

# The Architecture of High Performance Computers

# The Architecture of High Performance Computers

Roland N. Ibbett  
*Reader in Computer Science,  
University of Manchester*



Springer Science+Business Media, LLC

© Roland N. Ibbett 1982

Originally published by Springer-Verlag New York in 1982

All rights reserved. No part of this publication may be reproduced or transmitted, in any form or by any means, without permission.

*First published 1982 by  
The Macmillan Press Ltd  
London and Basingstoke*

Sole distributors in the USA and Mexico  
Springer Science+Business Media, LLC.

ISBN 978-1-4757-6717-9      ISBN 978-1-4757-6715-5 (eBook)  
DOI 10.1007/978-1-4757-6715-5

# Contents

## Preface

### 1 Introduction

1.1 Historical Developments	1
1.2 Techniques for Improving Performance	2
1.3 An Architectural Design Example	3

### 2 Instructions and Addresses

2.1 Three-address Systems - The CDC 6600 and 7600	7
2.2 Two-address Systems - The IBM System/360 and /370	10
2.3 One-address Systems	12
2.4 Zero-address Systems	15
2.5 The MU5 Instruction Set	17
2.6 Comparing Instruction Formats	22

### 3 Storage Hierarchies

3.1 Store Interleaving	26
3.2 The Atlas Paging System	29
3.3 IBM Cache Systems	33
3.4 The MU5 Name Store	37
3.5 Data Transfers in the MU5 Storage Hierarchy	44

### 4 Pipelines

4.1 The MU5 Primary Operand Unit Pipeline	49
4.2 Arithmetic Pipelines - The TI ASC	62
4.3 The IBM System/360 Model 91 Common Data Bus	67

### 5 Instruction Buffering

5.1 The IBM System/360 Model 195 Instruction Processor	72
5.2 Instruction Buffering in CDC Computers	77
5.3 The MU5 Instruction Buffer Unit	82
5.4 The CRAY-1 Instruction Buffers	87
5.5 Position of the Control Point	89

## **6 Parallel Functional Units**

6.1 The CDC 6600 Central Processor	95
6.2 The CDC 7600 Central Processor	104
6.3 Performance	110
6.4 The CRAY-1	112

## **7 Vector Processors**

7.1 Vector Facilities in MU5	126
7.2 String Operations in MU5	136
7.3 The CDC Star-100	142
7.4 The CDC CYBER 205	146
7.5 Future Developments	163

<b>References</b>	165
-------------------	-----

<b>Index</b>	170
--------------	-----

# *Preface*

Computer architecture has always formed an essential part of all single honours degree courses taught in the Department of Computer Science at the University of Manchester, and the material in this book is largely based on those parts of the course which are concerned with the design of high performance uniprocessors. One of the principal aims of the course has always been to encourage students to discover the 'whys' as well as the 'hows' of computer architecture. For this reason I have tried to place the computers described here in their proper historical context, in order to show both why and how they were developed. Similarly I have not tried to impose a uniform nomenclature on to the material, but rather to use the terms which came naturally to the designers and which are to be found in the original documentation.

I am indebted to many people for advice and assistance in the preparation of this book. At Manchester numerous colleagues joined in discussions about the various topics considered here, and the computer maintenance team kept MU5 in good order while it processed the text. Elsewhere computer maintenance engineers at various sites willingly answered obscure questions about the machines in their charge, and staff at Cray Research and Control Data Corporation vetted parts of the manuscript. Particular mention should be made of Chuck Purcell of CDC for much useful information, and I would like to thank Professors T. Kilburn, D. B. G. Edwards, F. H. Sumner and D. Morris for their constant encouragement and advice over the years. Finally a word of appreciation to my small son James for his 'help' with the manuscript.

Roland N. Ibbett

June 1982

# *1 Introduction*

Computer architecture has been defined in a number of ways by different authors. Amdahl, Blaauw and Brooks [1], for example, the designers of the IBM System/360 architecture, used the term to 'describe the attributes of a system as seen by the programmer, i.e. the conceptual structure and functional behaviour, as distinct from the organisation of the data flow and controls, the logical design and the physical implementation'. Stone [2], on the other hand, states that 'the study of computer architecture is the study of the organisation and interconnection of components of computer systems'. The material presented here is better described by this wider definition, but is particularly concerned with ways in which the hardware of a computer can be organised so as to maximise performance, as measured by, for example, average instruction execution time. Thus the architect of a high performance system seeks techniques whereby judicious use of increased cost and complexity in the hardware will give a significant increase in overall system performance.

## **1.1 HISTORICAL DEVELOPMENTS**

Designers of the earliest computers, such as the Manchester University/Ferranti Mark 1 (first produced commercially in 1951 [3]), were constrained by the available technology (valves and Williams Tube storage, for example, with their inherent problems of heat dissipation and component reliability) to build (logically) small and relatively simple systems. Even so, the introduction of B-lines and fast hardware multiplication in the Mark 1 were significant steps in the direction of cost-effective hardware enhancement of the basic design. At Manchester this trend was developed further in the Mercury computer, with the introduction of hardware to carry out floating-point addition and multiplication. This increased logical complexity was made possible by the use of semiconductor diodes and the availability of smaller and more reliable valves than those used in Mark 1, both of which helped to reduce power consumption (and hence heat dissipation) and increase overall reliability. A further increase in reliability was made in the commercial version of Mercury (first produced in 1957) by the use of the then newly

developed ferrite core store.

The limitations on computer design imposed by the problems of heat dissipation and component reliability were decreased dramatically in the late 1950s and early 1960s by the commercial availability of transistors, and the first generation of 'supercomputers' (Atlas, Stretch, MULTICS and the CDC 6600, for example [4,5,6,7]) appeared. These machines incorporated features which have influenced the design of succeeding generations of computers (the paging/virtual memory system of Atlas and the multiple functional units of the 6600, for example), and also highlighted the need for sophisticated software, in the form of an operating system, intimately concerned with activities within the hardware from which the user (particularly in a multi-user environment) required protection, and vice versa! In this book we shall follow the developments of some of the ideas from these early supercomputers into present day computer designs.

## **1.2 TECHNIQUES FOR IMPROVING PERFORMANCE**

Improvements in computer architecture arise from three principle factors

- (1) technological improvements
- (2) more effective use of existing technology
- (3) improved software-hardware communication

Technological improvements include increases in the speed of logic circuits and storage devices, as well as increases in reliability. Thus the use of transistors in Atlas, for example, not only allowed individual circuits to operate faster than those in Mercury, but also allowed the use of parallel adders, which involved many more logic circuits but which produced results very much more quickly than the simple serial adders used in earlier machines. Corresponding changes can be seen today as the scale of integration of integrated circuits continues to increase, allowing ever more complex logical operations to be carried out within one silicon chip, and hence allowing greater complexity to be introduced into systems while keeping the chip speed and overall system reliability more or less constant.

Making more effective use of existing technology is the chief concern of computer architecture, and almost all of the techniques used can be classified under the headings of either storage hierarchies or concurrency. Storage devices range from being small and fast (but expensive) to being large and cheap



(but slow); in a storage hierarchy several types, sizes and speed of storage device are combined together, using both hardware and software mechanisms, with the aim of presenting the user with the illusion that he has a fast, large (and cheap) store at his disposal.

Concurrency occurs in various forms and at various levels of system design. Low-level concurrency is typified by pipeline techniques, interleaved storage and parallel functional units, for example, all of which are largely invisible to software. High-level concurrency, on the other hand, typically involves the use of a number of processors connected together in some form of array or network, so as to act in parallel on a given computational task, and special software is generally required in order to organise the activities of these processors. In this book we shall be largely concerned with low-level concurrency techniques.

Communication between software and hardware takes place through the order code (or instruction set) of the computer, and the efficiency with which this takes place can significantly affect the overall performance of a computer system. One of the difficulties with some 'powerful' computers is that the 'power' is obtained through hardware features which can only be fully exploited either by hand coding or by the use of complex and time-consuming algorithms in compilers. This fact was recognised early in the Manchester University MU5 project, for example, where an instruction set was sought which would permit the generation of efficient object code by high-level language compilers. The design of this order code was therefore influenced not only by the anticipated organisation of the hardware, which was itself influenced by the available technology, but also by the nature of existing high-level languages. Thus the order code, hardware organisation, and available technology all interact together to produce the overall architecture of a given system. An example of such an interaction is given in the next section, based on some of the early design considerations for MU5 [8].

### 1.3 AN ARCHITECTURAL DESIGN EXAMPLE

At the time when the MU5 project was started (1966/7), the fastest production technology available for its construction was that being used by ICT (later to become ICL) for their 1906A computers. In this technology, MECL 2.5 small scale integrated circuits are mounted individually or in pairs on printed circuit modules, together with discrete load resistors, and up to 200 of these modules are interconnected via multi-layer platters. Platters are mounted in groups of six or nine within logic bays, with adjacent platters being

joined by pressure connectors. Inter-group and inter-bay connections are via co-axial cables. The circuit delay for this technology is 2 ns, but the wiring delay of 2 ns/ft (through matched transmission line interconnections) must also be taken into account. An average connection between two integrated circuits mounted on modules on the same platter involves a 1" connection on each module and a 4" connection between modules, involving 6" in all and giving an extra 1 ns delay. Some additional delay also occurs because of the loading of an output circuit with following input circuits, and so for design purposes a figure of 5 ns was assumed for the delay between the inputs of successive gates.

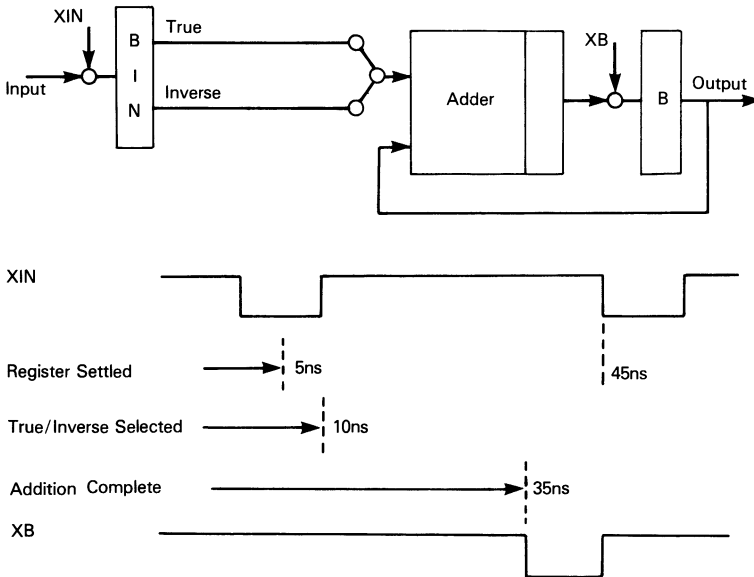


Figure 1.1 An Index Arithmetic Unit

A 32-bit fixed-point adder/subtractor constructed in this technology requires five gate delays through the adder, plus one gate delay to select the TRUE or INVERSE of one of the inputs to allow for subtraction, giving a total delay of 30 ns. This adder/subtractor can be incorporated into an index arithmetic unit as shown in figure 1.1. A 10 ns strobe XIN copies new data into register BIN, the output from which is steady after 5 ns. During the addition, information is strobed into a register which forms part of the adder, and the 10 ns strobe XB, which copies the result of the addition into the index register B, starts immediately after this internal adder strobe has finished. The earliest time at which the next XIN

strobe can start is at the end of XB, so that the minimum time between successive add/subtract operations in this unit is 45 ns.

This time is very much less than the access time to the main store of MU5, a plated-wire store with a 260 ns cycle time. (A suitable replacement semiconductor store, using 16K MOS RAMs, would actually be even slower.) Thus, in the absence of some additional technique, there would be a severe mis-match between the operand accessing rate and the arithmetic execution rate for index arithmetic instructions. An examination of the operands used in high-level languages, and studies of programs run on Atlas, indicated that over a large range of programs, 80 per cent of all operand accesses were to named scalar variables, of which only a small number was in use at any one time. Thus a system which kept these operands in fast programmable registers would be able to achieve high performance, a technique commonly used in existing computers. However, this is exactly the sort of hardware feature which causes compiler complexity, and which the designers of MU5 sought to avoid.

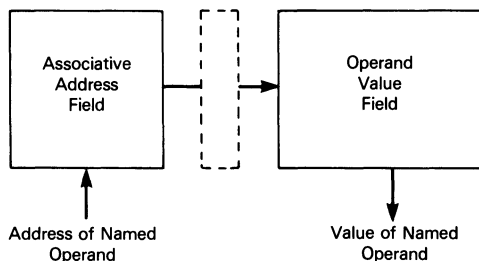


Figure 1.2 The MU5 Name Store

The alternative solution adopted in MU5 involves the identification within each instruction of the kind of operand involved, and the inclusion in the hardware of an associatively addressed buffer store for named variables, the Name Store. This store has to operate in conjunction with the main store of the processor, thus immediately introducing the need for a storage hierarchy in the system. Having observed this implication, further discussion of storage hierarchies will be left until Chapter 3; there are additional implications for the system architecture to be considered here, based on the timing of Name Store operations.

The Name Store consists of two parts, as shown in figure 1.2; the associative address field and the operand value field. During the execution of an instruction involving a

named variable, the address of the variable is presented to the associative field of the Name Store, and if a match is found in one of its 32 associative registers, the value of the variable can be read from the corresponding register in the Name Store value field. The time required for association is 25 ns, and similarly for reading. Thus, in order for the index arithmetic unit to operate at its maximum rate, the association time, reading time and addition time for successive instructions must all be overlapped (by introducing buffer registers such as that shown dotted in figure 1.2). In making provision for this overlap, however, another architectural feature has been introduced - the organisation of the processor as a pipeline. Further discussion of pipelines will be left until Chapter 4, but it should now be clear that there is considerable interaction between the various facets of the architecture of a given computer system.

## *2 Instructions and Addresses*

An important characteristic of the architecture of a computer is the number of addresses contained in its instruction format. Arithmetic operations generally require two input operands and produce one result, so that a three-address instruction format would seem natural. However, there are arguments against this arrangement, and decisions about the actual number of addresses to be contained within one instruction are generally based on the intuitive feelings of the designer(s) in relation to economic considerations, the expected nature of the implementation, and the type of operand address and its size. An important distinction exists between register addresses and store addresses, for example; if the instruction for a particular computer contains only register addresses, so that its main store is addressed indirectly through some of these registers, then up to three addresses can be accommodated in one instruction. On the other hand, where full store addresses are used, multiple-address instructions are generally regarded as prohibitively expensive both in terms of machine complexity and in terms of the static and dynamic code requirements. Thus one store address per instruction is usually the limit (in which case arithmetic operations are performed between the content of the store location and the content of an implicit accumulator), although some computers have variable-sized instructions and allow up to two full store addresses in a long instruction format. In this chapter we shall introduce examples of computer systems which have three, two, one and zero-address instruction formats, and discuss the relationships which exist between each of these arrangements and the corresponding hardware organisation.

### **2.1 THREE-ADDRESS SYSTEMS - THE CDC 6600 AND 7600**

The Control Data Corporation 6600 computer first appeared in the early 1960s, and was superseded in the late 1960s by the 7600 system (now renamed CYBER 70 Model 76). The latter is machine code compatible upward from the former, so that the basic instruction formats of the two systems are identical, but the 7600 is about four times faster than the 6600. The 6600 remains an important system for students of computer

architecture, however, since it has been particularly well documented [7], and an understanding of its organisation and operation provides a proper background not only for an appreciation of the design of the 7600 system, but also that of the CRAY-1, which can be seen as a logical extension of the 6600/7600 concepts from scalar to vector operation. The design of all these machines will be discussed in more detail in Chapter 6.

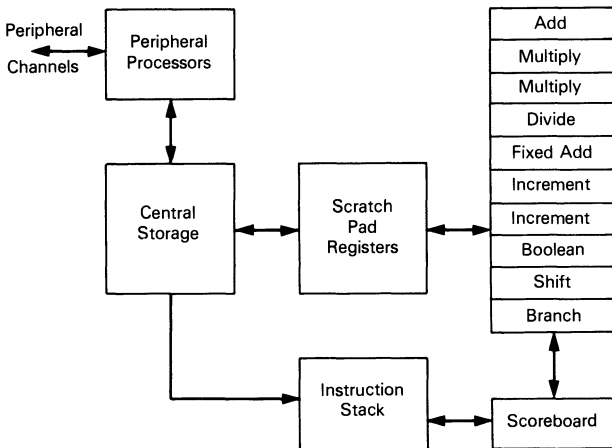



Figure 2.1 CDC 6600 Processor Organisation

The 6600 was designed to solve problems substantially beyond contemporary computer capability and, in order to achieve this end, a high degree of functional parallelism was introduced into the design of the central processor. An instruction set and processor organisation which could exploit this parallelism, while at the same time maintaining the illusion, at least, of strict serial execution of instructions, was therefore required. A three-address instruction format provides this possibility, since successive instructions can refer to totally independent input and result operands. This would be quite impossible with a one-address instruction format, for example, where one of the inputs for an arithmetic operation is normally taken from, and the result returned to, a single implicit accumulator. Dependencies between instructions can still occur in a three-address system, of course; for example, where one instruction requires as its input the result of an immediately preceding instruction, and the hardware must ensure that these are strictly maintained. This would be difficult if full store addresses were involved, but the use of three full store addresses would, in any case, have made 6600 instructions

prohibitively long. There were, in addition, strong arguments in favour of having a 'scratch-pad' of fast registers in the 6600 which could match the operand processing rate of the functional units.



F denotes the major class of function  
 m denotes a mode within the functional unit  
 i identifies one of 8 registers within X, B or A  
 j identifies one of 8 registers within X, B or A  
 k identifies one of 8 registers within X B or A  
 K 18-bit immediate field used as a constant or Branch destination

Figure 2.2 CDC 6600 Instruction Formats

Thus the overall design of the 6600 processor is as shown in figure 2.1, and the instruction formats are as shown in figure 2.2. There are three groups of scratch-pad registers, eight 60-bit operand registers (X), eight 18-bit address registers (A), and eight 18-bit index registers (B). 15-bit computational instructions take operands from the two X registers identified by j and k, and return a result to the X register identified by i. Operands are transferred between X registers and Central Storage by instructions which load an address into registers A1-A5 (thus causing the corresponding operand to be read from that address in store and loaded into X1-X5), or into registers A6 or A7 (thus causing the content of X6 or X7 to be transferred to that address in store). The contents of an A register can also be indexed by the contents of a selected B register, and the B registers can also be used to hold fixed-point integers, floating-point exponents, shift counts, etc. The 30-bit instruction format (identified by particular combinations of F and m, as are the register groups) is used where a literal ('immediate') operand is to be used in an operation involving an A or B register, or in control transfer ('Branch') instructions, where the value in the K field is used as the jump-to, or destination, address.

The issuing of instructions to the functional units and the subsequent updating of result registers is controlled by the Scoreboard (section 6.1), which takes instructions in sequence from the Instruction Stack (section 5.2.1). This in turn receives instructions from Central Storage. The order code of the central processor does not include any peripheral handling

instructions, since all peripheral operations are independently controlled by the ten Peripheral Processors, which organise data transfers between Central Storage and input/output and bulk storage devices attached to the peripheral channels. This relieves the central processor of the burden of dealing with input/output, a further contribution towards high performance.

## 2.2 TWO-ADDRESS SYSTEMS - THE IBM SYSTEM/360 AND /370

The design objectives for the IBM System/360 [1,9] (introduced in 1964) were very different from those for the CDC 6600. Here the intention was to produce a line of processors (initially six) with compatible order codes, but a performance range factor of 50. Furthermore, whereas the CDC 6600 was intended specifically for scientific and technological work, System/360 machines were required to be configurable for a variety of tasks covering both scientific and data processing applications. The former are dominated by floating-point operations, where fast access to intermediate values requires the provision of one or more registers close to the arithmetic unit(s) in the central processor. Data processing applications, on the other hand, frequently involve the movement (and fairly simple manipulation) of long strings of data, and direct two-address storage to storage operations are much more appropriate than operations involving three store addresses or operations involving the use of intermediate values held in registers. Furthermore, in machines at the low performance end of the range, logical registers are normally implemented within main storage, and the extra store accesses required for these registers would actually slow down this type of operation.

Store addressing is itself a major problem in the design of a compatible range of computers; to quote from IBM, in turn quoting from DEC [10,11], 'There is only one mistake...that is difficult to recover from - not providing enough address bits...'. In the design of the IBM System/360 large models were seen to require storage capacities in the millions of characters, whereas small models required short addresses in order to minimise code space requirements and instruction access time. This problem was overcome by the use of index registers to supply base addresses covering the entire 24-bit address space, and small displacement addresses within instructions. This can be seen in figure 2.3, which shows the five basic instruction formats: register to register operations (RR), register to indexed storage operations (RX), register to storage operations (RS), storage and immediate operand operations (SI), and storage to storage operations (SS). All operand addresses consist of a 4-bit B field



specifying which of the 16 general registers provided is to be used as a base, and a 12-bit displacement which is added to the base value to form the store address. In the case of the RX format, a further general register is specified by the X field, for use as a conventional index register allowing access to an element within an array, for example. In the RS format this field becomes R3, used in the specification of the number of registers involved in the simultaneous transfer of a multiplicity of register contents to or from store. A typical instruction using the RR format, for example, involves the addition of the content of the register specified by R2 to the content of register R1, the result being returned to R1.

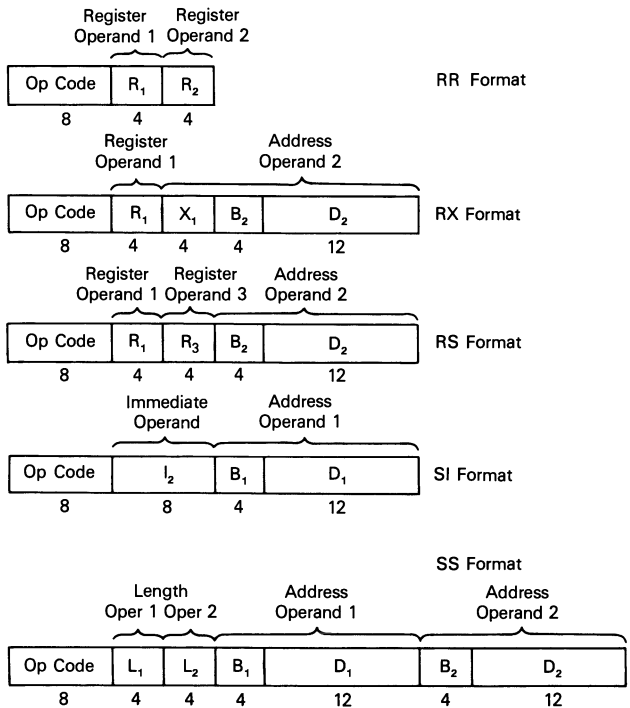


Figure 2.3 IBM System/360 Basic Instruction Formats

The number of registers to be provided was determined from an extrapolation, at the time when the System/360 was designed, of contemporary technological trends. These indicated the probable availability of small high-speed storage, and the design therefore incorporated what was then considered to be a substantial number of logically identifiable registers. These registers were to be physically

realised in core storage, local high-speed storage or transistors, according to the model. There are, in fact, two sets of addressable registers, sixteen 32-bit general purpose registers used both as index registers and fixed-point accumulators, and four 64-bit floating-point registers. Both sets occupy the same logical address space, the operation code determining which type of register is to be used in a given operation. The choice of 16 addressable register locations also allows the specification of two registers in a single 8-bit byte, the smallest addressable unit of information in the System/360. The separation of the two sets, using operation code information, allows instruction execution in larger models to be carried out in a separate unit from, and concurrently with, the preparation of subsequent instructions, thus leading to higher performance (section 4.3).

Store addresses generated by the central processor of almost all System/360 machines are 'real' addresses, each referring directly to a physical location in main store. (The same is also true of CDC 6600 addresses.) A move away from this situation occurred with the introduction of the System/370 in 1970 [10]. System/370 machines have essentially the same order code as System/360 machines, with minor modifications including, for example, the introduction of some one-address instructions. Addresses in System/370 machines are all 'virtual' addresses, however, which do not refer directly to physical locations in main store, but are mapped on to these locations by a dynamic address translation mechanism. Such a system was first introduced on Atlas, and will be described in some detail in Chapter 3.

### 2.3 ONE-ADDRESS SYSTEMS

Few high performance computers have strictly one-address instruction formats, although these are fairly common among mini and microcomputers where economy of bits in instructions is vital. The DEC PDP-10 [12], for example, has a single store address within each instruction (figure 2.4), but also specifies one of sixteen identical, general purpose registers, rather than assuming a single implicit accumulator. 'One-and-a-half-address' might be a more appropriate designation for this type of instruction, and clearly the IBM System/360 RX format (figure 2.3) would fall into this class. Instructions in MU5 (section 2.5), on the other hand, and its commercial derivative, the ICL 2900 Series [13], can be regarded as strictly one-address. There are only a few registers, all serving different dedicated purposes, and the register specification is therefore properly part of the function code, rather than addressing information.

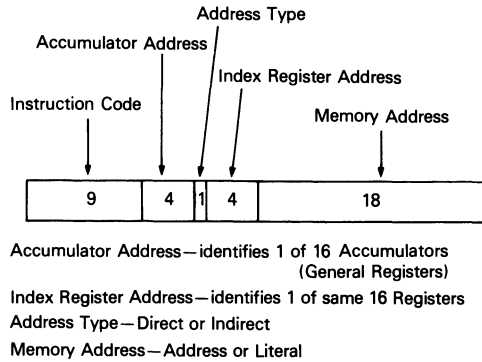


Figure 2.4 DEC PDP-10 Instruction Format

The Atlas computer [4] had a single full-length (48-bit) accumulator, closely associated with the main floating-point arithmetic unit, and so its A-codes were strictly one-address, with double B-modification (figure 2.5). The processor was organised in a manner which allowed operands to be accessed from the 2  $\mu$ s cycle-time core store at a rate matching the execution rate of the main arithmetic unit (of the order of 0.5M instructions per second), thereby avoiding the need for additional full-word buffer registers. This organisation (figure 2.6) involved having a separate 24-bit fixed-point B-arithmetic unit closely associated with a 128-line B-store, separate access mechanisms to the individual units of the core store (interleaving, described in more detail in section 3.1), and extensive overlapping of operations (pipelining).

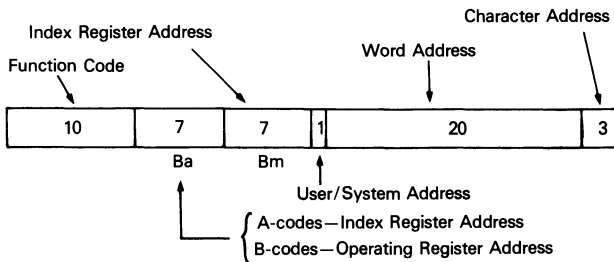


Figure 2.5 Atlas Instruction Format

B-codes in Atlas were of the one-and-a-half-address type, however, since they involved operations on one of the 128 B-lines specified by the Ba field in the instruction, Bm being used to specify a single modifier. The B-store was implemented as a fast core store (0.7  $\mu$ s cycle time) of 120 24-bit words,

together with eight flip-flop registers which served special purposes. These included the floating-point accumulator exponent, for example, and the three control registers (or program counters) used for user programs, extracode and interrupt control. The existence of these three control registers allowed the machine to switch rapidly from one context to another, while their inclusion within the addressable B-store avoided the need for separate control transfer functions. Extracodes were a set of functions which gave access to some 250 built-in subroutines held in the high-speed (0.3  $\mu$ s) read-only fixed store. These included the more complex mathematical functions such as sin, log, etc., peripheral control orders, input/output conversion routines, etc. The extracodes made use of 30 of the B-lines, leaving 90 for use as index registers or fixed-point accumulators in user programs.

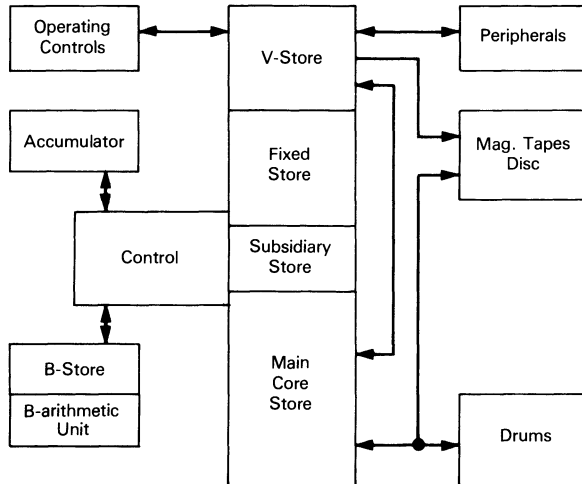


Figure 2.6 The Atlas Computer System

User addresses in Atlas (defined by a 0 in their most significant address bit) referred to a 1M word virtual address space, and were translated into real addresses through a set of page registers (as will be described in Chapter 3). The real store in the prototype machine consisted of a total of 112K words of core and drum storage, combined together through the paging mechanism so as to appear to the user as a 'one-level' store. Thus a user's total working store requirements could not exceed this amount, but code and data could be placed in any convenient part of the virtual address space. This led to an informal segmentation of the address space, a concept developed formally in MULTICS [6]. Formal

segmentation, or '2-dimensional addressing' is used in MU5 and the ICL 2900 Series (section 2.5), and has appeared more recently in minicomputers such as the Data General ECLIPSE MV/8000 and microcomputers such as the Zilog Z8000.

