecture, 258
- Routing, networks, 280–282
- S**
- Scalability, PACS architecture, 256
- ScanHead CR, 204–205
- Scanners
  - charge coupled device (CCD) scanners, 196, 525
  - laser scanners, 197, 525
- Scheduling
  - hospital information system (HIS) method, 14
  - and RIS, 13–14
  - Web-based method, 14
- Screen-film, compared to computed radiography (CR), 199–200, 202
- Screensaver, password protection, 315
- Secondary image capture, 150
- Sectors, hard drive, 321–322
- Secure sockets layers (SSLs), 381
- Security
  - administrator accounts, 315
  - default shares, disabling, 316
  - e-mail, authorized-only delivery, 314
  - encrypted user authentication, 314
  - “everyone” group, avoiding, 315
  - new technology file system (NTFS), 316
  - operating system patches, 312–313
  - passwords, 313–314
  - physical security, 314–315
  - screensaver password protection, 315

- speech recognition system, 474
- standards related to, 158, 167
- virtual private networks (VPNs), 300
- viruses, 316–317
- virus protection, 313
- for Windows-based systems, 315–316
- Servers, 250, 304–310
  - clustered failover model, 307–308
  - DICOM servers, 252, 305–307
  - grid servers, 334
  - hardware, 309
  - load-balancing model, 305–307
  - multiserver model, 304, 306
  - operating systems, 309–310
  - single server model, 304–305
  - Web servers, 304–305, 378–382
- Service class providers (SCP), 191
- Service class users (SCU), 191
- Service process, queuing analysis of, 87
- Set Partitioning in Hierarchical Trees (SPIHT), 233
- Showstoppers, 131
- Signal-to-noise ratio (SNR), 363
- Silicon field-effect transistor (FET), 212
- Simple mail transfer protocol (SMTP), 314
- Simple network management protocol (SNMP), 327
- Site visits, scope of, 67, 70
- Slots, graphics cards, 358–359
- Small computer system interface (SCSI), 323, 331–332
- Smoothing, 107
- Snappy Appliances, 264
- Software, 392–398
  - basic functions of, 180–181
  - image management, 399
  - image manipulation tools, 393, 396–397
  - navigation process, 393–396
- Solid-state diode laser scanner, 197
- Space requirements, assessment of, 42
- Spatial frequency processing algorithms, 207–208
- Spatial resolution
  - computed radiography (CR), 201
  - file size trade-off, 201
- SpeechMagic, 471
- Speech Q, 470–471
- Speech recognition system, 110–111, 467–482
  - application component, 473
  - capture of report, 15
  - companies offering, 470–471
  - cost-savings, 481
  - customizable fields, 475
  - desktop integration, 476
  - development of, 468–469, 472
  - implementation of, 478, 481
  - input devices, 471, 478
  - language model, 472–473
  - macros, 475
  - necessity of, 4, 111, 468–469
  - operational checks, 480–481
  - operational issues, 480
  - radiologist training, 478–479
  - reports, 110–111, 475, 482
  - resistance to use, 469
  - RIS-HIS interface, 476–478
  - speech engine, 471–472
  - system security, 474
  - technical support, 479
  - templates, 475
  - voice navigation, 479
- Speed
  - abnormal findings information, 114
  - and emergency teleradiology, 536–538
  - workstations, 392
- Stack view mode, 107, 402
- Standards
  - International Organization for Standardization (ISO), 270–272
  - teleradiology. *See* Legal/regulatory issues
- Star typology, PACS networks, 297–298
- States
  - licensure requirements, 158–163
  - malpractice insurance coverage, 165–166, 543–544
  - teleradiology initiatives, 163
- Static mode, 107
- Statis softcopy interpretation, 107
- Storage, 253–254
  - American College of Radiology (ACR) standard, 166–167, 190–191
  - archiving, 252–253



