

BASIC RESEARCH IN
INFORMATION SCIENCE
AND **TECHNOLOGY**
FOR **AIR FORCE NEEDS**

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

BASIC RESEARCH IN
INFORMATION SCIENCE
AND **TECHNOLOGY**
FOR **AIR FORCE NEEDS**

Committee on Directions for the AFOSR Mathematics and Space Sciences
Directorate Related to Information Science and Technology

Board on Mathematical Sciences and Their Applications

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

THE NATIONAL ACADEMIES PRESS
Washington, D.C.
www.nap.edu

THE NATIONAL ACADEMIES PRESS 500 Fifth Street, N.W. Washington, DC 20001

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

This study was supported by Contract No. F1ATA04295M001 between the National Academy of Sciences and the Air Force Office of Scientific Research. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the organizations or agencies that provided support for the project.

International Standard Book Number 0-309-10031-3

Copies of this report are available from The National Academies Press, 500 Fifth Street, N.W., Lockbox 285, Washington, DC 20055; (800) 624-6242 or (202) 334-3313 (in the Washington metropolitan area); Internet, <http://www.nap.edu>.

Copyright 2006 by the National Academy of Sciences. All rights reserved.

Printed in the United States of America

THE NATIONAL ACADEMIES

Advisers to the Nation on Science, Engineering, and Medicine

The **National Academy of Sciences** is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Ralph J. Cicerone is president of the National Academy of Sciences.

The **National Academy of Engineering** was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. Wm. A. Wulf is president of the National Academy of Engineering.

The **Institute of Medicine** was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Harvey V. Fineberg is president of the Institute of Medicine.

The **National Research Council** was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Ralph J. Cicerone and Dr. Wm. A. Wulf are chair and vice chair, respectively, of the National Research Council.

www.national-academies.org

**COMMITTEE ON DIRECTIONS FOR THE AFOSR MATHEMATICS
AND SPACE SCIENCES DIRECTORATE RELATED TO
INFORMATION SCIENCE AND TECHNOLOGY**

ALAN J. McLAUGHLIN, MIT Lincoln Laboratory (retired), *Chair*
RUZENA K. BAJCSY, University of California at Berkeley
ELWYN BERLEKAMP, University of California at Berkeley
PHILIP A. BERNSTEIN, Microsoft Corporation
ROGER W. BROCKETT, Harvard University
VINCENT CHAN, Massachusetts Institute of Technology
STEPHEN CROSS, Georgia Institute of Technology
EDWARD FELTEN, Princeton University
OSCAR GARCIA, University of North Texas
W. DAVID KELTON, University of Cincinnati
KLARA NAHRSTEDT, University of Illinois at Urbana-Champaign
PRABHAKAR RAGHAVAN, Yahoo, Inc.
RONALD W. SCHAFER, Hewlett-Packard Laboratories

Staff

SCOTT WEIDMAN, Director, Board on Mathematical Sciences and
Their Applications
BARBARA WRIGHT, Administrative Assistant

**BOARD ON MATHEMATICAL SCIENCES
AND THEIR APPLICATIONS**

DAVID W. McLAUGHLIN, New York University, *Chair*
TANYA STYBLO BEDER, Tribeca Investments, LLC
PATRICK L. BROCKETT, University of Texas at Austin
ARAVINDA CHAKRAVARTI, Johns Hopkins University School of
Medicine
PHILLIP COLELLA, Lawrence Berkeley National Laboratory
LAWRENCE CRAIG EVANS, University of California at Berkeley
JOHN E. HOPCROFT, Cornell University
ROBERT KASS, Carnegie Mellon University
KATHRYN B. LASKEY, George Mason University
C. DAVID LEVERMORE, University of Maryland
ROBERT LIPSHUTZ, Affymetrix, Inc.
CHARLES M. LUCAS, AIG
CHARLES MANSKI, Northwestern University
JOYCE McLAUGHLIN, Rensselaer Polytechnic Institute
PRABHAKAR RAGHAVAN, Yahoo, Inc.
STEPHEN M. ROBINSON, University of Wisconsin-Madison
EDWARD WEGMAN, George Mason University
DETLOF VON WINTERFELDT, University of Southern California

Staff

SCOTT WEIDMAN, Director
BARBARA WRIGHT, Administrative Assistant

For more information on BMSA, see its Web site at <http://www7.nationalacademies.org/bms/>, write to BMSA, National Research Council, 500 Fifth Street, N.W., Washington, DC 20001, call at (202) 334-2421, or e-mail at bmsa@nas.edu.

Acknowledgments

This report has been reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise, in accordance with procedures approved by the NRC's Report Review Committee. The purpose of this independent review is to provide candid and critical comments that will assist the institution in making its published report as sound as possible and to ensure that the report meets institutional standards for objectivity, evidence, and responsiveness to the study charge. The review comments and draft manuscript remain confidential to protect the integrity of the deliberative process. We wish to thank the following individuals for their review of this report:

C. William Gear, Princeton University,
Eric Horvitz, Microsoft Research,
John W. Lyons, U.S. Army Research Laboratory (retired),
Debasis Mitra, Bell Laboratories,
S. Shankara Sastry, University of California at Berkeley,
William Scherlis, Carnegie Mellon University, and
Sheila E. Widnall, Massachusetts Institute of Technology.

Although the reviewers listed above have provided many constructive comments and suggestions, they were not asked to endorse the conclusions or recommendations nor did they see the final draft of the report before its release. The review of this report was overseen by William H. Press, Los Alamos National Laboratory. Appointed by the National Research Council, he was responsible for making certain that an inde-

pendent examination of this report was carried out in accordance with institutional procedures and that all review comments were carefully considered. Responsibility for the final content of this report rests entirely with the authoring committee and the institution.

The committee thanks members of the AFOSR staff of the Air Force Research Laboratory's Information Directorate; staff of the Air Combat Command; Thomas Cruse, Chief Technologist of the Air Force Research Laboratory; and Shankara Sastry and Janos Sztipanovits of the Air Force Scientific Advisory Board for their helpful discussions and inputs to this study.

Contents

EXECUTIVE SUMMARY	1
1 INTRODUCTION	17
2 BACKGROUND	22
Overview of Air Force Goals That Rely on IS&T Research, 27	
The R&D Response: Current Directions, 29	
3 BASIC RESEARCH FOR AIR FORCE NETWORK SYSTEMS AND COMMUNICATIONS	35
Types and Characteristics of Communication and Network Services Needed in the Future, 35	
Technical Challenges Posed by Future Air Force Networks and Communications Systems, 37	
Challenges for Future Air Force Communications Systems, 37	
Challenges for Future Air Force Networks, 41	
Recommended Basic Research Areas in Support of Air Force Networks and Communications, 48	
Satellite Communications and Data Networking, 49	
Radio Communications and Networking, 50	
Free-Space Optical Networks, 51	

4	BASIC RESEARCH FOR AIR FORCE SOFTWARE	53
	Coevolution of Air Force Concepts of Operations and System Architectures, 55	
	Software Behavior Envelopes, 58	
	Evolvability Throughout the Life Cycle, 59	
5	BASIC RESEARCH FOR AIR FORCE INFORMATION MANAGEMENT AND INTEGRATION	61
	Background, 61	
	Major Information Management Challenges for Air Force IS&T, 63	
	Recommended Basic Research in Information Management and Integration, 67	
6	BASIC RESEARCH FOR HUMAN INTERACTIONS WITH AIR FORCE IS&T SYSTEMS	71
	Challenges Posed by Human Interactions with Air Force IS&T Systems, 72	
	Scope of the Challenge, 72	
	Interfaces for Air Force Decision Makers, 74	
	Machine Learning to Support HSI, 76	
	Simulation as a Design and Training Tool for HSI, 77	
	Basic Research Recommendations for HSI for Air Force IS&T Systems, 77	
7	PRIORITIES IN BASIC IS&T RESEARCH FOR THE AIR FORCE	80
	A Model for Air Force IS&T Basic Research, 80	
	Recommended Basic Research Priorities in IS&T, 84	
8	FUNDING MECHANISMS	87
9	FUTURE CONSIDERATIONS	91
	Program Organization, 91	
	Recruitment of Program Managers, 92	
	A Mechanism for Fostering Experimental Research in IS&T, 93	
	APPENDIXES	
A	Meeting Agendas	101
B	Acronyms	105

Executive Summary

The U.S. Air Force, like the other services, is transforming itself into a new type of force with capabilities appropriate for an emerging array of new threats. The Air Force roadmap for transformation, part of the *U.S. Air Force Transformation Flight Plan*,¹ describes the desired new capabilities, and it is readily seen that advances in information science and technology (IS&T) underpin most of them. For example, the three main new capabilities are information superiority, precision targeting (or strike), and improved battlespace awareness. The first requires secure and survivable command and control systems; methods for sharing, tailoring, and distributing vast amounts of information; decision aids; and offensive and defensive cyber warfare. Precision strike implies the ability to place munitions with minimal error anyplace required to achieve a military objective, and also the ability to perform rapid damage assessment. And improved battlespace awareness requires the ability to fuse and convey information so that decision makers can fully understand the plan of action and its execution in real time and be able to rapidly assess and anticipate necessary changes to the plan.

In order to refocus its program of basic research in IS&T to better support these Air Force goals, the Air Force Office of Scientific Research (AFOSR) asked the National Research Council to establish a committee charged with the following task:

¹Available at http://www.dtic.mil/jointvision/af_trans_flightplan.pdf. Referred to in this report as the *Air Force Flight Plan*.

The study will create a vision and plan for the IS&T-related programs within the AFOSR's Mathematics and Space Sciences Directorate. Based on the spectrum of Air Force IS&T needs and the context in which the Mathematics and Space Sciences Directorate operates, the committee will do the following:

- Identify which of the Air Force's IS&T needs seem to call for AFOSR-sponsored R&D;
- Recommend a program of 6.1 research² in IS&T that is not being done elsewhere (or is not readily applicable to Air Force situations) and that covers the most critical or broadly useful topics that fit within the purview of the Mathematics and Space Sciences Directorate;
- Develop rough estimates of the funding needed to make credible progress in this program of IS&T-related research, with a prioritization that defines what could be adequately covered with flat funding, a 10 percent decrease, a 10 percent increase, and a 25 percent increase. Recommend how the directorate might transition from its current program to the envisioned one under these various budget scenarios; and
- Recommend an appropriate balance of funding mechanisms for the directorate's IS&T-related research, choosing among the various mechanisms currently in use in the directorate.

This report is the outcome of that committee's study.

The committee learned about Air Force goals from a variety of sources, including printed reports, briefings at the Air Force Research Laboratory (AFRL) and the Air Combat Command, and discussions with senior Air Force leaders in research and development (R&D). From these sources, the committee concluded that most of the capabilities desired by the Air Force cannot be attained without continued IS&T R&D. This is because IT pervades most, if not all, envisioned Air Force systems and is often the principal enabler of system capability, yet IT is still an immature engineering discipline requiring much work to assure predictable results when a system requires IT-related innovation. Furthermore, nearly all of those capabilities require some advances that are unlikely to be developed commercially or by the other services and therefore will require targeted R&D by the Air Force itself. Moreover, nearly all of that Air Force-specific R&D must include ambitious basic research, because significant gaps exist in the knowledge base upon which the desired capabilities will be built.

²In the Department of Defense (DOD), funding lines are assigned numbers, and 6.1 is the line for basic research. Within the Air Force, the AFOSR is in charge of all 6.1 funding, most of which is used to support peer-reviewed academic research. Funds for applied research and development (R&D) are designated 6.2 or 6.3.

The committee, echoing what is already understood within the AFOSR R&D establishment, identified (1) access to disparate data and information, (2) their fusion and appropriate distribution, and (3) conversion of information into knowledge as the necessary building blocks for attaining the desired capabilities. These building blocks, like most of the Air Force's desired capabilities, rely on team-focused, network-enabled systems—that is, interlocking systems made possible by networks that enable the teams to work together. The committee concluded that research to develop those building blocks is the most important Air Force need, one that will persist as long as the Air Force relies on network-enabled systems, and from its initial store of ideas about which kinds of research would be relevant to Air Force IS&T, the committee identified four that underpin team-focused, network-enabled systems of any kind: research in networks and communications, software, information management, and human-system interactions (HSI). The committee's vision for AFOSR's IS&T program is captured in Figure ES-1. Distributed research and experimentation environments are discussed in Chapter 9 and some grand challenges are proposed in Chapter 7. Then, the committee summarizes the research it recommends in each of these areas.

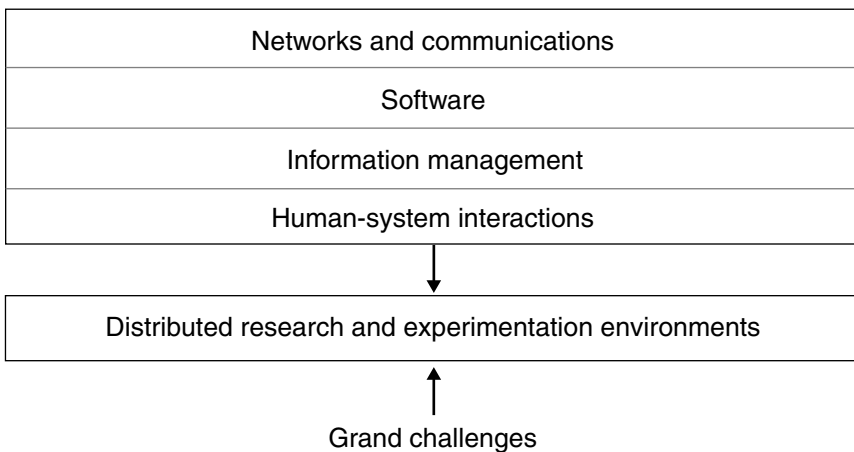


FIGURE ES-1 A vision for Air Force IS&T research: Team-focused, network-enabled systems are created by the four research areas shown. The concerted efforts in the four areas, which also affect one another, are to be focused by grand challenges identified by the AFOSR and by experiments conducted in distributed research and experimentation environments.

NETWORKS AND COMMUNICATIONS

Air Force applications must contend with communication modalities that are not encountered in commercial and civilian settings. For example, satellite channels have unusually long-delay data rates and randomly fading dispersive channel characteristics. Classical communication and information theories do not incorporate an element of adversarial attacks. Radio channels, especially those associated with mobile platforms, have rapidly changing link capacities and connectivity, with disconnections and dropouts that can last minutes or more. In contrast with this dynamism, traditional layer 3 (Network Layer) and layer 4 (Transport Layer) protocols assume fairly stable underlying substrates that change, if at all, over the course of minutes—that is, much more slowly than most transmissions. These traditional protocols often yield low throughputs and poor quality service when applied to defense systems; in some cases, they do not work at all despite valiant efforts to provide patches. Thus, the main challenge of Air Force communications is to provide assured connectivity between networks (albeit at varying rates) under difficult channel conditions, including during adversarial attacks.

Another Air Force communications challenge is how to recognize when multiple sensors have collected related observations so that redundancy can be removed or complementary data fused. This is essential in order to stay within network bandwidth capacities, especially in difficult communication environments. More generally, the theory of networks has not matured to a point where one can predict how well protocols developed heuristically in one application setting will perform on a communication network built on radically different communication modalities. To deal with the new and complicated modalities of importance to the Air Force, fundamental tools must be developed to help understand how networks might perform in new environments and to optimize architectures. It is simply too costly to develop these architectures and protocols ad hoc and then experiment with the communication links in the field.

Bandwidth will always be in high demand in the battlespace, so there is a need for a network management system that is able to translate high-level guiding principles into network actions such as routing and media access control priorities in a timely fashion without a human in the loop. Currently, asset management is done manually, and it is far from responsive or optimal. Because it will not always be possible to ensure that no nodes are compromised, the network should be designed to sense dead or malfunctioning network elements and route around them. In addition, the network should have an architecture that confines such damage to a local area and does not allow it to propagate across the network. When the network senses outside attacks, it should be able to locate the real entry points and then defend against and remove these attacks.

In response to these challenges, the committee recommends that AFOSR pursue basic research in the following topics of importance for Air Force networks:

- Robust protocols, addressed with new mathematical tools for network dynamics analysis.
- Error-free, end-to-end delivery, requiring better methods for performance prediction.
- Throughput, delay deadlines, and congestion control, all based on network coding.
- Network performance optimization, building on dynamic (convex and nonconvex) programming, game theory, and control theory.
- Policy-based network management, requiring means of monitoring, resource allocation, and making performance guarantees for subsets of users.
- Robust architectures, perhaps based on Byzantine robust networking.
- Network architecture and protocols for unmanned air vehicles (UAVs) and other air vehicles.

For sensor networks in particular, the committee recommends the following basic research topics:

- Real-time embedded processing.
- Embedded control systems.
- Minimization of power consumption, addressed through energy-efficient routing, Transport Layer protocols, and energy-efficient process management.
- Programming and support tools for large-scale networks.
- Energy-efficient coding schemes for information distribution.
- Techniques for real-time dynamic resource allocation.
- Energy-aware compilers and schedulers.
- Source compression and correlation methods for multiple sensors.

The Air Force communication systems that operate on these networks require basic research in the following areas:

- Unifying methodologies for modulation, coding, beam-forming, and scheduling optimization.
- Information theory extensions for dynamic self-adaptive communications.
- Wireless architectures for exploiting node-to-node cooperation.
- Ultrawideband (UWB) communication: air-ground, air-air, air-space.

- Dynamic exploitation of channel characteristics for increased capacity, reliability, and spectrum efficiency.
- Design of systems with performance guarantees for difficult channels, including channels under attack.
- Integrated design/optimization of networks plus communications systems, being conscious of the vulnerability to cross-layer adversarial attacks.

SOFTWARE

Network-enabled systems are by definition dependent on complex software because of the great number of possible states of the networks. The systems that require such software transcend a range of Air Force applications, from intensive human-machine systems (e.g., command and control, air operation centers) to embedded applications (e.g., avionics systems). Increasingly these applications are connected by networks into a system of systems and, in fact, the distinction between enterprise and embedded systems blurs as the focus is increasingly on the interconnectedness of all such systems.

Rather than focusing on large-scale code development—a challenge that is being researched by others—the committee recommends that AFOSR focus on a set of important software engineering issues that are key to successful Air Force network-enabled systems but that have received limited attention. This recommended set of issues centers on how to understand what to build and how to ensure that its behavior is relatively predictable and acceptable, both during design and in operational use. Three important questions emerge:

- How do we discover and understand what is needed?
- How can critical nonfunctional attributes (those that are desired or necessary but ancillary to the software's primary functionality) be implemented in a predictable fashion?
- Can the resulting software, once fielded, evolve to satisfy new needs discovered as it is used?

To address the first of these, the committee recommends a program of research aimed at the coevolution of Air Force concepts of operations and system architectures. This program extends the philosophy of software development models such as iterative development that support rapid prototyping of a software system so that end users can experiment with the system to see if it satisfies their needs. The prototype then becomes an explicit representation of the requirements. Current research in execut-

able architectures and in engineering tools for the design and analysis of functional and nonfunctional attributes provides a basis for this program.

The committee recommends research into the following:

- Methods to support rapid composability.
- Semantic extensions of current modeling languages to enhance composability and representation and reasoning of behavior.
- Development of tools that enable the construction of executable versions of models in system modeling languages.
- Methods that support experimentation, operational assessment, and the use of initial architecture representations in exercises. An example might be scripting languages that allow end users to explore early versions of software and help encode their preferences into the final architecture.
- Approaches that allow user tailoring, definition, and exploration of new processes and automated learning based on past problem solving.
- Experimentation and demonstration of these research approaches in domains of relevance to the Air Force.

To address the second question, the committee recommends that AFOSR support a new line of research, extending model-based software research funded by the Defense Advanced Research Projects Agency (DARPA) to build up an understanding of software behavior envelopes. Dynamic analysis of the nonfunctional attributes (e.g., scalability, interoperability, survivability, security, energy awareness) of software could define the performance envelope of a network-enabled system. It would be valuable to know the extent to which software could be modified, by developers or end users, and stay within the desired envelope. This topic would be a new area of research for the software community, but there is related work on which to build, as explained in Chapter 4 of this report.

Once a software architecture has been defined and the performance envelope explored, a logical third capability would be one that supports the continued evolution of complex software within its fielded context. While most other software engineering research focuses on developing new software-intensive systems, in fact the larger challenge is to learn how to maintain and upgrade the huge amount of Air Force software that has already been fielded. Thus, important research areas include methods to infer the architecture of legacy software systems, to identify software components within that architecture, to parallelize legacy system software and applications, and to migrate that architecture and components to new and improved architectures, possibly within a new computing environ-

ment. Since network-enabled systems will involve many legacy systems as well as new systems, it is imperative that software be designed so that it can be evolved in an affordable manner throughout its life cycle. The committee recommends that AFOSR support research to improve the evolvability of software-intensive systems. The following specific lines of research, which could build on readily available commercial frameworks, are recommended:

- Our ability to conduct dynamic, model-based analyses to analyze nonfunctional attributes needs to be improved.
- In order to improve component integration, research is needed to accelerate the development of abstract design-component systems and code-component-based systems, addressing automated discovery, composition, generation, interoperability, and reuse across hundreds of systems.
- Research in security is needed in support of the goal of measurable, available, secure, trustworthy, and sustainable network-enabled systems.
- To attain assured reliability with hard time-deadlines, methods are needed for modeling and analyzing integrated reliability, availability, and schedulability of components and systems in realistic conditions derived from user-specified scenarios.
- All participating components of the overall system need to be energy-efficient: (1) network energy on network interface and communication protocols of ad hoc networks, (2) processor energy and process management for scheduling various applications, (3) memory/storage energy and memory/storage management, and (4) display energy.
- Research is needed into novel integration of methods for verification and validation, such as integration of informal methods (e.g., software testing and monitoring) with formal verification (i.e., model checking and theorem proving) and abstract interpretation and static program analysis techniques. The ability to validate scalability, adoptability, usability, and measurement is also important, and some fundamental breakthroughs have occurred in the past 5 years that have led to a rapid rise in industry adoption and interest.

INFORMATION MANAGEMENT

One ramification of the ubiquitous deployment of IT in the Air Force is that both human and automated decision makers are now often faced

with voluminous multimedia data from which they must create knowledge. Even the first step in knowledge creation—the integration of raw data that are in different formats and managed by different data management technologies—is challenging, but future Air Force capabilities will require much more.