Atlas addresses with a 1 in the most significant bit were system addresses rather than user virtual addresses, and referred to the fixed store (containing the extracodes), the subsidiary store (used as working space by the extracodes and operating system), or the V-store. The latter contained the various registers needed to control the tape decks, input/output devices, paging mechanism, etc., thus avoiding the need for special functions for this purpose. This technique of incorporating peripherals into the address space was subsequently used in a number of computer systems, notably the PDP-11.

#### 2.4 ZERO-ADDRESS SYSTEMS

If we consider an Algol program statement of the form

```
RESULT := (A + B) * ((C + D) / (E + F))
```

then this could be compiled straightforwardly into one-address instructions as

```
LOAD A
ADD B
STORE TEMPO
LOAD C
ADD D
STORE TEMP1
LOAD E
ADD F
REVERSE DIVIDE TEMP1
MULTIPLY TEMPO
STORE RESULT
```

We note that two temporary variables have to be allocated and that these are re-used in the opposite order to that in which they were created. This corresponds to a last-in, first-out, or push-down stack arrangement, and where a multiplicity of base registers is provided, this stack is usually implemented by using one of the base registers as a stack pointer. Indeed, on the PDP-11, the auto-increment and auto-decrement addressing modes are specifically designed for this purpose.

An alternative arrangement is for the accumulator itself to be implicitly treated as the top-of-stack location, and for successive load orders to push previously loaded operands down

into the stack automatically. Arithmetic orders then operate between the two most recently loaded operands, and leave their result on the stack, so that no address specification is required. These are therefore zero-address instructions, and systems which use this arrangement are frequently known as stacking machines.

Stacking machines offer particular advantages to compiler writers, since the algorithm for translating from Algol into Reverse Polish notation, and thence into stacking machine code, is very simple. The expression in the example above becomes, in Reverse Polish

A B + C D + E F + / \* RESULT =>

Thus operands carry over into Reverse Polish without any relative change of position, and the rule for arithmetic operators is that the operator follows the two arithmetic values on which it is designated to work. Burroughs computers from the B5500 upwards are based on this system [14]. The push-down stack hardware consists of two or three registers at the top of the stack and a contiguous area of memory that permits the stack to extend beyond the registers.

Instructions in these machines are 12 bits long (packed four to a word), and fall into four categories: literal calls, operand calls and address calls (all of which involve pushing words into the stack), and operators. The operand calling mechanism is complex, since the actions which occur depend on the type of object found in store at the operand address (formed by treating part of the instruction or the top-of-stack word as an index, and adding it the content of a base register). The full word length is 51 bits, 48 bits of operand, instructions, etc., together with a 3-bit tag field which specifies the kind of information contained within the word. As examples, a word accessed by an operand call may be an operand (in which case it is simply stacked), a data descriptor (in which case the word addressed by the descriptor is stacked), or a program descriptor (in which case the program branches to a subroutine). Operators also use the tag field and treat words according to their tagged nature, so that there is only one ADD operator, for example, which takes any combination of single and double precision, integer or floating-point operands, and produces the appropriate result. Some operators require particular word types as operands, however, for which hardware checks are carried out, and an interrupt generated in case of error.

The designers of both the System/360 and the PDP-10 gave serious consideration to the use of a stacking machine, but

both rejected it on a number of grounds. The IBM team argued, for example, that the performance advantage of a push-down stack derived principally from the presence of several fast registers, rather than the way in which they were used or specified. Furthermore, they claimed that the proportion of occasions on which the top-of-stack location contained the item actually needed next was only about one half in general use, so that operations such as TOP and SWAP would be needed to copy submerged data to the active position and assist in clearing out-of-date submerged data. As a result, the static code requirements would be comparable with that obtained from a system with a small set of addressable registers.

The DEC team were concerned about high cost and possible poor performance in executing the operating system, LISP and FORTRAN. Burroughs' commitment to the essentially Algol-orientated stacking architecture was largely an act of faith, based on a belief held in several quarters in the early 1960s, that FORTRAN would disappear. It did not, has not, and probably will not for many years to come. It has been a constant source of difficulty for computer designers, and on Burroughs machines most FORTRAN programs execute more slowly than corresponding Algol-like programs.

## 2.5 THE MU5 INSTRUCTION SET

During the design of the MU5 processor the prime target was fast, efficient processing of high-level language programs in an interactive environment. This requirement led to the need for an instruction set which satisfied the following conditions

- (1) Generation of efficient code by compilers must be easy
- (2) Programs must be compact
- (3) The instruction set must allow a pipeline organisation of the processor leading to a fast execution rate
- (4) Information on the nature of operands should be available to allow optimal buffering of operands.

Algol and FORTRAN were taken as being typical of two major classes of programming languages, distinguished by the use of recursive routines (or procedures) and dynamic storage allocation in the former, and non-recursive routines and compile time storage allocation in the latter, but having the following features in common

- (a) Each routine has local working space consisting of named variables of integer (fixed-point) and real

(floating-point) types, and structured data sets (arrays, for example) which are also named

- (b) A routine may refer to non-local names either of an enclosing routine or in a workspace common to all routines
- (c) Statements involve an arbitrary number of operands which may be real names, array elements, function calls or constants.

In order to reconcile item (c) with conditions (1) and (4), it was decided that there would be an address form corresponding to each of the different operand forms. However, it was felt that to have more than one such operand specification per instruction would conflict with conditions (2) and (3) (although subsequent simulation experiments (section 2.6) have shown that this is not necessarily true), and the use of addressable fast registers was also rejected. There were two arguments against registers. The first was the desire to simplify the software by eliminating the need for optimising compilers. The second was the desire to avoid the need to dump and restore register values during procedure entry and exit and process changes. The alternative arrangement proposed for MU5 was a small associatively addressed Name Store, as described briefly in section 1.3, and which will be considered in more detail in Chapter 3.

The choice of instruction format was therefore limited to zero or one-address. A zero-address system was rejected on two main grounds. Firstly, the hand-coded library procedures which would be used for input/output handling, mathematical functions, etc., almost all required more code in zero-address than in one-address instructions. This was because the main calculation, the address calculations and control counting tended to interfere with each other on the stack. Secondly, execution of statements such as

$$A := B + C$$

would actually run more slowly on a stacking machine with the proposed hardware organisation for MU5. This statement would be coded in one-address and zero-address instructions as follows

ACC = B	STACK B
ACC + C	STACK C
	ADD
ACC => A	STORE A



If the operands, A, B and C, come from main store, then operand accessing would dominate the execution times of these sequences, and both would be approximately the same. However, using the Name Store, the access times to these operands and to the stack would be the same, and the zero-address sequence would involve six stack accesses (the ADD involves two reads and a write) in addition to the operand accesses.

The instruction format eventually chosen for MU5 represents a merger of single address and stacking machine concepts. All the arithmetic and logical operations involve the use of an accumulator and an operand specified in the instruction, but there is a variant of the load order (\*=) which causes the accumulator to be stacked before being re-loaded. In addition, there is an address form (STACK) which unstacks the last stacked quantity. Thus the expression

$$\text{RESULT} := (A + B) * ((C + D)/(E + F))$$

would compile into

```

ACC = A
ACC + B
ACC *= C
ACC + D
ACC *= E
ACC + F
ACC /: STACK
ACC * STACK
ACC => RESULT

```

The function '/:' is 'REVERSE DIVIDE'; if the operand to the left of an operator is stacked, it subsequently appears as the right hand side of a machine function (as the dividend in this case). Thus, for the non-commutative operations '-' and '/', the reverse operations '-:' and '/:' have to be provided. In pure stacking machines (such as Burroughs), all subtractions and divisions are implicitly reverse operations.

The instruction set is summarised in figure 2.7. There are two basic formats, one for computational functions and one for organisational functions (control transfers, base register manipulations, etc.). They are distinguished by the 'type of function' bits and use four and six bits respectively to specify the function. The remaining nine or seven bits specify the 'primary' operand. The 'k' field indicates the kind of operand, and determines how the 'n' field is interpreted. Where the k field indicates the Extended Operand Specification, the instruction is extended by an additional 16 bits in the case of a name, or by 16, 32, or 64 bits in the

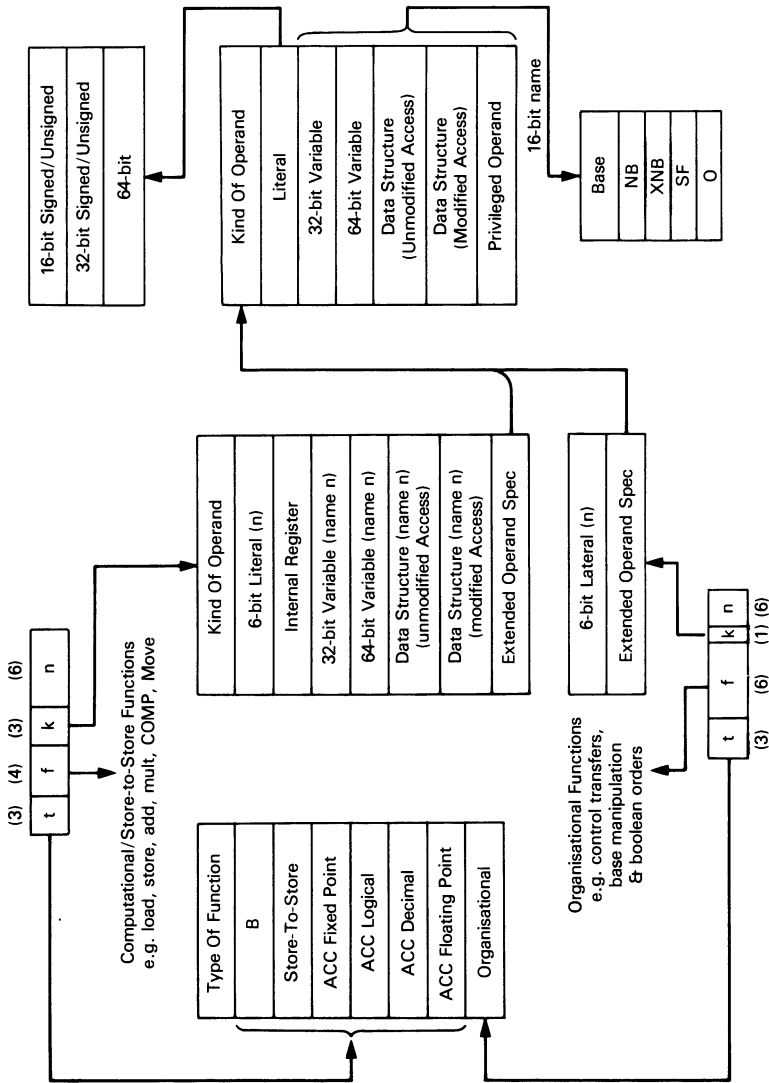


Figure 2.7 The MU5 Instruction Set

case of a literal.

The formal segmentation of the virtual address space, and the unique identification of store words belonging to different processes are achieved by splitting the address into three separate fields (thus producing 3-dimensional addressing), a 4-bit Process Number, a 14-bit Segment Number, and a 16-bit address defining a 32-bit word within a segment. Scalar variables are normally held in a single segment of a process at addresses formed by adding the 6 or 16-bit name in the instruction to the content of the Name Base register (NB). The name segment number and the process number (both held in special registers in the processor) are concatenated with these addresses to form the full virtual addresses required for store accessing. The value held in NB is unique to the name space associated with each procedure, and is altered at each procedure change. Access to names of other procedures and to common variables is made using the Extra Name Base (XNB), while using zero as base allows absolute addressing into the Name Segment (to access the Name Base values of different procedures, for example). The Stack Front register (SF) points to the area immediately beyond the name space of the current procedure, so that stacked variables are contained within the same address space as the names. SF is automatically incremented or decremented during the execution of a stacking function or a STACK operand access respectively.

Access to array elements is achieved through a descriptor mechanism. When a data structure element is specified by the k field, the corresponding primary operand is a 64-bit descriptor. This descriptor is sent to a D-unit within the processor, together with the content of the B register if the k field specifies modification. The D-unit interprets the descriptor and makes a 'secondary' store access for the element. It is this detachment of secondary addresses from instructions, and the use of names within instructions, which allow the full 30-bit virtual address space available to a process to be referenced by 16-bit instructions.

There are two main types of descriptor

String descriptor:	Ts	Length	Origin
Vector descriptor:	Tv	Bound	Origin
	8	24	32

String descriptors describe strings of bytes. The length field gives the number of bytes and the origin the byte address of the first byte. Strings are normally used in conjunction with

string processing orders (which may be two-address operations involving the use of two descriptors). These orders provide facilities for the manipulation of data records in languages such as Cobol, PL/1 and Algol 68 (move, compare, logically combine, etc.).

The size of element in a vector is specified within the type bits ( $T_v$ ) and may be between 1 and 64 bits. If the operand access is modified, the modifier must be greater than or equal to zero and less than the bound. This check is carried out automatically by hardware with virtually no time penalty, and removes the need for the costly software bound check which is often required, particularly during program development. The descriptor accessing mechanism will be discussed in more detail in Chapter 7 in relation to vector processing techniques.

## 2.6 COMPARING INSTRUCTION FORMATS

Comparisons of instruction formats are normally based on rather short sequences of hand-coded instructions which show the numbers of instructions required to obey various program statements. The expression

$$x := (a * b) + (c * d)$$

for example, could be coded in zero, one, two and three-address instructions (assuming all addresses are store addresses), as follows

<u>0</u>	<u>1</u>	<u>2</u>	<u>3</u>
LOAD a	LOAD a	TEMP = a	TEMP = a * b
LOAD b	* b	TEMP * b	x = c * d
*	=> TEMP	x = c	x = x + TEMP
LOAD c	LOAD c	x * d	
LOAD d	* d	x + TEMP	
*	+ TEMP		
+	=> x		
=> x			

Clearly the number of instructions to be obeyed decreases with increasing numbers of addresses in the instruction format, but if we assume that all function codes are 8 bits long and all addresses 16 bits long, then for this example the code space requirements are 18, 21, 25 and 21 bytes respectively for the zero, one, two and three-address instruction versions. Comparisons such as these tend to be misleading, however, (especially when operands may, or may have to be in registers rather than store locations) and they do not give an accurate

indication of the performance to be expected when complete programs are considered. A realistic comparison requires programs to be run on systems which have different numbers of addresses in their instruction formats, but are otherwise similar. Such machines do not exist in practice, but the ISPS architectural research facility developed at Carnegie-Mellon University [15] allows measurements to be made during the (simulated) running of programs on software models of different computer architectures, and hence allows for proper comparisons to be made.

ISPS is an implementation as a computer language of the ISP (Instruction Set Processor) notation, originally proposed by Bell and Newell [16] for documentation purposes. The process of using the ISPS facility involves writing a description in the ISPS language of the actions of each instruction, and the registers and store locations on which they operate. This description forms the input to the ISPS compiler, which runs on a PDP-10 computer and produces a PDP-10 program which is a simulator for the architecture under investigation. Programs and data can be loaded into the memory of this model, and the programs can then be executed. During execution the simulator keeps count of the number of bytes of information transferred in and out of the registers and store of the model, and the number of times labelled statements and procedures in the original description are executed. Various measurements can be derived from these counts, in particular the number of instructions obeyed by the model during the simulated execution of a program, and the number of instruction bytes transferred from memory (the dynamic code requirement).

This system has been used to evaluate both a two-address and a three-address version of the MU5 instruction set (designated MU5/2 and MU5/3 and having minimum instruction sizes of 24 and 32 bits respectively) against the one-address version [17]. Standard test programs were used in the evaluation, the Wichmann synthetic Algol scientific benchmark [18] and the Computer Family Architecture (CFA) test programs developed at Carnegie-Mellon [19]. The results of these experiments are summarised in figures 2.8 and 2.9.

	Ratio to MU5				
	MU5	MU5/2	MU5/3	MU5/2	MU5/3
Code Space (in bytes)	1914	2013	1804	1.05	0.94
Instructions Executed	9633	7121	4934	0.74	0.51
Bytes Executed	21282	22220	20592	1.04	0.97

Figure 2.8 Results for CFA Programs

The static code space requirements for the CFA programs run on MU5/2 and MU5/3 were not found to be significantly different from those for MU5, whereas the numbers of instructions executed were. The two-address version required roughly three-quarters as many instructions to be obeyed as the one-address version, and the three-address version only half as many. However, the two-address and three-address instructions are each proportionately longer than the one-address instructions, and counts of the numbers of instruction bytes executed showed that MU5/2 required 1.04 times as many bytes as MU5 and MU5/3 0.97 times as many. Clearly, neither of these figures differs significantly from the one-address figures, although the differences which do exist are in close agreement with the static code requirements.

	Ratio to MU5				
	MU5	MU5/2	MU5/3	MU5/2	MU5/3
Code Space (in bytes)	1411	1671	1720	1.18	1.22
Instructions Executed	107662	90527	75183	0.84	0.70
Bytes Executed	338854	386900	392082	1.14	1.16

Figure 2.9 Results for Wichmann Benchmark

For the Wichmann benchmark the code space requirements were found to be higher for both MU5/2 and MU5/3 with ratios of 1.18 and 1.22 respectively, and the numbers of instruction bytes executed were also found to be relatively higher than for the CFA programs, being 1.14 times higher for MU5/2 and 1.16 times higher for MU5/3. The numbers of obeyed instructions for the two and three-address versions again showed significant decreases relative to the one-address version, although these were less dramatic than for the CFA programs, with ratios of 0.84 and 0.70 respectively. The Wichmann benchmark contains many relatively long program statements with typically four arithmetic operators and an assignment operator (which can be coded more efficiently in one-address instructions than in two or three-address instructions) whereas the CFA programs are more typical of programs generally and contain a much higher proportion of short statements.

For both sets of measurements it is clear that the static and dynamic code requirements are approximately the same for all three systems, while the numbers of instructions obeyed show decreases as the number of addresses in the instruction format increases. In terms of overall system performance, however, these differences could only be translated into significant improvements if the instruction and operand

accessing rates could be improved correspondingly, since we can assume that the arithmetic speeds would be the same in each case. Thus the instruction byte accessing rate would need to be around 1.5 times higher for the two-address system and nearly twice as high for the three-address system. Furthermore, one three-address instruction clearly requires three operands, and there would be a significant engineering problem in providing three operands at a rate sufficient to maintain the same instruction execution rate as that of a one-address system. Given the same instruction and operand accessing rates, all three systems would have very similar performances, and we would be back where we started, with the intuitive feelings of the designer deciding the ultimate choice.

## 3 *Storage Hierarchies*

Storage hierarchies have existed for as long as stored program computers themselves. The Ferranti Mark 1 installed at Manchester University in February 1951 had a main store made up of eight Williams Tubes each containing 32 40-bit words and a drum backing store having a capacity of 3.75K words. Initially users were required to organise their own store transfers between a selected drum track and a selected pair of Williams Tubes, but the Mark 1 Autocode system introduced in 1954 carried out these tasks automatically on behalf of users and so made the two levels of storage appear to them as a one-level store. This arrangement was possible because the Mark 1 was a single user machine, and its performance was acceptable because the time required to transfer one drum track was roughly the same as the time for a programmed floating-point addition. Later developments in both architecture and technology led to the need for more sophisticated systems and in this chapter we shall consider the virtual memory and paging systems used in Atlas, the Cache stores used in some models in the IBM System/360 and System/370 ranges, and the MU5 storage hierarchy. Before dealing with these systems, however, we shall introduce the technique of store interleaving, used to obtain increased performance in these and many other computer systems.

### 3.1 STORE INTERLEAVING

Until parallel arithmetic hardware units were introduced in the late 1950s, the speed of random access storage devices was not a limiting factor on processor performance. The time for a floating-point addition in Mercury, for example, was 180  $\mu$ s, while the core store cycle time was 10  $\mu$ s. In Atlas, on the other hand, the time for a floating-point addition was 1.6  $\mu$ s, while the core store cycle time was 2  $\mu$ s. Thus, in the absence of some special technique, the time required to access both an instruction and its operand would have amounted to 4  $\mu$ s, leaving the floating-point unit idle for much of its time. Part of the solution to this problem was to overlap arithmetic and store accessing operations for successive instructions, a development which we shall consider in more detail in Chapter 4. In addition, however, the core store was made up of a



number of independent stacks, each with its own access circuits, so that provided successive accesses were to different stacks, one access would be able to proceed immediately after another without being held up for a complete store cycle. Had it been possible to ensure that instructions and operands always occupied different stacks, then given the appropriate degree of overlap, instructions could have been executed at a rate of one per core cycle. Even in the absence of this possibility, however, benefit could still be gained from the independent access facilities of the stacks. The addressing of words in the four stacks of the 16K word version of Atlas installed at Manchester University was arranged in the following manner

<u>STACK 0</u>	<u>STACK 1</u>
00000	00001
00002	00003
:	:
08190	08191
<u>STACK 2</u>	<u>STACK 3</u>
08192	08193
08194	08195
:	:
16382	16383

Since instructions more often than not follow in sequence, and since successive instructions occupied adjacent stacks in this arrangement, instructions could be accessed in pairs in Atlas, so that two instructions were accessed (and subsequently buffered) in one core cycle. Assuming a relatively small program occupying stacks 0 and 1, then the two operands required could come from any of the stacks, giving different limitations on performance. In a sequence of instructions in which the two operands accessed by each pair of instructions were all in either stack 0 or stack 1, for example, then the minimum time for the execution of the pair would be 6  $\mu$ s (2  $\mu$ s for the instruction pair access plus 2  $\mu$ s for each operand); for a sequence in which alternate operands were in stacks 0 and 1, the minimum execution time would be 4  $\mu$ s, and for a sequence in which alternate operands were in stacks 2 and 3, the minimum execution time would be 2  $\mu$ s (with the instruction pair and operand accesses fully overlapped). Assuming a random distribution of operands, however, we must consider all possibilities, as in the following table, and take an average to get an indication of the expected performance.

<u>OPERAND PAIR DISTRIBUTION</u>		<u>EXECUTION TIME</u>
0	0	6
0	1	4
0	2	4
0	3	4
1	0	4
1	1	6
1	2	4
1	3	4
2	0	4
2	1	4
2	2	4
2	3	2
3	0	4
3	1	4
3	2	2
3	3	4

The total of the execution times for these sixteen possibilities is 64  $\mu$ s, corresponding to an average of 4  $\mu$ s for the execution of a pair of instructions or 2  $\mu$ s per instruction. For sequences of computational instructions this figure was in fact achieved in Atlas. The overall average instruction time was higher, but this was due to the effects of control transfer instructions, as we shall see in Chapter 5, and not to store clashes. Thus, by having a store made up of four independent stacks with two-way interleaving of addresses, the execution time of computational instructions was improved by a factor of two.

In Atlas the single floating-point arithmetic unit was used not only for addition and subtraction, but also for the longer operations of multiplication and division (which took 4.9  $\mu$ s and between 10.7 and 29.8  $\mu$ s respectively), and the average rate at which this unit could execute arithmetic operations was well matched to the average store accessing rate. In the CDC 6600, on the other hand, the use of multiple functional units meant that instructions could be issued at the much higher maximum rate of one every 100 ns. The cycle time of main storage modules (or 'banks' in CDC terminology) exceeded the basic processor clock period by a factor of 10, however, and since the 24 computational registers offered a relatively small amount of buffering, a much higher level of store interleaving was required to sustain the instruction execution rate. This same ratio of processor clock period to store cycle occurs in the CDC 7600 (27.5 ns and 275 ns respectively) and in order to minimise the effects of store clashes on overall performance, the 60-bit word main stores of both these machines are 32-way interleaved.

Because of the factor of 10 difference between the processor clock and store cycle time, however, a maximum of 10 storage banks can be in operation at one time, and even this situation only arises during block copy operations between the Small Core Memory (SCM) and Large Core Memory (LCM) in the 7600 and between Central Storage and the Extended Core Storage in the 6600. In random addressing of the 7600 SCM for instructions, program data and input-output channel data, an average of four SCM banks in operation at one time is more normal. This means that there is a high probability that any one request will be to a free storage bank, and will be serviced immediately, and so much of the benefit of this storage arrangement comes not so much from interleaving per se, but from the fact that the store is made up of a large number of relatively small independent modules. The interleaving of these modules becomes important when many successive accesses are to sequential addresses, during instruction accessing, for example, and more importantly, during the processing of long vectors. Thus interleaving has been extended to even higher levels in the CDC CYBER 200 Series of vector processors. The 32-bit word storage banks of the CYBER 205, for example, are 256-way interleaved, and evidence for the necessity for such a high level of interleaving comes from the fact that reducing this interleaving to 128-way can cause a reduction in performance of nearly 40 per cent for some types of operation.

### **3.2 THE ATLAS PAGING SYSTEM**

In early computer systems peripheral transfers took place under direct program control using input/output instructions which normally required several milliseconds for their completion. By the time the Atlas computer was developed, however, the disparity between processor and peripheral speeds had become so great (of the order of 1000:1) that it was essential that peripherals be operated on a time division basis, with each peripheral interrupting the normal main program when it required attention. Since the interrupt routines controlling these peripherals were required to be resident in store concurrently with the main user program, and since the system was extended to allow several user programs to co-exist in store, it was no longer feasible either to optimise programmed transfers of information between the core and drum stores (so as to eliminate the drum access time of 6 ms) or even to allow users to program transfers within a shared physical store when they had no knowledge of other users' requirements. Consideration of these problems led to the idea of providing users with a virtual memory and to the development of a paging system to map virtual addresses on to the real physical storage devices. Thus the integration of the

Atlas core-drum combination into an apparently one-level store was a function of the machine itself, rather than a software facility, as on Mark 1 and Mercury.

### 3.2.1 Real and Virtual Addresses

In a virtual addressing system the programmer is allowed to use all available store addresses, but these virtual addresses do not refer directly to real store addresses, and each user programmer can assume that he or she is the sole user. In the hardware the virtual address space is considered to be split up into a series of contiguous pages of some convenient size (512 words in Atlas), which may be represented as follows

PAGE	0	1	2	3	4	5	6	7	8	. .
	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑
WORD	0	512	1024	1536	2048	. .				

Not all virtual pages contain information, since virtual addressing allows users to locate their code and data in areas which they find convenient, and which need not be contiguous. Thus the core and drum in Atlas might have contained, at some stage in the execution of a particular user program, the following virtual pages

CORE BLOCK	0	1	2	3	4	5	. .		
VIRTUAL PAGE	1	4	7	**	40	0	. .		
DRUM BLOCK	32	33	34	35	36	37	38	39	. .
VIRTUAL PAGE	12	13	3	9	2	6	5	41	. .

where \*\* indicates an empty core block. If the program called for a virtual address in page 7, for example, the hardware would have been required to translate this into a request to a real address in block 2 of the core store. This required the existence of a page table relating virtual and real addresses, as follows

<u>Virtual</u>	<u>Real</u>
0	5
1	0
2	36
3	34
4	1
5	38
6	37
.	.
.	.
.	.

### 3.2.2 Page Address Registers

In developing the hardware to implement the translation between virtual and real addresses, the designers of Atlas made the following observations

- (1) Every virtual address generated within the processor has to be translated into a real address by a table look-up, and this involves an extra contribution to the store access time.
- (2) If the number of virtual pages is large, the table is large and therefore slow to access.
- (3) The demanded virtual addresses must be in the core store most of the time if the system is to operate efficiently; a page on the drum must be moved into core before the processor can access individual words within it. Thus a small table containing only the entries corresponding to pages in core can be looked at initially, and the full table accessed only when the small table does not contain the required page address.

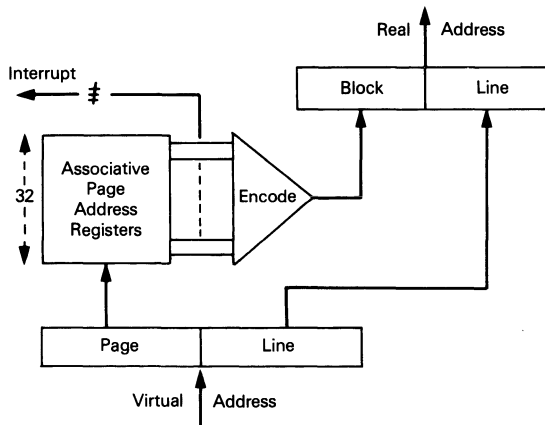


Figure 3.1 The Atlas Page Address Registers

The system of Page Address Registers used to implement this table is shown in figure 3.1. To locate a page, the 11 PAGE digits of the required address were presented to an associative store containing one register for each of the 32 blocks of core store. These registers were implemented in such a way that a check for equivalence between the presented PAGE bits and the contents of each register took place in parallel,

thereby minimising the address translation overhead.