- Storage (*cont.*)
- content addressable storage (CAS), 264, 334–335
  - direct attached storage (DAS), 253–254, 328
  - file system, 327–328
  - formats, 320–321
  - grid storage, 264, 334
  - hard drives, 321–324
  - hierarchical storage management (HSM), 254–255, 339–340
  - image manager, 53
  - mean time to failure, 324–325
  - mirrored system, 100–101
  - Moore's law, 230, 320
  - National Electrical Manufacturers Association (NEMA) format, 149, 190–191
  - near-line storage, 51–53
  - network attached storage (NAS), 253–254, 263–264, 332–334
  - optical storage, 336–337
  - redundant array of inexpensive disks (RAID), 20–21, 51, 100, 254–255, 325–326
  - request for proposal (RFP) description, 51
  - storage area network (SAN), 253–254, 263, 328–332
  - and switching vendors, 267
  - tape storage, 337–338
  - terminology (powers of 10), 320
- Storage area density, 320
- Storage area network (SAN), 253–254, 328–332
- advantages/limitations, 330–331
  - and Fibre Channel, 330–331
  - set-up diagrams, 329, 331
  - storage companies, 263
- Storage Management Initiative Specification (SNIA), 327
- StorageTek, 265
- Strategic plan, 28–41
- documentation requirements, 28–29
  - downtime contingency plan, 43
  - future demands, analysis of, 39–41
  - goals/objectives, PACS alignment to, 29
  - hospital roles, contributions of, 30–31, 41
  - impact analysis, 32–37
  - implementation plan, 31–35
  - management of system, 42–43
  - market share analysis, 37–39
  - PACS readiness assessment, 30–31, 41–42
  - technical, 41–43
- Structured reporting, 111–112
- Subpixels, 355
- Sunk costs, 127
- Surface rendering, 3-D imaging, 452–453
- Switching
- layer 3 switching, 300
  - networks, 283, 287, 289
- Sybase, 20
- System architecture, request for proposal (RFP) description, 50
- System reliability, request for proposal (RFP) description, 60–61
- System throughput, request for proposal (RFP) description, 57
- T**
- T1 line, 293, 296–297, 529
- Talk Technology TalkStation, 470
- Tape storage, 337–338
- capacities, 338
  - mean time between failures, 338
- Teaching files. *See* Digital teaching files (DTFs)
- Technologist Q/A workstation, request for proposal (RFP) description, 56
- Teleradiology, 523–547
- accuracy of, 530–532
  - applications of, 533–536
  - emergency manpower, 538–539
  - image acquisition, 524–526
  - image compression, 526–529
  - image interpretation, 530
  - image transmission, 529
  - interpretation site, layout, 524
  - job satisfaction, 541–542
  - licensure/credentialing, 542–543
  - operations section, request for proposal (RFP), 50
  - patient care, 540–541

- radiologist/clinician relationships, 539–540
  - regulations. *See* Legal/regulatory issues
  - reverse teleradiology, 535
  - speed, importance of, 536–538
  - Templates, speech recognition system, 475
  - Thick client application, 253
  - Thin client application, 253, 404
  - Thin-film-transistor (TFT), 212–213
  - 3-D imaging, 447–463
    - average-intensity projection (AveIP) images, 452
    - benefits of, 448–449
    - computed tomography (CT) scans, 449–450
    - cost of service, 458–460
    - image compression limitations, 243
    - imaging service staffing, 457–458
    - insurance reimbursement, 460–463
    - issues related to, 109
    - maximum-intensity projection (MIP) images, 452–453
    - Medicare coverage, 458–463
    - minimum-intensity projection (MinIP) images, 452
    - multidetector computed tomography (MDCT), 450–451
    - and multidetector CT (MDCT), 450–451
    - in PACS hospital environment, example of, 453–457
    - post-processed images, creating, 451–457
    - surface/volume rendering, 452–453
    - 3-D exam modifier, 455–457
    - volumetric navigation, 107, 109
    - volumetric viewing, 452
    - voxels, 184
  - Time-motion studies, 400
  - Tivoli, 265
  - Token ring cards, 275
  - Topology, networks, 275–276
  - Tracks, hard drive, 321
  - Training, request for proposal (RFP) description, 58–59
  - Transaction rule, HIPAA, 158
  - Transfer communication
    - protocol/Internet (TCP/IP) protocol, 275–276, 278, 374–375
  - Transformations
    - digital mammography display, 368
    - in irreversible (lossy) compression, 233
  - Trojan horse, 317
- U**
- Ultrasound, image compression, 242
  - Undo, 416
  - Uninterruptible power supply (UPS), 310–311
  - Universal Serial Bus 2 (USB-2), 178
  - Unshielded twisted pair (UTP) cable, 274, 285
  - Uptime
    - definitions, 60
    - request for proposal (RFP) description, 60
  - User interface, 407–417
    - early studies, 407–408
    - error recovery, 416
    - hanging protocols, 415
    - mental model as basis, 408–409
    - studies related to, 399–402
    - worklists, 410–415
- V**
- Variable costs, 135–136
  - Vendors
    - contract negotiation, 70–71
    - evaluation/selection process, 65–70
    - final decision-making about, 67, 70
    - and itemized pricing, 64
    - maintenance/support, 61–62
    - request for proposal (RFP) response form, 63–64
    - request for proposal information/selection description, 62–63
    - switching and stored data, 267
    - warranty period, 61
  - Veritas, 308
  - Versatility, and PACS architecture, 256
  - ViaVoice, 471
  - Video cards. *See* Graphics cards