The Air Force faces major open questions on how to manage and share information in a distributed system. The “publish-subscribe” paradigm is one that is being explored at the AFRL. That concept includes (1) a common repository where information is “published” and (2) “subscription” information for various users that defines which posted information their systems will download from the common area. The publish-subscribe concept has been shown to scale to hundreds of thousands of participants within stable network environments. However, an Air Force publish-subscribe system must work in an unstable wide-area network environment such as a battlespace network; it must in many cases weed out information that is outdated or redundant; its subscription rules must be more sophisticated than those available today, including having enough “intelligence” to take context into account; and the system must be trustworthy even if an adversary has gained access to publish or subscribe. These challenges are examples and not comprehensive. Moreover, they are not unique to publish-subscribe. Similar challenges accompany alternative infrastructures for information management. It is clear, therefore, that much research in fields such as distributed computing, database systems, security, and data mining must be accomplished before the Air Force can field a dependable information management system.

More generally, the Air Force needs to understand information at a more abstract level. It needs a model and architecture for situation understanding and a means of incorporating situation modeling, model-based processing, situation projection, and top-down management of situation understanding in order to explore topics in information fusion. It also needs a scientific basis and technologies for multisensor fusion for air and ground targets. Some of these topics are extensions of ongoing work in intelligence, surveillance, and reconnaissance (ISR) methods. An even bolder question would be, How can a computer understand data and information in context? In principle, background understanding of a mission or related intelligence could help a computer interpret information from the battlespace—for example, to help identify objects in video or image data. If such context-dependent processing were possible, perhaps information-understanding algorithms could be embedded in sensors and networks to enable rapid data assessment and rapid situation assessment.

The committee recommends the following basic research in support of Air Force information management:

- Query-processing techniques for large-scale sensor networks.
 - Where to place query functionality vs. limited power, bandwidth, etc.
 - Coping with mobile sensors, unreliable sensors, high data rates.
- Techniques for processing and managing semistructured content.
 - For data modeling, for querying and routing, for execution.
- Fusion of uncertain, inconsistent data and querying of incomplete information.
- Mechanisms for determining the certainty of answers as a function of the certainty of raw data.
- Multilevel representation of multimodal signals (video, images, hyperspectral, etc.).
 - For efficient transmission, storage, manipulation, multimodal data mining, and machine learning.

HUMAN-SYSTEM INTERACTIONS

The committee focuses on HSI to encompass not only human-computer interactions but also the coordinated and purposeful interactions of several or many humans with complex systems and the interactions of teams of humans mediated through systems. The committee recommends an AFOSR focus on HSI because it is essential to the successful operation of complex systems and to the accomplishment of network-enabled operations. An ultimate goal of HSI research would be to enable machines (or algorithms) to perform more of the complex data manipulation, correlation, computation, and data reduction—and even some decision-making—leaving humans to perform the most critical judgments that cannot be accomplished by algorithms or that rely on extrinsic knowledge. Furthermore, HSI should help humans to interact with one another in cooperative tasks where multiple humans are part of the system.

In the Air Force, there are many situations where one or more humans interact with one or more IS&T systems. This includes systems that are distributed not only among different platforms but also, perhaps, across geographical and organizational boundaries, most often with strict security and service reliability constraints such as near-real-time or time-critical services. Complexity increases if the humans and the systems interact with one another in ways that are not connected with the task being analyzed. What sorts of information, architecture, and format should be used to achieve desired effects, and how can designers and users estimate the uncertainties and internalize the context and caveats associated with each option? Assuming the right information is available at the right time and in the right form (e.g., text, images), what techniques will enable the user to make the best use of it? How can what-if simulations be considered and

evaluated? Such complex capabilities might require integrated and synchronized multimodal interfaces (visual, aural, and/or haptic) to capture the high dimensionality of a system of sensors and actuators in the battlefield.

Research into HSI should shed light on the usability of the (same) information in a battlefield command-and-control situation relative to the perspective (rank) of the user and the granularity (detail of the information). In other words, one must understand and characterize the most likely and useful level of complexity for each potential user, from the warfighter to the commander, so that the complexity and amount of information can be optimized for battlefield decision-making—not a paucity of data, but not data overload either.

The importance of HSI research is also driven home by the Air Force emphasis on influence operations, which are meant to alter adversaries' attitudes and perspectives so as to achieve U.S. goals without resorting to the tools of traditional warfare. Influence operations require fundamental research into behaviors and how they can be affected. To this base of knowledge must be added knowledge on interpretation and presentation, personnel training, and modeling and simulation, building on what is known about cultural and behavioral factors to carry out influence operations. As an example, characterization and recognition of normal and abnormal behavior would, in general, help in surveillance at all levels. Characterizing which actions, postures, and so on signify worrisome behavior requires ongoing research in the social sciences, and the ability to automatically recognize such behavior in sensed data is an ongoing challenge for IS&T.

The committee recommends that AFOSR pursue basic research in the following areas of importance to HSI:

- Tools for improved human interactions with automated reasoning and inference systems under constraints.
- Automated diagnosis and decision support, automated learning. Enable user navigation of systems involving complex and noisy data and decision systems.
- Learning-theory-based techniques for predictive modeling and anticipatory behavior involving cultural factors.
- Combination of heuristic and optimization techniques for complex searches, with adaptability to different levels of detail to avoid information overload of the warfighter.
- Trade-offs between power usage in sensors and displays and choices regarding the range of visual items, human attention, and control.
- Fundamental requirements and metrics in designing, implementing, and experimenting with complex, interactive, time-critical information systems.

- Enhanced, interactive, mixed-modality models, experiments, and testbeds for more integrated real-time human/system/sensor synergy and database decision support relevant to Air Force goals.

In particular, HSI research of importance to information usability and influence operations would include:

- Simulation of urban and human environments.
- Behavioral models of individuals, groups, and organizations.
- Fundamental attributes of information operations testbeds and experimental metrics for evaluating effectiveness.
- Decision support techniques for addressing partial-solution approximations based on evolving, nonstatic information.

Note that some of the HSI research falls squarely in the domain of psychology or sociology. AFOSR already has programs that are joint between IS&T and psychology, and the committee recommends that this interface continue to be strengthened and broadened.

PRIORITIES FOR AFOSR IS&T RESEARCH

The committee recommends that AFOSR prioritize its IS&T research in networks, communications, information management, software, and HIS, as shown in Table ES-1. With the current funding available for IS&T (the column headed “Stable”) the committee recommends that networks, communications, and HSI research merit the highest priority, while information management and software research portfolios would be better able to weather any forced reductions in the level of effort. The committee is not saying that the latter two research areas are less important to the Air Force. Rather, it is the committee’s judgment that if cutbacks are required, reductions in those programs would do the least harm in limiting future options. If the overall IS&T funding dropped by 10 percent, the committee would give software the lowest priority only because other organizations, and commercial enterprises, are doing some related research. If overall funding increases by 10 percent, the priority for information management research should be raised a notch. Finally, if overall IS&T funding were to increase by 25 percent, the committee recommends a balanced portfolio drawn from the particular research recommendations earlier in this summary. See also the footnotes to Table ES-1 for additional interpretative notes.

Because all of the major research areas listed in Table ES-1 contribute synergistically to the future fielding of team-focused, network-enabled systems, progress toward that vision is dependent on a balanced research effort across all five areas. As implied by Table ES-1, the overall funding

TABLE ES-1 Relative Priorities Under Four Funding Scenarios

IS&T Topic	10% Reduction	Stable	10% Increase	25% Increase
Networks	H	H	H	H
Communications	H	H	H	H
Information management	M	M	H	H
Software	L	M	M	H
Human-system interactions	H	H	H	H

Note: "H" means the general topic is a high priority, and its funding should be protected or increased. "M" means the general topic is of medium priority for AFOSR support, given the contributions by other players, not that it is of medium importance to the AFOSR. "L" means that funding in that area should be sacrificed so that a critical level of effort can be supported in other areas. "L" does not mean that the topic is not of importance to the Air Force, only that if resources are tight, it is a reasonable candidate for cuts because other organizations are contributing to the topic and/or the challenge is so great that a small AFOSR effort is unlikely to lead to significant progress. These priorities pertain to the five general research areas listed in the left-hand column as weighed only against one another, not against other programs funded by AFOSR's Mathematics and Space Sciences Directorate. The priorities are meant to show the committee's consensus on which of the areas to (de)emphasize if there are any changes in funding. The priorities take into account not only the importance of the research but also the relative need for Air Force-specific research. They reflect the committee's general sense of what can be meaningfully accomplished within the funding scenarios posited, but the committee did not develop a detailed estimate of the resources required for each of the research topics in the left-hand column.

level for basic research in IS&T will not support such a broad, balanced effort unless there is a significant increase. Therefore, the committee recommends a significant increase in IS&T funding within AFOSR centered on research to support team-focused, network-enabled systems of Air Force interest.

The committee also recommends that AFOSR consider designating some topics as grand challenges as a means of focusing its IS&T research, motivating the academic research community, and connecting that research to Air Force goals. Topics designated as grand challenges would be ones for which there is a recognizable gap in the knowledge base that would be properly addressed by a cross-disciplinary community of basic researchers; the grand challenge will help give that community an identity and thus strengthen its coherence. These grand challenges should be defined in terms that are recognizable to the basic research community, but AFOSR should also be able to map the grand challenges to future Air Force technologies. The grand challenges are not part of, nor do they compete with, the AFRL's focused long-term challenges (which are more ori-

ented toward technologies), but they should link to them. Building a program around grand challenges quite naturally facilitates new interdisciplinary research communities: “interdisciplinary,” because the breadth of the challenges calls for varied expertise, and “naturally,” because the associated researchers are interested in the whole range of efforts addressing the grand challenge.

The committee recommends that AFOSR consider the following as possible grand challenges, but this list is by no means exhaustive:

- *Control of multiple UAVs.* Research to enable the control of multiple UAVs by one human in mixed manned-unmanned airspace, in contrast to today’s requirement for many humans for a single UAV in carefully deconflicted manned and unmanned airspace.
- *Taskable airborne network.* Research to enable cost-effective and rapidly deployable tactical intelligence networks in urban environments, where the nodes generally are sensors carried on UAVs or lighter-than-air vehicles and the networks are taskable by ground- and air-based commanders.
- *Mixed-reality training environments.* Research to enable training for air crews, command post staff, and commanders in an environment of such fidelity that it would be indistinguishable from the real world (and in fact would sometimes involve the real world—hence “mixed” rather than “virtual” or “augmented”). The computer tools used in such training environments should be the same as those used in the real world.
- *An automated Air Operation Center staff assistant.* Research to enable software that can learn from being told, much as human staff members learn on the job.
- *Rapid system integration.* Research to enable the rapid integration of IT-based systems, such as those belonging to different members of ad hoc coalitions. This research would encompass HSI, networks and communications, security, software, and information management.

FUNDING MECHANISMS

AFOSR’s current IS&T research is supported through a range of funding mechanisms, and the committee found that each of those mechanisms provides value and that the AFOSR program managers are doing a good job of making use of them. The committee does not recommend any hard-and-fast rules for balancing the various funding mechanisms; rather, it encourages continued flexibility and is comfortable with the current mix. The committee did observe, though, that it would be beneficial for AFOSR to increase the number of young investigators who are aware of Air Force

challenges. Therefore, it recommends that the Air Force consider establishing a mechanism for young investigator awards so as to raise its visibility within that group of IS&T researchers.

FUTURE CONSIDERATIONS

The committee observed that AFOSR's IS&T portfolio is difficult to pin down because it is distributed among various programs in two AFOSR directorates. The committee recommends that AFOSR identify IS&T as a major topic within the Mathematics and Space Sciences Directorate and, as IS&T investment increases, establish a separate directorate with that single focus. It also recommends that all human-system and human effects research be consolidated within that IS&T directorate because of the critical importance of HSI to the effectiveness of complex Air Force systems.

As the IS&T program grows, the committee sees an opportunity for AFOSR to try new mechanisms for recruiting program managers, especially by reaching out to the broader IS&T community. Mechanisms such as the Intergovernmental Personnel Act and the Experimental Personnel Hiring Authority can be very useful for bringing in both program managers and higher-level staff.

Finally, the committee urges AFOSR to work with other parts of the Air Force to establish testbeds that will allow researchers and Air Force users to experiment with prototype IS&T concepts and systems. Besides the inherent benefit of experimental science, such an approach would provide an intellectual crossroads between the scientific and operational community to support the scientific discovery process. The committee uses the label "distributed research and experimentation environment" (DREE) to describe a shared computation infrastructure that supports experimentation within a community of researchers.

The committee believes that DREEs would be useful for each area of research cited in this report. A DREE for information management, for instance, would enable the associated community—including universities, AFRL laboratories, and perhaps federally funded R&D centers—to create sample data sets and develop associated queries that illustrate how the data are to be integrated. A DREE related to network-centric systems would allow exercises from which concrete performance requirements could be generated; those requirements are difficult to identify otherwise. While exercises are ongoing, operational Air Force participants can clarify their real, not hypothetical, needs; IS&T applied researchers can investigate engineering issues with the prototype network; basic researchers in IS&T can experiment with fundamental changes (e.g., to communication protocols); and HSI researchers can instrument the experiments and learn from them.

The DREE approach to promoting experimental science should not be prohibitively expensive, because the necessary network infrastructure is rapidly falling into place and there is the possibility of leveraging investment in testbeds made by other AFRL directorates. For example, a research version of the Distributed Mission Training environment housed in AFRL's Human Effectiveness Directorate might support experimental science in areas ranging from control of UAVs to decision making in real-time environments. The committee's recommendation is that basic research funds not be used to establish DREEs, only to support the involvement of researchers.

Introduction

In October 2004, the National Research Council (NRC) was asked by the Air Force Office of Scientific Research (AFOSR) to recommend a basic research program to support Air Force information science and technology (IS&T) goals. The NRC established the Committee on Directions for the AFOSR Mathematics and Space Sciences Directorate Related to Information Science and Technology and charged it to “create a vision and plan for the IS&T-related programs within the AFOSR’s Mathematics and Space Sciences Directorate.

The committee was not charged with reviewing the current IS&T program, so nothing in this report should be construed as a criticism of AFOSR or its current program. Rather, this report was designed to be a *de novo* look at what should be included in the Air Force’s basic IS&T program. Many of the recommended research topics are well known to AFOSR staff, and they clearly overlap the current program. The identification in this report of a basic research need does not imply that the need has been overlooked, only that it should be one of AFOSR’s top IS&T priorities.

To accomplish its charge, the committee held three meetings. The first two were intended to inform the committee of the Air Force needs in IS&T, from which the committee would deduce the required basic research portfolio. (Complete agendas of the three meetings are included in Appendix A.) The first meeting was held at the Air Force Research Laboratory (AFRL) Information Technology Directorate in Rome, New York. This directorate is a counterpart to AFOSR charged with conducting 6.2 and 6.3 R&D in IS&T; as such, it is both a user of the results of

AFOSR-sponsored IS&T and a conduit for transitioning results between the research community and Air Force operating units. At that meeting, the committee also interacted with two members of the Air Force Scientific Advisory Board, who conveyed results from that board's 2003 review of Air Force IS&T R&D. The second meeting was held at Langley Air Force Base in Hampton, Virginia, and was hosted by units of the Air Combat Command (ACC). The ACC's Command, Control, Intelligence, Surveillance, and Reconnaissance center and ACC activities in information operations make it the primary customer among Air Force operational units for AFOSR's IS&T research. At that site visit, the committee received additional perspectives on the Air Force's IS&T needs. At its third meeting, in Washington, D.C., the committee held a discussion with Thomas Cruse, chief technologist for the AFRL, to explore AFRL's directions in IS&T. At all meetings, the committee had ample opportunity for frank discussions with AFOSR personnel who manage IS&T-related research portfolios, and on two occasions it had in-depth discussions with Brendan Godfrey, the director of AFOSR.

The committee also sent one member to attend a program review of the Partnership for Research Excellence and Transition (PRET) for advanced concepts in space situational awareness, held in January 2005, and another member for a site visit to the Mesa, Arizona, unit of AFRL's Human Effectiveness Directorate. The first visit was helpful in bringing out different perspectives about AFOSR funding mechanisms, and both visits covered potential research.

As part of its investigation, the committee examined a wide range of past reports—from the Air Force itself, the Defense Science Board, the Air Force Scientific Advisory Board, and the National Research Council—that bear on Air Force basic research into IS&T. For example, the committee examined Air Force-specific planning documents such as the Air Force Long-Term Challenges, the Air Force Mission and Vision Statement, and the Air Force Flight Plan. From all of these investigations, readings, and discussions, the committee first established a consensus on what IS&T research is needed to support the Air Force's goal and then filtered those findings based on the committee members' collective insight about what research is being done, or is likely to be done, in industry, academia, and elsewhere. The study focused on information available from AFOSR, AFRL, and the Air Force Scientific Advisory Board. As such, this report centers on IS&T research for Air Force operations and does not explore possible IS&T research to improve Air Force processes. The committee also filtered the research needs according to whether or not there were Air Force-specific questions to be addressed: If not, then there is no need for AFOSR to carry out research in that particular area.

Based on these inputs and its own expertise, the committee first

generated a list of more than a dozen IS&T research areas of importance to the Air Force. Then, echoing what is already understood within the Air Force R&D establishment, the committee identified three building blocks for attaining the desired Air Force capabilities: (1) access to disparate data and information, (2) fusion and appropriate distribution of the data, and (3) conversion of the information into knowledge. These building blocks, like most of the Air Force's desired capabilities, rely on team-focused, network-enabled systems—that is, interlocking systems made possible by networks that allow teams to work together. The committee concluded that research to develop those building blocks is the most important for the Air Force, and from its initial list of Air Force-relevant IS&T research areas, it identified four that underpin team-focused, network-enabled systems of any kind: research in networks and communications, software, information management, and human-system interactions (HSI). This committee vision for AFOSR's IS&T program is captured in Figure 1-1; "distributed research and experimentation environments" and "grand challenges" will be explained in Chapters 7 and 9. By building up the knowledge base in these four fundamental areas, AFOSR can help the Air Force move beyond more heuristic approaches for developing disruptive technologies such as network-enabled systems.

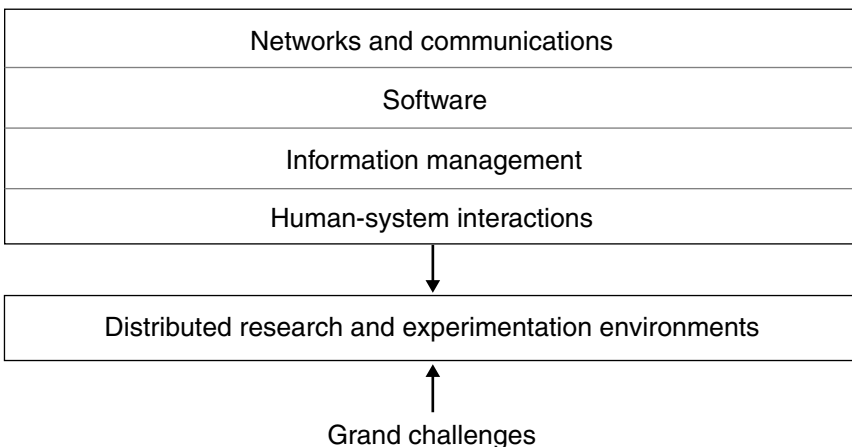


FIGURE 1-1 A vision for Air Force IS&T research: Team-focused, network-enabled systems are created by the four research areas shown. The concerted efforts in the four areas, which also affect one another, are to be focused by grand challenges identified by the AFOSR and by experiments conducted in distributed research and experimentation environments (see Chapter 9).

The committee believes this structure for IS&T research is generic in that it applies to a broad set of information-intensive challenges and embraces all of the high-priority topics for the Air Force. In developing this vision, the committee decided to recommend that AFOSR focus on the IS&T needs of team-focused, network-enabled systems. This choice of research areas is strongly suggested by the documents, briefings, and plans reviewed by the committee. Such a focused research program promises to be much more effective in achieving the goals of a network-enabled Air Force than would a piecemeal attack on disparate IS&T research topics. The committee recognizes that the inputs to network-centric operations and lower-level aspects of networks—for example, sensors, processors, information/surveillance/reconnaissance (ISR), and information assurance—also require continued AFOSR-sponsored research, but it believes the greatest research void when looking to the future lies with the system-oriented topics in Figure 1-1. It is at the interface to these systems that information and data emerge and must be dealt with in an effective way to enhance the decision superiority of the Air Force.

The basis for all network-enabled operations is a good understanding of networks and the communications that traverse them. As explained in Chapter 3, there are many differences between Air Force networks and commercial networks such as cell phone networks and the Internet. These differences mean that a large number of basic research challenges are Air Force-specific and are not being addressed outside the military. At present, there is no solid foundation of understanding to guide the design and management of these complex networks, and without that understanding building such a network would be akin to designing a new generation of fighter planes by piecing together components from the last.

On top of the networks and communications, the Air Force relies on a wide range of complex software. The challenge is to develop a capability to build complex software systems with predictable behavior. This is a long-standing challenge, and Chapter 4 recommends specific aspects that AFOSR should pursue.

Information management and interactions between human beings and systems are needed to create military value from the Air Force's emerging information dominance. For example, much basic research is needed before the Air Force can effectively use and coordinate all of the data currently available from both unmanned air vehicles (UAVs) and satellites. There are many unmet challenges in these areas, which are explored in Chapters 5 and 6. Included in the HSI discussion is information operations, which entails not only offensive and defensive cyberwarfare but also the emerging area of influence operations, which aims to effect military goals without necessarily damaging people or infrastructure. There is a great deal of research to be done on that topic.

The committee did not conduct a detailed audit of the present AFOSR IS&T programs other than to review the brief summaries that were provided. Such an audit was not within its mandate, and the committee took a blank-slate approach to defining an IS&T vision and a basic research program. By identifying the basic research challenges that underpin team-focused, network-enabled technologies rather than those suggested by specific Air Force or DOD systems and programs, the committee aimed to define an IS&T program that will remain relevant regardless of how particular technology foci might vary over the coming years.

Background

An NRC review of DOD's IS&T for air and space asserts the following:¹

The pervasiveness of information systems in military systems and in visions for future warfighting and the accelerating threat information technology poses require that [the Department of Defense] have a strong information systems S&T program. The committee found the opposite to be the case, however, in its review of the Air Force S&T investment. Since at least the mid-1990s, the Air Force S&T investment in information systems has declined steadily, despite recurring annual plans to increase it.

Later, on pages 29 and 30, that report notes as follows:

Future Air Force concepts are driven by information and information systems, which are becoming the "force multiplier" for the Air Force of the future. The increased emphasis on information and associated trends are reflected in Air Force warfighting concepts. These concepts include:*

- Dynamic aerospace command
- Joint battlespace infosphere
- Information operations
- Integrated aerospace operations
- The Expeditionary Aerospace Force
- Effects-based operations

*Defense Science and Technology Advisory Group, "Air Force strategy in infor-

¹National Research Council, Review of the U.S. Department of Defense Air, Space, and Supporting Information Systems Science and Technology Program, National Academy Press: Washington, D.C. (2001), p. 4.

mation," briefing by John A. Graniero at the Information Systems Technology (IST) Technology Area Review and Assessment (TARA), Air Force Research Laboratory, Rome, New York, March 13-17, 2000.