If an equivalence occurred between the required page and a page address stored in one of the associative registers, that register indicated a '1' while the others indicated '0'. Thus for virtual page 7 in our example, line 2 would have contained the value 7 and given equivalence when an address in page 7 was requested. The 32 outputs from the associative registers were encoded into a 5-bit binary number and concatenated with the 9 LINE digits from the required virtual address to produce the real address to be sent to the core store.

If no equivalence occurred, the required page was not in the core store, and the program could not continue until it was. A page fault interrupt was therefore generated and the assistance of the operating system (known as the Supervisor in Atlas) invoked. The Supervisor accessed the full page table in order to locate the required page, initiated the necessary transfer from the drum into an empty core block, and set up the corresponding Page Address Register. It also initiated the transfer of an existing page from the core back to the drum in order to maintain an empty core block in readiness for the next page fault.

The choice of page to transfer out of core was made by a 'learning program' which attempted to optimise these transfers such that pages currently in use were not chosen for rejection. The operation of this program was assisted by the provision of a 'use' digit in each page register, which was set every time the corresponding page was accessed. All 32 digits were read and reset at regular intervals (normally every 1024 instructions), and their values added to a set of running totals maintained and used by the learning program. This program was the first example of a page rejection algorithm. Page rejection algorithms now form an important part of most operating systems and will not be considered further here. The interested reader is recommended to reference [20], for example.

Each of the Page Address Registers in Atlas also contained another digit, the 'lock-out' digit, which prevented user programs from accessing a page even though it was core-resident. This served two purposes. Firstly, it prevented programs from accessing pages to which drum, magnetic tape or peripheral transfers were taking place; with the lock-out digit set, access to that page was only permitted if the address had been generated by the drum or tape systems, or as a result of a peripheral transfer. Secondly, it prevented interference between programs co-existing in the core store; whenever the Supervisor caused a change from one program to

another, it protected the pages belonging to the suspended program by setting their lock-out digits. In the next Manchester University computer (MU5), this inter-program protection is extended by the inclusion of a 4-bit process number in each paging register. The current process number is held in a register within the processor and updated by the operating system at a process change. The contents of this register are concatenated with each address generated by the processor, so that addresses can only produce equivalence in those paging registers that contain the appropriate process number.

The Atlas virtual memory/paging system, and the concept of dynamic address translation which it introduced, has had a profound influence on the design of many other computer systems, among which are the DEC PDP-10, the IBM System/370 series and the CDC STAR-100. The full implications of virtual memory and its attendant advantages and disadvantages are still subjects of controversy. An excellent survey of the issues involved can be found in reference [21].

### **3.3 IBM CACHE STORES**

#### **3.3.1 The System/360 Model 85 Cache**

A different form of storage hierarchy from that used in Atlas involves cache storage, first introduced into the IBM System/360 series with the Model 85 [22]. Here all program-generated addresses are real addresses referring to main (core) store locations, and the semiconductor cache store, which is invisible to the programmer, is used to hold the contents of those portions of main storage currently in use by the program. Thus it is provided solely as a means of improving performance, rather than as a means of solving the problems involved in multi-programming, and may be thought of as a real address paging system.

The cache mechanism operates by dividing both cache and main storage into logical sectors, each consisting of 1024 contiguous bytes starting on 1 Kbyte boundaries. During operation, cache sectors are assigned to main storage sectors in current use, and a sector address register associated with each cache sector contains the 14-bit address of the main storage sector to which it is assigned (figure 3.2). Since the number of cache sectors (16-32, depending on the configuration) is smaller than the number of main storage sectors (512-4096, depending on the configuration), most main storage sectors do not have a cache sector assigned to them. However, the localised nature of store accessing exhibited by most programs means that most processor accesses are handled

by the cache (which operates at the 80 ns processor rate) rather than by main storage (which has a 1.04  $\mu$ s cycle time).

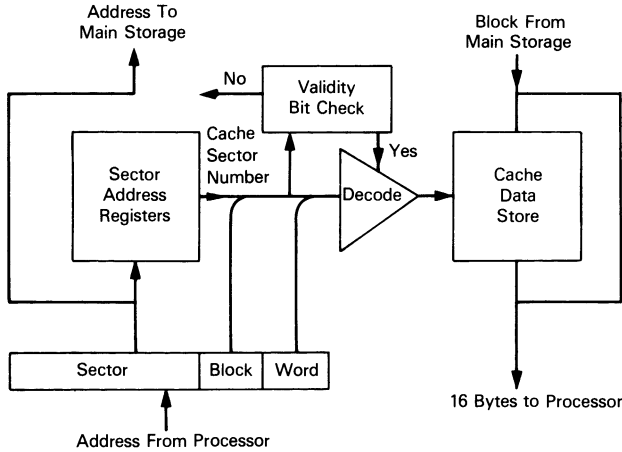


Figure 3.2 IBM System/360 Model 85 Cache Organisation

Each sector within the cache is made up of 16 blocks of 64 bytes and each block is marked with a 'validity' bit. Whenever a cache sector is re-assigned to a different main storage sector, all its validity bits are re-set, and the block containing the required store word in the new sector is accessed from main storage. The validity bit for this block is then set and the sector address register updated. Further blocks are accessed and their validity bits set as required.

The sector address registers constitute an associative store. Whenever an address is generated which requires an operand to be fetched from store, the sector bits within the address are presented for association. If a match occurs a 4-bit tag is produced indicating the cache sector address, and this is used in conjunction with the block bits to select the appropriate validity bit for examination. If a match occurs in the sector field, and the validity bit is set, the appropriate 16-byte word is read out in the next machine cycle and returned to the processor. Throughput is maintained at one access per cycle by overlapping association and reading for successive instructions.

If a match occurs in the sector field, but the validity bit is not set, the required block is read from main storage. Reading one 64-byte block involves one access to each of the four interleaved storage modules making up the main store, since each is 16 bytes wide. The delay experienced by the



processor as a result of these main storage accesses is minimised by always accessing the required 16-byte word first in the cycle of four, and by sending this word directly to the processor at the same time as loading it into the cache.

If a match does not occur, then a cache sector must be re-assigned to the main storage sector containing the failing address. A mechanism is therefore required to select a cache sector for re-assignment. The problem is similar to that involved in page rejection algorithms, except that in order to maintain performance, the algorithm must be implemented in hardware. The Model 85 cache implements a 'least recently used' algorithm by maintaining an activity list with an entry for each cache sector. Whenever a sector is referenced it is moved to the top of the list by having its entry in the activity list set to the maximum value, while all intervening entries are decremented by one. The sector with the lowest activity list value is the one chosen for re-assignment.

Re-assignment does not involve copying back to main storage values in the cache updated by the action of 'write-to-store' orders. Whenever such an order is executed, both the value in the cache and that in main storage are updated, a technique known as 'store-through'. Furthermore, if the word being updated is not held in the cache, the cache is not affected at all, since no sector re-assignment or block fetching takes place under these circumstances. While the store-through technique has the advantage of not requiring any copying back of cache values at a sector re-assignment, it also has the disadvantage of limiting the execution rate of a sequence of write orders to that imposed by the main storage cycle time.

Simulation studies of the Model 85 cache showed that the average performance of the cache/main storage hierarchy over 19 different programs corresponded to 81 per cent of a hypothetical system in which all accesses found their data in the cache. The average probability of an access finding its data in the cache for these programs (the hit rate) was 96.8 per cent. There do not appear to have been any subsequent measurements on real Model 85 systems to confirm or deny these predictions, but certainly a different arrangement was introduced in the System/360 Model 195 [23] and carried through into the System/370 Models 165 and 168. The average hit rate in the Model 195 cache was shown to be 99.6 per cent over a range of 17 job segments covering a wide variety of processing activities.

### **3.3.2 The System/370 Model 165 Cache**

The Model 85 cache is characterised by having a relatively

small number of fairly large contiguous units of buffer store. While this can perform well in a single program environment, it is unlikely to do so in a multi-access multi-programming environment where rapid switching between a number of store co-resident programs is essential. What is needed in these circumstances is a relatively large number of (necessarily) small contiguous units of buffer store. Smaller versions of the Model 165, for example, have 8-Kbyte cache stores consisting of 256 blocks each of 32 bytes. Accessing these blocks through a 256-word 19-bit associative store is not practicable in an 80 ns processor cycle, however, so cache and main storage are logically divided into 64 columns (figure 3.3).




Figure 3.3 IBM System/370 Model 165 Store Arrangement

Each column in the cache store contains 4 blocks and each column in main storage 512 blocks (in a 1-Mbyte version of the model). The block address registers constitute an address array made up of 64 columns, each containing four 13-bit

entries. Thus of the 19 bits of an incoming address used to access a block, the 6 low order bits are used to select a column, and the 13 high order bits are presented for association in the selected column of the address array. This arrangement is referred to as set associative mapping; two adjacent blocks from main storage contained within the cache will be referenced by identical entries in adjacent columns of the address array.

A validity bit associated with each block indicates whether or not an address giving a match on association refers to valid data; if not the required data is accessed from main storage. Apart from the factor of two difference in store widths, data paths and block sizes, the arrangements for transferring data between the cache and main storage are identical in the 370 Model 165 to those in the 360 Model 85. The cache replacement algorithm is also the same, except that in the 370 Model 165 a replacement array is required, containing a separate activity list for each of the 64 columns.

### 3.4 THE MU5 NAME STORE

A quite different approach to high-speed buffering was taken in the design of the MU5 computer. As we saw in section 1.3, the use of fast programmable registers in MU5 was rejected in favour of a small associatively addressed buffer store containing only named scalar variables and forming part of a one-level store with the Local Store of the processor [24, 8]. Simulation studies of this 'Name Store' indicated that a hit-rate of around 99 per cent would be obtained with 32 words of store, a number which it was technologically and economically feasible to construct and operate at a 50 ns rate, and which could be conveniently fitted on to two of the platters used in the construction of MU5.

Thus the address and value fields of the Name Store (figure 3.4) each occupy one of these platters, and form two adjacent stages of the Primary Operand Unit (PROP) pipeline, through which instructions move in a series of beats (section 4.1). At each beat a virtual address generated in the previous two stages of the pipeline is copied into the Interrogate Register (IN), and concatenated with the contents of the Process Number register (PN), for presentation to the address field of the Name Store. A full virtual address in MU5 consists of a 4-bit Process Number, a 14-bit Segment Number and 16 bits which identify a 32-bit word within a segment. Addresses presented to the Name Store do not contain the Segment Number, however, since it was assumed at the design stage that the Name Segment would always be zero, and only 15 of the word address bits are

used, referring to 64-bit operands. Where necessary, a 32-bit operand is selected from within a 64-bit word in a later stage of the pipeline.

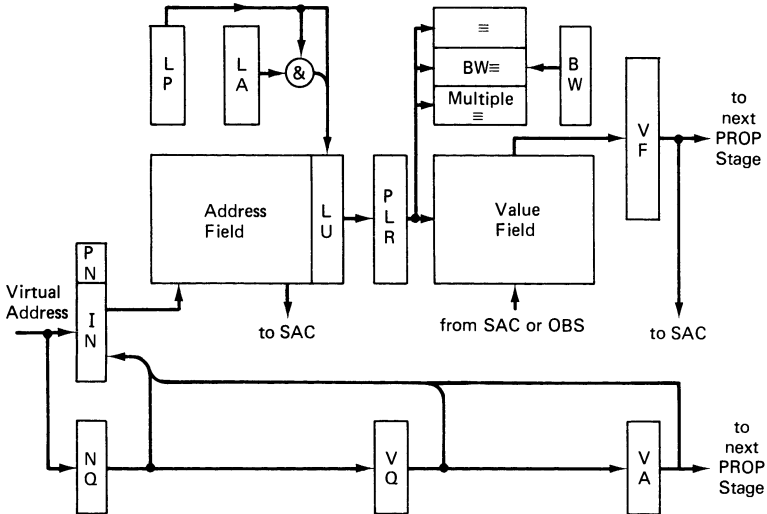


Figure 3.4 The MU5 PROP Name Store

### 3.4.1 Normal Operation

If the presented virtual address gives a match in the associative store, and the corresponding Line Used digit (equivalent to the 'validity bit' in an IBM cache store) is set to 1, then an equivalence has occurred, and on the next pipeline beat a digit is set in the PROP Line Register, PLR. The digit in PLR then selects a register in the Value Field, and the 64-bit word is read out and copied into the Value Field Register (VF) by the next beat. At the same time, checks are made to determine whether

- (1) an equivalence occurred
- (2) a 'B Write Equivalence' occurred
- (3) multiple equivalence occurred

The check for equivalence simply requires an OR operation on the digits in PLR. If no digit is set, however, this indicates non-equivalence, and the Name Store is updated by transferring a new word into it and discarding an old one. It would clearly be inefficient to enter software to organise

this one-word transfer, so the transfer is controlled directly by hardware. A 'B Write Equivalence' occurs if the line giving equivalence is the target line of an outstanding 'B => name' order, and this causes a pipeline hold-up (section 4.1.2) until the operand value has been returned from the B-unit (figure 3.5) and written into the Name Store. A multiple equivalence occurs if a hardware malfunction in the associative field causes more than one line to give equivalence. In this case an interrupt is generated.

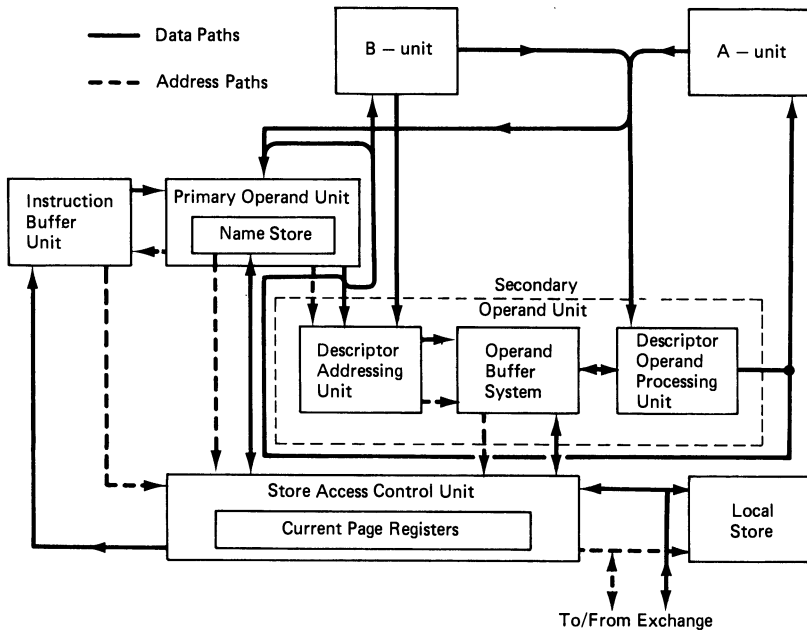


Figure 3.5 The MU5 Processor

When a Name Store entry is replaced the hardware must take into account the effect of store orders. To maintain the speed advantage of the Name Store, store orders only update the value of an operand in a Name Store, rather than operating on the store-through principle used in IBM cache stores. The Name Store contains values which would be held in registers in IBM machines, and it would clearly be inappropriate for these to be written back to the main store each time they were altered. Thus the old word may have to be copied back to the Local Store before it is overwritten. The decision concerning which line to replace requires the use of a replacement algorithm, and the effects of various replacement algorithms were studied by simulation before the Name Store was built. These varied from a simple cyclic replacement algorithm, requiring a minimum of additional hardware for its implementation, to a

multiple-use digit algorithm requiring, for a 32-line store, 32 5-digit counters. Very little difference in performance was found among these different algorithms, and the simple cyclic one was therefore chosen.

The actions which take place when a non-equivalence is detected also depend upon whether the instruction is an Accumulator (ACC) order destined for the A-unit and whether the required operand is already in an additional Name Store situated close to the A-unit. In order to improve the accessing of vector elements in MU5 (section 7.1), the Secondary Operand Unit (SEOP) incorporates an Operand Buffering System (OBS) in addition to the two parts of the D-unit (the Descriptor Addressing Unit and Descriptor Operand Processing Unit in figure 3.5). As a result there are some ten pipeline stages between the end of PROP and the A-unit through which all instructions destined for the A-unit must pass (in order to maintain the correct program sequence). Thus if a name held in the PROP Name Store were to be used to accumulate a total calculated by ACC orders in a small program loop, the order reading the total from the Name Store would have to be held up until the value calculated by the previous pass through the loop had been returned. The solution to this problem was the provision of 24 lines of Name Store in OBS.

The OBS Name Store is meant to keep names used by ACC functions and the PROP Name Store to keep those used by non-ACC functions. Thus for a non-ACC order the normal situation is for equivalence to occur in the PROP Name Store, while for an ACC order the normal situation is for a non-equivalence to occur in the PROP Name Store and equivalence to occur in the OBS Name Store. However, the same name might be accessed by both kinds of orders, and the hardware must guard against the possibility that a name is in the 'wrong' Name Store.

### 3.4.2 Non-equivalence Actions

When a PROP Name Store non-equivalence occurs a 'non-equivalence' digit is set in the function register associated with VF. This causes the normal control signals decoded for the instruction in the next pipeline stage to be overridden. As a result, a special order is sent through the pipeline to OBS carrying the non-equivalence address (including the Name Segment number). OBS accesses its Name Store, and if it finds equivalence, it returns the 64-bit store word to PROP. If it does not find equivalence, it makes an access via the Store Access Control Unit (SAC) to the MU5 Local Store on behalf of PROP, so that the 64-bit store word will be returned direct to PROP from the Local Store.

In PROP the occurrence of a Name Store non-equivalence causes the initiation of the next pipeline beat to be held up until the appropriate actions have been completed (figure 3.6). The first actions are the preparation of a line in the Name Store to receive the new address and store word and the copying of the new address, currently held in register VA, into the Interrogate Register (IN). The Name Store line to be replaced is indicated by the Line Pointer Register (LP), which contains one bit corresponding to each line in the Name Store and has, at any time, only one digit set to a 1. Thus LP is simply copied into register PLR in order to select the line for replacement. Two possible conditions are checked, however, and if necessary dealt with, before the line is ready to be overwritten with the new address and value. The first is that the selected line may be the target line of an outstanding 'B => name' order (section 4.1.2). If it is, LP is moved on to select the next line. This is done by first copying the content of PLR back into LP, and then copying LP into PLR. The outputs from PLR are routed back to the Line Pointer with a 1 digit shift, thus implementing the simple cyclic replacement algorithm. The second condition is that the contents of the selected line may have been altered by the action of a store order. This is checked using the appropriate digit in the Line Altered register. If it has, the virtual address and value are read out, and a write request is made to SAC. The selected line is then ready for overwriting.

The next action which occurs depends on whether the store word is returned from SAC or OBS, or on whether a CPR (Current Page Register) non-equivalence occurs, in which case no store word is returned, but the interrupt sequence is entered instead. If the store word comes from SAC, then when it arrives the address and value are written into the Name Store, the content of PLR is copied back into the Line Pointer in preparation for the next non-equivalence, and the corresponding bits in LU and LA are set to 1 and 0 respectively.

Although the actions needed to update the Name Store are complete at this point, the contents of the PROP Line Register and Interrogate Register no longer correspond to the orders in the pipeline, and must be restored. PLR is restored first by copying the address in register VQ into IN, and after a delay to allow for association, PLR is re-strobed. (Preserving a previous copy of PLR for use at this time would not be satisfactory, since it might be the line newly overwritten.) The address in register NQ is then copied into IN, and the actions are complete. These complications could have been avoided if a longer PROP pipeline beat time had been adopted in the design.

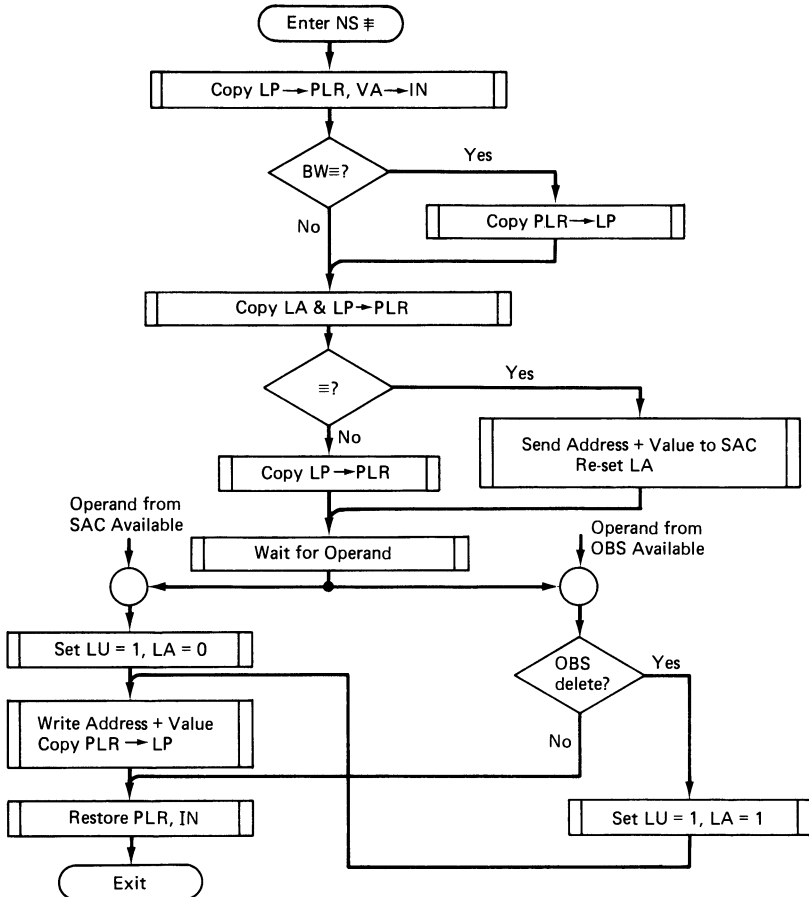


Figure 3.6 The Name Store Non-equivalence Routine

In the case where OBS indicates that the required store word is in its Name Store, then if all 64 bits are to be overwritten by a store order, the store word is not actually returned to PROP. In this case OBS simply deletes its copy and PROP assumes it to have arrived immediately, thereby reducing the non-equivalence delay time. In cases where the operand is not deleted from the OBS Name Store, then the only further action required in PROP is the restoration of PLR and IN before the routine is complete. (The clearing of a line earlier in the routine only copies its content back to the Local Store and sets it 'unaltered', so the Name Store is not disturbed.) If the operand has been deleted from OBS and returned to PROP, the actions are as for an operand from SAC.



### 3.4.3 Actions for ACC Orders

Although the normal situation for an ACC order using a named operand is for non-equivalence to occur in the PROP Name Store and equivalence to occur in the OBS Name Store, equivalence may be found in the PROP Name Store on some occasions. For example, a 32-bit variable required for an ACC order might be contained in the same 64-bit store word as a 32-bit variable already in use by B orders. Unless the ACC order finding equivalence is a store order, the operand is read out as for a non-ACC order and carried through to OBS in the same way as a literal operand, so that no access is made to the OBS Name Store. If the ACC order finding equivalence is a store order, however, (ACC Write Equivalence) then action is taken to delete the word from the PROP Name Store.

The actions for an ACC Write Equivalence are initiated in a similar way to those for a non-equivalence for a non-ACC order. These actions are similar to those involved in preparing a line for a non-equivalence and in restoring the pipeline at the end. The appropriate address and value are read out of the Name Store and sent to SAC, and the appropriate digits in the LU and LA digits are then set to zero, in order to mark the line empty. When the order reaches OBS its operand will not, of course, be found in the OBS Name Store, and the normal OBS non-equivalence actions will be initiated.

### 3.4.4 Name Store Performance

The management of the Name Stores of MU5 is clearly complex by comparison with the cache approach. This complication was thought to be acceptable because it was expected to give rise to a high 'hit-rate' and a high rate of instruction execution between 'misses'. Measurements of Name Store performance were made for a set of 95 programs containing both Fortran and Algol jobs ranging in complexity from simple student exercises to large scientific programs. For most programs it was found that 80 per cent ( $\pm 5$  per cent) of operand accesses were to named variables, that no more than 120 names were used in any one program, and that in all programs 95 per cent of name accesses were to fewer than 35 per cent of the names used. These figures confirmed the Atlas results which inspired the idea of using a Name Store, but the figures for hit-rates were not as good as had been anticipated. Thus although 96.1 per cent of name accesses found their operands in one or other Name Store, only 86 per cent of name accesses found their operands in the correct Name Store.

The comparatively high and largely unforeseen amount of

interaction between the two Name Stores was found to be due to the procedure calling mechanism used in MU5. Parameters are normally passed on by stacking into the PROP Name Store, but in many cases may subsequently be used with ACC orders. Conversely, it is possible for a particular word required to be in the OBS Name Store in one procedure to be required to be in the PROP Name Store in a subsequent procedure. It is by no means clear what is the best solution to this problem.

### 3.5 DATA TRANSFERS IN THE MU5 STORAGE HIERARCHY

Operand accesses which cannot be satisfied by the buffer stores within the MU5 processor cause requests to the Store Access Control Unit (SAC). SAC contains 32 Current Page Registers (CPRs) which translate the presented virtual address into a real address to be sent to the Local Store. Current Page Registers differ from Page Address Registers as used in Atlas, for example (section 3.2), in that they do not provide full cover of the addressable real store, and do not, therefore, correspond one for one with the pages in real store. Thus whereas in Atlas a real address was obtained by encoding the 32 Page Address Register outputs into a 5-bit binary number, this cannot be done with CPRs. Instead there is a real address register corresponding to each virtual address register, and a real address is obtained from the CPRs in exactly the same way as an operand is obtained from the Name Store. Unlike the Name Store, however, software action is invoked in the event of a CPR non-equivalence (frequently referred to as a page fault). Either a block of data must be moved into the Local Store and a CPR set up to address it, or a CPR may simply be set up to address an existing Local Store page.

The MU5 'one-level' store extends over two further levels of real store, a 2.5  $\mu$ s cycle-time Mass (core) Store and a Fixed-head Disc Store. Pages are brought into the Local Store from either the Mass Store or the Fixed-head Disc Store, but are rejected first to Mass and later to the Fixed-head Disc. The 'one-level' store contains all the working space and files needed by current terminal users and active background jobs. Most of the files, however, are stored on large capacity discs attached to a separate PDP-11/10 minicomputer.

#### 3.5.1 The Exchange

In order to provide a completely general and flexible interconnection scheme, allowing for data transfers between all these various stores and processors, an 'Exchange' was developed as part of the MU5 project [25, 8]. Logically it is a multiple-width OR gate operated as a packet switching system

at the star point of the interconnections. This configuration involves only a very short common path for transfers between the various units, allowing a much higher data rate than would be possible with a distributed highway or 'bus' system. Thus transfers can occur at a rate of one every 100 ns, and each can involve a 64-bit data word together with address and control bits.

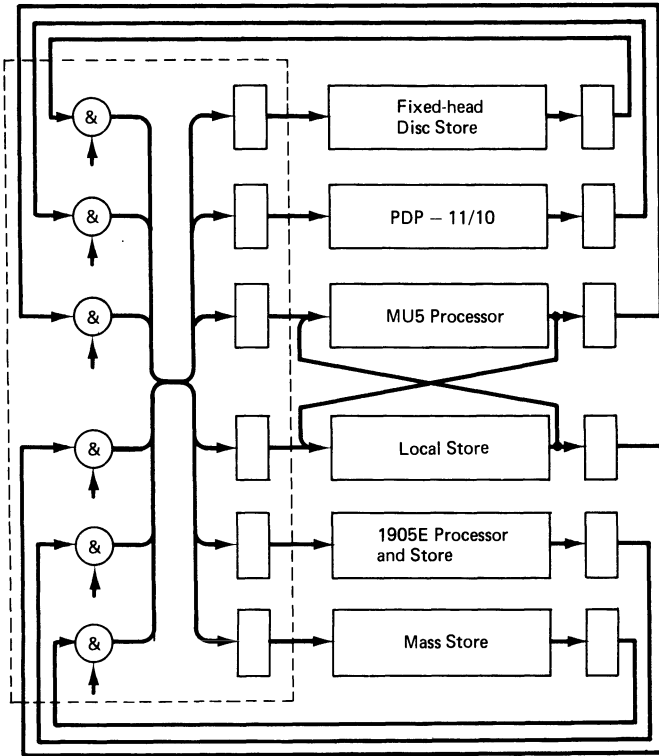


Figure 3.7 The MU5 Exchange System

Each unit attached to the Exchange (figure 3.7) provides a set of parallel inputs to the OR gate and each is connected, via its own buffer register, to the output of this OR gate. The Exchange operates by time sharing the OR gate between the units. Thus the transfer of a block of words from the Fixed-head Disc Unit to the MU5 Local Store, for example, involves a succession of 64-bit word transfers from the Disc, as 'Sending' Unit, to the Local Store as 'Receiving' Unit, with the OR gate connecting these units for the duration of each word transfer rather than for the whole duration of the block transfer. Other transfers can therefore be accommodated during this period, so that the 1905E computer, for example,

can make read requests to the Mass Store. Two transfers are required for a read request, one in which, in this case, the 1905E as Sending Unit sends the address and appropriate control information through the Exchange to the Mass Store as Receiving Unit, and a subsequent 'data available' request from the Mass Store as Sending Unit to the 1905E as Receiving Unit in order to return the data read out from the specified location in the Mass Store.