- Video Electronics Standards Association (VESA) mount, 358
- Video random access memory (VRAM), 359
- View box, PAC alternative to, 104–106
- Virtual private networks (VPNs), 300, 381
- Viruses, 316–317
  - operation of, 316–317
  - Trojan horse, 317
  - virus protection, 313
  - worms, 317
- Visual C, 179–180
- Voice recognition. *See* Speech recognition system
- Volumetric visualization. *See* 3-D imaging
- Voxels, 184
- VoxReports, 471
- W**
- Warranty period, request for proposal (RFP) description, 61
- Wavelet transform, 222, 233–234, 528–529
- Web-based image distribution, 374–382, 424–425
  - cost factors, 378
  - Web servers, 378–379
- Web-based PACS
  - functions of, 24–25
  - scheduling, 14
  - workstations, 25
- Web servers, 378–379
  - client device, 379–380
  - functions of, 304–305, 379
  - hospital example, 380–382
  - image distribution request for proposal (RFP) description, 55
  - needs analysis, 43–44
- Wide-area networks (WANs), 293–294
  - leased circuits, 293–294
  - virtual radiology department, 115–116
- Winchester Systems, 263
- Windows-based systems, security guidelines, 315–316
- Window shade display, 422–423
- Workflow analysis, 76–88
  - auditing workflow process, 117–118
  - for clerical personnel, 77–79
  - communication of abnormality, 114
  - interpretation process, 106–110
  - multi-institutional PACS, 115–116
  - PACS impact on physicians/clinicians, 112–114
  - queuing analysis, 85–88
  - for radiologic technologists (RTs), 79–83
  - reading room environment, 116–117
  - reporting methods, 110–112
- Workflow control. *See* Image management
- Workflow engineering, 74–92
  - and digital image environment, 75
  - hanging protocols/display, 104–106
  - industrialized approach to, 75
  - interactive workstations, 110–111
  - for new type of technologist, 85
  - pre-fetch strategies, 103–104
  - purpose of, 74
  - quality assurance, 89–91
  - structured reporting, 111–112
  - workflow analysis, 76–88
  - worklists, 101–102
  - workstation tools, 106–110
- Worklists, 101–102
  - advantages of, 193–194
  - contents of, 393, 410–415
  - DICOM standards, 102, 193–194
  - generation of, 102
  - modality worklist (MWL), 251–252
- Workstations, 24–25, 253
  - choosing, questions to ask, 418–420
  - clinical review, 25
  - and design of system, 402–406, 418–420
  - development of, 386–388
  - diagnostic, 24–25, 390
  - environment for set-up, 417, 420–422
  - ergonomic factors, 417, 420–422
  - functions of, 388–390
  - future trends, 422–425
  - grayscale adjustment, 393, 396
  - hardware, 391–392
  - lighting environment of, 420–421
  - for mammography, digital, 370
  - operation, studies of, 399–402

review, 390  
software, 392–398  
speed, 392  
thin and thick client, 253, 404  
user interface, 399–402, 407–417  
Web-based, 25, 253  
workstation tools, 106–110

Write once read many (WORM),  
337

## X

XPress, 203

X-ray dose-dependent noise, computed  
radiography (CR), 202