On pages 31 and 32, the NRC report notes that commercial IS&T will not, no matter the size and strength of the investment, address all of the IS&T needs of DOD:

Leveraging commercial information technologies is difficult, however, because industry rapidly changes direction to meet rapidly changing customer demands and because the time to market must be as short as possible. Fierce competition dictates limited, short research and development cycles and near-term investment strategies. Very little funding is being invested in basic research, which is usually outsourced to academia.* Industry's short-term needs cannot support the longer-range visions of the services. Although commercial technologies show promise in providing significant near-term capabilities, leveraging them could require much effort (and significant resources) to adopt, adapt, or reengineer them. . . . Another caveat about using commercial information systems is that they are becoming available to all nations and interest groups. If the services depend on commercial technologies for advancing the state of the art in their information systems, potential enemies may come close to achieving parity (or even asymmetrical superiority) with U.S. forces DOD needs to expand its basic research to explore the frontiers of science in search of new technological approaches for maintaining military superiority. . . . The committee believes that DOD should continue to explore the frontiers of science and that basic research has never been more important to DOD.

*National Research Council, *Trust in Cyberspace*, National Academy Press: Washington, D.C. (1999).

It is certainly desirable to use commercial technology wherever it fulfills Air Force needs, meaning when the Air Force requirements are in some sense close to those of the commercial sector. However, the Air Force, like all the services, has requirements that are not significant drivers in the commercial sector—requirements in security (e.g., being multilevel in many cases), in the complexity of its command and control systems, in the real-time requirements of some systems, and in the need for robustness and resilience in the face of broken network or communications links or of captured nodes.

That 2001 report ultimately reaches (p. 36) the following recommendation for Air Force IS&T research funding:

The committee believes that the Air Force should increase its science and technology (S&T) budget for information systems technology (IST). The basic research (6.1) program should support long-term air and space IST

needs, surpass previous-year levels, support a strong in-house program (with appropriate researchers), and compensate for limited long-term commercial investment.

The DOD investment in basic research in IS&T has not improved in the 4 years since that report was issued. While the committee did not develop a detailed estimate of the level of effort required for Air Force basic research in IS&T, its general sense is that it is small in relation to the challenges identified in Chapters 3 through 6.

The Air Force roadmap for transformation is described in the *Air Force Flight Plan*.² Though not intended as a primer on Air Force needs for IS&T, it describes deficiencies in current capabilities as well as desired new capabilities that require new information technologies. The three main new capabilities are information superiority, persistent precision strike, and battlespace awareness.

Assuring information superiority is a critical part of the Air Force mission. Issues requiring advances in IS&T include secure and survivable command and control systems, the rapid evolution and adaptation of command and control systems, decision aids based on rapidly changing operational needs, and increased capabilities in information operations. Information operations³ include all offensive and defensive actions necessary to ensure full access to timely and accurate information and to deny the same to adversaries.

Precision strike refers to the ability to place munitions with minimal error anywhere required to achieve a military objective. Persistence, a relatively new desired capability, refers to the ability to maintain precision targeting for long periods. Open goals for R&D revolve around decreasing the time from sensing to shooting to less than 10 minutes, providing actionable intelligence in a useful (i.e., decision-ready) form, and timely and accurate damage assessment, which is particularly difficult when nonkinetic weapons are used. A report from the National Intelligence Council⁴ notes as follows:

The use of precision-guided munitions will increase. The accuracy of these systems will improve. Targeting will remain the critical factor in determining success or failure of these missions. Having eyes on target . . . must be a core mission.

²Available at http://www.dtic.mil/jointvision/af_trans_flightplan.pdf.

³Available at <http://www.e-publishing.af.mil/pubfiles/afdc/dd/afdd2-5.pdf>.

⁴John B. Alexander, "The evolution of conflict through 2020: Demands on personnel, machines, and missions," CIA Conference on the Changing Nature of War (2004), pp. 9-12. Available at http://www.cia.gov/nic/NIC_2020_2004_05_25_intro.html.

Actionable intelligence is essential. In recent conflicts commanders at all levels have lamented the shortfalls of actionable intelligence. Operation Iraqi Freedom was initiated prematurely when it was reported that Saddam Hussein would be located at Dora Farms in Baghdad. A phone call from an agent in the area began the process. Because this was the beginning of the conflict the decision rested with President Bush. The reaction time was too long, and if Saddam ever was there, he left before the target could be hit.

By 2020 we should be able reduce the time from which a target is detected until it is hit to under ten minutes. Currently, for DOD elements, target identification and release must cycle back to the US for approval. This is simply too slow.

Battlespace awareness implies the ability of a commander and his or her staff to fully understand the plan of action and its execution in real time and to rapidly assess and anticipate necessary changes to the plan. The deficiencies that motivate current research are summarized from an Air Force report on information operations:⁵

Battlespace awareness information is often reactive in nature and rapidly loses relevance. Targeting decisions often are made too far away from the warfighter to effectively engage mobile targets. . . . It is still very difficult to integrate rapidly expanding data streams from multiple sources in a timely manner. . . . Commanders often do not have a clear, accurate real-time picture of the battlespace. . . . The military still cannot assess, plan, and direct air and space operations from anywhere or from multiple locations in near real-time, something the Air Force believes will be necessary in the future to give the commander the greatest flexibility to meet national tasking.

The Air Force needs take into account that our potential adversaries have access to the same commercial IT as we do and that those adversaries can field IT systems at commercial rates, typically much faster than possible under current DOD acquisition practices. Hence, Air Force IS&T research must be focused on capabilities that fill Air Force needs, will probably not be provided by the commercial sector, and will give superiority over the enemy.

In the abstract, to fully address Air Force needs, IS&T research must yield new capabilities that satisfy the four *S*'s—speed, security, scalability, and smartness. “Speed” means IT-based systems need not only to operate in real-time environments but also to be built and delivered much faster and to be extensible in the field by end users. Such systems must also be

⁵Available at <http://www.e-publishing.af.mil/pubfiles/afdc/dd/afdd2-5/afdd2-5.pdf>.

secure and scalable. Lastly, they need to be smart. Getting there requires a better approach for sponsoring and conducting IS&T research. The AFOSR, which is the Air Force's prime agent for the conduct of 6.1 research, is well aware of these issues.

As part of its investigation, the committee examined a 1996 NRC review of AFOSR programs in mathematical and computer sciences and found that some of the observations still pertain and are helpful for setting the context of the current study.⁶ That study reviewed the size, scope, and quality of the erstwhile AFOSR mathematical and computer sciences program. (Many of the elements of that program, and some of the program managers, are now part of AFOSR's IS&T program. The IS&T program, though, does not map directly onto any AFOSR unit. Most of the IS&T research is managed today within the AFOSR Mathematics and Space Sciences Directorate, but only about half of that unit deals with IS&T. In addition, AFOSR's research in human-computer interfaces is managed jointly with its Chemical and Life Sciences Directorate.) The following observations, paraphrased from the 1996 report (pages 1-3), appear to have withstood the test of time:

- AFOSR program managers in mathematical and computer sciences have a great deal of autonomy and compete for funds among themselves in a relatively cooperative way. They are self-motivated to promote their areas in order to gain an adequate share of the available funds.
- The majority of funds go to principal investigators (PIs) at universities, but the Air Force is successful at teaming these PIs with researchers at Air Force laboratories.
- Prior collaboration with Air Force laboratories is not essential for obtaining program support, but collaboration subsequent to an award characterizes the best and longest-lived research projects.
- During the early 1990s, these program managers were under increased pressure to fund projects with short- rather than long-term payoffs, although there appears to have been adequate flexibility to resist such pressure.
- In determining program priorities, emphasis should be given to research that clarifies Air Force problems, has an impact on operational procedures, and advances the state of the art in areas of interest to the Air Force.

⁶National Research Council, *Review of AFOSR Programs in Mathematical and Computer Sciences*, National Academy Press: Washington, D.C. (1996).

OVERVIEW OF AIR FORCE GOALS THAT RELY ON IS&T RESEARCH

The Air Force is already heavily dependent on IS&T, and its ongoing transformation is critically dependent on the successful fielding of new IS&T products. Many of those new products will be specific to military needs, and many require basic research. A good summary of Air Force IS&T goals is contained in the pamphlet "S&T Planning Review."⁷ This pamphlet presents six long-term challenges identified in 2001 to provide focus and motivation for Air Force science and technology over the next 20 to 50 years. The emphasis is on basic and applied research that will enable critical future Air Force capabilities:

- *Finding and tracking.* Provide the decision maker with quality information from anywhere in near real time. The needed capabilities include reliable assessment and monitoring, sensor placement and sustainment, and information systems.
- *Command and control.* Monitor, assess, plan, and direct aerospace operations anywhere from multiple locations in near real time. Capabilities to support this goal include monitoring and assessing global conditions and events, planning and executing military operations, and information assurance.
- *Controlled effects.* Rapidly create precise effects with the ability for quick retargeting. The needed capabilities include measured global force projection and dominant remote control.
- *Sanctuary.* Protect the total force from natural and manmade hazards and threats while permitting operation with lowest possible risks.
- *Rapid aerospace response.* Quickly respond to peacetime operations or crises. Required capabilities include rapid global reach, on-demand space surge, and aerospace power network.
- *Effective aerospace persistence.* Sustain the flow of equipment and supplies as well as the application of force as required. Capabilities include space awareness and control, space access, and operations.

It is clear that information systems pervade all of these capabilities and that sustained basic research in IS&T will be required to meet these challenges. It is also clear that there are Air Force-unique aspects to most, if not all, of this basic research.

More recently, the Air Force released the latest statement of its mission, which encompasses not only air and space operations but also informa-

⁷Available from the AFOSR.

tion operations.⁸ The latter includes the union of electronic warfare, both defensive and offensive; offensive and defensive cyber operations (that is, disruption of, or tampering with, computing and communications or defense against the same); and “influence operations,” actions intended to lessen an enemy’s resolve to fight.⁹ Influence operations was given special emphasis when the study committee met with the Air Combat Command, which is exploring a range of means to lessen adversaries’ will to fight, their negative impressions of the United States or U.S. forces, or otherwise reduce the need for traditional military action. The principal goal of all information operations is to ensure decision dominance over adversaries. Key challenges from an IS&T perspective include integration of, access to, and fusion of disparate data and information sources; conversion of information to actionable knowledge; basic modeling of human behaviors; and so on.

The ways the Air Force achieves its vision and accomplishes its mission are changing. This change is the process DOD calls transformation, and it relies in a central way on information superiority. The Air Force roadmap for transformation is described in *Air Force Flight Plan, 2003*, which on pages 51-56 expands the notion of information operations described in the vision statement to the notion of information superiority:¹⁰

Information superiority is a key enabler of the type of revolutionary change described by RMA [revolution in military affairs] advocates including effects-based operations and parallel warfare. . . . There are still many obstacles to achieving the full potential of information superiority under many circumstances today:

- There is still significant progress to be made in rapidly getting timely, accurate, and relevant intelligence from sensors-to-shooters (actionable intelligence in a usable format) in single-digit minutes.
- Battlespace awareness information is often reactive in nature and rapidly loses relevance. Targeting decisions often are made too far away from the warfighter to effectively engage mobile targets.
- It is still very difficult to integrate rapidly expanding data streams from multiple sources in a timely manner.
- Commanders often do not have a clear, accurate, real-time picture of the battlespace.
- The military still cannot assess, plan, and direct air and space operations from anywhere or from multiple locations in near real-time, something the Air Force believes will be necessary in the future to give the commander the greatest flexibility to meet national tasking.

⁸Available at <http://www.af.mil/library/posture/vision/vision.pdf>.

⁹Available at <http://www.e-publishing.af.mil/pubfiles/afdc/dd/afdd2-5/afdd2-5.pdf>.

¹⁰Available at http://www.dtic.mil/jointvision/af_trans_flightplan.pdf.

- Computer network and information systems are often vulnerable to attack.
- There is limited ability to disrupt adversary C4ISR [command, control, computing, communications, intelligence, surveillance, and reconnaissance] assets and information flow.
- “Tribal” platforms and procedures within the Air Force still must be integrated using information technology.
- New and planned C4ISR systems require a lot of additional bandwidth.
- Lack of data standards inhibits use and exploitation of Artificial Intelligence capabilities.
- The Air Force has not developed all necessary protocols for machine-to-machine interfaces.
- The Air Force lacks a scalable C4ISR system that can support operations across the spectrum of conflict.
- The Air Force still needs to evaluate its current systems and determine what they can contribute to its capabilities and what tools are necessary to transform those systems from a collection of platforms into a networked system that is greater than the sum of its individual parts.

THE R&D RESPONSE: CURRENT DIRECTIONS

The AFRL’s Information Directorate (AFRL/IF) is a major performer of applied research and development for Air Force IS&T, and the committee spent considerable time with researchers there to learn about their programs and their strategy for addressing Air Force IS&T needs. Northrup Fowler, chief scientist (at the time), told the committee on February 24, 2005, that the directorate’s main R&D thrusts were

- *Global awareness.* Methods to acquire, exploit, fuse, and reason about information.
- *Dynamic planning and execution.* Methods for rapidly exploiting knowledge of the battlespace and fostering better informed and more accurate decisions.
- *Global information enterprise.* Methods to move, process, manage, and protect information throughout the Global Information Grid (GIG) and provide assured information to decision makers.

The underlying science and technology focus areas identified by AFRL/IF are these:

- Information exploitation, which involves estimation and prediction of signals from electronic intelligence, imagery, audio, and speech processing.

- Information fusion and understanding, the process of combining information to estimate and predict the state of the battlespace.
- Information management, the harnessing of information resources and capabilities of an organization to accomplish its objectives.
- A framework for developing and demonstrating information management technology solutions.
- Advanced computing architectures, R&D focused on fundamental models of computation and on engineering techniques for the design of computing systems.
- Cyber operations, encompassing technologies for information assurance, computer network defense, and computer network exploitation.
- Connectivity, addressing the need for high-capacity C4ISR networking and enterprise management and control.
- Command and control, involving decision support, adversarial modeling, and battlefield simulation.

AFRL/IF and AFOSR have a joint strategic planning process to help them align 6.1 investments with AFRL/IF program needs. The process involves AFOSR program managers in strategic deliberations with AFRL/IF and input from AFRL/IF on AFOSR's 6.1 investments. The committee concluded that the AFRL/IF program, though certainly not exhaustive, is well matched to the long-term challenges outlined above, and it also concluded that the success of the AFRL/IF program is critically dependent on a number of basic research advances that are well within the purview of AFOSR.

Final background information was provided to the committee in a briefing by Shankara Sastry of the University of California at Berkeley, who chaired an IS&T review panel for the Air Force Scientific Advisory Board that issued a report on November 21, 2003. That review identified several key IS&T areas for investment: the scalability of the information management systems under investigation, security and real-time capability, mission management technologies for unmanned air vehicles in mixed manned/unmanned environments, and embedded software and systems. AFOSR's current response to these various challenges is suggested by these excerpts from its ongoing Broad Agency Announcement 2005-1,¹¹ which invites IS&T proposals in the following areas:

Dynamics and Control

The . . . program is interested in . . . guidance, navigation, and control of autonomous aerial vehicle systems and teams; image tracking and robust feedback control in high scintillation environments; . . . and . . . hybrid

control systems that can intelligently manage actuator, sensor, and processor communications in complex, spatially distributed systems. . . . Interest in control of complex, multi-scale, highly uncertain nonlinear systems is increasing. . . . Interest exists in the development of control concepts applicable to single and multiple unmanned aerial vehicles (UAVs) and micro air vehicles (MAVs). Areas of interest include cooperative/collaborative control of a team of UAVs conducting operations such as cooperative combat ISR, electronic attack, urban warfare, wide area search/attack, and persistent area denial. Real-time, adaptive acquisition, classification, prosecution and assessment of geographically dispersed targets is envisioned, requiring cooperation amongst UAVs. . . . A cooperative decision and control theoretic framework is of interest to address robust dynamic control of distributed UAVs executing multiple, strongly coupled tasks with a high degree of decentralization. A long-term goal of the Dynamics and Control program is control for intelligent autonomy, in order to achieve a higher level of autonomous control. . . . Providing UAVs and MAVs with faculties of wide field-of-regard perception will be a significant step toward the realization of autonomous control, and in this area research in vision-based guidance, navigation and control is of interest.

Optimization and Discrete Mathematics

Target tracking environments include multiple maneuvering targets in clutter, targets with low [signal-to-noise ratios], and cooperative tracking platforms. Computational complexity for real time applications is a key issue. . . . Research includes discrete event systems, especially as it relates to Air Force transportation, target tracking, command and control systems, and battlefield management.

Signals Communication and Surveillance

This . . . activity is concerned with the systematic analysis and interpretation of variable quantities that represent information, or convey information physically through a channel. Communications signals, enabling command and control, and surveillance images are of special importance An outstanding need in the treatment of signals is to develop resilient algorithms for data representation in fewer bits (compression), image reconstruction/enhancement, and spectral/frequency estimation in the presence of external corrupting factors. . . . These . . . hold promise in the detection and recognition of characteristic transient features, the synthesis of hard-to-intercept communications links, and the achievement of faithful compression and fast reconstruction for audio, video, and multi-spectral data.

¹¹Available at <http://www.afosr.af.mil/pdfs/BAA2005-1.pdf>, dated October 2004.

Software and Systems

The goal is . . . basic research needed to enable development of advanced computing science methods to support future Air Force needs in worldwide, 24-7 battlespace information management. . . . Research is . . . to meet . . . challenges including collection, control, and integration of the vast amounts of information flowing through battlespace information networks, protection of friendly information resources, and complexities in software and algorithm development in support of large information systems. Some specific areas of Infospheric Science research follow:

- Models of Information Flows
- Metrics for Information Flow
- Hierarchical Flow Models
- Information Dynamics
- Managing Massive Numbers of Triggers
- Information Pedigree/Certainty
- Stream Data Processing
- Automated Downgrading of Sensitive Information
- Preventing Self-Inflicted DoS Attacks
- Audit Data for Damage Assessment
- Steganography Detection
- Secure Code Composition
- Distributed, Assured Pipelines
- Application Layer Multicast Encryption
- Seamless Integration of Wireline-Wireless Networks
- Network Monitoring, Measurement, and Inferencing
- Ad Hoc Wireless Networking
- Middleware
- Joint Battlespace Infosphere (JBI) System Stability
- Dynamic System Management
- JBI Information Metadata and Structure
- Evolvable Components

The need to collect, integrate, and disseminate information from widely disparate sources will be crucial. . . . Information extraction from all sources of data is . . . of interest. . . . Network protection . . . detection of intrusion, forensics, and an active response and recovery from an attack on information, are needed. Basic research that anticipates the nature of future information system attacks is critical. . . . Research on effective security policies across large, heterogeneous infospheres is of high interest. Techniques to automatically detect deceptive data or information are of interest . . . [as are] mathematical approaches for the specification, design, and analysis of distributed software systems. Approaches for overcoming the increasing computational complexity of these systems are essential.

Artificial Intelligence

Mathematical foundations of information fusion must be established—robust, integrated fusion architectures for handling increasing diversity

of input sources. . . . Research is sponsored into how to make the best use of uncertain information; share and disseminate information; increase the accuracy, speed, and economy of the recognition and identification process; and aid the intelligence analyst. . . . Research is needed to develop large-scale intelligent systems that can address practical Air Force needs. . . . Means are sought to scale up those methods that work for small knowledge-based systems. . . . Formalisms need to be developed for the representation of and reasoning with uncertainty, in handling corrupt information, identifying deceptive information, and effectively using experiences. . . . We seek means to combine numeric and symbolic inference methods. . . . The program seeks to develop technology that will support decision-making. . . . Intelligent agents are needed [that are] capable of gathering information, reducing data to a manageable amount of essential information, and cooperating with other agents to solve problems. Research is also needed to combine artificial intelligence methods with operations research tools to overcome inefficiencies in solving some mission-critical Air Force problems (e.g., scheduling in a distributed, dynamic environment).

To cover this range of research, AFOSR has an annual budget of some \$25 million for IS&T: approximately \$16 million for academic grants, \$7 million for multidisciplinary university research initiatives (MURIs), and \$2 million for in-house research at AFRL/IF.¹² Although it did not examine the current IS&T portfolio in any depth, the committee's sense is that the resources are not adequate for the broad scope of the challenge identified earlier. While many of these topics are clearly important, the committee also found gaps and recommended changes of emphasis, which are reflected in Chapters 3 through 7.

In the next chapters, the committee offers recommendations for the top-priority basic research in IS&T to help the Air Force meet its IS&T goals. The recommendations fall into four very general areas: research to enable network-centric warfare, including networking, communications, security, real-time embedded systems, and other topics; research to support improvements in software development, including the ideas developed in Chapter 4; research in information management, including integration of complex databases, publish-subscribe systems, signal processing, and information fusion; and research into the interface between humans and IS&T, including human-computer interfaces, human-system integration, decision-support tools, and information fusion.

¹²Figures supplied by Clifford Rhoades, head of the AFOSR Mathematics and Space Sciences Directorate, in discussions with the committee on February 27, 2005, in Rome, N.Y.

3

Basic Research for Air Force Network Systems and Communications

The Air Force is becoming a network-enabled paradigm, wherein many of its capabilities will be generated through, and dependent on, the integrated efforts of multiple components. This approach to operations is expected to result in greater agility and attendant tactical advantages. However, as is the case with any untested concept, there is a need for technology that enables the analysis and execution of the new paradigm. In this case, the scope of change is extremely large, requiring reanalysis of force structures, doctrine, acquisition options, command-and-control systems, training, and long-range planning, not to mention the considerable challenges involved in engineering, constructing, and managing the actual networks. Much of the planning and analysis will depend on the research described in this chapter, which will provide the necessary conceptual and technical foundation. Network-centric warfare also is critically dependent on software, the subject of Chapter 4, on the effective distribution and management of information throughout the network, which is the prime topic of Chapter 5, and on the effective employment of that information, which is considered in Chapter 6.