The requests from different Sending Units arrive at the Exchange completely asynchronously, and much of the control logic within the Exchange is concerned with scheduling transfers through the OR gate on a priority basis (section 3.5.3). A substantial proportion of these transfers are paging transfers between the Local and Mass Stores. Since these are both passive units, the transfers are organised by a Block Transfer Unit (BTU). The BTU is activated by the MU5 processor, which sends to one of its four channels the starting address for the transfer in each store, together with the block size and a start command. The processor is then free to continue computation, while the BTU generates the necessary requests, via the Exchange, to the Mass and Local Stores to carry out the transfers. At the end of the block transfer the BTU activates a processor interrupt.

The width of the data path through the Exchange OR gate is 8 bytes (plus one parity bit per byte) corresponding to the width of the data paths within the MU5 processor, and also exceeding the minimum width necessary for some of the units to be able to communicate at all. The Fixed-head Disc, for example, has an effective data rate of 0.5  $\mu$ s/byte, while the storage modules constituting the Mass Store have a cycle time of 2.5  $\mu$ s, and both these devices and the data path between them must therefore be capable of dealing with at least 5 bytes per transfer for communication to be possible.

The address path contains 27 bits (including 3 byte-parity bits and one bit which distinguishes user and system addresses). The control field contains some information which is copied directly through the OR gate from the Sending Unit to the Receiving Unit, some information which is copied through the OR gate and is also used by the Exchange control logic, and some information (the Unit Number) which is transmogrified by the control logic before being sent to the Receiving Unit. In all, 14 control digits pass between units via the Exchange, making the total width of the OR gate 113 bits, and some additional control signals pass between the Exchange control logic and each of the units. The Strobe Outwards (SO) signal (figure 3.8), for example, is sent from a Sending Unit to the Exchange to initiate a transfer, and timed

to arrive as soon as the data, address and control information in the Sending Unit output buffer have become valid at the Exchange (allowing for cable length, etc.). This output buffer is required because, at the time SO is sent, either the Receiving Unit may not be free, or a higher priority transfer may be in progress, and the Sending Unit therefore has no means of knowing when the transfer will actually occur.

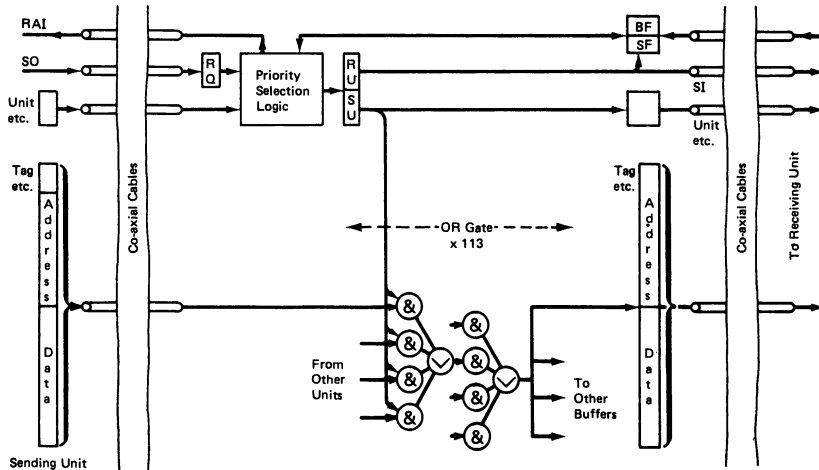


Figure 3.8 Exchange Control and Data Paths

The timing control logic of the Exchange itself is governed by a free running oscillator, so that the Exchange operates synchronously, at a rate of one transfer per 100 ns. Each transfer requires two 100 ns periods or 'slots' for its completion, one for the actual transfer through the Exchange OR gate, and a previous one in which the Exchange control logic determines which of the incoming requests to service. Within the Exchange these two activities are overlapped for successive transfers. When a request has been selected for servicing in one time slot, the information from the Sending Unit is gated into the OR gate in the next time slot by the appropriate decoded output from the Select Unit Register (SU in figure 3.8). The output signals from the OR gate then propagate to the input buffers of all the Units but only the buffer corresponding to the Receiving Unit of the current transfer is strobed, at the end of the slot. A Strobe Inwards pulse (SI) is then sent to the appropriate Receiving Unit, thereby completing the transfer as far as the Exchange is concerned. The Receiving Unit, on receipt of SI, deals with the data in its buffer at its own convenience and then returns a signal to the Exchange indicating that its buffer is free to

be overwritten by a further Exchange transfer.

### **3.5.2 The Exchange Priority System**

In any time-shared communication system to which a multiplicity of source units is connected there has to be some mechanism for establishing the order in which simultaneous requests are serviced. In the case of the MU5 Exchange the units can be classified into four categories, Peripheral Processing Units (PPUs), Central Processing Units (CPUs), Stores (Mass and Local), and the Block Transfer Unit. PPU's have highest priority since they are generally concerned with organising transfers which involve crisis time devices. CPUs normally maintain an intense traffic to and from their own local stores, via dedicated highways, and only make occasional requests to stores via the Exchange. Apart from PPU transfers, most of the store transfers are paging transfers between the Mass and Local Stores organised by the Block Transfer Unit. Since this Unit can control up to four block transfers simultaneously, it can easily saturate the Mass Store, and although CPUs have a crisis time extending to infinity, it would be unreasonable to hold up their requests for the duration of a block transfer. Thus CPUs have the second highest priority and the Block Transfer Unit has the lowest.

Exchange priorities are assigned in inverse proportion to the Unit Number of the port to which a device is attached, so that the Fixed-head Disc, for example, attached as Unit 0, has highest priority. The priority of a request is normally determined by the priority of the Sending Unit, but in the case of a 'data available' request occurring in response to a read request from a crisis time PPU, the priority of the request is determined by that of the Receiving Unit. All transfers to or from the Fixed-head Disc, for example, are organised by the Disc's own block transfer unit, and must all have top priority. Failure to transfer a request to or from the Disc within its 4  $\mu$ s crisis time would cause the whole block transfer to have to be aborted and re-started.

## 4 Pipelines

In any computer the execution of a single instruction requires various activities to be performed, such as instruction accessing, instruction interpretation, operand accessing and arithmetic. If separate hardware units carry out these activities their operations can be overlapped to give an increased rate of completion of instructions. This technique, first introduced in computers such as Atlas and Stretch, has become known as 'pipeline concurrency'. In a pipeline computer several partially completed instructions are in progress concurrently, and although the time to complete any one instruction is still limited by the sum of the times for the various activities, the rate at which instructions progress through the pipeline is only limited by the time for an individual activity. In Atlas and Stretch the number of concurrent operations was of the order of four. In more recent designs the pipeline concurrency principle has been extended to several tens of instructions and used in both arithmetic and instruction processing units. In this chapter we shall consider the design of the MU5 Primary Operand Unit as an example of the latter and the Texas Instruments Advanced Scientific Computer (TI ASC) as an example of the former. In addition we shall consider some techniques used in the IBM System/360 Model 195 to overcome problems which arise in pipeline designs.

### 4.1 THE MU5 PRIMARY OPERAND UNIT PIPELINE

We saw in sections 1.3 and 3.4 how processor performance considerations led to the adoption of a pipeline structure for the MU5 Primary Operand Unit (PROP). PROP is concerned with accessing the operand specified directly by an instruction (the primary operand) and routing the instruction, together with its primary operand, to the appropriate following unit for execution or further processing. If the primary operand is a named variable or literal, for example, the instruction is ready for execution at the end of PROP, whereas an instruction involving a data structure must be sent to the Secondary Operand Unit (SEOP) where the accompanying descriptor, accessed as a primary operand by PROP, is used to specify the secondary access for the data structure element.

Figure 4.1 shows the basic hardware required in the Primary Operand Unit to implement the primary operand accessing facilities of the instruction set described in section 2.5 and the various stages of operation involved in processing a typical instruction. Instructions are received from an Instruction Buffer Unit (IBU) into registers DF (function) and DN (name). The first action is the decoding of the instruction to select the content of the appropriate base register (NB, XNB or SF) and the name part of the instruction. For access to a 32-bit variable, the name is shifted down one place relative to the base and the least significant digit is used later to select the appropriate half of the 64-bit word obtained from the Name Store.

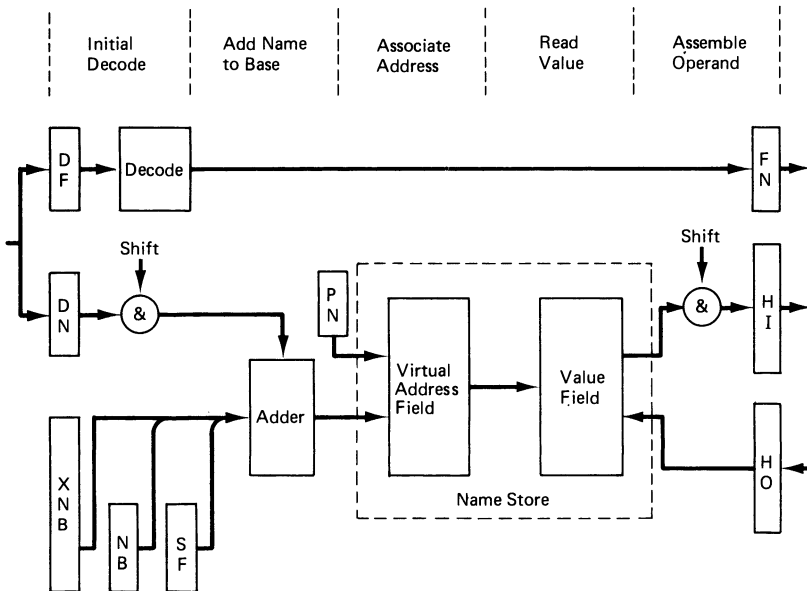


Figure 4.1 The Basic Components of the Primary Operand Unit

In the second stage the name and base are added together to form the address of a 64-bit word. This address is concatenated with the 4-bit Process Number (PN) and presented to the Associative Field of the Name Store in the third stage. If the address is found, access is made in the fourth stage to the 64-bit word in the Name Store Value Field. (If the required address is not in the associative store, the non-equivalence actions described in section 3.4.2 must be initiated.) The fifth stage of processing is the assembly of the operand. A 32-bit integer, for example, may be taken from



either half of the 64-bit store word, but must always appear at the least significant end of the data highway when presented to a succeeding unit. Registers FN and HI form the input to the highway which links PROP to other units in the processor (figure 3.5) and register HO is connected to one of the outputs of this highway in order to receive operands resulting from store orders.

The pipelining of the five stages in PROP is achieved by staticising the information obtained at the end of each stage in a buffer register (figure 4.2). An important aspect of the design of a pipeline is the timing of the strobes to the buffer registers, because some outputs from a stage (the function bits, for example) will derive directly from its inputs. With edge-triggered or master/slave flip-flops the various registers could be strobed simultaneously, but in the D-type flip-flops used in MU5 the information on the D input is copied through to the output as long as the clock is held low and a different technique must therefore be used. Thus the result obtained at the end of any one stage is only copied into its buffer register when the result of the following stage has itself been staticised. The strobes used to copy information into the buffer registers are therefore staggered, as shown in figure 4.3.

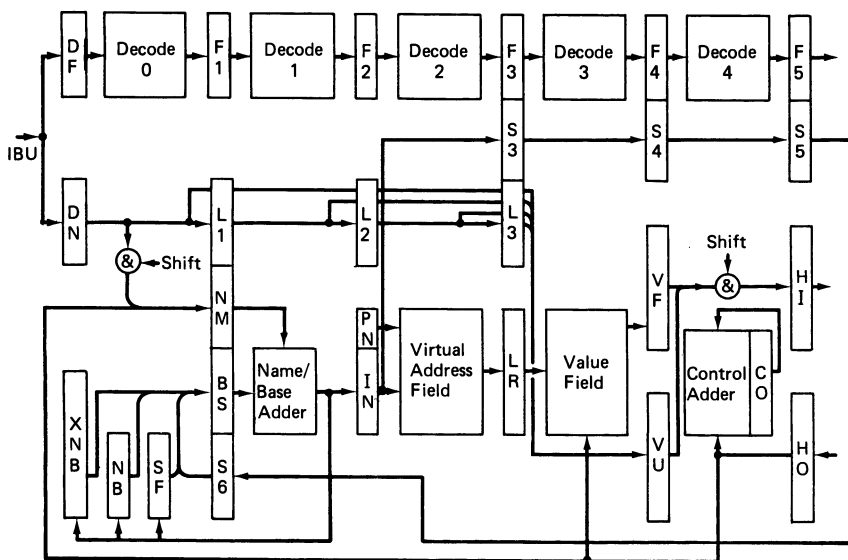


Figure 4.2 The Complete Primary Operand Unit

The shaded portions of figure 4.3 show the progress of one instruction through the PROP pipeline. It is first copied into DF and DN (function and name respectively) and the Decode 0 logic carries out the decoding of the instruction necessary to control the first stage. The decoding logic of figure 4.1 is spread out in the pipelined version into separate decoders for each stage. In many cases, however, the necessary decoding cannot be carried out in sufficient time to control the action of a given stage. In these cases it is carried out in the previous stage, and the various control signals appear as additional function digits, along with the original function, as inputs to the stage requiring them.

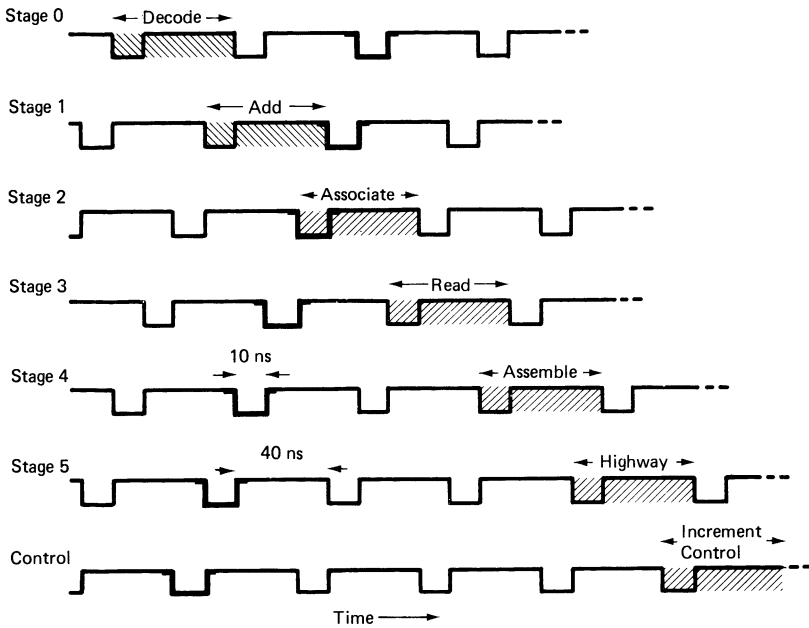


Figure 4.3 The Basic PROP Timing Diagram

The next pipeline strobe is timed to arrive no earlier than when the outputs of the first stage have settled and are ready to be strobed into the registers F1 (function), NM (name) and BS (base). The addition of name and base now takes place and when the next pipeline strobe arrives, the result is copied into IN, the Interrogate Register. The output of IN is concatenated with PN, the Process Number, to form the input to the associative field of the Name Store. The result of the association is then copied into the PROP Line Register (PLR), the output of which accesses the line in the Value Field containing the required operand. The Value Field output is

copied into the Value Field register (VF), and thence, after assembly, into the Highway Input register (HI).

Once an instruction has reached HI, PROP must wait until it has been accepted by another unit before taking any further action. When it has been accepted the Control Register is updated for that instruction and all other instructions in the pipeline move along one stage. Instructions therefore proceed through PROP in series of 'beats', the rate at which these beats occur being determined by the maximum operating rate of PROP and the acceptance rates of the succeeding units. The generation of a beat is initiated by the setting of a 'data gone' flip-flop (figure 4.4) which when any other necessary conditions at the end of the PROP pipeline have been satisfied, allows a 20 ns pulse to propagate through the pipeline delay chain. The pulses from the chain drive 10 ns pulse-forming modules at each stage, and the delays in the chain are adjusted to produce a 10 ns stagger between stages. The progress of one pulse through the pipeline is shown heavily drawn in figure 4.3.

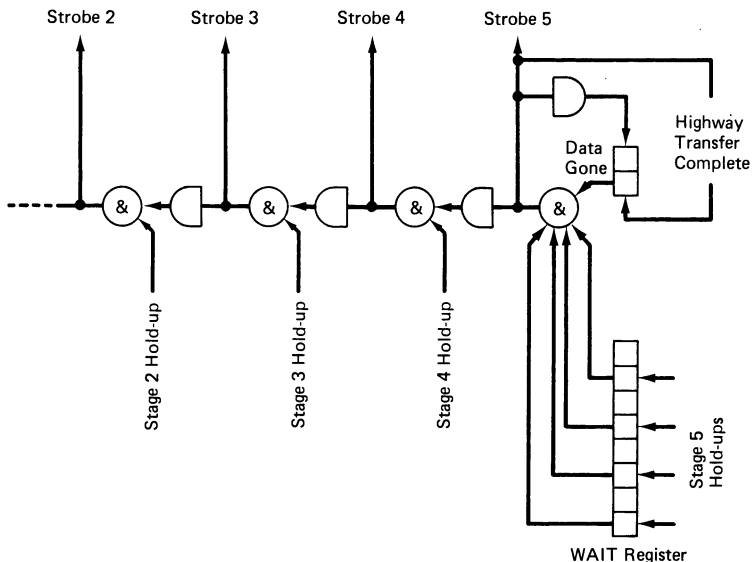


Figure 4.4 The Pipeline Delay Chain

Some problems arise as a result of the physical dimensions of the Processor and the layout of the platters within it. Thus it is not feasible to locate all registers pertaining to one stage in close proximity either to each other or to the timing control logic. As a result, all the registers of one stage cannot be strobed simultaneously, since 'far' strobes

would have to be sent out in advance of 'near' strobes by up to 20 ns. Alternatively, all 'far' registers could be strobed late (the strobes and data must all travel the same distance) but some control signals must travel back from the far registers to the near registers, and the double delay would slow the pipeline down. The problem is overcome in practice by deriving far strobes from the earliest level of fan-out in each section of the delay chain, and by designing for only three levels of logic in paths which involve data travelling from near to far registers. Thus the 50 ns within each stage is typically made up as follows

Input Buffer Settling Time	5 ns
Operation Within Stage	30 ns
Inter-platter Cable Delays	5 ns
Output Buffer Strobe	<u>10 ns</u>
TOTAL	50 ns

#### 4.1.1 Synchronous and Asynchronous Timing

PROP is an example of a synchronous pipeline in which the operation time of each stage is the same and the time relationship between the buffer register strobes is fixed. However, it is asynchronously connected to other units in the processor, and a 'handshake' system between PROP and the B-unit, for example, determines when the next PROP beat is to be generated once a B-order has entered the PROP output registers. Figure 4.5 shows the hardware used to implement a typical handshake system between two units and a schematic timing diagram of its operation. Information is passed from one unit to another when the sending unit has the data available and when the receiving unit is not busy. For a given rate of operation, the cable length  $L$  cannot exceed a certain maximum. Assuming a delay of 5 ns per circuit and a co-axial cable delay of 1.8 ns per foot, for example, the maximum distance between units for a 50 ns operating rate is approximately 8 feet. This might seem long, but in a large system such as MU5 cable lengths between units can easily be of this order, and the cable length between PROP and the B-unit is in fact almost twice this figure. A double handshake loop which relies on the presence of an input buffer in the B-unit is therefore used to link these two units.

In MU5 all interactions between the separate functional units are asynchronous. This type of operation allows information transfers to take place as and when required and offers greater speed than a completely synchronous system (where transfers only take place at fixed times), particularly when the units concerned are not heavily used or where the

different operations carried out within a particular unit take significantly different times (add, multiply and divide, for example, in the floating-point unit). Where a number of communicating pipeline stages are heavily used, however, the time penalty incurred by the handshake becomes a dominant factor, and a synchronous system such as that used in PROP may be preferred. In many high performance systems (the CRAY-1 for example) the central processor is completely synchronous.

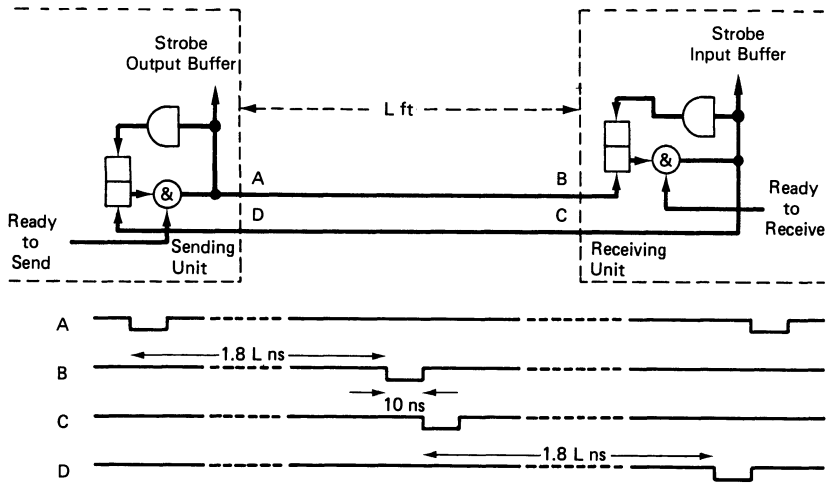


Figure 4.5 An Asynchronous Handshake System

Asynchronous systems can also give rise to additional problems, particularly when a unit can accept requests from a number of sources. In MU5 this problem occurs, for example, at the input to the Store Access Control (SAC). Here three different functional units may request a store cycle at any time, and in the event of a clash, SAC must decide which request to accept. Each incoming request pulse is staticised on a flip-flop, and the outputs of these three flip-flops drive a combinational priority logic circuit. The outputs from this circuit, only one of which can be a 1, form the inputs to a set of decision flip-flops. These flip-flops are strobed when SAC is free to accept a request and sufficient time has elapsed for the priority circuit to have settled after receipt of the first request.

Because a second request may occur a short time after the first request, however, it is possible for the inputs to the decision flip-flops to change state just before the end of the strobe, leaving the outputs somewhere between a 0 and a 1 level. Under these conditions the time taken for the flip-flop

outputs to reach proper logic levels may be long compared with the normal propagation delay, and it is possible for subsequent circuits to operate inconsistently. Sufficient time must therefore be allowed for the decision flip-flops to settle to a constant level, or failure of the control circuits may occur [26]. For ECL flip-flops with a propagation delay of 2.2 ns, over 30 ns settling time must be allowed in order to maintain this failure rate at 1 per month. In MU5 a tunnel diode flip-flop with a decision time of the order of 100 ps is used in critical circuits giving an MTBF of 136 years for a settling time of 3 ns.

#### 4.1.2 Variations Among Instruction Types

In describing the operation of the PROP pipeline, it was assumed that each stage of the pipeline contained one 16-bit order. However, circumstances arise which prevent full utilisation of each stage. For example, the order code allows for multi-length orders and some orders require multiple accesses (those involving the stack mechanism, for example). These can normally be dealt with by a purely logical control mechanism which does not affect the main pipeline timing, but may involve the creation of gaps (unused stages) in the pipeline. These gaps, or 'dummy orders', are distinguished by means of an additional function digit which, when set, inhibits the actions of each stage, including the Control Register updating.

The stacking mechanism is used for storing and retrieving partial results and for procedure links and parameters. For example, during the evaluation of expressions such as

$$a = b*c + d*e$$

partial results are stacked by the use of the '\*=' (stack and load) function. They are later unstacked by use of the operand form STACK (section 2.5). Stacked operands are therefore contained in the MU5 Processor storage system in exactly the same way as names, their addresses being generated relative to the Stack Front register (SF), which points to the most recently stacked operand within the Name Segment. Thus SF is advanced by both the '\*=' function and functions concerned with procedure entry (STACK and STACK LINK), and all these functions require two operand accesses to be made. Hence they are divided into two phases.

For the STACK function, for example, an access is first made for the specified operand followed by an access to the stack, while for the '\*=' order the first access is to the stack, in order to store the content of the specified

register, and the second is for the operand. For the stack writes the name/base adder is used to create the address SF+2 and at the same time SF is updated to this new value. The two phases of these orders are distinguished by extra digits carried through the pipeline with the function. These digits override the normal operand accessing mechanisms when access to the stack is required and also prevent the incrementing of the Control Register when the first phase reaches HI.

The implementation of this stack mechanism within a pipeline gives rise to additional problems in relation to control transfer orders. An order implicitly changing SF does so while there are still several orders ahead of it, but not past the Control Point (the point in the pipeline at which the Control Register is updated), and therefore not yet completed. Any one of these orders could be a control transfer order requiring that the partially processed orders behind it in the pipeline be abandoned (section 4.1.3). Should this situation occur, the SF Register may contain an incorrect value. The correct SF value could be maintained by preventing overlap in such situations, but this would cause serious degradation of the performance of the pipeline. The alternative solution adopted is to allow the SF register to change as and when required and to carry along the pipeline with the order the new value of SF created by it (registers S3, S4 and S5 in figure 4.2). When the Control Register is updated for the order, the value in S5 is copied into S6. Therefore when a control transfer occurs, the value in S6 is correct and is used to restore SF.

In a non-pipelined computer this situation would not arise since each instruction would be completed, and all register values would have been updated, before the next instruction could begin. A particular consequence of this is that an order which writes the contents of a programmable register to store in a non-pipelined computer can proceed without delay, since the register contents are guaranteed to be correct and available at the start of the order. This is not generally true in a pipelined machine where operand accesses are made several pipeline stages ahead of the programmable register into which the results of the operations involving those operands will be loaded. For example, in MU5 the Name Store is several stages ahead of the B-unit, so that when a store order of the type 'B => name' accesses the Name Store there can be several orders ahead of it in the pipeline which have themselves not yet reached the B-unit and operated on the B register content. Thus the B register value cannot be made available immediately, and in the absence of any additional technique, the 'B => name' order would have to wait at the Name Store stage of the pipeline until the B register content

could be guaranteed to be correct and available. This would involve a delay equivalent to at least four pipeline stages, and because these orders constitute 5-10 per cent of all orders, special action is taken to avoid this hold-up. (Store orders involving registers within PROP (NB =>, etc.) or SEOP (DR =>, etc.) occur much less frequently and no special action is taken for them. 'ACC => name' orders are dealt with separately, but in a similar manner to 'B => name' orders, in the Secondary Operand Unit pipeline.)

Whenever a 'B => name' order enters Stage 4 of the PROP pipeline, the content of the PROP Line Register (PLR in figure 4.2) is preserved, for later use, in an additional register BW, and the order proceeds normally through the pipeline without causing a hold-up and without impeding orders following, except as described below. Execution of a '=>' order in the B-unit involves sending the result to the H0 register in PROP. When it arrives the PROP pipeline is held up before the next beat is generated, and the information held in register BW is used to select the appropriate line in the Value Field of the Name Store so that the value in H0 can be written into it. The additional information needed to select one half of the line for over-writing is held in the F2 Function Register, together with a 'B => outstanding' digit which indicates that the BW Register is in use. When the action of writing into the store has been completed, the 'B => outstanding' digit is re-set and the pipeline is re-started.

While the 'B => outstanding' digit is set, two pipeline hold-ups can occur, one at Stage 2 and one at Stage 4. These hold-ups are typical of a number which can arise in the pipeline because some necessary information (such as the B register value for the 'B => name' order) is not immediately available. Apart from hold-ups which arise from the fifth stage, they cannot be detected in time for the next beat of the pipeline to be inhibited and therefore operate independently of the beat generating logic, by simply preventing subsequent beats from propagating back beyond the stage from which they arise (figure 4.4), and by causing dummy orders to be propagated forwards. The hold-up at Stage 2 occurs if a second 'B => name' instruction enters that stage, while the 'B => outstanding' digit is set, because the BW register can only deal with one outstanding order at a time. This hold-up prevents subsequent beats from propagating back beyond the input registers to Stage 3 (F3 etc.) and causes dummy orders to be copied into Stage 3. The hold-up at Stage 4 occurs if an instruction tries to access the same line in the Name Store as that indicated by BW. This is an example of the 'read after write' problem found in almost all instruction processing pipelines [27] and is a further consequence of the



fact that the B register value is not available for writing into store when the 'B => name' order reaches the Name Store. The hold-up prevents subsequent beats from propagating back beyond the input registers to Stage 5 (F5, etc.) and causes dummy orders to be copied into Stage 5. Both these hold-ups are automatically released when the 'B => outstanding' digit is re-set, or if the contents of the pipeline are discarded by the action of a control transfer before the 'B =>' order has left the end of PROP. If a control transfer occurs after a 'B =>' order has left the end of PROP, then the store updating action must still be carried out since the Control Register will have been incremented for this order.

Hold-ups arising from the fifth stage are normally those involving complex actions within PROP and possibly interactions with another unit (as in the case of a Name Store non-equivalence, for example), and which require the pipeline to be stopped. The need for one of these hold-ups is indicated by the setting of one or more bits in a 'WAIT' Register as the instruction is copied into F5 (figure 4.4). On completion of the highway transfer the 'data-gone' flip-flop is set as usual, but the next beat is prevented from being released by the presence of the digit in the WAIT Register. Instead, a hardware routine is entered appropriate to the most significant digit in the WAIT Register. When the routine is completed the corresponding WAIT digit is re-set and either the beat is released or another hardware routine is entered appropriate to the next most significant digit in the WAIT Register. A typical WAIT routine is that used for base register manipulation. This manipulation is carried out using the same adder as that used for address calculations (figure 4.2). If the order is of the type 'NB +', then the base forms one input to the adder, via the same route as that used for address calculations, and the operand forms the other input, in place of the name. The adder output is routed back to the inputs of all the base registers and when the addition is complete the appropriate one is updated. These orders must be completed before a succeeding instruction can be allowed into register DF, since it may use the content of NB to form an address. In these cases the decoding logic in Stage 0 of PROP sets a 'no-overlap' digit which prevents further instructions from being copied into DF until the base manipulation order has been executed.

#### 4.1.3 Control Transfers

In computers with a fixed instruction length the Control Register (or Program Counter) needs only to be incremented by 1 at the completion of each instruction and a simple counter may be used for this purpose. In computers such as MU5,

however, where the instruction length is variable, a full adder is required to accommodate the correspondingly variable increment. In a simple implementation this adder would be shared with other operations, but in a high performance system a separate dedicated adder is required. This same adder can therefore be used to execute control transfer instructions. Thus for a relative control transfer in MU5 the Control Adder in PROP (figure 4.2) is used to add the operand to the current value in the Control Register (CO), while for absolute transfers the operand passes through the adder with the Control Register input to the adder held at zero.

Control transfers (or branch instructions) are a major problem in pipelined instruction processors, since if the branch is taken all the instructions pre-fetched and partially executed in the pipeline must normally be discarded. This can not only cause serious delays in instruction execution, but also requires that the designer ensure that no irrecoverable action is taken during the partial execution of instructions ahead of the Control Point. We have already seen an example of the latter in the implementation of the stacking mechanism in PROP. Various techniques have been used in high performance computers to overcome the delay problems caused by control transfers and we shall examine these in detail in Chapter 5. We note here, however, that the MU5 Instruction Buffer Unit incorporates a prediction mechanism which, following the first execution of a particular control transfer, attempts to supply the correct sequence of instructions to the pipeline behind subsequent occurrences of that control transfer, so that if the branch is taken there is no delay in continuing instruction execution. This system offers a performance advantage, because unconditional transfers always branch and a high proportion of conditional control transfers are loop closing instructions which branch more often than not.

When a control transfer instruction reaches the end of PROP an 'out-of-sequence' digit accompanying the following instruction is examined. If set this digit indicates that the instruction is the start of a predicted 'jump-to' sequence, and if the control transfer is unconditional or conditional and the condition is met, instruction execution can proceed as soon as the Control Register has been updated. This also happens if the out-of-sequence digit is not set and the condition for a conditional transfer is not met. If the out-of-sequence digit is not set following an unconditional transfer or a conditional transfer for which the condition is met, however, or if the out-of-sequence digit is set and the condition for a conditional transfer is not met, then all instructions in the pipeline before the Control Point (those in PROP and in the Instruction Buffer Unit) must be discarded,

by the setting of their dummy order bits, and a long delay is incurred while the new jump-to sequence of instructions is fetched.

#### 4.1.4 Performance Measurements

The overall performance of a pipelined instruction processing unit is affected by the relative frequencies of occurrence of different instruction types and the additional delays which they incur. Measurements of these frequencies have been made for MU5 [28] using a hardware performance monitor, for a number of benchmark programs and the results are shown in columns three and five of table 4.1 for program execution and compilation respectively. Column two gives the time required in excess of the 50 ns beat time for the execution of each type of order, and columns four and six give the net additional contribution of each type of order to the execution time of the average order.

Table 4.1 PROP Performance

Type of Order	Execution			Compilation	
	Excess Time (ns)	Frequency (%)	Net Time Added (ns)	Frequency (%)	Net Time Added (ns)
Multi-length	50	56	27.9	45.7	22.8
B Store	100	6.2	6.2	9.1	9.1
Base Manipulation	450	5.1	23.0	8.4	37.8
Predicted Control Transfer	100	8.7	8.7	8.0	8.0
Unpredicted Control Transfer	1350	4.5	60.7	10.5	141.8
Name Store NEQ	1180	3.5	41.3	11.9	140.4
Total net time added			<u>167.8</u>		<u>359.9</u>
Average instruction time			<u>217.8</u>		<u>409.9</u>

For both execution and compilation the longest delays are due to unpredicted control transfers and Name Store non-equivalences. The significant difference in performance between these two type of activity is due largely to the data-dependent nature of the compilation process which involves many alternative processing sequences. Name Store delays are largely a consequence of the organisation of the storage hierarchy (section 3.4), whereas the delays caused by control transfers, B-store orders and base manipulation orders are directly attributable to the difficulties inherent in processing sequentially dependent instructions in a pipeline, and those due to multi-length orders are consequences of the particular organisation chosen for the IBU-PROP interface. In each case, however, improvements could be made which would reduce these delays. Many of the decisions taken during the design of MU5 were based, inevitably, on inadequate information about instruction usage, and in a commercial environment changes almost certainly would have been made to subsequent models in the light of experience gained from what was essentially a prototype design.

#### 4.2 ARITHMETIC PIPELINES - THE TI ASC

Although instruction and operand accessing are highly pipelined in MU5, no attempt is made to pipeline the activities of the main floating-point arithmetic unit. Floating-point addition, for example, consists of four distinct operations (exponent subtraction, mantissa shifting, mantissa addition, normalisation) which can be pipelined in a very straightforward manner. In a one-address system such as MU5, however, no benefit would be gained from this arrangement, since one of the input operands to an addition is the content of the Accumulator, and this cannot be used as an input to any operation until the previous operation has been completely finished. For successful operation of a pipelined arithmetic unit each instruction must reference at least two and preferably three operands, and even when this requirement is satisfied the possibility of dependencies between successive instructions requires complex control mechanisms and may cause hold-ups in instruction execution. Examples of these control techniques are the data forwarding arrangement use in the IBM System/360 Model 91 (section 4.3) and the Scoreboard used in the CDC 6600. The ideal arrangement is for one instruction to cause two independent operand streams to be combined in the arithmetic unit to form a third independent result stream. This is achieved in vector processing systems such as the CDC Cyber 200 Series (section 7.2), the CRAY-1 (section 6.3) and the Texas Instruments Advanced Scientific Computer (TI ASC).

Texas Instruments has traditionally been a supplier of instrumentation to the oil industry, and it was this industry's need for a powerful seismic processing capability that led to the development from 1966 onwards of the Advanced Scientific Computer [29]. A significant feature of this type of processing is the frequent use of triple-nested indexing loops, and an important characteristic of the ASC is the provision of three levels of indexing within a single vector instruction.

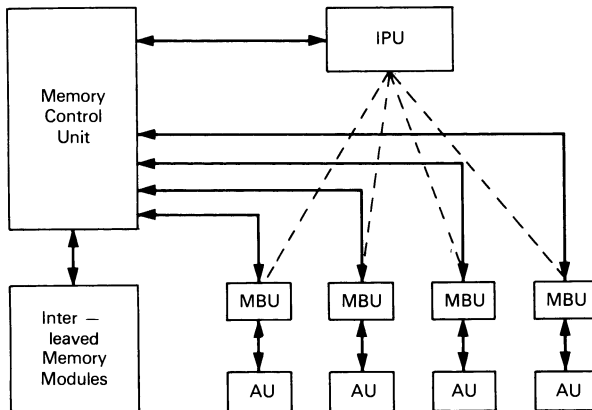


Figure 4.6 TI ASC Central Processor Structure

The basic structure of the central processor of the ASC is shown in figure 4.6. The Memory Control Unit allows two-way traffic between eight 256-bit wide processor ports and eight interleaved memory modules (plus an optional memory extension), with each processor port having full access to the 24-bit address space. The central processor design is such that one, two, three or four pipelined arithmetic units (AU) can be provided. Each AU is linked to the Memory Control Unit through its own Memory Buffer Unit (MBU), the primary function of the latter being to supply, from memory, a continuous stream of operands to the arithmetic unit, and to return to memory a continuous stream of arithmetic results. Addresses for sustaining this continuity are computed in the MBU using parameters supplied through the Vector Parameter File at the start of a vector instruction, and double buffering (in units of eight 32-bit words) is provided for each of the operand streams. The Vector Parameter File contains a starting address field for each of the three vectors involved in a typical vector instruction, a vector length, inner and outer loop counts, and increments to be added to each of the three starting addresses at the completion of each inner and outer

loop. Information is supplied to the Vector Parameter File by the Instruction Processing Unit (IPU), which is mainly concerned with initiating operations in the MBU-AU vector processing pipelines, but also acts as a conventional scalar processor dealing with instruction fetching and decoding, address generation, control transfers, etc.

The arithmetic units are 64-bit parallel fixed and floating-point units split into two halves of 32 bits each. Double length operations are carried out using both 32-bit halves in parallel, while single length operations use only the most significant half of an AU. The ASC differs from the CDC Cyber 200 Series in this respect, since the 64-bit arithmetic units in the latter split into two identical 32-bit units which can work in parallel. Numbers are represented in the ASC in IBM format. Thus fixed-point numbers are represented as 2's complement integers, while floating-point numbers are represented in sign and magnitude form with a base 16 (hexadecimal) exponent represented by an excess 64 binary number, as shown in figure 4.7.




Figure 4.7 ASC Floating-point Formats

Each AU is made up of eight distinct sections, each of which performs a separate arithmetic or logical operation (figure 4.8). Each section can be connected to any other section to allow the correct sequence of operations to be executed for a particular instruction, with the appropriate configuration being established at the start of each vector instruction. In any given configuration the various sections form a pipeline into which a new pair of operands can, in principle, be entered at each 60 ns clock, and after a start-up time, corresponding to as many clock periods as there are sections in use, result operands emerge at a rate of one per clock period. At the end of a vector instruction there is a similar run-down time between the entry of the last operand pair and the emergence of the corresponding result.

Floating-point addition, for example, requires the use of

the Receiver Register, Exponent Subtract, Align, Add, Normalise and Output sections, connected as shown by the solid line in figure 4.8. Pairs of operands from the MBU are first copied into the Receiver Register, the cable delays between the MBU and AU effectively forming a complete stage in the overall pipeline arrangement. The Exponent Subtract section then performs a 7-bit subtraction to determine the difference between the exponents of the two floating-point operands, or in the case of equal exponents uses logic to determine which of the fractional mantissae is larger (this logic is also used by those instructions which test for greater than, less than or equal to, in order to avoid duplication of hardware).

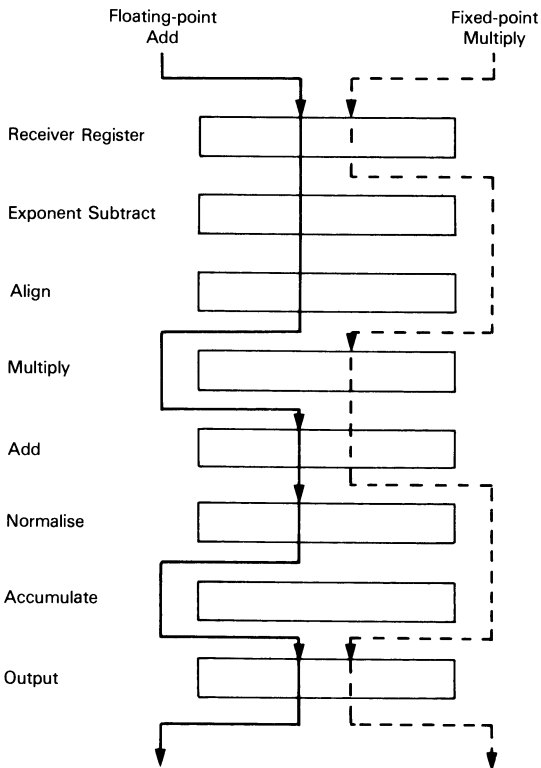


Figure 4.8 ASC Arithmetic Pipeline

The exponent difference is used in the Align section to shift right the mantissa of the operand with the smaller exponent. In one cycle any shift which is a multiple of four can be carried out, which is all that is required for floating-point numbers represented in base 16. (Fixed-point

right shifts require two cycles, one shifting by the largest multiple of four in the shift value, and a second in which the result of the first is re-entered and shifted by the residue of 0, 1, 2 or 3.)

Having been correctly aligned, the fractional parts of the two floating-point numbers are added in the Add section, and the result is passed on to the Normalise section. This section closely resembles the Align section in that floating-point operations only require one cycle, while the fixed-point left shifts which it also carries out require two. The major difference between these two sections is that Align receives information concerning the length of shift required in floating-point operations, while the Normalise section has to compute the shift length by determining which four-bit group contains the most significant digit. It also contains an adder to update the exponent value when a normalisation shift occurs. The results of all arithmetic operations pass through the Output section before being returned to the Memory Buffer Unit. The partitioning of the arithmetic unit into these various sections is primarily intended to give high throughput of floating-point addition and subtraction. Each section is capable of operating on double length operands so that vector double length instructions can proceed at the clock rate. Double length multiplication, and all divides (which are performed by an iterative technique) proceed more slowly.

The dashed line in figure 4.8 shows the interconnection used for fixed-point multiplication. The Multiply section can perform a 32 by 32-bit multiplication in one clock period, so that the results of both fixed-point and single-length floating-point multiplication are available after one pass through the multiplier. Because a carry-save addition technique is used, the output of the Multiply section consists of a 64-bit pseudo-sum and a 64-bit pseudo-carry. These must be added in the Add unit to produce the true result. Double-length multiplication requires three separate 32 by 32-bit multiplications to be performed and these can therefore proceed at a rate of only one every three clocks. After passing through the Add section the three separate results are added together in their proper bit positions in the Accumulate section.

The Accumulate section is similar to the Add section and is used in all instructions which require a running total to be maintained. An important example of this type of instruction is the Vector Dot Product, which is used repeatedly, for example, in matrix multiplication. Pairs of operands are multiplied together in this instruction and a single scalar result, equal to the sum of the products of the pairs, is



produced. Because the running total is maintained in the arithmetic unit, the read-after-write problems which occur in scalar implementations of this operation are avoided in the ASC. In MU5, for example, the sequence of instructions required to produce, on the stack, the vector dot (or 'scalar') product of two vectors defined by descriptors VEC1 and VEC2, each of length LIMIT, is

```

      B = 0
      ACC = 0
L1:  ACC *= VEC1[B]
      ACC * VEC2[B]
      B CINC LIMIT
      ACC + STACK
      IF /=, -> L1

```

where [B] implies modification of the descriptor Origin by the content of the index register B, and the function '\*' is processed in two stages, the first stacking the content of the specified register (the floating-point accumulator in this case) and the second re-loading it with the specified operand (the Bth element of VEC1 in this case). Thus there are only two Accumulator instructions between the instruction which writes a new value from the Accumulator into the top-of-stack location held in the Name Store (section 3.4.1) and the instruction (ACC + STACK) which reads it out again. In the MU5 pipeline the Name Store is several stages earlier than the Accumulator, and a similar mechanism to that used for B write orders (section 4.1.2) is used to prevent ACC write orders from creating an immediate hold-up. The hardware must guard against an operand being read out before it has been updated by an outstanding write order, however, and in so doing it does cause a hold-up in the passage of instructions through the pipeline. Exactly the same read-after-write problem occurs in pipelined, multi-address register machines such as the high-performance models in the IBM System/360 series; in the case of Models 91 and 195, however, a 'data forwarding' mechanism was developed in order to avoid this hold-up.

#### 4.3 THE IBM SYSTEM/360 MODEL 91 COMMON DATA BUS

The System/360 Model 91 was developed in the mid 1960s with the primary aim of providing the highest performance capability that advanced design philosophy and System/360 circuit technology extensions could achieve within a balanced development schedule [30]. Performance was taken to mean general computer availability and high-speed execution of general problem programs, and a factor of between one and two orders of magnitude over the earlier 7090 was achieved, depending on the nature of the problem. Only a small number of

Model 91s was actually produced, but most of its architectural features were carried over into the commercially more successful Model 195.

One of the problems facing the designers of high-speed computer systems is the difficulty of achieving the fastest possible execution times for a particular technology in universal execution units. Circuitry designed to carry out both multiplication and addition, for example, will do neither as fast as two units each limited to one kind of operation. Thus not only did the Model 91 contain separate fixed-point and floating-point execution areas (figure 4.9), but the floating-point area contained separate add and multiply/divide units capable of concurrent operation. Both units were pipelined; the add unit as a two-stage pipeline capable of starting a new operation in each 60 ns clock cycle, and the multiply/divide unit capable of starting a second operation three cycles after the start of a previous multiply or twelve cycles after a previous divide.

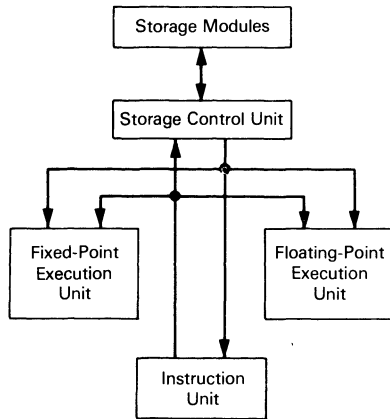


Figure 4.9 IBM System/360 Model 91 Central Processor

Independence of fixed-point and floating-point operations is guaranteed in all System/360 machines by the separation in the instruction format of the two sets of general accumulator registers (section 2.2). Within the floating-point unit, however, the hardware must preserve essential instruction sequence dependencies while allowing the greatest possible overlap of independent operations. This same problem arises in a number of machine designs, particularly the CDC 6600, with which we shall deal in some detail in Chapter 6. The control mechanism used to solve this problem in the Model 91 shows some similarities to the technique used in the CDC 6600,

although the organisation of the data paths is quite different.

The organisation of the Model 91 floating-point unit [31] is shown in figure 4.10. Instructions are prepared for this unit by the Instruction Unit pipeline and entered in sequence, at a maximum rate of one per clock cycle, into the Floating-point Operand Stack (FLOS). Instructions are taken from the FLOS in the same sequence, decoded, and routed to the appropriate execution unit. The Instruction Unit maps both storage-to-register and register-to-register instructions into a pseudo-register-to-register format, in which the equivalent of the R1 field (figure 2.3) always refers to one of the four Floating-point Registers (FLR), while R2 can be a Floating-point Register, a Floating-point Buffer (into which operands are received from store), or a Store Data Buffer (from which operands are written to store). In the first two cases R2 defines the source of an operand; in the last case it defines a sink.

The most significant feature of this floating-point system is the Common Data Bus (CDB). The CDB is fed by all units which can alter a register, and itself feeds the floating-point registers, the store data buffers and all units which can have a register as an input operand. The latter connections allow data produced as the result of any operation to be forwarded directly to the next execution unit without first going through a floating-point register, thus reducing the effective pipeline length for read-after-write dependencies, as found, for example, in the scalar product loop. The running total in this loop would not actually appear in a floating-point register in the Model 91 until the last execution of the loop.

The operation of the CDB is controlled by the use of tags. A tag is a 4-bit number generated by the CDB control logic to identify separately each of the eleven sources which can feed the CDB. Thus there are six floating-point buffers, three parallel 'reservation stations' (containing input buffer registers) associated with the adder, and two parallel reservation stations associated with the multiplier/divider. Tag registers are associated with each of the four floating-point registers, with the source and sink input registers of each of the five reservation stations, and with each of the three store data buffers. There is also a busy bit associated with each of the floating-point registers. This bit is set whenever the FLOS issues an instruction designating the corresponding register as a sink, and is re-set when a result is returned to the register.

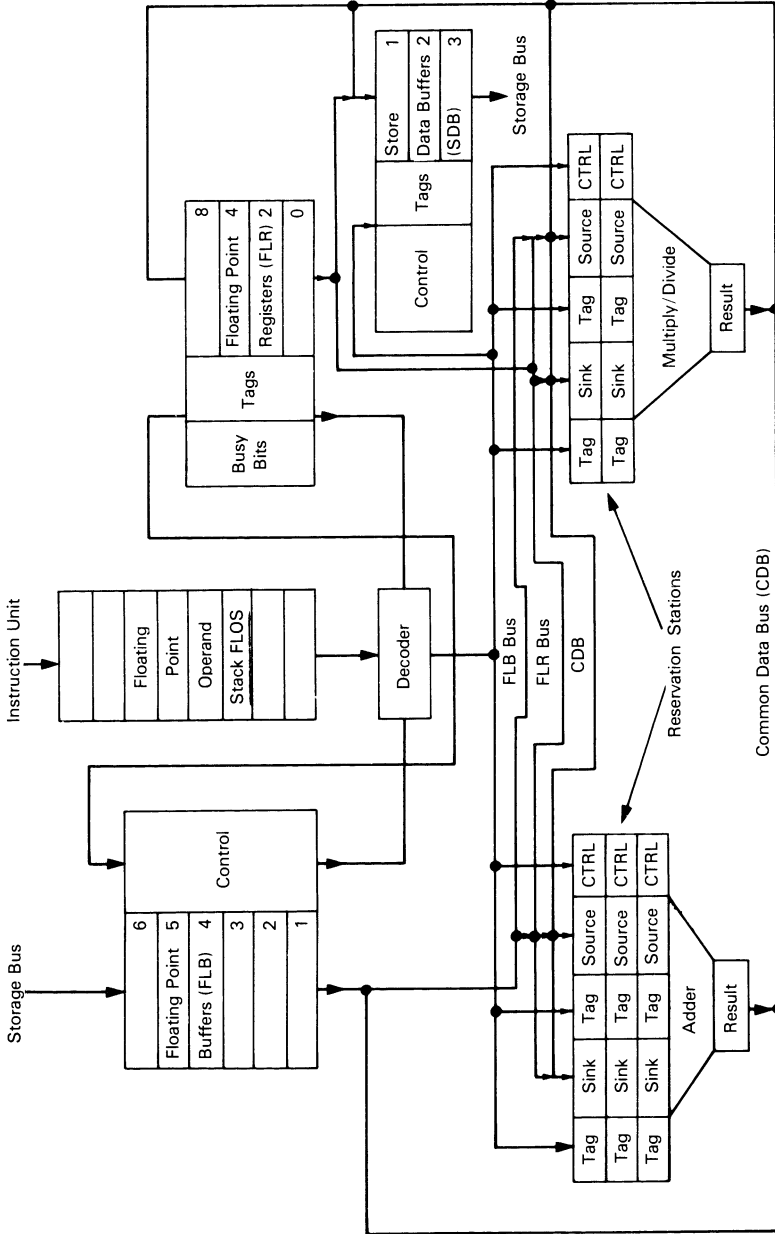


Figure 4.10 The IBM System/360 Model 91 Floating-point Unit

Whenever the FLOS decodes an instruction it checks the busy bit of each of the specified floating-point registers. If the bit is zero, the content of the register is sent to the selected reservation station via the Floating-point Register (FLR) Bus. On issuing the instruction the FLOS sets the busy bit of the designated sink register, and enters into its tag register the tag number of the selected execution unit. If the FLOS finds a busy bit set, however, it does not transmit the register contents to the selected reservation station, but instead transmits the current value of the corresponding tag register, and enters into that tag register the appropriate new tag number. Thus the tag register of a busy floating-point register identifies the last unit (in proper program sequence) which will produce a result destined for that register.

Whenever a result appears on the CDB, the tag corresponding to its source is broadcast to all destinations. Each active reservation station (selected but awaiting a register operand) compares its sink and source tags with the CDB tag. If a match occurs (a sink is also a source in the System/360 two-address instruction format), the reservation station takes the data from the CDB. In a similar manner, the CDB tag is compared with the content of the tag register associated with each busy floating-point register. All busy registers with tags matching that on the CDB are set to the value on the CDB and their busy bits are re-set.

Issuing an instruction in this system only requires that a reservation station be available for whichever execution unit is required. If a source register is awaiting the result of a previously issued, but as yet uncompleted instruction, or if a floating-point buffer register is awaiting an operand from store, the tag associated with that register is transmitted instead to the reservation station, which then waits for that tag to appear at its input. Thus it is the reservation stations which do the waiting for operands, rather than the execution circuitry, which is free to be engaged by whichever reservation station fills first. Execution of an instruction starts when a reservation station has received both operands. It may receive one or both operands from either the CDB or FLR bus in the case of a register-to-register instruction, while in the case of storage-to-register instructions the source operand is transmitted by the FLB bus. Both the FLR and FLB busses could have been incorporated, in principle, into the CDB, but in practice this would have caused a reduction in performance due to conflicts over the common facility. The use of the CDB and the associated tagging mechanism has been shown to reduce the execution times of the inner loops of programs used to solve partial differential equations, for example, by about one-third.

# 5 *Instruction Buffering*

In dealing with operand accessing in earlier chapters we considered various techniques used to overcome the disparity between processing speed and main store accessing rate. This problem also impinges on instruction accessing, since for efficient operation instructions must also be supplied to the processor at a rate matching its execution rate. In the case of instruction accessing, however, the problem is ameliorated by the fact that most instructions are obeyed sequentially and the main store word size is normally such that one word fetched from main store can contain several instructions. Furthermore, with an interleaved store, successive accesses for sequential instructions reference each stack in turn and are not held up by cycle time effects. Thus store requests can be made in advance of the corresponding instruction being required and the replies buffered until they are needed for execution. This pre-fetching technique is used in almost all high performance pipelined processors. A significant proportion of instructions cause control transfers, however, and each such transfer requires a request to be made to the store for a new sequence of instructions. Thus although the accessing rate for instructions can normally be matched satisfactorily to the processing rate, the access time for the first instruction of a new sequence can result in a long delay to the processor. Techniques for overcoming this problem rely on the fact that the cause of many control transfers is a branch back from the end to the start of a loop of instructions, and loop catching buffers are incorporated into a number of processors. In this chapter we shall consider the instruction accessing and buffering techniques used in the IBM System/360 Model 195, the CDC 6600 and 7600, MU5 and the CRAY-1.

## 5.1 THE IBM SYSTEM/360 MODEL 195 INSTRUCTION PROCESSOR

The organisation of the System/360 Model 195 Central Processor is very similar to that of the Model 91 (figure 4.9) and, as in the case of the Common Data Bus discussed at the end of Chapter 4, the instruction buffering technique used in the Instruction Processor of the Model 195 derives directly from that used in the Model 91 [32]. The Model 195 Instruction

Processor is concerned with fetching and buffering instructions from storage, fetching the operands which those instructions specify, issuing instructions to the appropriate execution units, handling interrupts, and executing all branching (control transfer), status switching and input/output instructions.

Instructions fetched from store are buffered in an eight-doubleword (64-bit) Instruction Stack (figure 5.1). The instruction fetching mechanism is controlled by three registers, the Instruction Register (IR) which addresses the instruction currently being decoded, the Upper Bound Register (UB) which points to the most recent doubleword brought into the stack, and the Lower Bound Register (LB) which points to the earliest doubleword in the stack. During normal operation the stack contains the current instruction doubleword, some doublewords ahead of the current instruction and a copy of some instructions which have already been issued.