TYPES AND CHARACTERISTICS OF COMMUNICATION AND NETWORK SERVICES NEEDED IN THE FUTURE

Figure 3-1 shows the range of components that will be networked in the future Air Force. The following capabilities are key to achieving the desired functionality, but off-the-shelf technologies are not yet adequate:

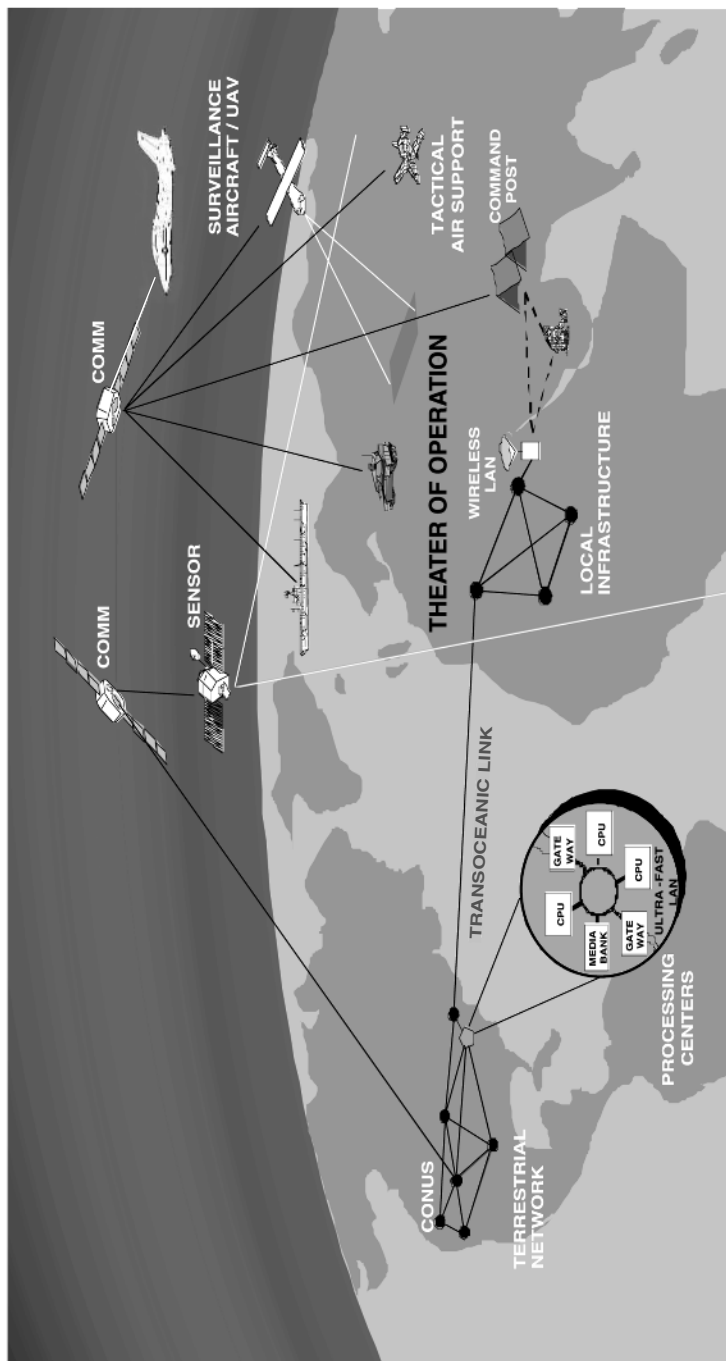


FIGURE 3-1 Schematic showing the range of components that will be networked in the future Air Force. Robust dynamic tactical networks support 10,000 to 500,000 communicating devices in theater. They go beyond traditional “tactical,” which is oriented to the voice and to human users, and beyond traditional ad hoc networks.

- The network must provide robust data and circuit services to tens or hundreds of thousands of fixed and mobile users with different service levels. Some of the service challenges include guaranteed rates, communication over difficult channels, hard time-deadlines, reliable message delivery over unreliable networks, security, and policy-driven resource allocation (explained below).
- Sensors in the networks must be able to relay data while maintaining coherency, and there must be high compression of correlated sensors to reduce communication needs.
- Laser communication must be available between satellites and aircraft and to terrestrial sites.
- Satellite services must be available to mobile and fixed users. Those services must (1) maintain connectivity at all times and track important users, (2) have low probability of detection (LPD), low probability of intercept (LPI), and antijamming (AJ) capabilities, (3) be secure, and (4) support very small as well as large terminals.
- Wireless radio services must be available in the field without pre-existing infrastructures, and they must be secure, power-efficient, have LPD and LPI, have AJ capabilities, and be able to maintain connectivity with important users.
- The satellite, radio, and wireless networks must be integrated with a high-bandwidth, affordable, and secure terrestrial fiber network as one supernetwork.
- The network must be an integrated one with multilevel security, data delivery within deadlines, and the capability for rapid reconfiguration and adaptation.

With the possible exception of power efficiency and some aspects of security, all of these characteristics are primarily driven by defense needs, and they will not be advanced by the commercial sector.

The remainder of this chapter describes some of the technical challenges of future communication modalities planned for the Air Force and the performance of networks when they are connected together in military settings. It concludes with recommendations of particularly important basic research areas that typically cut across multiple challenges.

TECHNICAL CHALLENGES POSED BY FUTURE AIR FORCE NETWORKS AND COMMUNICATIONS SYSTEMS

Challenges for Future Air Force Communications Systems

Defense applications have to contend with communication modalities that are not encountered in commercial and civilian settings. Satellite

channels have unusually long-delay data rates (sometimes known as delay-bandwidth products) and randomly fading dispersive channel characteristics. Radio channels, especially those associated with mobile platforms, have rapidly changing link capacities and connectivities, with disconnections or dropouts that can be on the order of minutes or more. In contrast to this dynamism, traditional layer 3 (Network Layer) and layer 4 (Transport Layer) protocols assume fairly stable underlying substrates that change, if at all, over the course of minutes—that is, much more slowly than most transmissions. Recent applications have shown these traditional protocols often yield low throughputs and poor quality of service when applied to DOD systems. In some cases, these protocols do not work at all despite valiant efforts to provide patches. Thus, the main challenge of Air Force communications is to provide assured connectivity between networks (albeit at varying rates) under difficult channel conditions, including during adversarial attacks.

A cost-effective investment strategy is to make maximum use of commercial technologies. However, not all the services listed above can be supported by commercial technologies. In many cases, commercial architectures only need to be modified, but in some cases completely new designs are needed.

Source Compression for Correlated Sensors

One communications challenge is the need for source compression methods when multiple sensors collect correlated observations, whether or not this was done in a coordinated manner. This is critical in order to fit the large volumes of sensed data in modest network capacities, especially in difficult communication environments. Particularly challenging is the identification and compression of highly correlated data in the absence of fine coordination among sensing platforms. This will require both lossless and information-lossy compression and new approaches to measuring signal fidelity tailored to battlefield signals. Another challenge lies in creating joint source-network coding: The separation theorem, which for point-to-point communication channels allows one to decouple source and channel coding, does not apply to networks of Air Force interest, where the data volume itself may cause packet loss due to congestion. It is important to explore the joint consideration of coding for both source compression and network transport (from layer 1 to layer 4).

Design of Communications Systems

A better understanding of capacity limits is needed in order to design the best architectures for Air Force communications systems, which

should perform at the theoretical limit for difficult channels and in the presence of adversarial attacks. The open-air satellite, radio, and optical channels planned for the Air Force exhibit fading, dispersion, and interference (benign and adversarial) not typically encountered in well-known classical channels, especially when the channel is broadband. For these channels, there are many different ways to design the communication system, and good system design will depend on understanding the channel effects and the fundamental limits of communication performance optimized over the class of possible architectures. Our current understanding of the fundamental limits of these channels is only rudimentary, and the space of communication designs has yet to be fully characterized and optimized. There are also open questions related to the capacity limits of multiple-access channels with or without coordination. Multiple-access techniques for open-air channels, such as satellite and battlefield radio channels, are extremely important for good communication performance. Classical techniques must be reconsidered in a context that includes adversarial jamming, interference, coupling, and little or no coordination with some key inputs such as the satellite channel. Classical information theory articulates the limits of communication capacities with no specific attention paid to the transaction delay of the session. In fact, most results are asymptotic, taken as the delay becomes large. In contrast, some critical Air Force applications are very delay sensitive, and a theory for the fundamental capacity limits of a communication system with delay constraints needs to be developed. In some applications power efficiency is also required.

Classical communication and information theories do not incorporate an element of adversarial attacks, although some results related to LPD, LPI, and AJ communications were developed 40 years ago. The technology assumptions underlying those results have changed, of course, and none of the results were obtained with networking in mind. For instance, networks using open-air interfaces present a new set of vulnerabilities that can be easily exploited by adversaries. It is important to understand the set of possible technologies and techniques available to adversaries and to design the network with those in mind rather than treating them as an afterthought. Game-theoretic and optimization techniques will likely be relevant approaches, as will methods for authentication and cryptography.

Reliable and Efficient Free-Space Optical Communications

Free-space optical communications in space is relatively well developed, though its capabilities for terminal spatial tracking systems and the weight and power demands of the systems can be improved using more

sensitive receivers. The current sensitivity of the best receivers (direct or heterodyne detection) available in prototype form requires 2-10 received photons per bit. With the advent of solid-state, single-photon detectors, a photon counting receiver can achieve 5-10 bits per received photon. A fundamental investigation into the best combination of modulation and coding under technology constraints could lead to large payoffs for Air Force space and aircraft systems.

More problematic is free-space optical communication over the turbulent atmospheric channel where fast fading (a few multiseconds) presents difficulties in the design of modulation/coding and receivers. These systems are important for aircraft and ground applications. For supersonic platforms, the bow shock presents even bigger challenges than atmospheric turbulence owing to its speed, which is three orders of magnitude higher. The current mainstream receiver uses adaptive optics to compensate for phase-front distortion, which gives rise to fading and other undesirable channel effects. While its performance is adequate, the solution is costly and heavy. The Air Force should pursue new system techniques, such as spatial-, temporal-, and frequency-diversity receivers in both coherent and incoherent forms, to mitigate the turbulence effect.

Energy-Efficient Communications

The Air Force needs more energy-efficient communications for two reasons. The first is that for operations in a hostile environment, keeping the transmitter power low is good for LPD and LPI. The second is that some operations may be conducted with ground mobile radios, where battery conservation is very important. The first problem can and should be addressed as part of the research outlined above aimed at developing capabilities for efficient communications for difficult channels and under attack. The second includes efficient signal processing and higher-layer network protocols, which will be addressed below in the discussion of sensor networks.

Dynamic Information Theory for Network Applications

Recent developments in the new applications of communication networks often require that information be processed in a distributed fashion in real time. For example, in a sensor network, a relay node often needs to forward a received signal without being able to decode the embedded message reliably; when controlling a dynamic system, decisions often need to be made based on noisy observations. Conventional information theory, which is based on the concept of reliable communications and

thus assumes large delays and extensive coordination, is often not applicable to these new problems. The theoretical understanding of how to handle information without the error-free guarantees is therefore important for future developments of highly dynamic and distributed network applications. The study of communication systems with partial or list decoding using a layered coding structure could be the first concrete step in this direction.

Challenges for Future Air Force Networks

The theory of networks has not matured to a point where one can predict how well the protocols developed heuristically in one application setting will perform on a communication network built on radically different communication modalities. To deal with the new and complicated modalities of importance to the Air Force, fundamental tools must be developed to help understand how networks may perform in new environments and to optimize architectures. It is simply too costly to develop these architectures and protocols ad hoc and then experiment with the communication links in the field. In addition, there are many possible ways to configure a communication system. If a communication system is going to be used as part of a network, it should be designed jointly with the network and not independently. For example, end-to-end reliable data delivery can be a function of the communication system (using diversity receivers and error-correcting codes), or it can be a function of the network (using diversity path routing and automatic repeat requests at various layers such as the Data Link Control Layer and the Transport Layer).

High-Bandwidth, Affordable, and Secure Global Fiber Network

The Air Force needs a global broadband fiber network to interconnect its multiple modalities and act as the high-speed backbone. Current networks, such as the Global Information Grid (GIG), use commercial technologies with wavelength-division multiplexed optical transport and electronic packet switching. This technology is both costly (it does not scale well to very high data rates) and insecure. For terrestrial networks there is a burning need to create economical new optical transport mechanisms such as optical flow switching (dynamically set up with short-duration connections on demand) with lightweight protocols and also a security architecture for both the Physical Layer (the transport mechanisms) and the higher layer protocols.

Network Layer (Routing) Design

Routing over the Internet today is purposely set to change rather slowly to prevent oscillations and overreactions to sudden rate surges. In contrast, most of the channels used by the Air Force can experience rapid changes in connectivity and link capacities due to random channel effects, mobility, and adversarial attacks. Thus, most routing protocols in use today cannot keep up with such changes nor can they provide effective congestion and flow control. To create better protocols it is important first to understand the fundamental effects on the Network Layer of these fast dynamics. For example, mobile networks (especially those that support high data rates and time-deadline services) are not well served by conventional commercial protocols, which are mostly designed for static or quasi-static network topologies. Many of these dynamics appear first in the Physical Layer and permeate up the protocol stack. The mobility aspect calls for the joint design of the Physical Layer and the higher layers, including routing and the Transport Layer. The development of fundamental mathematical tools is a prerequisite to the understanding and solution of this very important problem.

Transport Layer Design

The task of providing error-free, end-to-end delivery of messages in a network is jointly shared by the Data Link Control Layer (layer 2) and the Transport Layer (layer 4). For random channels (especially those in mobile networks), these two layers interact in ways that are stochastic and not yet well understood; the interactions often result in drastically reduced throughput. Researchers have not yet succeeded in predicting the performance of specific implementations and determining the fundamental limits of these protocols when coupled with random channels. The answer might come from control theory and stochastic system analysis.

Current Transport Layer protocols assume a stable Physical Layer communication infrastructure, so packet losses are typically interpreted as buffer overflow at routers due to congestion. In the random channels with which the Air Force has to deal, packet losses can also be due to path fades or intentional interference by an adversary. Some of these effects can be very fast owing to the mobility and agility of electronic attacks. The Transport Layer Protocol will react poorly to these effects, resulting in very low network throughput. This is plaguing many current DOD programs. For the difficult channels DOD and the Air Force deal with, fundamental research is needed to address this large and critical performance gap. Joint design of the Physical Layer and upper Network Layers must be done in these cases.

Network Coding and Diversity Routing for Difficult Channels

For networks using difficult channels, such as channels that change quickly compared with network coordination time, it may be impossible to solve the network routing problem with traditional techniques. Network coding is a new technique that uses almost no buffering, no route computation, and no flow control. The simplicity of this technique could be the answer to these problems (especially for mobile networks that have no infrastructure and dynamic topologies that change too quickly). However, the area of network coding is just beginning to be defined, and there are many theoretical problems in throughput, delays, and congestion control that need to be solved. This topic is explained further later in this chapter.

Because some packets on Air Force networks have hard time-deadlines, those networks must be designed with an understanding of the limits and trade-offs of network throughput, delay, and available resources. Diversity routing, for instance, is one possible technique for increasing performance, albeit at the expense of more network resources. Some network topologies are better for this purpose than others, so there is a need to optimize network performance as a function of topologies and protocol design at the higher layers.

Policy-Driven Efficient Network Resource Allocation

There may be a misconception that with modern fiber and radio technologies, bandwidth is plentiful and should not be a bottleneck to network performance. This is far from true in a tactical battlespace, where there are several limitations to bandwidths and quality of service. First, the amount of asset that can be deployed is limited. Second, difficult channel conditions can drastically reduce rates; in some situations, expensive relay assets must be put in place to provide minimal but critical connectivity. Finally, in adversarial attacks, the number of surviving routes may be few and signaling bandwidths may have to be used to provide AJ via band spreading, reducing the actual data rates. These factors all point to the reality that high-performance tactical network assets will always be precious and in demand. To allocate these precious resources fairly and provide the best operational support to our forces, the network management system must be able to take into account external policy on priorities, which would change from time to time depending on the important missions at the moment and requirements that arise in the field. The network management system should be able to translate these high-level guiding principles into network actions such as routing and media access control priorities in a timely fashion without a human in

the loop. Currently, assets are managed manually, and the process is far from responsive and optimal.

There are risks associated with a network management system that is agile and responsive. To guard against the network going into undesirable states, it needs to be closely monitored for unusual behavior, and fallback procedures and network states must be implemented to ensure minimal critical network service performance. The fundamental research that must be done to support such a vision includes multicommodity resource allocation in a competitive environment, game-theoretic approaches to deal with adversarial attacks, and inference techniques to assess network states in the presence of noise and intentional masking.

Networks Designed to Function Under Adversarial Attacks

There will be many nodes in a tactical network and the span can be wide, with connectivities back to the continental United States from anywhere on the globe. Since it will not always be possible to ensure that no nodes are compromised, the network should be designed to sense dead or malfunctioning network elements and route around them. In addition, network failures often result from operator errors, so the network should have an architecture that confines such damage to a local area and does not allow it to propagate across the network. When the network senses outside attacks, it should have the capability of first locating the real entry points and defending and removing the attacks. Because sensing and other protection techniques can fail, the network should be designed to recognize such failures and be able to continue to function, probably at a lower performance level. This feature is necessary because it may be impossible to avoid rogue nodes or network operator errors. Techniques such as Byzantine robust networking can provide sabotage-resistant routing and defend against compromised trusted network elements. There is some parallel here to fault-tolerant computing, except in this case the adversarial factor must be included in the analysis. Of particular importance are the techniques for refreshing and distributing cryptographic keys for dynamic narrowcast groups. This is especially difficult when a central authority is not readily connected to the population. A provably secure consensus agreement and key exchange mechanism must be devised.

In particular, jamming of open-air communication systems (e.g., satellite links and wireless radios) is a standard military technique to bring about denial of services. Only a few military radio communication systems today have been designed with any AJ capability. Of these, only the satellite system Milstar has adequate AJ capability to deal with jammers that employ modern technologies. Milstar succeeds by deploying spread spectrum and antenna nulling techniques. However, Milstar was designed

for voice communications. When data networking protocols are put on top on this physical layer medium to provide data services in new DOD networks, such as those described by the Transformational Communication Architecture, new vulnerabilities have surfaced that significantly weaken the system's AJ capability. For example, the adversary need only jam a few bits per Internet Protocol (IP) packet to fool the Transport Control Protocol (TCP) into believing that there is congestion at some routers downstream. In response, TCP will begin closing its transmission window (reducing the number of packets released in flight), reducing the effective throughput of the system to less than 1 percent. Counteracting this new attack requires a combination of techniques from the Physical Layer to the Transport Layer: spread spectrum, nulling, rerouting and diversity routing, and changes to TCP. There are several other known network weaknesses, some of which would appear to be correctable, although history tells us that ad hoc changes to overcome vulnerabilities often open (or overlook) other vulnerabilities. What is needed is a systematic fundamental look at networking, perhaps with a solid mathematical foundation, to provide some assurance of protection. In particular, vulnerabilities to cross-layer attacks (which might be more effective than traditional within-layer attacks) should be examined and addressed.

Cross-Layer Network Design and Optimization

More generally, there would be value in rethinking the network architecture across the layering boundary for military networks, because the difficult channels encountered by DOD are a reality that cannot be avoided. It is easy to state the obvious—namely, that when network layering structures are broken and the architecture is optimized without boundary constraints, the network will perform better. However, the unstructured problem may become so unwieldy that any insight and hope of arriving at a good architecture is lost. There should be a wholesale reevaluation of the functions of each traditional layer to see if there is a better way to group these functions in view of new communication modalities. Even if it is highly unlikely that one could find a layering structure that totally decouples the network architecture, breaking down the architecture problem into weaker interacting subsystems would be a great advance over today's architecture.

Optimization is a fundamental discipline very important to network design and multiagent control. It relies on tools such as mathematical programming for optimization in static contexts, dynamic programming for optimizing systems that evolve over time, and game theory for optimizing in the presence of competing interests. A key characteristic of the Air Force's communication networks is their decentralized nature. They

are being constructed from the interconnection of different systems and users with different objectives and performance measures, and their operation relies on varying degrees of collaboration and competition among these entities. Thus, there is a need to understand the optimal allocation of resources in the presence of users with heterogeneous service requirements, which would seem to necessitate the use of game-theoretic and economic market models for resource allocation and network control.

Challenges in this area might include (1) developing game-theoretic and artificial intelligence models for resource allocation problems in different types of networks, (2) studying the optimality/efficiency properties of such multiagent control schemes, (3) developing new theoretical tools in optimization and game theory to analyze convex and nonconvex optimization problems and games that arise in these settings, (4) devising efficient computational methods with established convergence/rate of convergence properties, and (5) developing tractable approximate dynamic programming methods that generate suboptimal control strategies.

Some of the difficult network design problems can benefit from new mathematical analysis tools. In the past few years, ideas developed in statistical physics seem to have had a strong impact in many areas of electrical engineering and computer science. In particular, “replica” and “cavity” methods arising from physics have led to better (heuristic) understanding as well as the development of novel algorithms for hard questions arising in combinatorial optimization, error-correcting codes, statistical inference, and the like. It would be profitable to study these connections rigorously as well as to identify their strengths and weaknesses. This would also lead to a healthy shift in the engineering thinking of large-scale systems such as a globally connected defense network.

Network Coding

Traditionally, network schemes have considered transmission of information in networks in terms of flows—in essence, networks are used in a manner akin to any transportation system. Recently, a radically different approach has been proposed. Rather than considering information as a good to be distributed, one could actively make use of the fact that information is composed of algebraic entities, such as bits, that can be manipulated. One could then consider combining the information of different packets in the interior of the network in order to increase throughput, even in the case of a lossless network. This type of network routing has significant potential application in military networks, where links are unreliable to begin with and there are also adversarial attacks. The main theoretical underpinnings of this area should be developed by merging algebraic theory, networks, and stochastic algorithms to estab-

lish the fundamental performance characteristics of network coding. A strong focus should also be maintained on randomized distributed algorithms for multicast network coding, which has shown promise for network operation in preliminary industry demonstrations.

Some specific basic research challenges in network coding would include the following:

- Concurrently explore the theoretical foundation of network coding and applications of that theory.
- Bring network coding to the area of network security by developing the ability of network coding to perform algebraic manipulations in the interior of the network.
- Consider the economic implications of network coding. In particular, the pricing of resources in network coding differs significantly from traditional network and other pricing mechanisms based on additive resources. Network coding builds on the use of entropy rather than volume as a pricing unit.
- Develop the theory and implementation of network coding in lossy networks for data dissemination and for creating altogether new means of performing distributed storage.
- Bring network coding to wireless applications. The dynamic and lossy nature of wireless networks renders the use of network coding, particularly for robustness, very effective.

Energy-Efficient Sensor Networks

This is a topic of considerable interest in the private sector, so there are opportunities for leveraging developments there. However, there remain significant research challenges in the area of networked sensor-based embedded systems (for instance, to control swarms of UAVs). Within sensor networks of importance to the Air Force, the following characteristics are common: (1) sensors vary widely in size, computing resources, and sensing information, ranging from very simple sensors like temperature sensors to acoustic sensors and surveillance cameras capturing and distributing video signals, (2) sensors are deployed in very large numbers to ensure adequate coverage, (3) sensors generally perform only a single task (such as sensing temperature or other environment characteristics), but they operate at very high sensing rates to ensure accuracy in environment monitoring, object identification, coordinate location, or other tasks, (4) sensors are deployed in difficult environments where energy sources are not easily accessible and where external forces can damage them, and (5) sensors need to communicate with the environment (other sensors or other IT devices) to distribute information.