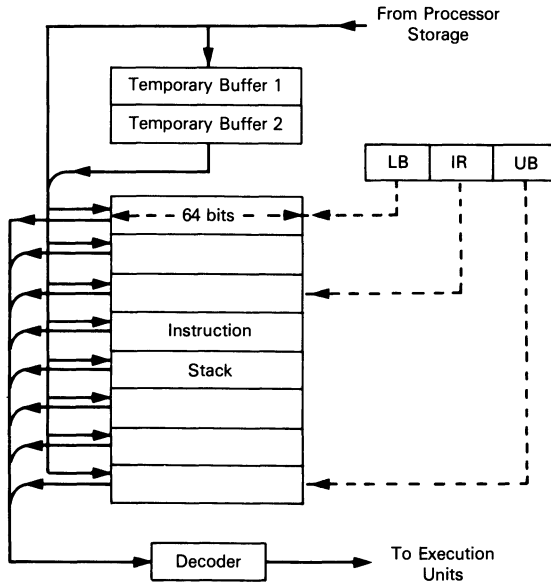


Figure 5.1 The IBM System/360 Model 195 Instruction Buffer

### 5.1.1 Sequential Instruction Fetching

Pre-fetching of instructions is controlled by the UB register. When instruction fetching is initiated following an interrupt, for example, the Instruction Stack is declared empty and the

main storage address of the first instruction doubleword is loaded into UB and LB. The instruction fetching mechanism associated with UB then accesses this doubleword and loads it into the location in the Instruction Stack addressed by the three least significant doubleword address bits in UB. Initially this location is also addressed by IR, which selects each instruction in sequence for decoding and processing. After an instruction has been decoded and passed to the next stage in the processor pipeline, IR is incremented by the number of half-words in that instruction and the next instruction selected.

Once the first instruction access has been sent to store, the instruction fetching mechanism increments UB and continues to make sequential store accesses until prevented from doing so either because the address in UB is seven doublewords higher than that in IR (and any further accesses would cause instructions not yet decoded to be overwritten), or because the Instruction Processor has detected a condition giving rise to a change in the instruction sequence (a branch instruction or an interrupt, for example).

During normal operation the instruction fetching mechanism continually attempts to increment UB and fetch instruction doublewords from store, while the instruction decoding mechanism continually increments IR as instructions are decoded and passed along the processor pipeline. Once IR has been incremented beyond the address in LB, instructions in the first doubleword fetched into the stack can be overwritten with new information. Provided IR remains ahead of LB, then when incrementing UB would cause its three least significant doubleword address bits to match the corresponding bits in LB, both these registers are incremented together. Thus at each instruction access the oldest doubleword in the stack is replaced by the latest doubleword fetched from store.

Use of this pre-fetching mechanism allows a continuous sequence of instructions to be supplied to the processor at a rate approaching one per machine clock cycle, and thus roughly matching the instruction execution rate. (Although the processor pipeline was designed to execute instructions at a rate of one per clock cycle, instruction dependencies, storage conflicts and the frequency of operations requiring multi-cycle execution combine to reduce the average rate to about half this figure.) When a new sequence of instructions is required as a result of the branch being taken in a branch instruction, however, the start-up delay is of the order of six clock cycles, and in the absence of some additional technique the average performance of the processor would be seriously degraded. Conditional branches cause even further



problems since the branch decision depends on the outcome of a previously issued, but not necessarily completed arithmetic instruction, and an additional delay may be incurred in awaiting this outcome. This problem is discussed further in section 5.5. In the Model 195 two techniques are used to ameliorate the problems caused by branches, one involving the establishment of a 'Conditional Mode' of operation, and the other a 'Loop Mode'.

### 5.1.2 Conditional Mode

Conditional branch instructions interrogate a 2-bit Condition Code at their point of execution in order to determine whether or not the branch is to be taken. The Condition Code is set by a variety of instructions, but only the last of these issued before a conditional branch must be allowed to affect its outcome. This is accomplished by tagging at decode time each instruction which will set the Condition Code. At the same time a signal is forwarded through the pipeline to remove the tags from any previously issued but uncompleted instructions. Only a tagged instruction may set the Condition Code, at which point its tag is removed, and a conditional branch instruction can only execute when there are no outstanding tags in the processor.

Since in general the Condition Code will not be valid when a conditional branch is decoded, the hardware always assumes this to be the case and establishes Conditional Mode. In Conditional Mode further sequential instruction accesses are inhibited, but rather than hold up further activity entirely, processing of the remaining instructions in the Instruction Stack proceeds as far as possible (until a further branch is decoded or the pipeline becomes full, for example), with the instructions being marked as 'conditional'. Conditional instructions are decoded, their operand fetches are initiated, and they are forwarded to the relevant execution units in the normal way. The conditional tag inhibits the execution units from actually completing them, however, and once the first such instruction reaches the point of execution, further processing is held up until the Condition Code is set and the branching action determined. If the branch is not taken, the conditional tags are re-set and the pipeline is re-started without further delay.

If the branch is taken, the conditional instructions must be abandoned and a fresh start made with a new sequence. The delay incurred in refilling the pipeline from the decoder onwards is unavoidable, but the delay in accessing the first instruction at the target address of the new sequence is minimised in the Model 195 because the hardware assumes at the

start of Conditional Mode that either outcome is equally likely and fetches the first two instruction doublewords at the branch target address immediately. These two doublewords are loaded into the two Temporary Buffers shown in figure 5.1, in order that the Instruction Stack remain unaffected if the branch is not taken. Clearly these instruction fetches will have been made unnecessarily on many occasions, and since instruction accesses have priority over operand accesses on the store address path, some performance degradation can occur due to interference with operand accesses for the conditional instructions. This disadvantage is more than offset, however, by the advantage gained, when the branch does occur, of the access time for the target instructions having been overlapped with the wait for the Condition Code. In the case of an unconditional branch to an instruction not in the Instruction Stack, there is, of course, no need to wait for the Condition Code to become valid. As in the conditional case, the target instruction sequence is requested immediately, but unless the execution unit pipelines are also held up (as a result of divide operations, for example) the six clock cycle start-up delay inevitably causes a gap to occur in the instruction processing sequence.

The primary purpose of the whole conditional philosophy was the circumvention of storage delays, and in retrospect the designers felt that the complications of the system, which involves numerous interlocks throughout the processor, would become increasingly difficult to justify as storage access times decrease.

### 5.1.3 Loop Mode

Without the use of branch target instruction pre-fetching in Conditional Mode, the time lost when the branch is taken would be roughly equal to the sum of the time spent waiting for the Condition Code to be set and the storage access time. With pre-fetching the time lost becomes equal to only the greater of these two, but even so, where the branch is closing a short loop of instructions, this loss can severely limit overall processor performance. Thus for short loops a different philosophy is adopted whereby the entire loop is contained within the Instruction Stack and storage accesses are avoided altogether until the program exits from the loop. Clearly, the longer the loop, the smaller the proportion of time lost as a result of the branch, and the choice of eight doublewords as the capacity of the stack represents a compromise between hardware cost and performance in Loop Mode.

Loop Mode is entered whenever a branch backwards is taken to a target address within eight doublewords of the current

instruction. The Instruction Stack is immediately re-initialised to contain the appropriate eight doublewords, after which instruction fetching ceases and the address path to store is fully available for operand fetching throughout execution of the loop. Loop Mode is controlled by two additional registers, one containing the loop target address (SLT) and the other the value of IR corresponding to the loop closing instruction (SLCIR). Once in Loop Mode the address of any branch instruction being decoded is compared with that in SLCIR, and if it is the same the branch is made immediately to the target address held in the other. Thus the role of Conditional Mode is reversed, since it is assumed that the branch will be taken, and instructions are therefore decoded from the target path rather than the straight through path. Furthermore, no fetches are made to the Temporary Buffers in Loop Mode.

Loop Mode is normally turned off because an exit is taken from the loop. This can happen in a variety of ways. If the branch closing the loop is not taken, for example, IR will run off the end of the instructions held in the stack and require a store access. Alternatively some other branch within the loop may be taken to a target outside the stack, or the address in SLCIR may be invalidated. This can happen if the base or index register specified in the instruction which caused SLCIR to be set up is altered. A record of which of these registers is involved is kept with SLCIR and a check made against this record if any instruction in the loop alters a fixed-point register.

## **5.2 INSTRUCTION BUFFERING IN CDC COMPUTERS**

### **5.2.1 The CDC 6600 Instruction Stack**

The organisation of the CDC 6600 Central Processor [7] was discussed in section 2.1 in connection with the 6600 instruction set. We noted that instructions are fetched from Central Storage and placed in an Instruction Stack before being decoded and issued to the appropriate functional unit under control of the Scoreboard. The Instruction Stack itself consists of eight 60-bit registers (I0-I7 in figure 5.2) which operate as a push-up stack and which can contain instruction loops. Programs are initiated in the 6600 by an 'Exchange Jump' in which the contents of all addressable registers in the central processor are interchanged with the contents of a designated store area. Following such an Exchange Jump the new contents of the program address register are used to access the first instruction word. This word is received from Central Storage into an Input Register and then loaded into the bottom register of the Instruction Stack.

Instruction words are made up of four 15-bit 'parcels' and as the first instruction word enters the bottom register of the stack (I0), the first two parcels within the word (starting from the left) are transferred into a series of instruction registers within the Scoreboard control logic (section 6.1). At the same time a further instruction fetch is initiated. Two parcels are taken to allow for long format (30-bit) instructions. If the first parcel is a short format (15-bit) instruction, the second parcel is ignored and in the next processor minor cycle the second and third parcels are taken from I0. When a long instruction is encountered an extra minor cycle is spent skipping over the second half, so that dealing with one complete instruction word never takes less than four 100 ns minor cycles. This matches the rate at which instructions move into the stack. Whenever a new instruction stack fetch is initiated, the contents of the stack ripple upwards one register every half minor cycle, with the topmost location (I7) being overwritten first. At the end of four minor cycles the contents of I0 are moved up and I0 is then ready to receive a new instruction from the Input Register.

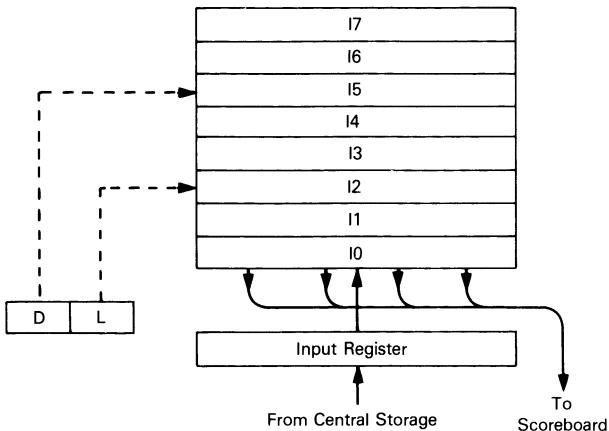


Figure 5.2 The CDC 6600 Instruction Stack

In practice the rate at which instructions could be accessed from Central Storage turned out to be longer than anticipated at the design stage, which meant that a new instruction word would not actually be available until after a total of eight minor cycles. Since the rate at which instructions are issued to the functional units cannot often be maintained at one per minor cycle, however, the overall effect of this delay on performance is not quite so bad as might be imagined, and when executing loops which can be

contained within the stack, no Central Storage accesses for instructions are required at all.

Information about the contents of the stack is contained in two registers, the Depth (D), which measures the number of valid instruction words in the stack, and the Locator (L), which specifies the location in the stack of the instruction word currently in use. During execution of a loop held entirely in the stack the instructions remain in fixed locations and the program address register can point to any one of the stack registers within a distance D from the bottom. D is re-set to zero whenever a branch out of the stack is taken, and is incremented by one for every new instruction word brought in. When the stack is full, D remains equal to seven.

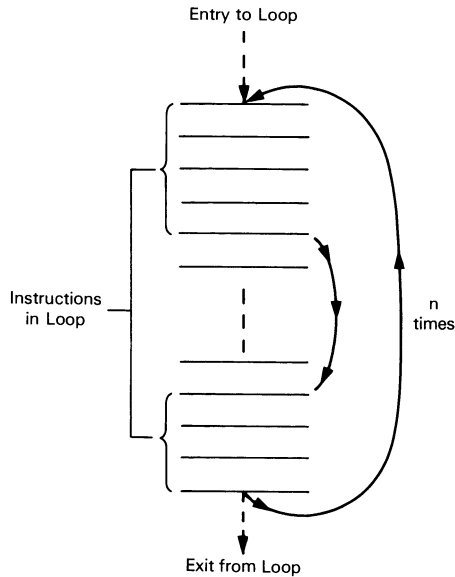


Figure 5.3 A Non-contiguous Instruction Loop

When a conditional branch is decoded a test for 'jump within stack' is made. This involves subtracting the current program address from the branch address. If the absolute value of the result is less than seven words, and if the values in D and L indicate that the branch is to a location within the stack, no further store accesses are made for instruction words until instruction parcels are again taken from IO. Thus a branch may jump forwards or backwards within the stack and loops may be held in the stack in various forms.

A very similar Instruction Stack was used in the STAR-100 computer, CDC's first commercially produced vector processor. The STAR-100 had a much longer instruction format than the 6600 so that its Instruction Stack was larger, being made up of sixteen 128-bit registers, but it used essentially the same control mechanisms. The main drawback of both the 6600 system and that used in the IBM System/360 Model 195 is that where the total number of instructions being obeyed in a loop will fit into the stack, but the code is actually made up of a number of non-contiguous segments (as in figure 5.3, for example), the loop may not be caught in the stack. With machine code programming this situation can normally be avoided, but it is a common occurrence in compiler generated code and the increasing emphasis on high-level language programming has caused processor designers to seek alternative solutions. The CDC 7600 and CYBER 205, for example (successors, respectively, to the 6600 and STAR-100), both use associatively addressed buffers, with the CYBER 205 Instruction Stack again being correspondingly larger. Further discussion of the CYBER 205 will be left until Chapter 7.

### 5.2.2 The CDC 7600 Instruction Stack

The CDC 7600 [33] was designed to be machine code upward compatible with the 6600, but to provide a substantial increase in performance. The ways in which this was achieved will be discussed in more detail in section 6.2. For now we will note that the organisation of the 7600 central processor is very similar to that of the 6600. It contains nine parallel functional units, a scratch pad of eight X registers, eight A registers and eight B registers, and an Instruction Stack. The 7600 Instruction Word Stack is made up of twelve 60-bit registers, however, compared with the eight used in the 6600, and each register also has its own 18-bit associative address register in an Instruction Address Stack (figure 5.4).

The Instruction Stack is filled two words ahead of the instruction currently being executed, thus giving a greater degree of pre-fetching than was possible in the 6600, and hence overcoming the storage access delay for sequential instructions. Furthermore, instructions are obeyed from a Current Instruction Word register, rather than from the bottom stack register, and a complete 60-bit word is transferred from the Instruction Stack into this register whenever the word address changes in the program address counter. This transfer can be made from any of the twelve registers in the Instruction Word Stack, allowing a considerable degree of flexibility in pre-fetching and loop catching. Whenever a new word is required in the Current Instruction Word register the address in the program address counter is compared with the

entries in the Instruction Address Stack, and if a coincidence occurs for any of these entries, the content of the corresponding register in the Instruction Word Stack is transferred into the Current Instruction Word register.

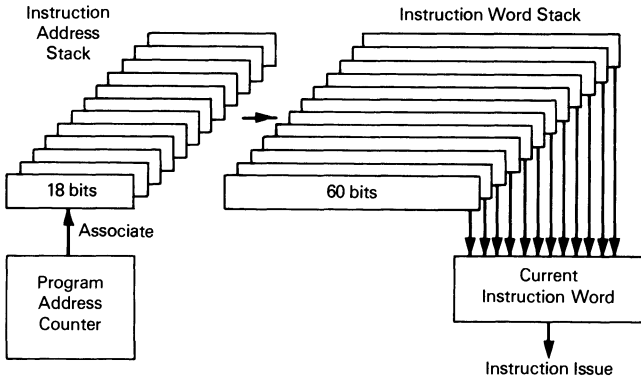


Figure 5.4 The CDC 7600 Instruction Stack

When obeying sequential code the required word will normally be in one of the bottom two registers. When a branch instruction is executed and the branch taken, the required word may already be in one of the top ten registers, obviating the need for a store access, and giving improved performance. If the required word is not in the stack, the first two words at the target address are immediately requested from store and instruction execution continues when the first of these is received. Whenever an instruction word is received from store all the entries in the Instruction Word Stack and the Instruction Address Stack are simultaneously moved up one position, with the new address and instruction word being entered at the bottom of the stack and the oldest entry being lost. Entries in the stack are only invalidated by the execution of a subroutine call or Exchange Jump, and not by normal branch instructions, so that a program may branch back and forth between short sequences of non-contiguous code held in the stack.

Although this stack is larger than that of the 6600, it is still relatively small, and considerable effort is frequently required to reduce the amount of code in program loops in order that they may fit into it. A quite different scheme from any of those considered so far is required if loops of unrestricted size are to be accommodated. Such a scheme is to be found in the MU5 computer.

5.3 THE MU5 INSTRUCTION BUFFER UNIT

The need for sequential instruction pre-fetching, and the disruptive effects of control transfers, become increasingly apparent as the degree of pipelining is increased. In the MU5 computer the Primary Operand Unit (section 4.1) has a maximum execution rate of one 16-bit instruction per 50 ns, while the plated-wire Local Store has a 260 ns cycle-time and is four-way interleaved. The Local Store can therefore supply successive 128-bit words at 65 ns intervals and so there is no problem in supplying sequential instructions at the required rate. Because the access time is much longer than 65 ns, however, instruction requests have to be sent out well in advance of their being required by the Primary Operand Unit, and buffered in an Instruction Buffer Unit. This Instruction Buffer Unit (figure 5.5) contains three 128-bit buffer registers which constitute its Data Flow section. The Data Flow control logic unpacks instructions from the first buffer and assembles them in the second and third buffers ready for PROP to take them as required. The necessary store requests are made by the Store Request System, which issues store addresses formed by a counter at a rate matched to that at which instructions are taken from the Data Flow by PROP.

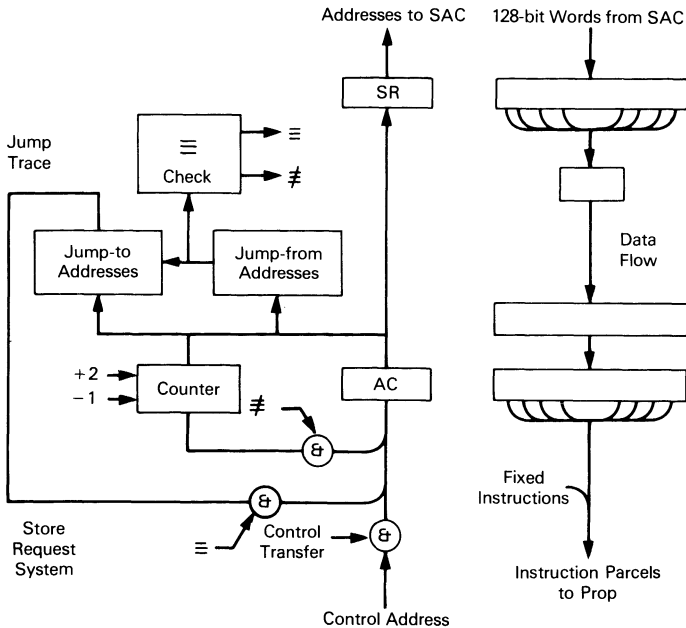


Figure 5.5 The MU5 Instruction Buffer Unit



This system operates satisfactorily until a control transfer occurs as a result of either an unconditional control transfer instruction, or a conditional control transfer instruction for which the condition is met. Then all the pre-fetched instructions must be abandoned, and the correct new instruction cannot be sent to PROP until the store has been accessed, using the new control address, and the instruction has passed through the Data Flow. As a result the total time between the execution of the control transfer and the first instruction of a new sequence is 950 ns.

Measurements made during the execution of a number of benchmark programs run on MU5 have shown that control transfers constitute around 13 per cent of obeyed instructions [28]. Thus without some attempt being made to overcome the effects of these control transfers, the average instruction execution time would be

$$50 * 87/100 + 950 * 13/100 = 158.5 \text{ ns}$$

which represents a reduction by a factor of over three from the peak execution rate.

A system was initially considered which had buffer registers containing the first few instructions at the destination or 'jump-to' addresses of recently obeyed control transfers. Access to these instructions was to have been via an associative search on their addresses. The pre-fetching mechanism would proceed normally until a control transfer occurred, and the destination address would then be presented to the associative store. If a match was found, the instructions in the corresponding buffer register would be read out and sent to the Primary Operand Unit (PROP). If no match was found, one of the set of associative and buffer registers would be updated when the instructions had been obtained from store.

Simulation studies of this technique showed that only eight lines of store would be needed to trap 80 per cent of jump instructions and that increasing the number of lines to sixteen would only produce an extra 1 per cent improvement. The problem in implementing this scheme was the width of the buffer store. In order to allow the pre-fetching mechanism to catch up after a control transfer, each line of the buffer would need to hold up to 950 ns worth of instructions. At 50 ns per 16-bit instruction the number of bits needed in each line would have amounted to over 300.

In order to retain the advantage obtained by using an associatively addressed store, without incurring the cost of

buffering large amounts of data, the system actually used in the MU5 Instruction Buffer Unit (IBU) involves an eight-line associatively addressed 'Jump Trace' store which attempts to predict the outcome of an impending control transfer. Whenever a new instruction address is generated by the IBU it is presented to the associative 'jump-from' address store before being sent to the Local Store via the Store Access Control Unit (SAC). If an equivalence is found, this address is replaced by the corresponding 'jump-to' address, so that pre-fetching of the new sequence takes place instead. When the control transfer instruction which gave equivalence in the trace is sent to PROP, it is accompanied by a bit indicating that the instructions following it are 'out of sequence'. This bit is used in PROP to determine the action after execution of the control transfer. If the following instructions have been correctly predicted, execution of instructions continues uninterrupted. If the instructions are not out of sequence, but should have been, a request is made to SAC for the instructions at the 'jump-to' address, and at the same time a line in the Jump Trace is loaded with the 'jump-from' and 'jump-to' addresses. Thus when the 'jump-from' address re-occurs within the IBU, the instructions at the 'jump-to' address are automatically pre-fetched.

Simulation studies of this system indicated that about 75 per cent of control transfers could be trapped using an eight-line store, and that, as before, increasing the number of lines in the store did not significantly improve the performance. The apparent drop in performance as compared with the first system considered occurs because the prediction mechanism sometimes predicts a transfer which does not occur. No attempt is made to correct the Jump Trace when a predicted branch does not occur, however, since the drop in performance is more than offset by the fact that the prediction mechanism allows useful overlapping of instructions to continue in PROP when the prediction is correct.

### 5.3.1 Sequential Instruction Fetching

The Store Request System (figure 5.5) is responsible for initiating requests for 128-bit words from SAC at the required intervals. Two different types of request may occur according to circumstances, ordinary requests and priority requests. Priority requests are made whenever PROP signals a discontinuity in the instruction stream, that is, when the instructions following the one currently being executed must be replaced by a different sequence. Following such an event, the instruction address received from PROP is loaded into both the Advanced Control Register, AC, and the Store Request Register, SR, and a priority request is made to SAC. AC is

then incremented at 40 ns intervals (it operates at a slightly faster rate than PROP) and whenever a carry across the 128-bit word address boundary occurs, the new address is copied into SR and an ordinary request is made to SAC. In addition, each new address generated in AC is checked against the contents of the associative jump-from field of the Jump Trace. If an equivalence occurs the corresponding jump-to address is read out and used instead. The Store Request System then continues to make ordinary requests starting from this new address.

Ordinary requests normally continue to be made until an instruction sequence discontinuity arises. This happens when an interrupt occurs, for example, or when the IBU has sent an incorrect sequence of instructions to PROP. This can occur either because an unpredicted control transfer instruction causes a jump, or because a predicted control transfer does not cause a jump. In either of these cases the priority mechanism is re-invoked when the address of the required instruction is sent from PROP. In the case of an interrupt two fixed (hard-wired) instructions are read from within the IBU. The first of these preserves essential linking information in store and the second causes a control transfer to the start of the appropriate interrupt routine. This again invokes the IBU priority mechanism.

The operation of the Store Request System can also be temporarily halted as a result of interlocks incorporated into the IBU to ensure that no loss of information occurs as a result of asynchronous operation. These interlocks simply cause time delays before the next ordinary request can be sent. For example, the rate of issuing of requests from the IBU to SAC is normally geared to the maximum rate at which PROP can process instructions. When PROP executes an instruction which requires a long interval of time for its completion, however, a hold-up occurs and this hold-up propagates back through PROP to the IBU Data Flow buffers. Since the IBU is obliged to accept data from SAC as soon as it becomes available, sufficient space must be maintained in these buffers to receive it. In order to meet this condition and to be able to maintain the maximum throughput rate, IBU ordinary requests may be held up within SAC until the IBU can guarantee to accept the requested instructions.

### 5.3.2 The Jump Trace

Although the Jump Trace is in principle a fairly straightforward system, a number of factors complicated its implementation. For example, although most control transfers use a literal operand (which is invariant in a machine like MU5 which uses 'pure procedure' code), some do not, and in

these cases the jump-to address will vary on some subsequent executions of the instruction. No attempt is made in MU5 to take advantage of cases where this variation occurs infrequently, however, since the designers felt that the extra hardware complication needed to check the correctness of the predicted jump-to address would not be cost-effective. Instead, the problem of variable jump-to addresses is avoided by only loading the Jump Trace for those control transfers which use a literal operand.

A further problem arises from the variable instruction length. The Control Register always addresses the first parcel of a multi-length instruction, but clearly jump-from addresses must always correspond to the last. This address must therefore be computed specially, since it is not otherwise required. However, the value in the Control Register after the execution of a conditional control transfer will either be the address of the first parcel of the next instruction in sequence (the jump-from address + 1), or some quite different address (the jump-to address). The first alternative can be generated immediately, since it requires no operand, whereas the second must await the arrival of the operand and the result of the condition. The Control Adder associated with the Control Register in PROP can therefore perform two cycles with no loss of performance, and when an unpredicted control transfer using a literal operand occurs, both values are sent to the IBU. The first is loaded into register AC (figure 5.5) and then decremented by the address counter before being used to load a line in the jump-from field. The second is loaded into AC and thence into the jump-to field and also into SR to be sent to SAC as a priority request. The line used in the Jump Trace is selected by a Fill Pointer according to a cyclic replacement algorithm and as each line is overwritten a 'use' digit associated with it is set. These use digits are normally only re-set, and the Trace thereby cleared, at a process change.

The technological problems which occurred during the implementation of the Jump Trace were mainly concerned with timing. All instruction addresses are presented to the jump-from field as they are generated by the Advance Control Register (AC), and the association proceeds in parallel with the operation of the counter. This allows the 40 ns address generation rate to be sustained. However, if an equivalence is found, the next address is not that generated by the counter but that read out of the jump-to field. This action requires an additional 40 ns. To sustain an average 40 ns rate, advantage is taken of the fact that the addresses generated in AC are sequential, and association is performed simultaneously on two addresses differing only in their least significant

digit. Thus the associative field is actually 30 bits wide with the 31st bit of the full virtual instruction address being held in a separate non-associative flip-flop. This bit is then used to determine whether or not an equivalence is genuine. When a genuine equivalence does occur, the corresponding jump-to address is read out and copied into AC and into SR.

**5.4 THE CRAY-1 INSTRUCTION BUFFERS**

The need for instruction pre-fetching and loop catching techniques arises from the impracticability of having a large high-speed random access store close to the central processor, and all the instruction buffering techniques with which we have dealt so far have been aimed at circumventing this problem. With the advent of MSI and LSI storage technologies, however, different hardware organisations have been made possible. An example of this can be seen in the instruction buffering mechanism used in the CRAY-1 [34], a machine with which we shall deal in more detail in Chapter 6.




Figure 5.6 CRAY-1 Instruction Buffers

The overall structure and instruction set of the CRAY-1 are largely derived from the CDC 6600 and 7600, with instructions

being executed by a set of parallel functional units. The processor obtains all its instructions directly from a set of instruction buffers, however, and does not send a stream of instruction requests to the main store. These buffers are organised as shown in figure 5.6. Each of the four buffers holds 64 consecutive 16-bit instruction parcels, and if an instruction request cannot be satisfied from within these buffers, a full 64-parcel block of instructions is transferred from main store into one of them.

The CRAY-1 uses a 22-bit instruction address and the first instruction parcel in a buffer always has an address starting on a 64-parcel address boundary. Any one buffer is therefore defined by the 16 most significant bits of a parcel address, and for each buffer there is a 16-bit starting address register containing this value. At each clock cycle the high order bits of the program address counter are compared with the contents of these registers, and if a match occurs the required instruction parcel is selected from within the appropriate buffer either immediately, if the buffer concerned is the same as the one which supplied the previous parcel, or after a two clock period delay if a change of buffers is involved.