The embedded sensor area is relatively new, although some work has been done in the development of operating systems for sensors and in protocols for the Physical, Media Access Control (MAC), Network, and Transport Layers. Related work exists in vehicular sensor networks, applied to sensor information such as Global Positioning System (GPS) positions, video, and others. Efforts are starting in programming and debugging large-scale sensor networks. Many relevant results exist in the use of sensors for location identification, but their integration with other sensing platforms has still not happened. Energy is an issue, and extensive work has been done on minimizing the consumption of ad hoc sensor networks. However, significant research challenges specific to Air Force applications remain in the area of sensor networks:

- Development of scalable large-scale sensor network topologies and energy-efficient architectures, including algorithms, operating systems, protocols, and internetworking management for the collection of information, taking into account constrained resources such as energy, communication bandwidth, and processing capacity.
- Design/development of efficient algorithms and protocols when sensors are part of a vehicular area network and need to communicate with one another under high-mobility scenarios.
- Development of efficient programming support and tools for programming large-scale systems—that is, developing efficient languages, compilers, and debuggers for programming large-scale networks of sensor nodes.
- Development of efficient memory systems and caches on sensor nodes to capture the state of the environment.
- Development of energy-efficient encoding schemes to distribute information at low data rates.
- Development of location identification sensors (e.g., radio-frequency identification), their associated processing algorithms, and their integration into the overall computing/communication environment.
- Investigation of security and privacy to protect sensing information.
- Incorporation of fault tolerance to ensure reliability of processing at sensing nodes and during information dissemination.

RECOMMENDED BASIC RESEARCH AREAS IN SUPPORT OF AIR FORCE NETWORKS AND COMMUNICATIONS

Some of the challenges described above need more investment by the Air Force than others. In this section the committee recommends a few particularly important basic research areas that typically cut across mul-

multiple challenges. The committee believes that the most valuable results will be achieved when communications and network challenges are considered together, so it has formulated recommendations that target these topics jointly. Because these recommendations are for basic research, the research programs should be tailored to the abstractions that arise from these applications rather than the specific applications themselves.

Satellite Communications and Data Networking

Satellite communications will be an important means of providing linkages between disconnected pockets of radio clusters and theaters to the outside world. Very high frequencies (>20 GHz) are necessary to support high data rates in benign situations and provide good AJ capability with wide-bandwidth, spread-spectrum modulation. Thus far, all military satellite systems have been designed for circuit operations having relatively few users (~1,000s), and the systems are highly inefficient for the transmission of bursty data traffic. These inefficiencies are exacerbated by the unique characteristics of the satellite channels. For example, satellite systems often have longer propagation delays and higher bit error rates than their terrestrial wired counterparts, while the open-air interface for satellite channels (and the high sunk cost of those channels) argues for dynamic sharing of resources. To accommodate data communication for many (~100,000) bursty computer users, the architectures of the Air Force's satellite communications systems will have to undergo a radical change. At the Physical Layer, the communication system should adapt to weather-induced impairments with power and variable rate control; over 10 dB of efficiency gain could be realized. To serve a large population of bursty data communication users, an efficient and rapid-response MAC Protocol needs to be adopted for the unique properties of the satellite channel. Rapidly changing channel conditions, especially under open-air attacks, will require a fresh look at the routing and flow/congestion control algorithms at the Network and Transport Layers. Since this satellite network will be interconnected to other networks, issues of routing at the Network Layer, reliable packet delivery at the Transport Layer, and network management and control will also need to be addressed. The corresponding research efforts should center on communication theory, experimentation, and practice, with emphasis on heterogeneous networks involving satellites, aircraft, UAVs, and terrestrial wireless and optical systems. Examples of some high-priority research topics include the following:

- Design of architectures and protocols for satellite communications systems, including design of satellite constellations, efficient

resource allocation algorithms, Link Layer Protocols for dealing with special properties of space systems, and efficient channel-sharing mechanisms (i.e., MAC Protocols).

- Development of hybrid space-terrestrial network architectures, including space-ground network architectures and interfaces, and the design of protocols for internetworking over heterogeneous networks that are efficient, reliable, and able to assure quality of service.
- Network architectures and protocols for air vehicle systems, including architectures for reliable communications between autonomous air vehicles (e.g., UAVs) for the purpose of delivering time-critical control information.
- Development of analytically tractable models of wideband channels and exploitation of channel characteristics to increase transmission availability, reliability, and/or spectral efficiency.
- Development of new lossless and lossy techniques for joint compression and aggregation of correlated sensor signals.
- Efficient design and performance analysis of transmission systems, receiver algorithms, diversity methods, and impairment mitigation techniques for air-ground, space, and inter- and intravehicle communications.
- Reliable and robust ultrawideband communications for high-speed wireless access, inter- and intra-airplane and -spacecraft communications, UAV communications, and special operations in harsh environments.

Radio Communications and Networking

The development of wireless applications and wireless networks in recent years has motivated some significant improvements in communication theory and information theory. However, maintaining good performance at the Physical Layer for difficult radio channels in the battlefield will be very challenging. Channel conditions and topologies can change rapidly, and the network must efficiently manage its communications over these shared, dynamic, unreliable resources. High-priority research goals in this area include the following:

- Development of a methodology that unifies the rich collection of new results in the wireless field. While numerous new techniques have been introduced in recent years to address different aspects of wireless applications (e.g., multiple input, multiple output (MIMO), beam-forming, energy-efficient ad hoc networks), translating these results into fundamental principles requires that we

first merge them into a unifying framework. Designs for modulation and coding, antenna shaping, and processing should be optimized across multiple nodes and with resource scheduling and routing and adversarial attacks all incorporated as key elements in the problem.

- Extension of the basic setup of information theory to allow more flexible communications. The primary goal of information theory is to enable reliable communication over long blocks at a prescribed rate. Wireless applications, however, call for breaking such operational scenarios. Many of these applications require communication over dynamic links, with varying qualities and performance requirements. The goal is thus to extend information theory to address the dynamics of information transmission and thereby develop rateless, blockless, self-adaptive communication schemes with delay as a prominent constraint.
- Development of cooperation in wireless environments and the architecture of wireless networks. While most of the current wireless networks adopt the cellular architecture, new self-organizing networks promise much better flexibilities, making many more applications possible. A basic research goal in this area would be to study cooperation between wireless nodes with limited information exchange between them and, further, to develop new wireless network architecture that takes advantage of such local cooperation.
- The emergence of new applications with diverse requirements has made traditional wireless systems suboptimal, often leading to inefficient use of spectrum, energy, and other resources. It is important to understand how the communications paradigm is evolving and to develop methodologies for the design of flexible, efficient networks. Of particular interest is how decentralized systems—where decisions are made with local information only—can achieve globally optimal use of resources. Also, the research efforts should incorporate factors that are traditionally overlooked, such as the competitive equilibrium and mechanism design for the resource allocation problem, in order to allow the design of fundamentally different systems based on the paradigms of cognitive/software-defined radio and opportunistic communication.

Free-Space Optical Networks

Free-space optical networks are important for Air Force applications connecting satellites, aircraft, and ground terminals. The Physical Layer properties of fading and phase distortion for optical links passing through atmospheric and aircraft turbulence create serious problems for protocols

for higher network layers, often resulting in very low throughputs and long delays for data delivery. It is important to consider optics as a network system rather than an isolated communication link. As described in the section above on Transport Layer Design, the Physical Layer and the network architecture must be codesigned. High-priority areas of research include the following:

- Understand the interaction of network protocols with the optical link and quantify the performance shortcomings.
- Jointly design the Physical Layer and higher network layers, making judicious use of spatial, temporal, frequency, and route diversity.
- Internetworking with other modalities, including using them as backup for guaranteed message delivery.

Basic Research for Air Force Software

Network-enabled systems are by definition dependent on complex software because of the great number of possible states of the networks. Building such software is very challenging. The systems that require such software transcend a range of Air Force applications, from intensive human-machine systems (e.g., command and control, air operation centers) to embedded applications (e.g., avionics systems). Increasingly these applications are connected by networks into a system of systems, and, in fact, the distinction between enterprise and embedded systems blurs as the focus is increasingly on the interconnectedness of all such systems. This chapter develops recommendations for basic research to enable more successful development of network-enabled systems, and those recommendations therefore apply to enterprise systems as well as to embedded systems.

The Air Force's experience with software over the past 30 years or more has been one of grappling with more and more lines of code to implement increasingly complex functionality in a software-intensive system that itself may contain many interconnected subsystems (e.g., a fire control system and a navigation system on a fighter aircraft). This software complexity has contributed to the challenges encountered in developing and fielding operational systems such as the F-22 aircraft and space systems such as the Space-Based Infrared System (SBIRS).¹ The

¹Department of Defense, Report of the Defense Science Board Task Force on Defense Software (2000). Available at <http://www.acq.osd.mil/dsb/reports/defensesoftware.pdf>.

advent of network-enabled systems is in some sense an extrapolation of the old challenge: the development, integration, and sustainment (maintenance) of many interconnected, large, software-intensive systems. And the Air Force expects to continue to develop and sustain systems that are unprecedented in terms of their size and complexity. Therefore, a basic research program in software must address how to develop and sustain such large families of software-intensive systems.

Rather than focusing on large-scale code development—a challenge that is being researched by others—the committee recommends that AFOSR focus on a set of important software engineering issues that are key to successful Air Force network-enabled systems but that have received limited attention. This recommended set of issues centers on how to understand what to build and how to ensure that its behavior is relatively predictable and acceptable, both during design and in operational use. Little is predictable about an end product's behavior when and where it is needed, and this knowledge is critical to understanding both how to engineer a new system and how to sustain and upgrade an existing system. In order to address this challenge, the committee recommends revisiting a concept articulated in 1987 by the Air Force Scientific Advisory Board: Pursue research that helps imbue unprecedented software development projects with some of the attributes that enable success when there *is* a precedent. Quoting from that report:²

The main philosophy behind risk reduction [for software] is to make unprecedented systems as precedented as practical.

Precedented systems are those for which (1) the requirements are consistent and well understood, (2) the system architecture (both hardware and software) is known to be adequate for the requirements, and (3) the acquisition and development teams have worked together to develop a similar previous system. Violation of one or more of these elements of definition causes the system to be unprecedented. Consequently, certain risks arise that require an acquisition-development strategy for their mitigation.

A primary assertion . . . is that precedence is one of the most important elements in the timely development of quality software. In many cases, it is more dominant than many of the more narrowly focused issues that receive more attention, such as language, tools, hardware instruction set, architecture, etc. Simply stated, the concept of precedented systems means that the learning curve applies to software development, as it does in all other areas of human endeavor.

²Air Force Scientific Advisory Board, *Adapting Software Development Policies to Modern Technology* (1987). This report is summarized in R.J. Sylvester and M. Stewart, "Unprecedented systems," *Encyclopedia of Software Engineering*, J. Marciniak, ed., John Wiley (1994).




FIGURE 4-1 Logical flow of software development in support of Air Force operations.

That report went on to note that engineering and management techniques have been developed to address the predictability of software and its behaviors, but only in domains that are fairly well understood and where the team developing and acquiring the software-intensive system has worked together on a similar system. With network-enabled Air Force systems, such experience in software development and maintenance does not yet exist, at least not with actual systems in operational settings (as opposed to tightly controlled and scripted exercises). That gap can be addressed by research into three major questions:

- How do we discover and understand what is needed?
- How can critical nonfunctional attributes (those that are desired or necessary but ancillary to the software's primary functionality) be implemented in a predictable fashion?
- Can the resulting software, once fielded, evolve to satisfy new needs discovered as it is used?

Figure 4-1 shows the logical flow from one question to the next and also indicates an appropriate research theme for each question. These themes, which leverage related research already under way at AFOSR and other research funding organizations, are explored next, and specific research topics are recommended.

COEVOLUTION OF AIR FORCE CONCEPTS OF OPERATIONS AND SYSTEM ARCHITECTURES

One of the persistent problems in software development is requirement uncertainty. This is particularly problematic with the complex unprecedented systems envisioned by the Air Force. The software management literature is replete with methods to instill disciplined approaches to

developing and managing (e.g., the Software Engineering Institute's Capability Maturity Model Integration),³ but an even more powerful methodology is needed for the Air Force's next-generation systems; it will require far more than process discipline to avoid repeating recent problems encountered with Air Force software systems.

The committee recommends that AFOSR sponsor research to extend the philosophy of software development models, such as iterative development, that support rapid prototyping of a software system so that end users can work with the system to see if it satisfies their needs. Through such interplay, the prototype then becomes an explicit representation of the requirements. To extend this research for, say, a network-enabled C4ISR system, a software model of the proposed system could be constructed and used in simulation-based exercises. The exercises would serve the dual purpose of discovering and validating user requirements, which help build Air Force concepts of operations (CONOPS) while supporting the design and analysis of the architecture. This is consistent with the architecture-first approaches advocated in best practices (e.g., using the Unified Modeling Language (UML) to enable architecture design that is motivated by cases supplied by end users) and the DOD standards for operational, system, and technical views of systems. This model for requirements definition has the additional benefits of building stronger communication between researchers and the user communities and of providing infrastructure to better support team-based problem solving involving humans and automated decision aids, which is particularly valuable for next-generation enterprise systems.

Current research in executable architectures and in engineering tools for the design and analysis of functional and nonfunctional attributes provides a basis for enabling this coevolution. Architecture development methods typically rely on standards-focused specifications and policy documents (e.g., those for the GIG, JTA, DODAF, DII COE) and on environments based on industry standards such as real-time CORBA. Those methodologies bring together operational views, which describe how the system will be used from the users' perspective and which might be expressed by the software or system engineer in a formal language such as UML; technical views, which describe the building codes or standards with which systems will be constructed; and the system view, the actual architecture of the system to be fielded. The system view satisfies the needs as defined in the operational architecture and adheres to the constraints imposed by the technical view.

³Capability Maturity Model Integration is a service mark of Carnegie Mellon University.

Technologies now exist for modeling an architecture, the most popular being UML. Tools and methods exist for evaluating the architecture from the perspective of scalability, real-time performance, survivability, and other nonfunctional attributes; for example, the Architecture Tradeoff Analysis Method (ATAM)⁴ is a commonly used analysis technique for architecture assessment. Some work has been done on executable and model-based architectures, which enable early execution in a simulation environment. These analysis methods typically work on static representations. A new approach in the software engineering community, the Model Driven Architecture (MDA), is designed to document an organization's business and business rules and the structure of its information infrastructure so as to create an expression of the business need for software that can guide IT projects. The power of the MDA approach is that the people who know the business will develop and maintain the business model, while the software engineers will develop and maintain the design and documentation of the software. These ideas are also being explored by other communities. For example, the workflow management community has developed an XML-compliant standard to support the explicit definition of business processes and the separation of these processes from underlying applications.⁵ In addition, the Object Management Group (OMG) recently released a draft version of a new language called the System Modeling Language.⁶

These recent advances would be extended through the recommended research into methods to allow the coevolution of Air Force CONOPS and system architecture. The committee recommends a general goal of providing executable versions of systems earlier in the development process to ensure that the operational view (the CONOPS) is accurately, effectively, and consistently represented by the system and technical views. To achieve that goal, the committee recommends research into the following:

- Methods to support rapid composability.
- Semantic extensions of current modeling languages to enhance composability along with representation of and reasoning about behavior.
- Development of tools that enable the construction of executable versions of models in system modeling languages.
- Methods that support experimentation, operational assessment, and the use of initial architecture representations in exercises. An

⁴ATAM is described at http://www.sei.cmu.edu/ata/ata_init.html.

⁵For example, see <http://www.wfmc.org/standards/docs/tc003v11.pdf>.

⁶See <http://www.sysml.org>.

example might be scripting languages that allow end users to explore early versions of software and help encode their preferences into the final architecture.

- Approaches that allow user tailoring, definition, and exploration of new processes and automated learning based on past problem solving.
- Experimentation and demonstration of these research approaches in domains of relevance to the Air Force.

SOFTWARE BEHAVIOR ENVELOPES

Air Force systems need to be scalable, interoperable, survivable, secure, evolvable, and capable of predictable real-time performance and energy awareness, among other things. These are “nonfunctional” attributes: characteristics of the software that, while necessary for the success of the system, are ancillary to the software’s ability to execute its function(s). The degree to which such attributes can be achieved in a system requires trade-off analysis during design and monitoring during execution. To this end, the committee recommends that AFOSR support a new line of research, extending DARPA-funded, model-based software research to explore the concept of a software behavior envelope. Dynamic analysis of nonfunctional attributes based on embedded models could be viewed as the behavior envelope of a network-enabled system. The extent to which the software could be modified within the envelope, by developers or end users, provides a measure of system elasticity. The analogy is to the flight envelope of an aircraft: Stay within the envelope and performance is safe and assured; wander outside the envelope, and tragedy could follow.

This topic would be a new area of research for the software community, but there is related work upon which to build. One related area is the approach called rate monotonic analysis (RMA), which is commonly used for real-time scheduling in operating systems. RMA was credited with enabling the rapid debugging and repair of the Mars Rover during its 1997 mission.⁷ Another avenue of research is qualitative modeling.⁸ This is worth reexamination because of the huge increases in computing power over the past 20 years, the lack of which was a limiting factor in those earlier efforts. A third related area comes from a recently completed DARPA program, the Model-Based Integration of Embedded Software

⁷See <http://www.cs.cmu.edu/afs/cs/user/raj/www/mars.html>.

⁸For example, see D.S. Weld and J. de Kleer, eds., *Readings in Qualitative Reasoning about Physical Systems*, Morgan Kaufmann: Los Altos, Calif. (1990).

(MoBIES) program, which was largely executed by AFRL's Information Directorate.⁹ A fourth area of related research would build on reference frameworks for composable systems (e.g., the work of David Garlan and colleagues at Carnegie Mellon University) and the work at the Software Engineering Institute in predictable assembly from certifiable components.¹⁰ Another relevant line of work is that by Janos Sztipanovits and colleagues at Vanderbilt University.¹¹ The union of these directions provides a good starting point for research into how to include executable versions of intended system behavior into actual systems, thus providing a means to establish a software performance envelope. This recommended program of research would leverage and extend ongoing research called out in AFOSR's BAA 2005-01 and research supported by other DOD laboratories, the National Science Foundation,¹² and other agencies.

EVOLVABILITY THROUGHOUT THE LIFE CYCLE

Once a software architecture has been defined and the performance envelope explored, a logical third capability would be one that supports continued evolution of complex software to continue to adapt it to its fielded context. While most other software engineering research focuses on developing new software-intensive systems, in fact the larger challenge is to learn how to maintain and upgrade the huge amount of Air Force software that is already fielded. Thus, important research areas include methods to infer the architecture of legacy software systems, to identify software components within that architecture, to parallelize legacy system software and applications, and to migrate that architecture and components to new and improved architectures, possibly within a new computing environment. The products of the research outlined in the previous section, for predicting software performance envelopes, could also be applied during such a migration to end users to tailor the software in a new operational environment.

⁹Mark Lewis, Air Force chief scientist, described the opportunities and benefits of AFOSR continuing, and thus completing, research projects initiated by DARPA in his October 2005 address to the Air Force Scientific Advisory Board.

¹⁰See <http://www.sei.cmu.edu/pacc>.

¹¹K. Chen, J. Sztipanovits, A. Ledeczi, and S. Neema, "Toward a semantic anchoring infrastructure for domain-specific modeling language," Proceedings of EMSOFT'05, ACM Special Interest Group on Embedded Software, September 2005.

¹²For example, see the NSF-sponsored study at <http://www.cs.virginia.edu/~sullivan/sodsis.html> or a related NSF program at http://www.cise.nsf.gov/funding/pgm_display.cfm?pub_id=13078.

Since network-enabled systems will involve many legacy systems as well as new systems, it is imperative that software be designed in a way that it can be evolved in an affordable manner throughout its life cycle. The committee recommends that AFOSR support research to improve the evolvability of software-intensive systems. The following specific lines of research, which could build on readily available commercial frameworks, are recommended:

- Improvements to our ability to conduct dynamic, model-based analyses to analyze nonfunctional attributes.
- In order to improve component integration, research to accelerate the development of abstract design-component systems and code-component-based systems, addressing automated discovery, composition, generation, interoperability, and reuse across hundreds of systems.
- Research in security in support of the goal of measurable, available, secure, trustworthy, and sustainable network-enabled systems.
- To attain assured reliability with hard time-deadlines, methods for modeling and analyzing integrated reliability, availability, and schedulability of components and systems in realistic conditions derived from user-specified scenarios.
- Energy-efficient components of the overall system: (1) network energy on network interface and communication protocols of ad hoc networks, (2) processor energy and process management for scheduling various applications, (3) memory/storage energy and memory/storage management, and (4) display energy.
- Research into novel integration of methods for verification and validation, such as integration of informal methods (e.g., software testing and monitoring) with formal verification (i.e., model checking and theorem proving) and abstract interpretation and static program analysis techniques. The ability to validate scalability, adoptability, usability, and measurement is also important, and some fundamental breakthroughs in the past 5 years have led to a rapid rise in industry adoption and interest.

Basic Research for Air Force Information Management and Integration

BACKGROUND

The primary reason for developing a network-centric Air Force is to gain the ability to rapidly collect and distribute information and thereby create information superiority and, from it, military superiority. This is a very challenging objective. While the Air Force has greatly increased its ability to collect information, the capability does not yet exist for quickly managing that information so as to create knowledge where it is needed while avoiding information overload.

Technology to manage and integrate information is critical to the Air Force for its current operations and its long-term challenges. The need for this technology arises in finding and tracking, command and control, decision support for situation assessment, counterintelligence, and public affairs. For example, finding and tracking require integrating data from large numbers of sensors and fusing them into actionable information. Command and control and situation assessment have similar requirements—for example, to tie together sensors and databases into a publish-subscribe environment.

More concretely, the Air Force will need an integrated battlespace management system that combines all of these capabilities in order to significantly improve its ability to conduct effective joint and coalition warfare. Such a system is dependent on several key IS&T advances: information exchange with complex event processing; transformation of multiple sources of information into a common representation and then into knowledge; and distributed collaboration via shared knowledge.

Building such a system is made more difficult by many real-world factors that are well beyond the state of the art. For example, integration often needs to be done quickly between systems that were not designed to interoperate, such as those of partners in an ad hoc coalition. Efficient techniques for evaluating complex subscriptions may enable warfighters to identify targets more quickly, while they are still vulnerable to attack. Sensors may be moving and unreliable, so that integration needs to be more dynamic and flexible. As improvements are made in power sources for sensors and networking technology and the cost and size of sensors decrease, we can expect the number of sensors and the data rate per sensor to increase. The combination of these effects will greatly increase the overall data rate from sensors, seriously taxing the scalability of current stream-processing algorithms. And, inaccuracies in human sources of information must be factored in.