If no match occurs, instructions must be loaded into one of the instruction buffers before execution can continue. A two-bit counter is used to determine which buffer is to be loaded; this counter is incremented by one whenever a load operation occurs, thus implementing a cyclic replacement algorithm. The 64-bit main store in the CRAY-1 is an 8-way or 16-way interleaved bipolar semiconductor store having a 50 ns cycle time. During a block transfer all other store requests are inhibited, and sequential accesses can be made at a rate of one per 12.5 ns clock period. In the case of transfers to an instruction buffer, four storage banks can be accessed in parallel, giving access to 16 instruction parcels in one cycle and allowing all 16 banks in a 16-bank configuration to be accessed in four clock periods. Since the cycle time is also equal to four clock periods, the first four banks are then ready to accept a further request, and a complete block transfer to an instruction buffer occupies four cycles of each bank. The total time required to access the first group of instruction parcels is nevertheless quite long, and a 14 clock period delay is incurred whenever a buffer has to be loaded. This delay is constant regardless of the position of the first parcel required from the buffer, since a technique is employed similar to that in the IBM System/360 Model 85 cache, whereby the first group of 16 parcels delivered to the buffers is always the one required immediately by the processor. Subsequent groups arrive at a rate of 16 parcels per clock

period and fill the buffer circularly.

When a branch is taken the new value in the program address counter is compared with the contents of the buffer starting address registers in exactly the same way as it is following execution of instructions in sequence. If a match occurs the required instruction is selected from the appropriate buffer, and if not a block transfer is initiated. Non-contiguous segments of code within a loop, or even separate subroutines may be held concurrently in separate buffers. The buffer contents are only invalidated, by having their starting addresses set to all ones, when an Exchange Jump occurs.

### 5.5 POSITION OF THE CONTROL POINT

In any computer executing instructions in a sequential manner, there is a register (called variously the Program Counter, Instruction Address Register, or Control Register) that contains the address of the instruction currently being executed. In a pipelined processor this Control Register points to an instruction at one particular stage of the pipeline, the Control Point. All instructions before the Control Point are being decoded, having their operands fetched, and so on, in such a way that they can be abandoned at any time, without causing any irreversible change to the state of the machine. Those after the Control Point have had the Control Register incremented past their address, and so must be guaranteed to complete their execution, even in the event of an interrupt.

Depending on the position of the Control Point, the pipeline can be described as 'early control', with the Control Point at an early stage in the pipeline (often the first), or 'late control', with the Control Point several stages further on. The position of the Control Point depends on a number of factors including the instruction set, hardware organisation, and the nature of program addresses (real or virtual, for example). Thus the parallelism of the functional units in machines like the CDC 6600 indicates an early Control Point, while the pre-processing required for an MU5 virtual address is easier to perform with a late control pipeline. Within these guidelines, there is a degree of flexibility in the position of the Control Point, and designers have to take other factors into consideration.

The effects of control transfers are particularly important. Control transfers can be divided into conditional control transfers, where the branch only occurs if some condition within the processor is satisfied, and unconditional control transfers, which always branch. Clearly, a conditional

control transfer instruction may be successful or unsuccessful, depending on the state of the condition when it is executed, and the relative proportions of successful conditional, unsuccessful conditional, and unconditional control transfers are important to the performance of the instruction fetching strategy used in the processor.

Control transfer instructions are normally executed at the Control Point stage in the pipeline. Since control transfers are often followed down the pipeline by wrong sequences of instructions, the delay incurred when these instructions have to be discarded increases in proportion to the number of stages in the pipeline between the main store and the Control Point. A late control pipeline would therefore seem to imply long delays for control transfer instructions. However, a closer examination reveals a more complicated situation.

The decoding of any instruction in the pipeline normally takes place before the final execution of some preceding instructions. This produces an information gap, or 'Gulf of Ignorance', between instructions entering the pipeline and those completing execution. This particularly affects conditional control transfers, since their execution depends on a value or condition determined by some previous instruction, or, in the case of machines like the CDC 6600, by the control transfer instruction itself. If this value is calculated by an instruction within the Gulf of Ignorance, then the control transfer is 'unresolvable' and cannot proceed until the required instruction completes execution. 'Resolvable' control transfers are those which are unconditional or which depend on the outcome of a long-passed instruction, outside the Gulf of Ignorance.

Thus the proportion of control transfers that are resolvable is very important, as a resolvable control transfer can be obeyed as soon as it is decoded. Unresolvable control transfers must normally be held at the Control Point until the required condition or result has been evaluated, usually at, or near, the end of the pipeline in one of the arithmetic units, and a gap will occur in the flow of instructions through the pipeline. The length of this gap is determined by the number of pipeline stages between the Control Point and the later stage where the arithmetic is performed. (It is this gap which the Conditional Mode of the IBM System/360 Model 195 is designed to overcome, in cases where the outcome of the conditional transfer has been correctly predicted. When it has not the Instruction Address Register at the Control Point has to be wound back to take account of the instructions issued but not actually obeyed.)



Once the condition has been resolved, the control transfer can be obeyed. Unless the appropriate sequence is already following the control transfer down the pipeline, however, a further delay will be incurred while the new sequence of instructions is requested from store (or the instruction buffer). This delay will be equal to the instruction access time plus the time to pass through the pipeline from the store to the Control Point. Thus, for a conditional control transfer followed by the wrong sequence of instructions, the delay subsequently incurred depends on the numbers of pipeline stages both between the main store and the Control Point and between the Control Point and the arithmetic unit, that is, the total pipeline length. However, if the correct instruction sequence following the conditional control transfer has been supplied to the pipeline, then this delay is only proportional to the number of stages in the pipeline beyond the Control Point.

On at least some occasions conditional control transfer instructions will be followed by the correct sequence (when they are unsuccessful in a simple system), and in these cases the pipeline delays are reduced if the Control Point is as late as possible in the pipeline. In the absence of any special techniques for supplying correct instructions after unconditional transfers, however, the pipeline delays for unconditional transfers are reduced by placing the Control Point as early in the pipeline as possible. Clearly, the relative proportions of these instructions are important. Measurements made using the MU5 hardware performance monitor produced the results shown in figure 5.7 for the proportions of different types of control transfer instruction obeyed during compilation and execution of 10 Algol and 15 FORTRAN benchmark programs [36].

During Algol execution 14.0 per cent of all instructions obeyed are control transfers, made up of 6.0 per cent successful conditional, 2.4 per cent unsuccessful conditional and 5.6 per cent unconditional transfers. For FORTRAN execution 12.5 per cent of all instructions are control transfers, made up of 4.5 per cent successful conditional, 2.8 per cent unsuccessful conditional, and 5.2 per cent unconditional transfers. During both Algol and FORTRAN compilation the proportions of control transfers are higher, and the success rates of the conditional transfers are lower. Both of these differences arise from the nature of the tasks. The execution figures include large numbers of loops of numerical calculations ended by successful conditional control transfers, while the data-dependent nature of the compilation task involves many unsuccessful tests for particular values of the input data.

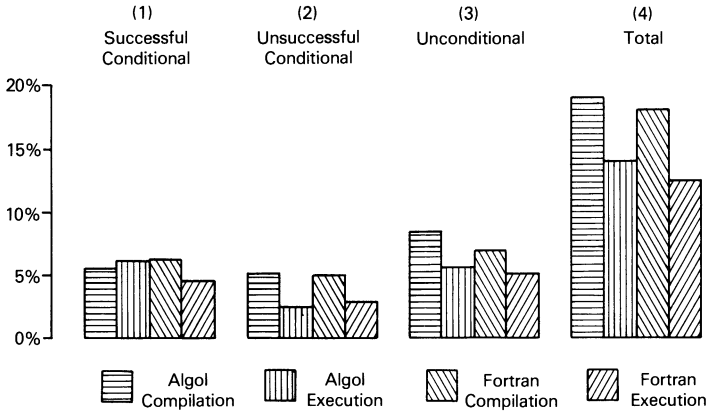


Figure 5.7 Percentages of Control Transfers

The relative numbers of predictable and unpredictable control transfers are also important in MU5, since only predictable control transfers (those which use a literal operand) cause a Jump Trace entry to be made following their first successful execution. The results in column 1 of figure 5.8 show that the per centage of these predictable control transfers (including conditional and unconditional) is fairly constant at over 85 per cent for all the types of program considered.

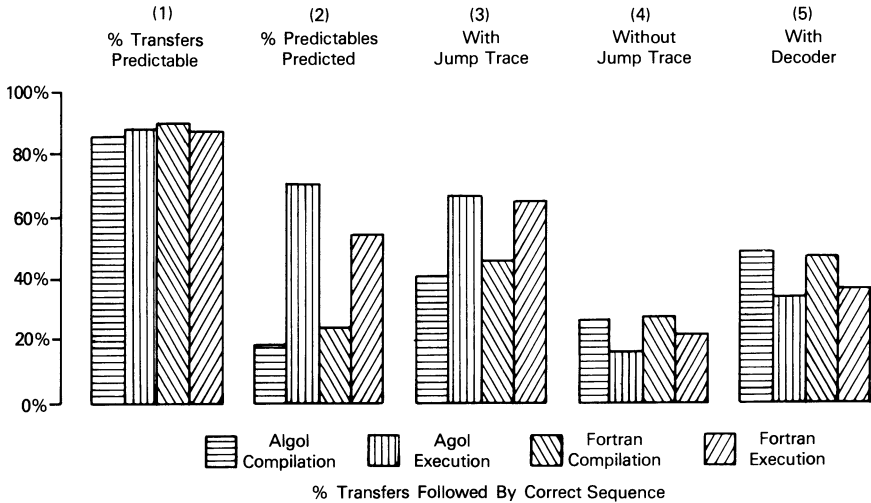


Figure 5.8 Control Transfer Prediction Rates

Column 2 of figure 5.8 shows the proportion of predictable transfers that are correctly predicted by the Jump Trace. The figures for execution are much better than for compilation, and again the nature of the tasks accounts for this difference. During execution, the Jump Trace appears to offer less advantage for FORTRAN than Algol programs, even though many of the benchmarks actually correspond to the same task carried out in each language. When the per centages of control transfers followed by correct sequences of instructions are considered, however, as in column 3, the lower success rate of conditional transfers in FORTRAN compensates for this effect. The figures in column 3 may be compared with those in column 4, which shows the per centages of control transfers what would be followed by correct sequences in the absence of the Jump Trace.

An alternative system that might seem to produce useful improvements in the numbers of correct sequences is one that decodes and executes predictable unconditional control transfers at an earlier stage in the pipeline than the Control Point. Such a scheme is used in the more powerful machines in the ICL 2900 series of computers. Using MU5 figures for the numbers of unsuccessful conditional control transfers and the predictable unconditional transfers, the per centages of control transfers followed by correct sequences using the decoder scheme alone, would be as shown in column 5 of figure 5.8. The effects for compilation are remarkably similar to those obtained using the Jump Trace, though clearly for execution this scheme is less attractive. Furthermore, large gaps will occur in the pipeline unless the decoding and execution of an unconditional control transfer occurs sufficiently far in advance of the Control Point for the new sequence to be accessed from store and supplied to the pipeline before the preceding instructions have all been processed.

A further set of measurements taken from MU5 concerns resolvability. Conditional control transfers in MU5 depend on the outcome of a previously executed COMPARE order, the result of which is held in a separate Test Register, and measurements of the numbers of conditional control transfers separated by a given number of instructions from the preceding COMPARE order were used to determine the per centages of conditional control transfers that would be resolvable for a given pipeline length (figure 5.9). Clearly, for a pipeline separation of more than one stage between the Control Point and the arithmetic unit carrying out the comparison, most conditional control transfers are unresolvable, even though the MU5 compiler writers attempted to separate comparison and control transfer orders wherever possible in the compiled code. If all

conditional control transfers are assumed to be unresolvable in MU5, their proportion of the total number of control transfers amounts to very close to 60 per cent for each type of program considered. This figure agrees very closely with that obtained by Flynn from measurements taken from an IBM System/360 machine [35]. Here the overall proportion of control transfers is higher than for MU5 at 9 per cent resolvable and 16.5 per cent unresolvable (27.5 per cent in all), but the proportion of control transfers which are unresolvable is also 60 per cent.

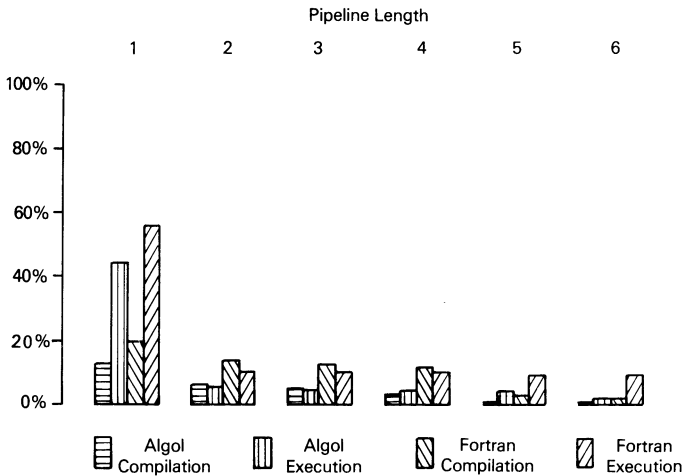


Figure 5.9 Percentages of Resolvable Control Transfers

The performance of pipelined processors is critically dependent on the effects of control transfers, as we saw in section 5.3, for example. What is important is the total length of the pipeline and simply designing for greater overlap within a processor will not necessarily improve performance. For a given pipeline length, however, performance can be improved by placing the Control Point as late as possible in the pipeline and supplying the pipeline with the correct sequence of instructions after control transfers as frequently as possible. This necessarily involves the prediction of the outcome of a control transfer, and this must normally be made as early as possible in order to overcome store access delays. A system such as the MU5 Jump Trace, which bases its prediction on the addresses used to pre-fetch instructions, seems to offer the best chance of success in this situation.

# *6 Parallel Functional Units*

In Chapter 4 we saw an example of the use of parallel functional units in the IBM System/360 Model 91. In this machine the parallelism was introduced as a means of enhancing the performance of a processor at the top end of a range of general purpose computers. In the case of the CDC 6600, performance was the principal criterion of the design, and as we saw in Chapter 2, the use of parallel functional units with an instruction set capable of exploiting this parallelism was a key feature. In this chapter we shall consider the design of the CDC 6600 central processor, with particular emphasis on the Scoreboard, the mechanism used to control the operation of these functional units. We shall then go on to consider some of the design modifications introduced in the CDC 7600, the successor to the 6600, and then see how limitations inherent in the 6600/7600 architecture led to the CRAY-1 design. We shall consider the design of this machine in some detail.

## **6.1 THE CDC 6600 CENTRAL PROCESSOR**

The overall design of the CDC 6600 central processor was introduced in section 2.1, and is shown again in figure 6.1. Instructions are taken in sequence from the Instruction Stack and issued by the Scoreboard to the appropriate execution unit. Each unit takes its input operands from among the 24 scratch-pad registers (eight 60-bit X (operand) registers, eight 18-bit A (address) registers, and eight 18-bit B (index) registers) and returns its result to one of these registers. The maximum rate of issue is one instruction per minor clock cycle (100 ns), while the units take typically 300 or 400 ns to complete their operations. Thus while the units can in principle all operate simultaneously, simultaneous operation of three or four units is more typical in practice. The issuing of instructions is not straightforward since instruction dependencies can require that some instructions be held up until the completion of other, previously issued, instructions. In order to maximise the amount of concurrent processing, the Scoreboard is designed to issue each instruction as early as possible, within the limits set by these dependencies, so as to allow the following instruction to be issued, hopefully to a different unit. Dependency

conflicts are resolved by the Scoreboard through control signals which link it to each unit and by the use of information buffered about the registers in respect of each unit, and about the units in respect of each register. Before discussing the details of this system, however, we shall consider the characteristics of the units themselves.

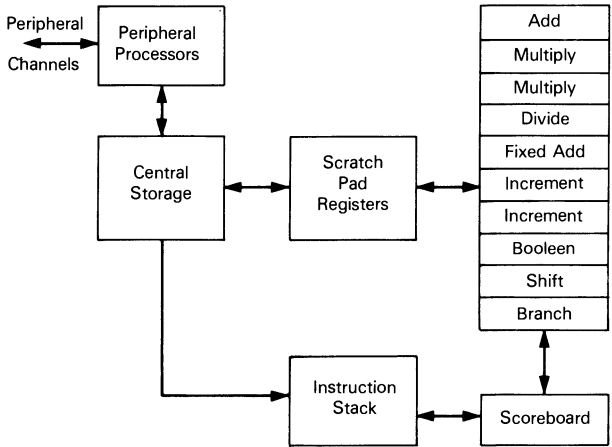


Figure 6.1 CDC 6600 Central Processor Organisation

### 6.1.1 Functional Units in the CDC 6600

#### Boolean Unit

The Boolean Unit performs logic operations of the type

$$X_i = X_j \& X_k$$

$$X_i = X_j \vee X_k$$

It operates on each pair of input digits in parallel and returns its result to the designated result register in 300ns.

#### Fixed Add Unit

The Fixed Add Unit performs the operations

$$X_i = X_j + X_k$$

$$X_i = X_j - X_k$$

It treats the input operands as 1's complement numbers and uses a 60-bit parallel adder which allows it to return its

result to the designated result register in 300 ns. It is also used as a partner to the Branch Unit when the latter is executing a conditional branch instruction which depends on a value in an X register.

### Shift Unit

The Shift Unit performs a variety of operations. The shifter itself is organised as a 6 logic level network capable of performing left circular and right arithmetic shifts to any position in 300 ns. It executes the operations Pack and Unpack, which respectively couple and separate the exponent and mantissa of a floating-point number using registers Xi, Bj, and Xk, and Mask, which forms a string of ones for use by the Boolean Unit in AND operations used for masking. The Shift Unit also executes the Normalise operation which takes 400 ns. In most computers the results of floating-point operations are automatically normalised, but this increases the time required for these operations. Normalisation is treated as an optional extra in the 6600, for which the user only pays the time penalty when the facility is used.

### Floating Add Unit

The Floating Add Unit also uses 1's complement representation for the 48-bit mantissae of floating-point numbers, but uses an 11-bit biased binary exponent (an exponent value of 0 is represented as 10000000000). The mantissae are assumed, unusually, to be integers rather than fractions. This has the advantage of allowing fixed-point integers to be converted to floating-point numbers by simply ORing in the exponent bias. Thus the floating-point number format is as shown in figure 6.2.




Figure 6.2 CDC 6600 Floating-point Format

Floating-point addition and subtraction can be carried out in 400 ns and the result may be a rounded or unrounded single-length result, or the upper or lower half of a double-length result.

**Floating Multiply Unit**

The 6600 boasts two identical Floating Multiply Units, each capable of producing a rounded or unrounded single-length or the upper or lower half of a double-length floating-point result in 1000 ns. This unit is a very interesting example of a high performance arithmetic unit, combining as it does the techniques of carry-save addition, multiplier pairs and split multiplier operation. However, we shall not be considering the details of arithmetic unit design in this book; the interested reader is referred to Thornton [7] for details of the 6600 arithmetic units, or to Gosling [37] for the design of arithmetic units generally.

**Floating Divide Unit**

Floating-point division is the longest arithmetic operation in the 6600, as it is in most computers, and requires 2900 ns to return a result to the designated result register.

**Increment Unit**

Each of the two Increment Units performs fixed-point addition or subtraction on 18-bit numbers in 300 ns. These units are used for indexing and for the loading of A registers (thereby causing transfers of operands between the corresponding X registers and Central Storage) in operations such as

$$\begin{array}{ll} A_i = A_j + K & B_i = A_j + K \\ A_i = A_j - B_k & B_i = A_j - B_k \end{array}$$

Either of these units may also be used as a partner to the Branch Unit when the latter is executing a conditional branch instruction which depends on a value in a B register.

**Branch Unit**

The result register of the Branch Unit is the Program Address register, which is overwritten by an unconditional branch or a conditional branch for which the condition is met. Determination of the condition is carried out by a partner unit, an Increment Unit in the case of a B value being involved or the Fixed Add Unit in the case of an X value. The Branch Unit itself performs the 'jump within stack' test (section 5.2.1) to determine whether or not the target instruction of a branch is in the Instruction Stack. Branch instructions take 800 or 900 ns to complete if the target instruction is in the Instruction Stack, and 1400 or 1500 ns if not. They also hold up the issuing of further instructions until they have completed.



### 6.1.2 Instruction Dependencies

Dependency conflicts between successive instructions in the 6600 are classified by Thornton [7] into three types, first order, second order and third order. The Scoreboard resolves these conflicts using the control signals which link it to each unit. The Scoreboard can send an Issue, Go Read and Go Store signal to each unit separately, and can receive from each unit a Request Release signal (figure 6.3). The Issue signal enters information into the unit identifying the source of its input operands and the mode in which it is required to operate, while Go Read causes the operands themselves to be copied into the unit. When a unit completes its operation it sends a Request Release signal to the Scoreboard, which responds with Go Store, a signal which allows the result of the operation to be copied out of the unit's Temporary Register and into the appropriate result register.

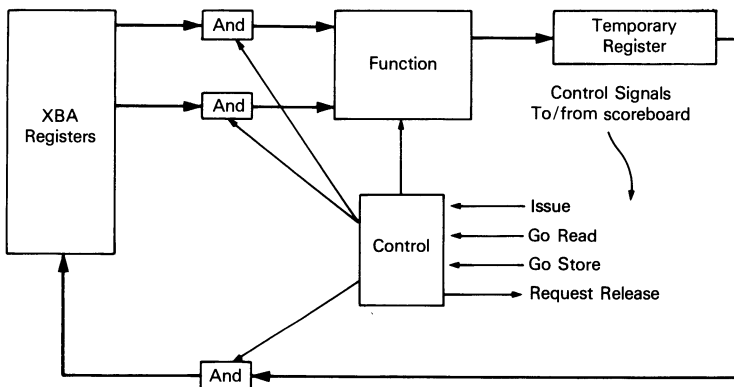


Figure 6.3 Control Schematic for a 6600 Functional Unit

#### First Order Conflicts

A first order conflict occurs whenever an instruction which is about to be issued requires the use of an arithmetic unit or a result register which is already in use or has been reserved by a previously issued, but as yet uncompleted instruction. Each instruction in the following pair, for example, requires the use of the Floating Add Unit

$$\begin{aligned} X6 &= X1 + X2 \\ X5 &= X3 + X4 \end{aligned}$$

while in the next example each instruction requires X6 as its result register

$$\begin{aligned} X6 &= X1 * X2 \\ X6 &= X4 + X5 \end{aligned}$$

Although this latter example is unlikely to arise in normal programming practice, it must nevertheless give the correct result. Without proper interlocks the add operation would complete first and the result in X6 would then be overwritten by that of the multiplication.

Both these conflicts are resolved by holding up the Issue signal for the second instruction until the first has completed. Issuing an instruction involves a sequence of four separate actions. Firstly the functional unit required by the instruction is reserved by the setting of its 'busy flag', and the operating mode, derived from the m field of the instruction, is entered into it. No subsequent instruction can be sent to this unit until its busy flag has been re-set at the completion of the instruction which reserved it. Secondly register designators F<sub>i</sub>, F<sub>j</sub> and F<sub>k</sub>, derived from the i, j and k fields of the instruction, are copied into the functional unit in order to identify the operand and result registers which it will use. Thirdly the Scoreboard copies into the current unit two numbers, Q<sub>j</sub> and Q<sub>k</sub>, taken from the identifier registers associated with these operand registers, and finally the identifier register associated with the result register is loaded with the unit number of the current unit. Thus a subsequent instruction which requires the content of this result register as an input operand will receive this unit number as a Q number, and a subsequent instruction which also requires this register as a result register will not be issued until the identifier has been cleared. The importance of Q numbers becomes clear when second order conflicts are considered.

### Second Order Conflicts

A second order conflict occurs whenever an instruction which is about to be issued requires the result of a previously issued but as yet uncompleted instruction. An example of such a conflict is the following

$$\begin{aligned} X6 &= X1 + X2 \\ X7 &= X5 / X6 \end{aligned}$$

Here the second instruction can be issued, but must not be allowed to start until the result of the first instruction has been entered into X6. This is achieved by holding up the Go Read signal. The first action within a functional unit at the start of any operation is the simultaneous copying into the

unit of the two input operands, and this can only occur when Go Read is set. Go Read is the logical AND of the Read Flags associated with each of the two input operands.

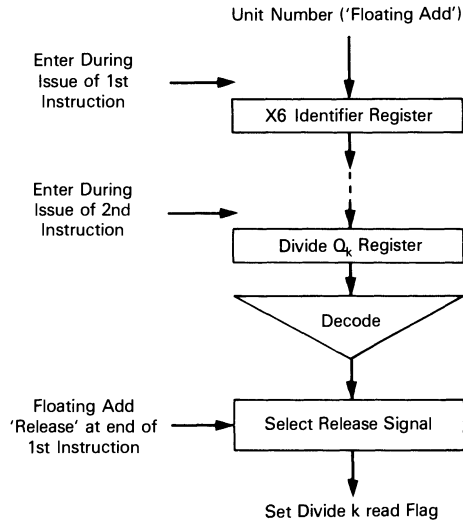


Figure 6.4 Second Order Conflict Event Sequence

The sequence of events for the example above is shown in figure 6.4. The unit number of the Floating Add Unit is entered into the identifier register of X6 when the first instruction is issued. When the second instruction is issued this unit number is copied from the X6 identifier register into the Qk register in the Divide Unit. When the Floating Add Unit completes its operation it sends its Request Release signal to the Scoreboard, and on receiving a Go Store response, sends a Release signal to all the other functional units. Because the Divide Unit has the unit number of the Floating Add Unit set in its Qk register, it detects this Release signal and sets the k operand Read Flag. Assuming that the Qj register contains zero (indicating that there was no outstanding result destined for X5 at the time when the divide instruction was issued), the Read Flag for the j operand would have been set immediately, and the divide operation therefore starts as soon as the Read Flag for the k operand is set.

### Third Order Conflicts

A third order conflict occurs when an instruction which has just completed its operation wishes to store its result in a register which is waiting to supply an input operand for a

previously issued, but as yet unstarted instruction. Such a conflict occurs in the following sequence

$$\begin{aligned} X3 &= X1 / X2 \\ X5 &= X4 * X3 \\ X4 &= X0 + X6 \end{aligned}$$

The third order conflict here is on register X4, and arises because of the second order conflict on X3. The second instruction can be issued immediately after the first, but is held up for its Go Read signal because X3 cannot be read until the Divide Unit completes its operation. The third instruction can likewise be issued immediately after the second, and since there are no result reservations on its input operands, the Floating Add Unit also receives Go Read immediately and starts its operation. The floating-point add operation completes in very much less time than division, however, and the Floating Add Unit is therefore ready to store its result in X4 before the Multiply Unit has read the current value in X4. Thus the Floating Add Unit sends its Request Release signal to the Scoreboard, but the Scoreboard holds up the Go Store response until after the multiplication has started.

This interlock is achieved through use of the Read Flags. When the multiplication instruction is issued there are no reservations on X4, and its Read Flag is set immediately. This Read Flag remains set until the X3 Read Flag is set at the end of the division operation, after which the multiplication is started and both Read Flags are re-set. Thus the fact that the value in a particular register is required as an input operand by a previously issued, but as yet unstarted, instruction is indicated by the presence of a Read Flag for that register. Conversely, if there are no Read Flags set for a particular register, that register can be described as 'All Clear' and a Go Store can be returned to a unit which is ready to send a result to it. The All Clear signal is produced by decoding the Fj and Fk operand designators for each unit and ANDing the decoder outputs with the corresponding Read Flag signal to produce an indication of whether or not the unit has a Read Flag set for a given register. For each register the signals from all the units are then combined to produce the All Clear signal indicating whether or not there are any Read Flags set for it. The Go Store signal for a particular unit is formed by decoding its Fi (result register) designator and combining together the results of AND operations between these decoded Fi signals and the corresponding All Clear signals for the registers.

In the example above the Fi designator for the Floating Add Unit identifies register X4, and until the multiplication

starts the Multiply Unit has a Read Flag set for X4, inhibiting the X4 All Clear signal. When the multiplication starts, this Read Flag is re-set, the X4 All Clear signal becomes true, and the Floating Add Unit receives Go Store. It then sends its result and Fi value along a data highway to the XBA registers where the Fi value is decoded to control entry of the result into the appropriate register, and the clearing of the result reservation in the corresponding identifier register.

### 6.1.3 Data Highways

The instruction times quoted in section 6.1.1 refer to the time between the issuing of an instruction and the availability of its result in a register. Included in these figures is the time required to transfer operands to a unit at the start of an operation and the time to transfer the result back to the appropriate X, B or A register at the end. These times add up to 100 ns typically, so that the 'highway' time amounts to as much as one-third of the total time required to execute short operations. Furthermore, because the Go Read signal may be held up for a particular instruction, the transfer of operands into the unit does not necessarily take place concurrently with the issuing of an instruction to that unit, and may be required to occur concurrently with the transfer of operands for a subsequent instruction issued to a different unit. Thus in order to maintain the maximum benefit of functional unit parallelism, a multiplicity of highways is required between the XBA registers and the functional units.