The Air Force also needs advanced data integration technology in support of nonkinetic operations—that is, those that do not rely on damage to physical objects. For example, public relations challenges might require quickly searching and correlating information across many data sources, only some of which may be owned by the Air Force and many of which have not been integrated at the time the search request is issued. Influence operations might rely on very fast unearthing of information about particular individuals (e.g., a leader of a threatening group) or groups (e.g., a tribe, clan, or religious community). Other nonkinetic operations might depend on quick searches of blueprints, building permits, inventories, cybercrime logs, and so on.

Commercial database and distributed systems technology can be expected to meet many Air Force needs. However, for many technical areas, the Air Force places a higher value than the commercial sector on obtaining highly innovative solutions or on obtaining them well before their widespread commercial availability. Because the Air Force uses many custom-developed application systems, off-the-shelf solutions are often insufficient, particularly for aspects of databases and distributed systems that enable interfacing with other systems: schema reconciliation, metadata, and so on. Rather, fundamental models of information integration are needed, models that can be tailored to unique Air Force requirements.

What is the current state of the art? Robust database system products exist to support the storage and management of structured information, and recent research is leading to versions of these products that can manage XML documents. Extract, transform, and load (ETL) tools are available to help clean and integrate data—for example, to store it in a data warehouse. Enterprise information integration (EII) tools are becoming available that enable querying heterogeneous databases, typically rela-

tional or XML, although this approach is much less mature than data warehousing.

Data integration projects, particularly for data warehouses, are now an important part of the IT work done in large enterprises. However, such projects are usually quite expensive. Simple ones typically require a year of elapsed time and several person-years of engineering effort, and complex projects require much more effort.

There is also a large market for application servers and related middleware products that support the integration of distributed applications. Application servers enable the integration of applications through remote procedure call (RPC). The newest incarnation of this technology is service-oriented architecture (SOA), where the RPC is Internet-enabled using the XML-based Simple Object Access Protocol (SOAP). Enterprise application integration (EAI) tools have a similar goal, but they use asynchronous message-passing and generally run behind the firewall. As in the database area, the deployment of such products is usually an expensive multiyear affair.

There is some redundancy of application development tasks appropriate to ETL, EII, EAI, and/or application servers. For example, data may need to be cleaned and reshaped in the same way for each of them. This redundancy increases the cost of integration and potential inconsistencies between functionally related solutions. However, there are technical impediments to eliminating this redundancy, such as incompatible execution models (e.g., streaming vs. RPC vs. batch processing). This area is ripe for research.

MAJOR INFORMATION MANAGEMENT CHALLENGES FOR AIR FORCE IS&T

A central issue in the ubiquitous deployment of IT is the efficient and effective use of the large variety and large quantities of data that can be acquired, transmitted, and stored. In the specific area of national defense, both human and automated decision makers must rapidly sort through voluminous multimedia data and fuse disparate sources to create knowledge. The first step is to integrate raw data that arrive in different formats and are managed by different data management technologies so that they can be correlated and searched using a common mechanism. There has been much progress on integrating classical formatted files with semi-structured data, such as XML, which is becoming a standard feature of commercial database systems. The integration of other information management functions into a common database infrastructure—such as information retrieval, natural language processing, geographical information systems, image and video stream processing, and data streams—is not as

far along. It is an open question how best to integrate these functions into a common database system architecture.

Similarly, the Air Force faces major open questions on how to manage and share information in a distributed system. The “publish-subscribe” paradigm is an increasingly popular one for enabling distributed information-system integration and which is being explored by the Air Force to work with DOD’s Global Information Grid. That concept includes a common area where information is “published” and “subscription” information for various users that defines which posted information their systems will download from the common area. While the publish-subscribe concept has been shown to scale to thousands of participants within stable local area network environments,¹ there are many Air Force requirements that will be met only with further research:

- An Air Force publish-subscribe system must work in an unstable wide-area network environment such as a battlespace network.
- An Air Force publish-subscribe system must scale to a very large number of users and data sources, while coping with a bursty load that generates network and server contention.
- A publish-subscribe system’s scalability depends in part on the amount of state information it needs to maintain about old events. Thus, in order to scale to the sizes contemplated by the Air Force, systems based on this paradigm must weed out published information that is outdated or redundant, even if it has not been replaced by a clear duplicate—except that in some cases, old information is still useful for some subscribers.
- The system must distinguish between subscription rules that do not fire because the rule’s conditions do not hold and those that do not fire because their data sources are unavailable. The former situation indicates that no action is required but the latter is ambiguous.
- A system that handles more expressive subscription rules can be more selective in identifying events of interest to battlefield commanders and can reduce the cognitive load on commanders by tracking the state of partially satisfied subscriptions. However, the expressiveness of subscription rules is limited by the speed with which they can be processed and the amount of memory needed to maintain the state of partially satisfied subscriptions—remember-

¹A. Carzaniga, D.S. Rosenblum, and A.L. Wolf, “Design and evaluation of a wide-area event notification service,” *ACM Transactions on Computer Systems* 19(3):332-383 (2001); Yi-Min Wang, Lili Qiu, Chad Verbowski, Dimitris Achlioptas, Gautam Das, and Paul Larson, “Summary-based routing for content-based event distribution networks,” *Computer Communication Review*, October 2.

ing, for example, that two of the three requested events have occurred.

- In effect, rules are continuous queries being posed on streams of data arriving from sensors and other data sources. Database-style query optimization will be needed to cope with high data rates and the large number of rules that need to be concurrently checked. It will also be needed to process queries over a combination of historical data and newly arriving data.
- The system must have some “intelligence” to take context into account when interpreting subscription rules. For example, a user’s information needs can vary tremendously depending on external conditions—for example, whether a battle is ongoing, whether other elements of the chain of command are online, and whether other information sources are operational.
- The importance of timeliness can vary greatly according to the information posted or the external conditions. Rules must be evaluated on input data streams at a rate that suits the timeliness requirements.
- The relative importance of certain rules can vary according to external conditions. When the system is overloaded due to a burst of activity in one area of the battlespace, performance should degrade gracefully—that is, the system must not thrash. The most important rules should be prioritized to access the limited system resources.
- Security of the published information is critical, and the system must be trustworthy even if an adversary has gained access to publish or subscribe. In an Air Force setting, it must also offer multi-level security.

The above requirements are examples and not comprehensive. Even so, it is a long list of challenges that spans many areas of computer science in addition to publish-subscribe concepts, such as database systems, security, and data mining. Moreover, in addition to publish-subscribe, other distributed computing paradigms might be relevant, such as peer-to-peer messaging, media delivery, grid services, distributed services for wireless networks, and overlay networks. Different paradigms offer different options for power management, net capacity, security, and other factors.

Information needs to be understood at a more abstract level. The Air Force needs a model and architecture for situation understanding and a means of incorporating situation modeling, model-based processing, situation projection, and top-down management of situation understanding in order to explore topics in information fusion. It also needs a scientific basis and technologies for multisensor fusion for air and ground targets. Some of these topics are extensions of ongoing work in intelligence, sur-

veillance, and reconnaissance methods. These sorts of analyses need to account for the uncertainty of information sources, sometimes due to known error characteristics of instrumentation but also due to hard-to-quantify uncertainties such as those inherent to human-supplied intelligence.

An even bolder question would be, How can a computer understand data and information in context? In principle, background understanding of a mission or related intelligence could help a computer interpret information from the battlespace. For example, it might help identify objects in video or image data. If such context-dependent processing were possible, perhaps information-understanding algorithms could be embedded in sensors and networks to enable rapid data assessment and hence rapid situation assessment. Source compression methods might also be possible when multiple sensors collect correlated observations, whether or not the collection was done with coordination. If the correlation is known, as is often the case (e.g., when data are bursty), redundancy can be compressed out. This is critical in order to fit the large volumes of sensed data into modest network capacities, especially in a difficult communication environment. One challenge is the compression of highly correlated data without fine coordination among sensing platforms.

Defense applications today involve a wide variety of information-bearing signals. Examples include audio and video communication signals, video and still light imagery, hyperspectral imagery, and radar imagery, to name but a few. Often, these different data acquisition modes will be employed simultaneously in the same system. To use such signals effectively, it is necessary to represent them in digital form in a hierarchy ranging from the basic uncompressed analog-to-digital form to a compressed digital form to a symbolic form. Continuing research will be required to find multilevel representations of multimodal signal sets so that they can be efficiently stored, transmitted, manipulated, and understood by digital techniques. Although much progress has been made in standardizing signal and data representations (e.g., MPEG7 and XML), additional AFOSR-sponsored research is needed to focus on Air Force signal representation problems. Given appropriate digital representations of multimodal signal sets, it is also important to find new ways to search for desired information and to aggregate that information into a convenient form. This means that basic digital signal processing topics such as signal enhancement, signal modeling, redundancy removal, and feature extraction will continue to be important fundamental research topics. This also requires advances in multimedia data mining and machine learning in the defense context. Combining information across different signal modes and fusing that information into robust multimedia "documents" with high semantic value is another topic of broad interest beyond just the Air Force, but there are also Air Force-specific aspects to this challenge.

Throughout information management and integration, there is a critical need for security, but many of the research challenges are military-specific. A very difficult challenge is how to enable the multilevel security needed in a fully functional publish-subscribe system. However, some of the most important issues there are policy ones, not technical ones.

There is a great need for research in defensive and offensive information operations, some of which has already begun at the AFRL. There is also a great need for better understanding of how humans react to information operations—that is, what actions will create the desired effects—and ways to evaluate options for influence operations, trade-offs, and courses of action. Damage assessment for influence operations is a special challenge so far unexplored.

RECOMMENDED BASIC RESEARCH IN INFORMATION MANAGEMENT AND INTEGRATION

Following are areas of information management and integration functionality that are important for future Air Force systems and are either beyond the state of the art and require basic research or are areas where progress has been slow and could be sped up by additional research initiatives:

- Large-scale sensor networks need to support queries, because it is often infeasible to send all sensor data to a central location to process queries. Research is needed to understand where to place query-processing functionality when nodes have limited power, limited computing capability, and limited network bandwidth and connectivity. Conversely, as network and power limitations are overcome, higher data rates can be expected, which will lead to optimization trade-offs different from those being investigated under today's limitations.
- Applications needed in the field will demand coping with dynamic, late-binding, and unpredictable structure in the data, or even no structure at all. Accordingly, dealing with semistructured content with unpredictable structure—for the data model, for querying and routing documents, and for efficient execution—is critical.
- Data arriving from different sensors and other sources may be inconsistent. Yet they need to be rolled up into a coherent view of the environment, thereby being turned into knowledge. Technologies are needed for fusing uncertain or inconsistent data and for querying databases containing incomplete or inconsistent information. Some of these technologies may be generic across arbitrary data types, while others may be specialized—for example, for

merging location information from mobile sensors. It is also important to be able to mine inconsistent data for patterns, to obtain hints about a coherent view without necessarily being able to construct such a view. It may be helpful to create and maintain an ontology of the information domain, to help direct the data mining or interpret the results.

- Database mechanisms are needed to track the lineage of all stored data, from its initial source through transformations and updates that were applied to it. Solutions exist for tracing schema information, for example, from the column of an output report to the input data sources and transformations of those inputs that generate it. However, much less is known about how to track lineage at the level of data instances, especially when arbitrary transformations, not just database queries, can be applied.
- Database system mechanisms are needed to deduce the degree of certainty of query results as a function of the uncertainty of the raw data, under appropriate probabilistic models. Uncertainty may arise, for example, from measurement errors, incomplete structure, imprecise knowledge of the location where a measurement was made, data cleaning of raw data, or conflicting data from different sources. Some of the challenges are to identify mathematical models of uncertainty that have general utility across a variety of applications, to develop efficient deduction algorithms for these models, and to find optimal ways of making these algorithms effective in the context of information management systems (e.g., in a database system, in middleware, or in application programs).
- New approaches are needed to reduce elapsed time and level of effort to integrate independently designed databases, such as the databases supporting each participant in a joint mission. Such approaches must be embodied in tools that can help an engineer design, develop, and test mappings between databases, find and reuse existing mappings in the same domain as the mapping to be generated, and evolve and customize mappings as data sources change and new ones are added. Here, too, the creation and maintenance of an ontology may be useful to organize a repository of schemas and mappings.
- The Air Force needs advanced distributed systems technology in support of battlespace information management. As discussed above, in the section on major information management challenges, publish-subscribe is one important distributed systems infrastructure, though other distributed computing paradigms should also be explored. Although there is much research on publish-subscribe systems for centralized infrastructure, many improvements in the

functionality of publish-subscribe systems are needed, which requires additional research. These include scalability to wide-area, unstable network environments; filtering outdated or redundant information; coping with intermittent availability of event streams; supporting more expressive subscription rules; improved optimization of rule execution; processing strategies that account for context, timeliness, and priority; and incorporating security guarantees. Architectural principles that enable us to create secure, scalable, reliable publish-subscribe platforms would also be valuable.

- Air Force distributed systems need to be interoperable, secure, massive in scale, highly available, dynamically reconfigurable, and easily managed. These system-level characteristics are notoriously more difficult to engineer than component-level capabilities. Moreover, some military requirements for these system-level characteristics are beyond today's state of the art.
- The data that flow through such distributed systems are subject to the same integration problems described above for information management. For example, identifying errors, cleaning the data, fusing data from multiple sources into a coherent view, and tracking data lineage are all relevant to data flowing through a distributed system. Some data must be moved from a real-time publish-subscribe system into a database to enable historical queries of (possibly recent) publications, to obtain answers to point questions, and to identify trends.
- System management is important, especially in an environment where information publishers, subscribers, and even major servers are unreliable. System managers must know who is connected, must be notified when users and information sources connect and disconnect, must be warned of impending performance and reliability problems, must be helped in understanding the impact of these problems on system behavior, and must be able to reconfigure the system dynamically in response to these and other changes. Again, techniques are needed to ensure these capabilities can be made available despite incompatible network architectures and security domains.
- Performance and scalability, which are always a major concern, are especially hard to manage in heterogeneous distributed systems that are frequently reconfigured as components and communications fail and recover. For example, quality of service for real-time streaming of audio and visual information needs to be maintained, while avoiding overload that interferes with concurrent high-priority interactive workloads.

- Given the dynamic nature of battlespace systems, testing for a broad range of contingencies is extremely challenging. Better models of large-scale systems are needed to allow virtual testing and analysis.
- Continuing research will be required to find multilevel representations of multimodal signal sets so that they can be efficiently stored, transmitted, manipulated, and understood by digital techniques. This means that basic digital signal processing topics such as signal enhancement, signal modeling, redundancy removal, and feature extraction will continue to be important fundamental research topics. However, the special nature of defense applications argues for continued support by AFOSR of work that leverages and extends signal-processing R&D of a more general nature.
- There is also a need for research on multimedia data mining and machine learning in the defense context. Research on basic signal representation and research on signal analysis must be closely intertwined. Combining information across different signal modes and fusing that information into robust multimedia “documents” with high semantic value and focused on information sources of interest in Air Force applications will allow AFOSR to leverage non-defense-based research and foster the creation of new knowledge with direct applicability to Air Force problems.

Basic Research for Human Interactions with Air Force IS&T Systems

The concept of human-system interactions (HSI) includes the traditional field of human-computer interactions (HCI) but also encompasses the coordinated and purposeful interactions of several or many humans with complex systems and the interactions of teams of humans mediated through systems. The committee recommends that AFOSR focus on HSI because it is essential to the successful operation of complex IT-based systems and to the accomplishment of network-enabled operations. There is strong historical precedence for this area of research. One example of clear importance to the Air Force is the use of cockpit simulators for training and rehearsing. However, the increasing pervasiveness of interactive complexity in IS&T systems argues for a broader interaction, much beyond the computer interfaces and human-machine synergy. Examples are interactions in real-time virtual and simulated environments and in the collaborative control of UAVs. Clearly, such capabilities are both team-focused and network-enabled. An ultimate goal of HSI research would be to enable machines (or algorithms) to perform most if not all of the complex data manipulation, correlation, computation, and automatic data reduction—and even some decision making—leaving humans to perform the most critical judgments that cannot be accomplished by algorithms. Furthermore, HSI should help humans to interact with one another in cooperative tasks where multiple humans are part of the system.

CHALLENGES POSED BY HUMAN INTERACTIONS WITH AIR FORCE IS&T SYSTEMS

Scope of the Challenge

In the Air Force, there are many instances where one or more humans interact with one or more IS&T systems. These include systems that are distributed not only among different platforms but also perhaps across geographical and organizational boundaries, most often with strict security and service reliability constraints such as near-real-time or time-critical services. Added complexity comes about because both the humans and the systems might be interacting with one another in additional ways that are disconnected from the task being analyzed. New paradigms such as publish-subscribe architectures are being investigated to cope with some of the challenges. What sorts of information, architecture, and format should be used to achieve desired effects, and how can designers and users estimate the uncertainties and internalize context and caveats associated with each option? Assuming the right information is available at the right time and in the right form (e.g., as text or images), what techniques will enable the user to make the best use of it? How can what-if simulations be considered and evaluated? Such complex capabilities might require integrated and synchronized multimodal interfaces (visual, aural, and/or haptic) to capture the high dimensionality of a system of sensors and actuators in the battlefield.

Research into HSI should shed light on the usability of the (same) information in a battlefield command-and-control situation relative to the different perspectives (ranks) of the users and the different degrees of granularity (detail). In other words, one must understand and characterize the most likely and useful level of complexity for each potential user, from the warfighter to the commander, so that the complexity of the information can be tailored to provide an optimal amount of information for battlefield decision-making—not a paucity of data, but not data overload either. Visualization (interactive visualization in particular) is an important topic. Because this subject is receiving much attention as part of DARPA's Command Post of the Future program, AFOSR should work to complement the DARPA investment.

The issue of how to ensure security pervades the topic of HSI—in particular because human behavior regarding security has a great influence on the effectiveness of embedded security protocols. User behavior can reinforce or undermine security systems, and we do not know enough about how the interface influences those behavioral choices. The most frequent cause of system failure is human error, and even under the best of circumstances (good training, no stress, well-designed interface system,

simple tasks), human operators have been shown to have a finite (non-zero) rate of placing a system outside its operational allowed boundaries.¹ Usability and error-free HSI design may work at cross purposes in certain cases. For example, the Therac-25 medical radiation device allowed operators to bypass error messages and apply treatment even if lethal dosages were being used. In interface design, unusual system circumstances must be prominently flagged. It is well known that a high level of stress increases the errors that may cause security or system failures. A modicum of system sensor display redundancy (from different sensors) also gives the user greater confidence in the validity of critical information. HSI research could lead to improved design and evaluation of IT-based systems to achieve better usability, safety, and security simultaneously.

Much of the research on influence operations goes beyond IS&T and requires fundamental research into behaviors and how they can be affected. To this base must be added research into information interpretation and presentation, personnel training, and modeling and simulation for influence operations, building on what is known about cultural and behavioral factors. As an example, characterization and recognition of normal and abnormal behavior, in general, would help in surveillance at all levels. Characterizing those behaviors requires ongoing research in social sciences, and automatic recognition of the identified characteristics is an ongoing challenge for IS&T.

Power consumption—concerns about which were noted in Chapter 3—may also be improved through HSI research into how users interact with mobile devices with small displays whose display component consumes a significant proportion of the device's energy. While hardware advances allow for the dimming of displays and the use of screen savers to lower power consumption, the possibility of selective dimming of display areas necessitates a better understanding of which display regions to dim and when and of which display effects may be acceptable to the users or allowed by different operations (e-mail, for example, will have different criteria on dimming than video streaming). Research efforts could be shared by industries that are also seeking power-efficient displays.

HSI challenges also arise when using commercial off-the-shelf hardware and software, which when merged into complex systems may present the user with a confusing, even contradictory, set of interfaces. HSI challenges certainly present in specialized equipment interfaces such as those used by the remote operators of UAVs. Such applications drive the development of context-aware, cross-modal analysis techniques that

¹Barry Kirwan, *A Guide to Practical Human Reliability Assessment*, Taylor and Francis, Ltd.: London (1994).

provide input to automated systems. In other cases, the analysis techniques can assist human decision makers in sorting through huge amounts of information. In such cases, data-mining and machine-learning techniques can simplify complex multidimensional data. How to represent such information in multimodal form to facilitate its use is an important area of research. New approaches for the display and rendering of multimodal signal sets, such that important context-dependent features are highlighted and subtle features are not obscured, will be necessary to provide context-aware, user-oriented interfaces that allow effective decisions to be made in complex situations. Such system interfaces must allow users to probe deeper layers to capitalize on uniquely human analysis capabilities.

Finally, two issues related to the sharing of information should be kept in mind when planning the HSI research program:

- DOD and the services have a range of geographically distributed R&D efforts in HSI and modeling. Because the HSI challenge is so huge, it is important that these resources and accomplishments be shared, striving for commonalities of icons, network services, standards, interoperability, response time, scalability, and so on.
- Some key concerns with respect to information sharing in battle-field situations include scalability, allocation and management of bandwidth, message and resource priority allocation policies, time order of the processing steps used in publish-subscribe systems, intelligence needed to establish logical and semantic relations on the information objects, potential deadlocks, etc. These concerns are closely related to the shift to network-enhanced operations, where multiple sites and organizations need to share common information on a timely basis.

Interfaces for Air Force Decision Makers

There is a general need within the Air Force for decision-support tools providing more effective interfaces between decision makers and the voluminous information available to them. Over the years, various software packages have been developed to support decision-making. One of the simplest and most useful is the spreadsheet, but others such as critical path methods and the Program Evaluation and Review Technique (PERT) have also found a place. However, most real-world problems involve uncertainty. Whether we are dealing with scientific issues, engineering problems, or capabilities of individuals, we make decisions based on imprecise assumptions and incomplete and time-varying data. Decision-making under uncertainty is usually studied as an offshoot of probability

theory, merged with ideas of risk. In principle, decision theory offers a set of structured procedures that assist decision makers.

Classical decision theory is most widely studied in a static setting in which the data are given and then analyzed and acted upon. In the real world—and especially in Air Force information operations and influence operations—the inputs to the decision often change before the solution can be implemented. Thus, one needs to pay attention to techniques that lend themselves to rapid, perhaps approximate, solutions and techniques for rapidly updating given solutions when small parts of the problem change. Stochastic programming methods could be useful in some probabilistic settings, although problems involving military logistics are often too large to be treated successfully using these methods. There is also the need to develop software that incorporates Bayesian inference, although a completely resolved Bayesian inference model might not be realistic given the computational resources likely to be required. As a consequence, statistically sound heuristics could be developed and explored to extend computational feasibility to much larger problems. As a model one might think of greatly extending the functionality of spreadsheet software by combining it with packaged statistical algorithms and an appropriate user interface designed to help users formulate questions and display the probabilistic information associated with the results. Part of any effective system of this type would be effective optimization algorithms, because every nontrivial what-if question requires optimization. (See Chapter 3 for some discussion of optimization, dynamic programming, and game theory in complex systems.) It would be desirable for this type of research to develop new classes of statistical models for complex phenomena that (1) are rich enough to capture relevant physical laws and other known constraints in large classes of problems of interest to the Air Force and (2) have exploitable structure that allows the development of new methods for information extraction and for the development of bounds on achievable performance for problems that previously defied practical solution.

Another example where better decision support is needed is in control systems that involve humans as essential parts of the decision chain. Technology—especially IT, collaborative work analysis, and communication technology—can provide such decision makers with better local and global awareness, e.g., about the weather and other environmental conditions, including other aircraft and machines in the environment. The research community has been quite successful in modeling and predicting the behavior of machines interacting with relatively well-understood or well-measured environments, including those with uncertainty and some bounded nonlinearity. But people in the loop make decisions that are highly nonlinear and not always predictable, and therefore their

inclusion creates a greater modeling challenge. In addition, the Air Combat Command told the committee that it was increasingly concerned with winning the local population's hearts and minds as the most effective means of achieving military goals, and in such a context, improved understanding of people can be more important than improved models of equipment. To address all of these people-dependent modeling challenges, there is a great opportunity to devise research experiments that instrument people and their environment with cameras, microphones, stress and pressure sensors, and other measurement devices, perhaps collecting data via wireless technology. Such measurements would offer objective observations about people's behavior under varied conditions, which could help in its modeling.

One specific need is HSI systems that make the control of UAVs more efficient. Currently, Global Hawks and Predators require as many as 25 people for each operation: at least three ground-based pilots, one for each 8-hour shift, and various support personnel to launch and capture UAVs, process collected data, etc. While some attention has been paid to the design of ground control units, significant improvements could probably be made using state-of-the-art knowledge of HSI design and collaboration. However, significant advances in intelligent control, HSI methods, and other areas of interactive-collaboration research will be needed to move toward the goal of having one person control a single UAV, with even more research being required before one person could control multiple UAVs.

Machine Learning to Support HSI

System components that adapt and learn from their environment or from coaching by a human supervisor would open an amazing number of doors, from intelligent control systems to intelligent decision aids. Machine learning is a critical technology for eventually enabling software-intensive systems to (1) adapt and extend their functionality in response to specific training and (2) make intelligent inferences either from general human guidance or through reasoning based on interaction with the system's domain.

One particularly important application area for intelligent systems would be immersive environments for training and exercise support. This will require advances in modeling and simulation, eventually yielding virtual worlds that can be used to support training, exercises, and critical decisions about system development and acquisition. All too often modeling and simulation are underused because technical difficulties prevent reuse and combination.

Simulation as a Design and Training Tool for HSI

Another challenge for human interactions with complex IS&T systems is the need for realistic, validated simulations that allow humans to design systems and learn about their management, operations, and vulnerabilities. Design options and capabilities of systems as complex as a network-enabled battlefield force or a publish-subscribe system with 100,000 users can only be explored via simulation. The software engineering and validation and verification aspects of such simulations are known to be challenging. Less well known is the need for the proper design and analysis of experiments run via simulation. Without attention to this need, complex simulations often lead to inaccurate or imprecise results. Moreover, the widespread use of simulation for Air Force training presents more challenges owing to the presence of a human in the loop. The simulation modeling of human behavior is not well developed or understood, though efforts along these lines (e.g., models of crowd behavior) are ongoing. In a training mission, the vagaries of human behavior make careful attention to statistical design and analysis even more critical.

Another area where the potential of simulations is underutilized is input/output sensitivity analysis. The simulation model is not only for obtaining predictions but also for investigating which inputs are important and which are not (often of equal interest). The same statistical design/analysis methods for “production” runs can be used for sensitivity analysis to elucidate things like main effects, interactions, and so on. This leads rather naturally to validation and verification, because validation relies on statistical comparisons between the outputs from the simulated system and those from the real system. A rigorous simulation procedure offers the potential for objective and rational validation. And a good simulation system can also be a good training tool for users.

BASIC RESEARCH RECOMMENDATIONS FOR HSI FOR AIR FORCE IS&T SYSTEMS

Some of the most valuable basic research efforts that would address the challenges in this chapter are listed below. The committee recommends that these topics be addressed in an interdisciplinary manner with goals that could enhance multiple capabilities (e.g., network security, simulation, information operations). Some of this research would seem to be appropriately conducted collaboratively with other branches of the military or with the National Science Foundation. The research should strongly articulate basic design principles, goals, and evaluation of human interfaces with complex IS&T systems. The committee recommends a focus on the following areas:

- Behavioral models for individuals, groups, and organizations.
- Specification of an information operations testbed and experimental metrics for evaluating its effectiveness in distributed environments that preserve realistic situation awareness.
- Better understanding of mixed-modality methods (systems that understand speech, gestures, sound, touch, etc.) for imparting information and effecting control. Research should aim to illuminate the minimum speech capability (or iconography, haptic stimuli, etc.) needed under varied conditions.
- Design of communication methods that combine text, speech, and visual modalities appropriate to the information being communicated.
- Understanding user needs that have implications for power usage (and energy efficiency) in displays, such as how users scan areas of a display.
- Exploration of the theoretical tools used to design, implement, and field complex, time-critical information systems. Topics include the development of large-scale deterministic, stochastic, convex and nonconvex, dynamic, and integer-constrained mathematical programming; algorithms for large-scale Bayesian analysis, including approximate techniques, partial evaluations, and update algorithms; game theory; multiscale modeling and multiscale analysis; and formal approaches to data mining, knowledge representation, and reasoning.
- Development of better tools for automated reasoning and inference under uncertainty. Such tools can draw on diverse techniques, such as data mining and knowledge discovery in databases; models that learn from data; information retrieval; automated diagnosis and decision support; and machine learning. The goal is to enable users to quickly navigate complex data sets, to set up and analyze decision trees or Bayesian networks, to understand the effects of decision rules, and so on.
- The paradox of automation is that the more automation is added to aid humans, the more work the humans have to do to understand both the automation and the task(s) which the automation is designed to support. To counter this, specific research is needed to enhance models of users' capabilities to interact with and understand systems via different modalities; to automate acquisition and enhancement of information fusion models; to create models of users that can be trained from instruction or in exercises with humans; and experimentation in domains of Air Force relevance.
- Interactive system components that can adapt and learn from their environment or from human supervision should be investigated.

Appropriate methods include indexing, retrieving, and reusing past problem-solving episodes; encoding problem-solving methods from instruction or observation; explanation-based methods; inferring or suggesting improved problem-solving methods; enhancing the efficiency of algorithms based on past problem solving; inductively inferring regularities using, for example, graphical models (Bayesian networks); reinforcing learning methods; understanding emergent behavior; transferring lessons learned in one domain to another domain; using multiagent teams to learn how groups collaborate to solve problems; integrating multiple machine learning techniques; developing toolkits and frameworks that integrate and facilitate adoption of learning methods and provide guidance on the best methods to use in different situations; and experimenting in domains of Air Force relevance.

- Research on the modeling and simulation of scenarios of importance to the Air Force, including more rapid and cost-effective construction of domain-faithful models of information and influence operations; support for the authoring of scenarios and their translation into the models and databases required for simulations; and software tools to allow users to tailor and understand the behavior and functionality of software components.

Priorities in Basic IS&T Research for the Air Force

It is the committee's consensus that the basic research topics described in Chapters 3 through 6 are of great importance for enabling the capabilities sought by the Air Force and that none of them are likely to be adequately explored by the private sector. Recognizing, too, that AFOSR's resources for IS&T are relatively limited, the committee recommends in this chapter how to prioritize among those research topics given different funding scenarios. As background for that discussion, it first presents its overall vision for Air Force basic research.

A MODEL FOR AIR FORCE IS&T BASIC RESEARCH

Given the limited availability of 6.1 funds for IS&T research, the committee believes it is best to give priority to those research areas that AFOSR gives a realistic chance of enabling significant new Air Force capabilities. This is not to say that all of AFOSR's basic research in IS&T should be tied to particular technology goals—in fact, such tie-ins should be an unusual exception—but that IS&T basic research should normally be motivated by the recognition of a gap in the knowledge base that underpins desired Air Force technologies. The committee suggests that these knowledge gaps be seen as “grand challenges,” which is a helpful way to motivate research and build research communities and a useful mechanism for communicating the relationship between basic research and future Air Force technologies. Topics designated as grand challenges would normally be ones that are properly addressed by a cross-disciplinary community of basic researchers, and such a designation gives that community an identity and

strengthens its coherence. Grand challenges should be defined in terms that are recognizable to the basic research community, but AFOSR should also be able to map the grand challenges to future Air Force technologies, strengthening the linkage between basic research and the rest of Air Force R&D. The grand challenges are not part of, nor do they compete with, the AFRL focused long-term challenges (which are more oriented toward technologies), but they should link to them.

Some examples of grand challenges in a general sense are included in a recent paper by well-known computer scientist Butler Lampson:¹

- Development of a computer that hears, speaks, and sees as well as a person;
- A system architecture that scales up by 10⁶; and
- An information system that can be used by millions yet only requires a part-time support staff.

Such grand challenges quite naturally facilitate new interdisciplinary research communities: “interdisciplinary” because the breadth of the challenges calls for varied expertise and “naturally” because the associated researchers are interested in the whole range of efforts addressing the grand challenges. The Computing Research Association has for several years held conferences that similarly seek to identify grand challenges for computer science more generally. Information about these meetings may be found at <http://www.cra.org/grand.challenges/>.

Figure 7-1 suggests how the grand challenges model would contribute to cohesion across the total Air Force R&D enterprise. The grand challenges would be defined by AFOSR as a clear response to AFRL and Air Force technology needs, but in terms of concepts that motivate the basic research community. They have the advantage of not biasing the research directions chosen to address a particular Air Force need: They define a challenge for the research community that relates to desired capabilities but do not present the research community with specific technology goals, as would be the case if the basic research program were motivated directly by current Air Force technology programs. In turn, the 6.1 research builds up a base of knowledge related to the grand challenges (the arrow on the lower right of Figure 7-1), and the rest of the Air Force R&D community can track progress on the grand challenges that support its technology goals. That transition (indicated by the arrow at the lower left of Figure 7-1) is aided by the computational laboratories recommended in Chapter 9.

¹B. Lampson, “Computing meets the physical world,” *The Bridge* 33(1) (2003).

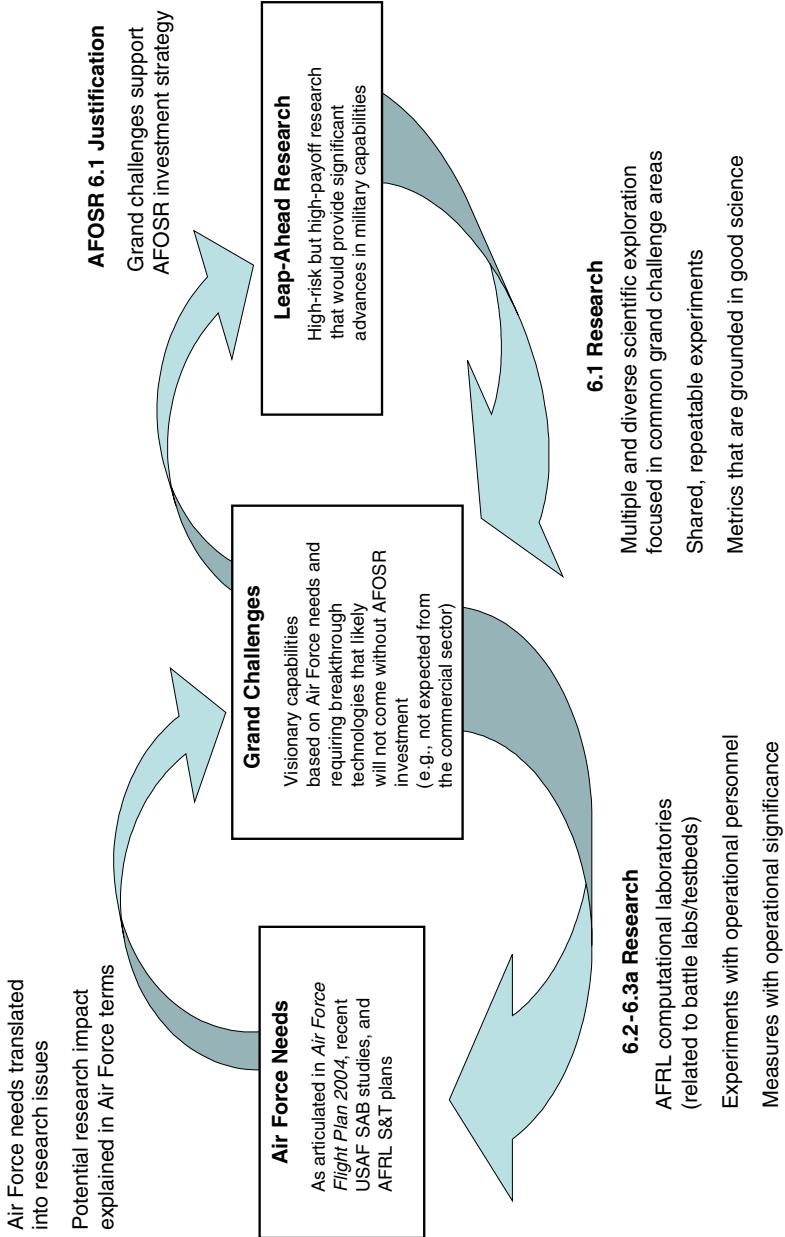


FIGURE 7-1 Mapping of Air Force needs to basic research through grand challenges.

The characteristics that are important in defining grand challenges for the AFOSR are these:

- They should be representative of important systems that the Air Force would like to have but that are well beyond the state of the art, like the examples in the next paragraph.
- The key component technologies that would make them realizable should be technologies where the Air Force wants research advances.
- They should be readily comprehensible to researchers who do not have domain-specific knowledge of the Air Force system.
- They should bound the research somewhat, so that researchers can aim toward specific applications rather than general capabilities.

The committee recommends that AFOSR consider the following as possible grand challenges, but the list is by no means exhaustive:

- *Control of multiple UAVs.* Research to enable the control of multiple UAVs by one human in mixed manned and unmanned airspace, in contrast to today's requirement for many humans for a single UAV in carefully deconflicted manned and unmanned airspace.
- *Taskable airborne network.* Research to enable cost-effective and rapidly deployable tactical intelligence networks in urban environments, where the nodes generally are sensors carried on UAVs or lighter-than-air vehicles and the networks are taskable by ground- and air-based commanders.
- *Mixed-reality training environments.* Research to enable training for air crews, command post staff, and commanders in an environment of such fidelity that it would be indistinguishable from the real world (and in fact would sometimes involve the real world—hence “mixed” rather than “virtual” or “augmented”). The computer tools used in such training environments should be the same as those used in the real world.
- *An automated Air Operation Center staff assistant.* Research to enable software that can learn from being told, much as human staff members learn on the job.
- *Rapid system integration.* Research to enable the rapid integration of IT-based systems, such as those belonging to different members of ad hoc coalitions. This research would encompass HSI, networks and communications, security, software, and information management.

It is envisioned that an interdisciplinary community would engage in collaborative research focused on each grand challenge. Researchers would describe repeatable experiments and metrics with which to assess

the research results. Greater synergy with 6.2-6.3a research programs could result in shared research environments or computational laboratories. Funding for 6.1 research should continue to be used to find solutions but not (normally) to develop or demonstrate particular technologies.

RECOMMENDED BASIC RESEARCH PRIORITIES IN IS&T

The committee recommends that in determining program priorities, AFOSR should emphasize research that clarifies Air Force problems, impacts operational procedures, and advances the state of the art in areas of interest to the Air Force. Prior collaboration with Air Force laboratories should not be essential to obtaining program support. AFOSR should also continue to monitor the emergence of possibly disruptive technologies or Air Force missions and adapt the research program and strategy accordingly: The rise of network-enabled systems and the emergence of a global war on terrorism have been recent disruptive changes, and they will not be the last.

The reader is reminded that the committee did not examine the entire portfolio of the AFOSR Mathematics and Space Sciences Directorate, which manages most of AFOSR's IS&T research, and it did not mean to imply that other topics covered by that directorate are less important. Rather, the topics funded by the programs in physical mathematics and applied analysis, computational mathematics, electromagnetics, and space sciences were not within the scope of this study, and it is not the committee's intent that IS&T topics should be emphasized at the expense of those topics in other programs. Even greater care is necessary in interpreting the priorities given here that overlap the current AFOSR programs in dynamics and control, optimization and discrete mathematics, high-performance computing and advanced architectures, and signals communication and surveillance, which overlap with IS&T only in part. Thus, any high-priority research identified by the committee in Chapters 3 through 6 that would be funded by one of those programs was not weighed against other topics covered by the same program, only against other IS&T topics. For instance, challenges in the control of complex ad hoc networks, distributed information systems, and multiple UAVs were considered by the committee in setting priorities, but it did not weigh in the importance of other challenges in control, such as control theory for avionics.

Recommendation. The committee consensus on how to prioritize Air Force IS&T research in networks, communications, information management, software, and HSI under different funding scenarios is encapsu-

TABLE 7-1 Relative Priorities Under Four Funding Scenarios

IS&T Topic	10% reduction	Stable	10% increase	25% increase
Networks	H	H	H	H
Communications	H	H	H	H
Information management	M	M	H	H
Software	L	M	M	H
Human-system interactions	H	H	H	H

NOTE: "H" means the general topic is a high priority, and its funding should be protected or increased. "M" means the general topic is of medium priority for AFOSR support, given the contributions by other players, but not of medium importance to the Air Force. "L" means funding in that area should be sacrificed so that a critical level of effort can be supported in other areas. "L" does not mean that the topic is not of importance to the Air Force, only that if resources are tight, it is a reasonable candidate for cuts because other organizations are contributing to the area and/or the challenge is so great that a small AFOSR effort is unlikely to lead to significant progress. These priorities pertain to the five general research areas listed in the left-hand column as weighed only against one another, not against other programs funded by AFOSR's Mathematics and Space Sciences Directorate. The priorities are meant to show the committee's consensus on which of the areas to (de)emphasize with any change in funding for their support. They take into account not only the importance of the research but also the relative need for Air Force-specific research. They reflect the committee's general sense of what can be meaningfully accomplished within the funding scenarios posited. The committee did not develop a detailed estimate of the resources required for each of the research topics in the left-hand column.

lated in Table 7-1. See Chapters 3 through 6 for detailed recommendations about research directions in each of these five areas.

With the current funding available for IS&T (the column labeled "Stable"), the committee recommends that networks, communications, and HSI portfolios merit the highest priority, while the information management and software research portfolios would be better able to weather any forced reductions in effort. The committee is not saying that the latter two research areas are less important to the Air Force. Rather, it is the committee's judgment that if cutbacks are required, reductions in those programs would do the least harm in limiting future options. If the overall IS&T funding dropped by 10 percent, the committee would give software the lowest priority only because other organizations, and commercial enterprises, are performing some related research. If overall funding increased by 10 percent, the priority for information management research should be increased. Finally, if overall IS&T funding were to increase by 25 per cent, the committee recommends a balanced portfolio drawn from the recommendations of Chapters 3 through 6.

Because all five major research areas listed in Table 7-1 contribute synergistically to the future fielding of team-focused, network-enabled systems, progress toward that vision is dependent on a balanced research effort across all five areas. However, overall funding would have to increase significantly to support such an effort, a conviction that is contained in the following recommendation:

Recommendation. The committee recommends a significant increase in IS&T funding within AFOSR, centered around research to support team-focused, network-enabled systems of interest to the Air Force.

Funding Mechanisms

The AFOSR is responsible for oversight and management of the Air Force program in basic (6.1) research. AFOSR program managers orchestrate the research program with universities, industry, other government organizations, and the AFRL technical directorates. The goal is to create revolutionary breakthroughs and to facilitate the transfer of research results to support Air Force challenges and needs.

In FY 2004, AFOSR managed funding for approximately 1,350 grants, cooperative agreements, and contracts. The total funding of \$394 million was distributed to approximately 380 academic institutions and industrial firms. As shown in Figure 8-1, the AFOSR FY 2004 budget authority included \$205 million for 6.1 defense research and \$104 million for university research initiatives (URI), both distributed in the form of extramural grants to universities, contracts to industry, and intramural grants to Air Force in-house laboratories. AFOSR also executed approximately \$85 million for other programs, such as the Small Business Technology Transfer (STTR) program and programs at DARPA.

One task given to the committee was to recommend an appropriate balance of funding mechanisms for IS&T research. The committee was briefed by John Tangney of AFOSR on funding mechanisms used by the AFOSR Directorate of Mathematics and Space Sciences, and it had the benefit of several reports,¹ excerpts from AFOSR BAA 2005-1, notes from

¹Among them, National Research Council, Review of the U.S. Department of Defense Air, Space, and Supporting Information Systems Science and Technology Program, National Academy Press: Washington, D.C. (2001).

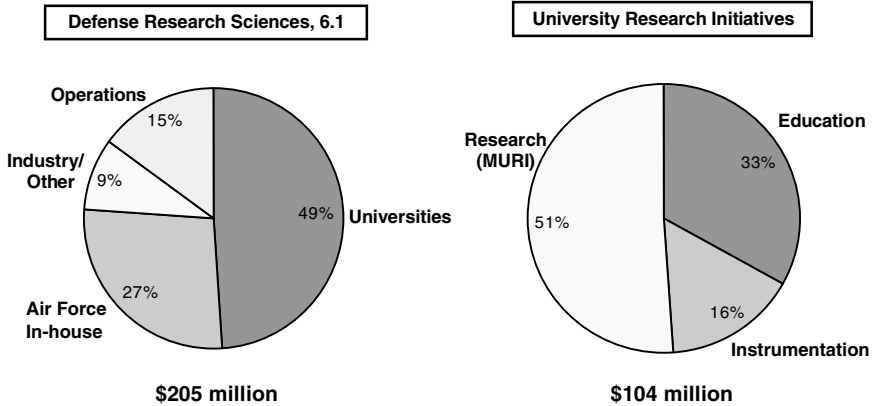


FIGURE 8-1 Funding for all of AFOSR in FY 2004.

a review of the Partnership for Research Excellence and Transition (PRET) for Advanced Concepts in Space Situational Awareness, which a member of the committee attended, and an overview briefing of AFOSR.

The current AFOSR IS&T program is executed through a variety of mechanisms:

- *Individual investigator grants at universities, at Air Force laboratories, and in industry.* A large majority of these grants are held by academic researchers. The committee agrees with that choice, because the academic setting is well suited for long-range, basic research and has a time-tested record of success. A small number of grants are held by investigators within AFRL, which is appropriate for topics that require in-depth understanding within the Air Force or that have not yet caught the attention of strong academic researchers.
- *Multidisciplinary university research initiatives (MURIs).* These longer-term group grants (up to 5 years) provide a good mechanism for investigating challenges requiring an unusual and/or concerted mix of efforts.
- *University centers.* These grants enable the Air Force to establish a critical mass of expertise focused on particular Air Force challenges at selected institutions. Those institutions in turn become a resource for in-house Air Force scientists and engineers.
- *Partnerships for Research Excellence and Transition (PRETs).* This mechanism brings together mixtures of academic, industrial, and Air Force researchers to jointly develop a new area of research (e.g., space situational awareness) while also establishing some path-

ways for technology transition. The industrial partners must contribute resources.

- *Small-business research, DOD Experimental Program to Stimulate Competitive Research (DEPSCoR), and grants designated for Historically Black Colleges and Universities or Minority Institutions (HBCU/MI).* These research mechanisms allow AFOSR to tap into a broader portion of the research community while, in the first case, also increasing the potential for technology transition.

Finding. The committee concluded that all of these mechanisms provide value of different sorts to the Air Force and that AFOSR program managers understand their pros and cons and use the mechanisms appropriately. Therefore, the committee concluded that the current choices appear to be good ones and that AFOSR staff should continue to be given the latitude to make choices.

Recommendation. The committee recommends that AFOSR consider augmenting these choices by implementing a young investigator award along the lines of the National Science Foundation's Faculty Early Career Development program or the Office of Naval Research's Young Faculty Investigator Program. Such a program in the AFOSR would help attract young science and engineering faculty of exceptional promise to work on Air Force research problems early in their careers. In particular, the Navy's program, which includes a basic award of \$75,000 per year for 3 years, coupled with possible equipment support and additional funding if the investigator receives support from another part of the Navy, is a very interesting model to consider since it includes incentives to connect to Navy laboratories and systems commands.

The committee also offers two related recommendations for consideration in planning and executing an expanded program in IS&T:

Recommendation. Using whichever funding mechanisms are appropriate, AFOSR should maintain a healthy balance between theory, experimentation, and transition efforts. This balance can be achieved only by providing freedom to program managers to resist pressures toward short-term goals, which in the long run will not serve the Air Force well because the Air Force long-term challenges are dominated by IS&T. This recommendation is almost a repeat of one from a 1996 review of AFOSR's mathematical and computer sciences program.²

²National Research Council, Review of AFOSR Programs in Mathematical and Computer Sciences, National Academy Press: Washington, D.C. (1996).

Recommendation. To accommodate the challenges of an expanded program in IS&T—which is called for because of the pervasiveness of IS&T in network-enabled Air Force operations—AFOSR should not only continue to build a strong, world-class base of PIs but should also develop highly interactive research communities of PIs with common goals and benchmarks, probably focused on the grand challenges suggested in Chapter 7. In some cases this will be possible through multidisciplinary activities (MURIs, PRETS, etc.). In others it will be essential to develop teams of researchers, perhaps on the scale of MURIs, to meet the challenges posed by the Air Force long-term goals.

Future Considerations

PROGRAM ORGANIZATION

Chapter 2 explained why the Air Force is growing increasingly reliant on IS&T. Consistent with this explanation is the committee's recommendation that AFOSR significantly increase the Air Force investment in basic research in IS&T, supporting such areas as decision-support systems for all levels of command, real-time control of air and ground systems with or without humans in the loop, human-system interfaces and interactions, and ubiquitous connectivity of commanders, deployed forces, and manned and unmanned systems.

Currently, AFOSR is organized into four directorates: Physics and Electronics, Aerospace and Materials Sciences, Chemistry and Life Sciences, and Mathematics and Space Sciences. The Mathematics and Space Sciences Directorate has a mix of programs that include topics such as dynamics and control, systems software and reliability, artificial intelligence, and information fusion. Nowhere in this breakout of topics does the term IS&T appear. In fact, as noted elsewhere in this report, the IS&T efforts are spread over multiple programs, such as those listed above. As it reviewed the available AFOSR information by topic and by budget, the committee had difficulty in determining the full scope of what could be categorized as IS&T research. While this might not hamper staff within AFOSR and AFRL, it could be an unnecessary hurdle for interfacing with outside organizations.

Recommendation. Given the Air Force's increasing reliance on IS&T, its role in enabling revolutionary concepts in military affairs, and the committee's recommendation to significantly increase the investment in basic research for IS&T, the committee recommends that AFOSR consider identifying IS&T as a major area within AFOSR. In fact, as the investment in that area increases over time, AFOSR should consider organizing a separate directorate for IS&T and related topics. Given the committee's recommendation regarding human-system interfaces and interactions, it would seem appropriate to also consider alignment of all human-system and human-effects-related research within this separate directorate.

RECRUITMENT OF PROGRAM MANAGERS

With the recommended increase in funding for IS&T and the corresponding increase in scope of the research investment, AFOSR will be faced with a staffing challenge. This challenge can be viewed as an opportunity to reach out to the IS&T research community to recruit new program managers. The DARPA model for recruiting and rotating program managers has quite successfully made use of the Intergovernmental Personnel Act and the more recent Experimental Personnel Hiring Authority. The former has been used to recruit candidates from universities and non-profit organizations such as federally funded research and development centers (FFRDCs), while the latter has been used to bring candidates from industry to government for a limited term to participate as program managers in exciting new areas of R&D. These tools have been used at the program manager and higher levels with great success. Aside from attracting the best and brightest from universities, nonprofits, and industry, these tools also rotate program managers out of the agency and back to endeavors of their choice when their term is completed. AFOSR then has an opportunity to both staff the recommended growth in IS&T and maintain the vitality of the area by bringing in new program managers as replacements for those who are rotating out.

Recommendation. AFOSR has an opportunity as well as a challenge to staff the recommended increase in funding for research in IS&T. The committee recommends that AFOSR take advantage of mechanisms such as the Intergovernmental Personnel Act to recruit new program managers and, possibly, directorate directors for some if not all of the expanded investment in IS&T. This approach would facilitate continued vitality in this important and growing area of research.

A MECHANISM FOR FOSTERING EXPERIMENTAL RESEARCH IN IS&T

The committee recommends that AFOSR encourage and support more experimental science in its IS&T program. While visiting the Air Combat Command in April 2005, the committee was briefed on an interesting approach to collaboration, the Information Operations Innovation Network. This network, while not a research effort, facilitates collaboration and coordination among technologists and operators in the development of experiments and supporting capabilities; it is used for carrying out technology demonstrations of new concepts in response to capability gaps. Because the specifications for team-focused, network-enabled systems are poorly defined both technically and operationally, there is a real need for joint experiments involving R&D staff and Air Force operations personnel working with early prototype systems. The experiments would then help to define the system requirements, and researchers would gain a tangible understanding of Air Force goals and emerging capabilities. The challenges posed by distributed information systems and their underlying science and technology require a new approach, and the committee recommends establishing team-focused experimental environments for R&D, data collection, experimentation, and demonstration of new concepts. Examples of issues that present a serious challenge to both researchers and Air Force operators include these:

- How to control distributed systems with or without humans in the loop.
- How to support decision-making by sharing information and intentions.
- How to fuse and manage heterogeneous information.
- How to test and evaluate techniques for offensive and defensive information warfare.

The committee uses the label “distributed research and experimentation environment” (DREE) to describe a shared computation infrastructure that supports experimentation within a community of researchers. As the name implies, DREEs could be virtual centers (geographically distributed), and they could involve a mixture of physical and simulated assets. In some cases, the scale or complexity of the experimental environment might be too great to emulate in a DREE. In such cases, modeling and simulation might enable experimentation and collaboration among distributed participants through micro world descriptions that are domain faithful but presented in an unclassified environment. Such modeling and simulation capabilities would enable a proactive discovery process through which Air Force IT challenges come into clearer focus. In addi-

tion to their technical goals, DREEs would themselves add to the Air Force's experience base in team-focused, network-enabled systems and contribute to interactions across various Air Force R&D communities.

Such collaborative experimental arrangements are not unprecedented in research. For instance, consider the following from the economics research community:¹

Computational laboratories (CLs) are computational frameworks that permit the study of complex system behaviors by means of controlled and replicable experiments. Agent-based computational economics (ACE) is the computational study of economies modeled as dynamic systems of interacting agents. ACE researchers generally conduct their studies in the context of CLs.

Research groups frequently cite the need for support to create shared data sets, testbed environments, and other infrastructure that will facilitate access to larger, scalable problems and community sharing. For example, consider the following excerpt from a 2003 report on strategic direction for artificial intelligence research at Cornell University:²

A common theme that emerged in all workshop sessions was the importance of data sets and testbeds. . . . Data sets and testbeds also have a powerful multiplier effect on research progress. A compelling, freely available data set may motivate hundreds of separate research studies, most by researchers with no connection (funding or otherwise) to the original producers of the data set. Particularly in areas where it is difficult to identify in advance the most promising technologies, the resulting breadth of voluntary effort can be crucial. . . . Support for the creation of freely available, sharable resources is likely to do more to move forward AI research in areas of interest to the Air Force than any other single action.

There are several examples of analogous testbeds already in use in the IS&T research community:

- Many researchers in robotics are participants in periodic RoboCups.³ These events are exciting, involve a large community of researchers, and provide a challenging shared domain for research in perception, adversarial planning, cooperative agent behavior, and machine learning.

¹Available at <http://www.econ.iastate.edu/tesfatsi/acedemos.htm>.

²Available at <http://www.cis.cornell.edu/iisi/SRDAI-workshop/srdai-iisi-report-2003.doc>.

³Available at <http://www.econ.iastate.edu/tesfatsi/acedemos.htm>.

- Similar benefits of sharing and leverage were seen in the DARPA spoken language experiments from the early 1990s. A large data set based on military messages was prepared each year, and this was used in a community-wide competition to benchmark progress in spoken language research. DARPA-funded research groups were expected to participate, and leading commercial groups joined in; competition was intense and progress was rapid. The DARPA investment was in creating and providing data sets that were shared with the research community. Sharing in this domain required no expensive testbeds or equipment, only shared data sets and performance metrics. Replication of experiments was encouraged and, in fact, conducted intensely. It created an understanding of how different approaches could achieve community-leading performance according to the metrics.
- In the domain of network services, PlanetLab⁴ has evolved as a community-supported, distributed multiuser platform that serves as a testbed for overlay networks. It currently consists of 578 machines hosted by 275 sites in more than 25 countries. Most of the machines are hosted by research institutions, although some are elsewhere (e.g., on Internet2's Abilene backbone). The key objective of the community-developed and community-shared software is to support distributed virtualization—that is, to allocate a slice of PlanetLab's networkwide hardware resources to an application. This allows an application to run across all (or some) of the machines around the globe, where at any given time multiple applications may be running in different slices of PlanetLab. The advantage to researchers using PlanetLab is that they are able to experiment with new services under real-world conditions and at large scale. An arrangement offering capabilities would be of great value to the Air Force as it tries to develop a better understanding of the design, operation, and management of large, complex networks and systems and of the incorporation of such systems into Air Force operations. Although PlanetLab is more limited in scope than what the committee has in mind, it is given here as an example that works.

These examples have three things in common:

- Each is directed at an exciting problem area.
- Each involves a research community that is willing, or encouraged, to share research ideas, experiments in which they investigate those

⁴Available at <http://www.planet-lab.org>.

ideas, and results. Ad hoc teams focused on similar problems spawn new ideas at the boundaries of their own idea sets.

- Each research community supports an infrastructure that can be shared, which may be as modest as data sets and repositories for reports and shared code.

The committee urges AFOSR to work with other parts of the Air Force to establish DREEs as analogous testbeds that will allow researchers and Air Force users to experiment with prototype IS&T concepts and systems. Besides the inherent benefit of adding this experimental component to other IS&T research efforts, such an approach would serve as an intellectual crossroads between the scientific and operational communities in support of the scientific discovery process. The approach also gives researchers something very concrete to work on, perhaps including real Air Force data and proposed system specifications, so they can measure the quality of their solutions. The committee was encouraged by a presentation of the AFRL chief technologist, who described how a C2 wind tunnel could provide a computationally based experimentation framework for exploring new technologies and their operational utility. This is an area of great interest to senior Air Force leadership, which has requested a study in 2006 by the Air Force Scientific Advisory Board on rapid, affordable experimentation.

The committee believes that DREEs would be useful for each of the main areas of research covered in this report. A DREE for information management, for instance, would enable the associated community—including universities, AFRL directorates, and perhaps FFRDCs—to create sample data sets and develop associated queries that illustrate how the data are to be integrated. A DREE related to network-centric systems would allow exercises that might generate concrete performance requirements, which are otherwise difficult to identify. While exercises are ongoing, operational Air Force participants could clarify their real, not hypothetical, needs; IS&T applied researchers could investigate engineering issues with the prototype network; basic researchers in IS&T could experiment with fundamental changes (to, for example, communication protocols); and HSI researchers could instrument the experiments and learn from them.

The DREE approach to experimental science should not be prohibitive in cost, because the necessary network infrastructure is rapidly falling into place and it may become possible to leverage the investments in testbeds by other AFRL directorates. For example, a research version of the distributed mission training (DMT) environment housed in AFRL's Human Effectiveness Directorate might support experimental science in areas ranging from control of UAVs to decision-making in real-time envi-

ronments. The committee's intent is that DREEs be established without expending basic research funds, which would only be used to support the involvement of basic researchers. Planning for such facilities should be coordinated with AFRL and the Air Combat Command and funded with 6.2 and 6.3 money.

Recommendation. The committee recommends that as AFOSR expands its investment in basic IS&T research, it work with its 6.2 and 6.3 partners to establish DREEs, which will not only speed the pace of research but also facilitate the transition of critical technology. The committee's intent is that 6.1 money not be used to establish these DREEs; they would be set up by the technical directorates of AFRL or by operational arms of the Air Force, which would use them to exercise early prototypes of complex, software-intensive systems and provide feedback and guidance to the research community. Basic researchers supported by 6.1 funds would contribute to those exercises and extract lessons for further research.

The concepts proposed for experimentation can start small and later scale up to fully supported DREEs, as deemed appropriate to support experimental research. Therefore, the committee recommends that AFOSR start small but soon. In its next broad agency announcement that deals with IS&T research, AFOSR could request an optional task—namely, an experiment that the proposer would create, publish, and share with other research groups. AFOSR could also consider holding summer workshops on DREE topics for doctoral students. For instance, AFRL's Human Effectiveness Directorate is interested in revolutionary approaches to DMT and would like a capability that is more like a free-play video game than the current tightly scripted scenarios. Students at a summer workshop could learn about Air Force needs in DMT, discuss and debate basic research issues that have to be resolved to advance the associated technology, and construct a DREE that they would take back to their respective universities for use in their research. In short, AFOSR, using only modest funding, could seed the development of a DREE and its associated community.

AFOSR should collaborate with other organizations to support the development and use of DREEs. For example, DOD's High Performance Computing Modernization Office (HPCMO), through its Defense Research and Engineering Network, supports over 4,300 scientists and engineers at government laboratories, universities, and industrial research laboratories. HPCMO sponsors an annual call for "challenge problems" and invests up to 25 percent of its funding to solve these challenge problems. AFRL might consider leveraging its own funding with some of HPCMO's "challenge problem" funding to create one or more DREEs focused on Air Force-specific IS&T needs.

Appendixes

A

Meeting Agendas

MEETING 1 ROME, NEW YORK

Thursday, February 24, 2005

Closed session

8:00 am Working breakfast

Open session

- 9:30 Car pool to AFRL Information Directorate (AFRL/IF)
- 10:00 Welcome and introductions (Alan McLaughlin, committee chair)
- 10:10 Charge to the committee and discussion of goals (Clifford Rhoades, head of Mathematics and Space Sciences Directorate, AFOSR).
- 10:30 Overview of Air Force IS&T priorities and Air Force long-term challenges (Robert Herklotz, AFOSR).
- 11:15 Overview of the AFRL/IF and its response to Air Force IS&T priorities and Air Force long-term challenges (Nort Fowler, chief scientist, AFRL/IF).

- 12:30 pm Lunch
- 1:15 Parallel sessions. Committee breaks into groups to learn more about AFRL/IF's major areas.
- 3:15 Break
- 3:30 Overview of IS&T-related research sponsored by the AFOSR Mathematics and Space Sciences Directorate (AFOSR/NM); summary of investments and funding mechanisms; discussion of AFOSR/NM's program-planning process and reviews (Clifford Rhoades, head, AFOSR/NM).
- 4:30 Discuss with Air Force staff the flow of knowledge between AFOSR/NM and AFRL/IF, AFRL's connections to the broader research community, the transition of advances from AFOSR and AFRL into Air Force technology, and other issues raised during the day.
- 5:30 Adjourn
- Closed session
- 7:00 pm Committee working dinner

Friday, February 25, 2005

Open session

- 8:00 am Working breakfast
- 8:30 Discussion with Brendan Godfrey, AFOSR director, on interplay between AFOSR and other AFRL directorates in addressing Air Force S&T goals and how the Air Force decides whether or not it needs to sponsor or conduct R&D in a given area, particularly for R&D related to IS&T and command and control.
- 9:15 Discussion with Shankar Sastry, University of California at Berkeley, and Janos Sztipanovits, Vanderbilt University, both members of the Air Force Scientific Advisory Board, on the recent review of AFOSR's IS&T research.

- 9:45 Discussion with Clifford Rhoades and other AFRL/IF and AFOSR/NM staff, as needed, about the information received on February 24 to identify gaps and how to fill them. Identify site visits, telephone calls, and background information that the committee should pursue.
- 11:00 Break
- Closed session
- 11:15 Committee discussion of preliminary impressions
- 12:15 pm Lunch
- 1:00 Report from Elwyn Berlekamp, who attended the January program review of AFOSR's research consortium on space situational awareness.
- 1:30 Develop initial points of committee consensus and identify topics for investigation.
- 3:00 Develop plans for site visits and identify any additional information needs.
- 3:30 Adjourn

**MEETING 2
HAMPTON, VIRGINIA**

Tuesday, April 26, 2005

Closed session

7:45 am Working breakfast

Open session

- 8:45 Committee will be picked up by a shuttle from the hotel lobby for travel to Langley AFB.
- 9:15 Welcome (Alan McLaughlin, committee chair)
Overview of the Air Combat Command (Janet Fender, ACC chief scientist, and Gen Maluda, head of ACC Communications and Directorate).

- 9:45 Presentation on communications and information systems (Col Kemp).
- 10:30 Presentation on the Air Force Information Warfare Center (John Kretzer).
- 11:15 Presentation on the Operations Directorate for Information Operations (Lt Col Lance).
- 11:45 Presentation on requirements, ACC Information Operations (Col Anninos and Lt Col Sowell).
- 12:15 pm Lunch with open discussion
- 1:00 Presentation on selected IS&T challenge areas identified by the C2ISR Center (Jon Vona).
- 3:00 Presentation on funding mechanisms used by AFOSR/NM (John Tangney), followed by general discussion with AFOSR/NM program managers.
- 4:00 Wrap-up discussions
- 5:00 Adjourn
- Closed session
- 7:00 Working dinner

Wednesday, April 27, 2005

- Closed session
- 8:00 am Working breakfast
- 8:30 Discussion of preliminary drafts
- Noon Lunch

12:45 pm Continue discussions. Develop plans for additional information gathering and writing. Issue assignments.

2:00 Adjourn

**MEETING 3
WASHINGTON, D.C.**

Monday, June 6, 2005

Closed session

8:00 am Working breakfast

8:30 Discuss progress since last meeting and plan for the current meeting.

9:45 Break

Open session

10:00 Discuss goals of study with Tom Cruse, chief technologist, AFRL. Brendan Godfrey, head of AFOSR, will also participate.

Closed session

11:15 Discuss Chapters 1 and 2 of the draft report

12:15 pm Lunch

1:00 Impressions of AFRL's Human Engineering Directorate (Steve Cross, based on his recent visit there).

1:30 Discuss Chapters 3, 4, and 5 of the draft report

3:30 Break

3:45 Discuss Chapters 7 and 8 of the draft report

5:00 Identify questions for tomorrow's discussion with AFOSR program managers.

- 5:30 Reception
- 6:30 Working dinner
- 8:00 Adjourn

Tuesday, June 7, 2005

Open session

- 8:00 am Working breakfast
- 8:30 Discussion with AFOSR program managers regarding their current program and the committee's open questions.

Closed session

- 9:30 Revisit topics from Monday's discussion
- 10:00 Break
- 10:15 Discuss Chapter 6 of the draft report; convergence on research priorities under different funding scenarios.
- Noon Lunch
- 12:45 pm Identify additional research topics and/or information-gathering trips needed to explore topics or develop background. Identify other issues and recommendations.
- 2:00 Detailed plans for completing the report
- 3:30 Adjourn

B

Acronyms

ACC	Air Combat Command
AFOSR	Air Force Office of Scientific Research
AFRL	Air Force Research Laboratory
AFRL/IF	AFRL Information Directorate
AFRL/NM	AFOSR Mathematics and Space Sciences Directorate
AJ	antijamming
BAA	broad agency announcement
C4ISR	command, control, computing, communications, intelligence, surveillance, and reconnaissance
DARPA	Defense Advanced Research Projects Agency
DEPSCoR	DOD Experimental Program to Stimulate Competitive Research
DMT	distributed mission training
DOD	Department of Defense
DoS	denial of service
DREE	distributed research and experimentation environment
EAI	enterprise application integration
EII	enterprise information integration
ETL	extract, transform, and load
GIG	Global Information Grid
GPS	Global Positioning System

HBCU/MI	Historically Black Colleges and Universities or Minority Institutions
HCI	human-computer interface
HPCMO	High Performance Computing Modernization Office
HSI	human-system interactions
IO	information operations
IS&T	information science and technology
ISR	intelligence, surveillance, and reconnaissance
IT	information technology
JBI	Joint Battlespace Infosphere
LPD	low probability of detection
LPI	low probability of intercept
MAC	Media Access Control
MAV	micro air vehicle
MIMO	multiple input, multiple output
MoBIES	Model-Based Integration of Embedded Software
MPEG-7	Moving Picture Coding Experts Group standard No. 7
MURI	multidisciplinary university research initiative
NRC	National Research Council
PERT	Program Evaluation and Review Technique
PRET	Partnership for Research Excellence and Transition
R&D	research and development
RPC	remote procedure call
S&T	science and technology
SOA	service-oriented architecture
SOAP	Simple Object Access Protocol
STTR	Small Business Technology Transfer
TCP	Transport Control Protocol
UAV	unmanned air vehicle
URI	university research initiative
UWB	ultrawideband
XML	extensible markup language