Figure 6.5 shows the data highways used in the 6600 Central Processor. The functional units are grouped together according to physical layout and performance criteria and each group is served by a separate highway. The two Increment Units, for example, can cause a traffic of one operand on a highway every two minor cycles, on average, while the two Multiply Units and the Divide Unit can together cause an average traffic of approximately one operand on a highway every four minor cycles, as can the Boolean Unit. Thus the requirements of the Multiply, Divide and Boolean Units taken together roughly match those of the Increment Units. A dual 18-bit highway is therefore used to supply operands to the Increment Units, while two dual 60-bit highways are used to supply each of the other two groups of units, and a single 60-bit highway is used to return operand values from the XBA registers to Central Storage.

The XBA registers are supplied with operands from Central Storage via a single 60-bit highway, and can receive results from the Multiply group of units via a 60-bit highway, from

the Increment Units via an 18-bit highway, and from the Add group of units via a dual 18-bit plus 60-bit highway. These latter highways would be used in an Unpack instruction, for example, where the mantissa of a floating-point number is returned to an X register, and the exponent to a B register.

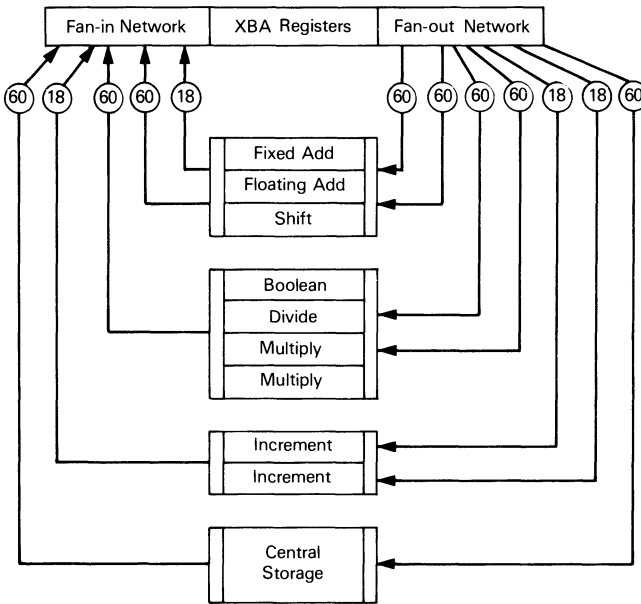


Figure 6.5 Data Highways in the CDC 6600

The highways form a considerable part of the processor hardware and represent a significant cost penalty to be offset against the performance advantage obtained from the functional unit parallelism. The Scoreboard control logic also represents a cost penalty, but the amount of hardware which it involves is rather less than that in an average functional unit.

### 6.2 THE CDC 7600 CENTRAL PROCESSOR

As we observed in Chapter 5 when discussing instruction buffering techniques, the CDC 7600 was designed to be machine code upward compatible with the 6600, but to provide a substantial increase in performance. This increase was achieved principally by a reduction of the minor cycle clock period from 100 ns to 27.5 ns (made possible by advances in technology), together with a corresponding reduction of the major cycle clock period from 1000 ns to 275 ns, and by architectural changes such as the modifications to the

Instruction Stack which we have already encountered, and the segmentation, or pipelining, of the functional units.

The overall design of the 7600 (figure 6.6) is in many ways similar to that of the 6600. Instructions are fetched from the Small Core Memory (SCM) and placed in the 12-word Instruction Stack before being copied into the Current Instruction Word (CIW) register for decoding and issue to one of the nine functional units. The control logic associated with the CIW register (section 6.2.2) serves the same purpose as the Scoreboard in the 6600, but is somewhat less complex. On early models of the 7600 the SCM consisted of either 32K or 64K words of 275 ns cycle-time core store made up from 16 or 32 independent 2K word banks. On current models the SCM consists of either 64K or 128K words of semiconductor memory made up from 32 or 64 independent 4K word banks. Although quite different internally these memories appear virtually identical to the rest of the system and because of the interleaving both can transfer sequentially addressed data at a rate of one word every 27.5 ns.

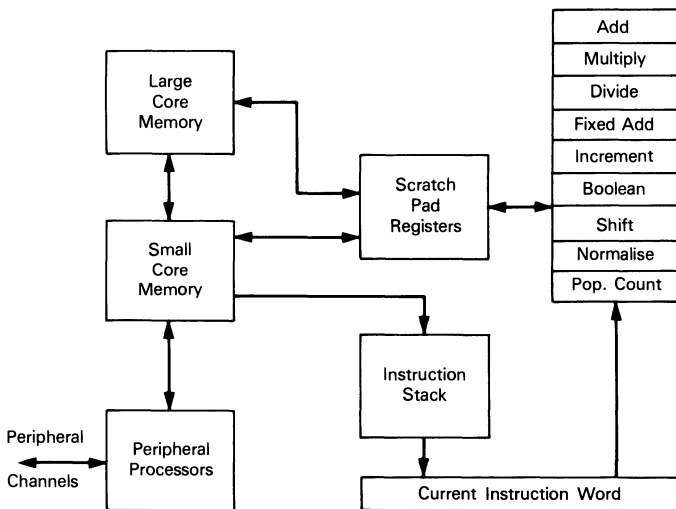


Figure 6.6 CDC 7600 Central Processor Organisation

All instruction accesses and most operand accesses are made to the SCM, and the SCM also communicates with the Large Core Memory (LCM) and with external devices through the Peripheral Processors. The LCM serves as a first level of backing store for the SCM, with blocks of code and data being exchanged between these two, when required, under direct program control. The LCM consists of 256K or 512K 60-bit words in four

or eight banks of 1760 ns, word-organised (2D) core store. Each bank operates independently and is eight words wide; once an access has been made to a particular bank the data read out remains available in an output buffer until a subsequent access is made to the same bank. In an eight-bank model this arrangement allows words to be accessed at a rate of 64 per LCM cycle. This time (1760 ns) is equal to 64 minor cycle clock periods, and block copy operations between the LCM and the SCM (in which LCM banks are accessed sequentially) can therefore proceed at the SCM maximum transfer rate of one word per minor cycle.

### **6.2.1 Functional Units in the CDC 7600**

Most of the functional units in the CDC 7600 perform identical functions to those in the 6600. The pipelining of 7600 units obviates the need for the duplication of units found in the 6600, however, and control transfers are dealt with by the control logic associated with the CIW register rather than in a separate Branch unit. Two extra units are included, since the normalise operation is carried out in a separate unit from the shift operations, and a separate Population unit is used to carry out the population count function. In the 6600 this function was implemented within the Divide Unit.

#### **Boolean, Fixed Add, Shift and Increment Units**

The Boolean, Fixed Add, Shift and Increment Units in the 7600 carry out similar operations to those in the equivalent units in the 6600, but each is pipelined into two stages. Input operands are copied into one of these units at the start of one minor cycle clock period, and a second pair of operands may be copied in at the start of the next minor cycle. At the end of this second minor cycle the result of the first operation is copied into the designated result register.

#### **Floating Add Unit**

The Floating Add Unit operates as a four-stage pipeline, but in comparison with the TI-ASC arithmetic pipeline the partitioning is expedient rather than elegant. Thus whereas the TI-ASC pipeline consists of a number of functionally distinguishable sections, the pipeline stages in the 7600 Floating Add Unit each contain as much logic as can be accommodated, in terms of logic gate delays, within one clock period. The mantissa shifting logic and the mantissa adder, for example, are each split between successive pipeline stages. It is interesting to note that although a new floating-point addition can be started at each minor cycle clock period, the time for any one addition occupies four



minor cycles, just as in the 6600.

### **Floating Multiply Unit**

In the 6600 Multiply Unit the mantissa multiplier was split into two parts, each of which carried out a 24-bit by 48-bit multiplication. The results of these two parallel operations were then combined to form a 96-bit product. The 7600 Multiply Unit contains the equivalent of only one of these half multipliers, and this is used twice in a two-pass mode of operation. In the first pass the lower 24 bits of  $X_j$  are multiplied by all 48 bits of  $X_k$  to form a partial product. This result is then shifted right 24 places and fed back into the multiplier to be combined with the result of multiplying the upper 24 bits of  $X_j$  by all 48 bits of  $X_k$  in the second pass. Thus although the Multiply Unit is pipelined in a similar fashion to the other functional units in the 7600, and clocked at every minor cycle clock period, a new multiplication can only be started in every alternate clock period. The total time for any one multiplication is 5 minor cycles, representing a factor of two improvement over the 6600 Multiply Units, each of which took 10 minor cycles to complete a multiplication.

### **Divide Unit**

Floating-point division is executed by a repeated subtract and test algorithm which cannot be pipelined. This operation therefore has a very much longer execution time than any other operation, amounting to 20 minor cycles. A second division can, however, be started 18 minor cycles after the start of a previous division.

### **Normalise Unit**

Normalisation is carried out in a separate functional unit in the 7600, rather than within the Shift Unit as in the 6600. This is because normalisation requires a three-stage pipeline, while shifting and mask generation can be carried out in a two-stage pipeline.

### **Population Unit**

The population count function counts the number of ones in the operand taken from  $X_k$  and returns this number as a result to  $X_i$ . In the 6600 this operation was carried out by hardware contained within the Divide Unit, but since it only requires two clock periods for its execution, a separate Population Count Unit, implemented as a two-stage pipeline, is used for this purpose in the 7600.

### 6.2.2 Instruction Issue

The introduction of pipelining into the functional units in the 7600 had important consequences for the instruction issuing mechanism. Thus whereas the Scoreboard in the 6600 issued instructions as soon as possible, and used the Go Read and Go Store interlocks to maintain proper sequential dependencies, the hold-ups which these interlocks caused could not be tolerated in the functional unit pipelines of the 7600. An instruction is therefore issued from the CIW register in the 7600 only when the conditions in the functional units and operating registers are such that the instruction may be carried through to completion without conflicting with a previously issued, but as yet uncompleted instruction. This condition occurs when the two signals 'Registers Free' and 'Enable Issue' are both in the TRUE state.

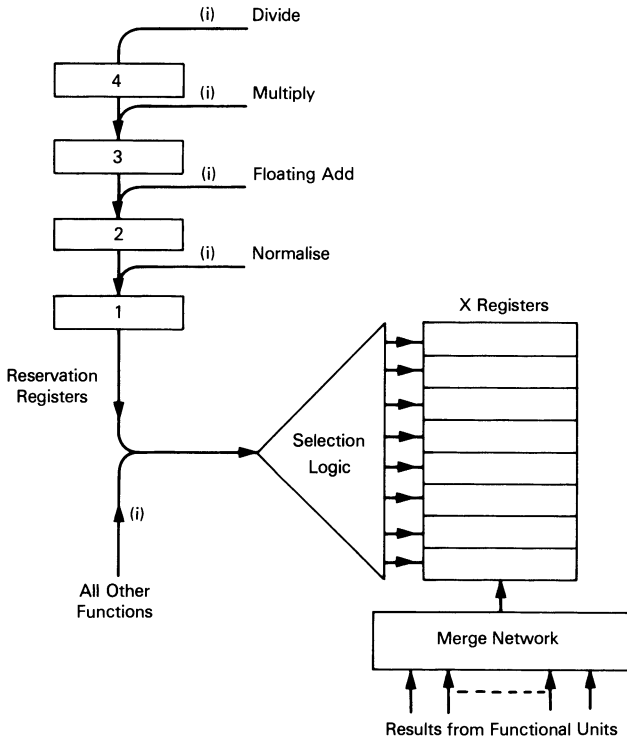


Figure 6.7 X Register Access Control Logic in the 7600

Associated with each of the X, B and A registers is a register busy flag. A busy flag is set to reserve a specific register when an instruction which will deliver a result to

that register is issued from the CIW register; the flag is re-set when the result appears and is copied into the register. Whenever an instruction in the CIW register is about to be issued, its operand and result register designators are used to select the busy flags associated with these registers. The states of these selected busy flags are then combined to form the Registers Free signal. This signal is only TRUE (allowing the instruction to issue) when there are no busy flags set for the selected registers; if a busy flag is set the issuing of the instruction is held up until the flag has been re-set. This mechanism prevents an instruction from reading or overwriting any register which is due to be updated by a previously issued, but as yet uncompleted, and possibly slower instruction.

The Enable Issue signal indicates that the states of the functional units, register destination paths and storage access path are such as to allow the current instruction in the CIW register to go to completion. The Multiply Unit, for example, cannot accept a new instruction until two minor cycles after a previous instruction, and the Divide Unit until after 18 cycles. The major problem to be overcome, however, is the occurrence of conflicts on the entry of results into the X registers, which can only accept one result per clock period. If all the functional units took the same time to produce their results this problem would not arise, of course, and does not arise in the case of the A registers, which only receive results from the Increment Unit. The X registers can receive results from any unit, however. Five of the units produce their results in two clock periods, one in three, one in four, one in five and one in 20 clock periods.

The entry of operands into the X registers is controlled by the X Register Access Control Logic, which contains four 4-bit Reservation Registers (figure 6.7). The normal situation is for a functional unit to produce a result two clock periods after receiving an instruction, and for its 3-bit *i* (result register) designator to be sent directly to the Selection Logic for use at the appropriate time. In the case of a normalise instruction, however, three clock periods are required, and instead of its *i* designator being sent directly to the Selection Logic, it is copied into Reservation Register 1, together with a fourth, 'valid' bit which indicates its presence. This causes the necessary extra clock period delay to be inserted into the path of the *i* designator en route to the Selection Logic. For a floating-point addition the *i* designator is entered into Reservation Register 2, and for multiplication it is entered into Reservation Register 3. For division the *i* designator is held in the Divide Unit until the 16th cycle in the divide sequence, at which point it is

entered into Reservation Register 4. At each minor cycle clock period information is copied through the Reservation Register chain towards the Selection Logic, with the input to any one of these registers being taken either from the previous register in the chain or from the CIW register when a new i designator is to be entered into it.

The Enable Issue logic uses the information held in the Reservation Registers to ensure mutual exclusivity of the two sources of information available to each Reservation Register. If a floating-point addition is issued in one clock period, for example, and the next instruction in sequence is a normalise instruction, then in the next clock period the presence in Reservation Register 1 of the valid floating-point add i designator inhibits the Enable Issue signal for the the normalise instruction and prevents its being issued until one clock period later. If it were not held up, it would produce a result in the same clock period as the floating-point add, and would also require entry into Reservation Register 1 for its i designator while that register contained the designator for the floating-point add. In the same way, a two-clock period instruction is prevented from issuing in the minor cycle immediately following the issue of the normalise instruction.

### 6.3 PERFORMANCE

Since the real work of the computer is carried out on floating-point numbers, with the fixed-point numbers being involved mainly in 'house-keeping' operations, it has become fashionable to gauge performance in terms of the rate of execution of floating-point operations, normally measured in millions of such operations per second (MFLOPS). In order to compare the architectures of the 6600 and 7600, however, it is more convenient to consider the number of floating-point operations which each can execute in one clock period (FLOPS/CLOCK). In the case of the 6600, floating-point addition/subtraction requires 4 clock periods, multiplication 10 and division 29. (Division may be ignored in this comparison; it always takes a long time, and computer designers have usually assumed that users will have the good sense to avoid using it as much as possible.) Thus, neglecting division, and allowing for the fact that there are two multiply units, the 6600 can perform 9 floating-point operations (5 additions and 4 multiplications) in 20 clock periods (corresponding to 0.45 FLOPS/CLOCK or 4.5 MFLOPS), with all three of the corresponding functional units being fully occupied. Since the Scoreboard is capable of issuing up to 20 instructions in this time, the right mix of instructions would allow this rate to be achieved. This rate also corresponds to the sum of the maximum rates of execution of

long sequences of additions (0.25 FLOPS/CLOCK) and multiplications (0.2 FLOPS/CLOCK) occurring separately.

Maintaining these rates requires the availability of the appropriate operands in the X registers, of course. Where these have first to be fetched from store, additional delays are incurred. However, the separation of the operand accessing and function execution facilities in the instruction set allows the possibility, at least, of programs (or compilers) organising appropriate pre-fetching of operands.

In the 7600 an individual floating-point addition or subtraction still takes 4 clock periods, as in the 6600, but because of the pipelining the maximum execution rate for these operations is 1 per clock period (1 FLOP/CLOCK). Individual floating-point multiplications take 5 clock periods, but as we have seen, the multiply unit is also pipelined and can perform multiplications at the rate of 1 every 2 clock periods (0.5 FLOPS/CLOCK).

The sum of these two rates produces a total maximum execution rate of 1.5 FLOPS/CLOCK. This rate cannot be achieved, however, since instructions can only be issued, and results entered into the X registers, at a rate of 1 per clock period at most. Thus the maximum floating-point execution rate is 1 FLOP/CLOCK, equivalent to 36.4 MFLOPS and made up, for example, of a sequence of additions (in which case the add unit is fully occupied and the multiply unit idle) or a sequence of alternate multiplications and additions (in which case the multiply unit is fully occupied and the add unit 50 per cent occupied). Even sustaining this rate for any length of time is virtually impossible, of course, since it does not allow for the execution of other instructions such as operand accesses and control transfers. However, being able to sustain the maximum rates for addition and/or multiplication for any period of time gives a performance bonus over the 6600 additional to the improvement in clock rate, and CDC claim that the overall performance of the 7600 is 15 million instructions per second.

The first CDC 7600 was delivered in 1969. By the mid 1970s technological advances offered the possibility of increasing the clock rate by a factor of about two, but in seeking to provide an increase in performance over that of the 7600 comparable with that which the 7600 had offered over the 6600, the designers (principally Seymour Cray) were faced with the problem of overcoming the instruction issuing bottleneck in the 7600 design. The solution was found in vector processing, and the architecture which resulted appeared commercially as the CRAY-1.

#### 6.4 THE CRAY-1

The problem of executing floating-point operations at a rate in excess of 1 per clock is overcome in the CRAY-1 processor [34] by the use of vector orders, which cause streams of up to 64 data elements to be processed as a result of one instruction issue. Vectors are contained in a set of eight V registers, each capable of holding 64 elements (each of 64 bits), and a typical vector instruction causes sets of operands to be taken from two V registers and the results to be returned to a third. In the following instruction sequence

```
V0 <- V1 + V2
V3 <- V4 * V5
```

the second instruction uses different registers and a different functional unit from the first and can be issued one clock period after the first instruction. Subsequent to the pipeline start-up delays in the functional units (each of which can carry out operations at a rate of one per clock period), a floating-point result will appear from both the adder and the multiplier in each successive clock period. Thus if performance is again estimated in terms only of floating-point addition and multiplication, the maximum floating-point execution rate is 2 FLOPS/CLOCK. Furthermore, around 60 other instructions can be issued before these units require further instructions to keep them busy. With a clock period of 12.5 ns, 2 FLOPS/CLOCK corresponds to 160 MFLOPS.

The other serious bottleneck in the CDC 7600 architecture is the entry of results into the X registers, which is also limited to one per clock period. In the CRAY-1 each vector register has its own input multiplexer circuitry for selecting results from among the seven functional units which can produce vector results and, correspondingly, each vector functional unit has its own input multiplexers for selecting vector register operands. Without these circuits the CRAY-1 would also be limited to 1 FLOP/CLOCK.

The overall design of the CRAY-1 processor is shown in figure 6.8. In addition to the eight 64-element V registers, there are eight 64-bit S (scalar) registers and eight 24-bit A (address) registers (corresponding to the X and A registers in the CDC 6600 and 7600), together with 64 B registers (each of 24 bits) and 64 T registers (each of 64 bits). The B and T registers are used in a different way from any of the registers in the 6600 and 7600, however, in that they act as buffer stores for A and S register values, respectively. The functional units take their input operands from the A, S and V registers only, and only return results to these registers.

The A registers are used primarily as address and index registers for scalar and vector memory references, but are also used for loop control, input-output operations and to provide values for shift counts. An A register can be loaded either from a B register or direct from memory, while the B register contents can be transferred to or from memory in block copy operations which proceed at a rate of one per clock period. The S registers contain scalar operands, which may be used in scalar operations in the same way that X register values are used in the 6600 and 7600, but an S register in the CRAY-1 may also supply a scalar value required for a vector operation. S register values may be transferred to or from memory or the T registers, the latter allowing intermediate results of complex computations to be held in fast buffers rather than main memory. T register values can be transferred to or from memory in the same way as B register values.

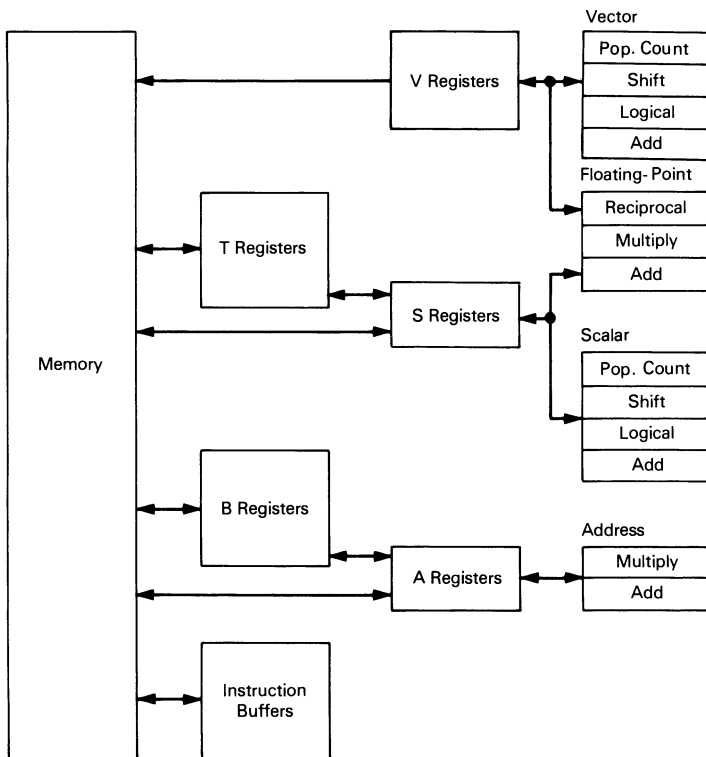


Figure 6.8 CRAY-1 Processor Organisation

The instruction format used in the CRAY-1 (figure 6.9) is very similar to that used in the CDC 6600 and 7600, except that the major function field (g) contains four bits rather than three, and instructions are therefore 16 or 32 bits long rather than 15 or 30. The extra function bit allows vector as well as scalar operations to be specified, and a typical vector instruction takes the form

$$V_i \leftarrow V_j + V_k$$

implying that successive elements of  $V_k$  are to be added to successive elements of  $V_j$ , and the results returned as successive elements of  $V_i$ . Instructions which cause the transfer of an operand between A and B or S and T register use the combined j and k fields to specify the B or T register, and the j and k fields are also combined to produce shift counts in shift instructions.





Figure 6.9 CRAY-1 Instruction Format

Instructions which use immediate (literal) operands use the 32-bit format and combine the j, k and m fields to produce a 22-bit literal value. Memory referencing instructions similarly combine the j, k and m fields to produce a 22-bit memory address, and also use the h field to specify an A register for indexing. (The memory itself is an 8-way or 16-way interleaved 50 ns cycle-time semiconductor store containing 0.5M, 1M, 2M or 4M words according to the configuration.) Branch instructions combine the i, j, k and m fields to produce a 24-bit memory address field, allowing any 16-bit instruction parcel within a 64-bit word to be specified.

In addition to the operating and buffer registers, the CRAY-1 processor also contains several additional registers which support the control of program execution. These include the program counter, datum and limit addressing registers, interrupt registers, a Vector Mask register (VM) and a Vector Length register (VL). The VM register contains 64 bits, one per element position in the vector registers. In merge operations each bit in VM is used to select the corresponding element of one or other source vector for copying into the



destination vector, while in test operations bits in VM are set according to whether or not corresponding elements in a source vector satisfy the chosen condition. The VL register contains a number in the range 0-64 and determines how many vector elements take part in an operation. In the case of an operation on a 150-element vector, for example, the hardware would be required to treat this as two successive 64-element operations (with VL = 64) followed by a 22-element operation (with VL = 22).

#### **6.4.1 Functional Units in the CRAY-1**

The thirteen functional units in the CRAY-1 can be classified into four groups according to the kind of operations they perform and the operating registers to which they are connected. These groups are address, scalar, vector and floating-point. Each unit is pipelined into single clock period segments, so that each can start an operation on a new pair of input operands in each successive clock period.

##### **Address Units**

The Address Add and Address Multiply Units perform 24-bit 2's complement integer arithmetic on operands obtained from the A registers and both return their results to an A register. The Address Add Unit (which performs both addition and subtraction) is the equivalent of the Increment Unit in the CDC 7600, and is similarly pipelined into two stages. The Address Multiply Unit is pipelined into six stages and produces as its result the least significant 24 bits of the integer product of two 24-bit operands. Address multiplication is frequently used in the handling of multi-dimensional arrays, and whereas the number format used in the CDC 7600 allowed integer multiplication to be carried out quite straightforwardly in the floating-point multiply unit, the use of a sign and magnitude, fractional mantissa in the CRAY-1 makes the use of this option much less desirable. Furthermore, the floating-point multiply unit in the CRAY-1 is frequently reserved for long periods by vector operations, and a separate address multiplier is therefore essential for performance reasons.

##### **Scalar Units**

The Scalar Add, Scalar Shift and Scalar Logical Units perform operations on 64-bit operands taken from S registers and each delivers a 64-bit result to an S register. The Scalar Add Unit performs 2's complement integer addition and subtraction and is pipelined into three stages. The Scalar Shift Unit performs single and double-length shifts on one or two S register

operands, the former requiring two clock periods and the latter three. The Scalar Logical Unit is the equivalent of the Boolean Unit in the CDC 6600 and 7600, but it produces its results in one clock period rather than the two required in the 7600.

The fourth unit in this group is the Population & Leading Zero Count Unit which takes a 64-bit operand from an S register and returns a 7-bit result, equal to the number of ones in the operand or the number of zeros preceding the most significant 1 in the operand, to an A register. The first of these operations requires four clock periods for its execution, and the second three.

### **Vector Units**

The Vector Add, Vector Shift and Vector Logical Units take operands from two V registers and return their results to a V register. Successive operand pairs are transmitted to a vector unit in successive clock periods, and after a start-up delay equal to the pipeline length of the unit, results are also copied back to the result register in successive clock periods. The Vector Add Unit performs 64-bit 2's complement integer addition and subtraction and is pipelined into three stages. The Vector Shift Unit is a four-stage pipeline which performs single-length shifts on individual elements of a V register or double-length shifts on consecutive pairs of V register elements. The Vector Logical Unit performs operations similar to those in the Scalar Logical Unit, but acts on operands taken from V registers rather than S registers, and is implemented as a two-stage pipeline. There is also a Vector Population Unit which operates on vectors in a manner similar to that in which the Scalar Population Unit operates on scalars.

None of the address, scalar or vector arithmetic units detects overflows; CRAY-1 users are supposed to know what they are doing and to write error-free programs. Floating-point out-of-range errors are detected, however, in the Floating-point Units.

### **Floating-point Units**

The Floating-point Add, Floating-point Multiply and Reciprocal Approximation Units perform floating-point arithmetic for both scalar and vector operations. For scalar instructions the operands are obtained from S registers and the result returned to an S register, while for vector instructions the operands are obtained from a pair of V registers or a V register and an S register, and results returned to a V register. When

executing vector instructions successive operand pairs are transmitted to a unit in successive clock periods, and results are similarly obtained.




Figure 6.10 CRAY-1 Floating-point Format

The Floating-point Add Unit performs addition or subtraction of 64-bit operands in floating-point format (figure 6.10) and always produces normalised results. This is a departure from the 6600 and 7600 tradition where normalisation was an 'optional extra' which had to be paid for in extra instructions. The consequence is a longer pipeline in the CRAY-1 Floating-point Add Unit, involving six stages rather than the four used in the 7600.

The Floating-point Multiply Unit is pipelined into single clock period segments, like all other functional units in the CRAY-1, and in contrast to the CDC 7600 Multiply Unit, which has a two clock period segment time. The CRAY-1 Multiply Unit uses a multiply pyramid to produce the mantissa product, and requires seven clock periods to produce any one result, rather than the five required in the 7600.

The CRAY-1 Floating-point Multiply Unit also participates in division. The standard 'subtract and test' algorithm used in most machines to implement division cannot easily be pipelined, and a Newton-Raphson iteration algorithm is therefore used in the CRAY-1, similar to that used in the IBM System/360 Model 91 and in MU5. In the CRAY-1, however, division is not implemented directly, but instead an approximation of the reciprocal of the divisor is formed in the Reciprocal Approximation Unit, and a separate instruction must be used to multiply the result by the dividend in the Floating-point Multiply Unit. Furthermore, the Reciprocal Approximation Unit produces a result accurate to only 30 bits (in a 14-stage pipeline), and in order to produce a result accurate to 47 bits (which is still one bit short of the full mantissa), an additional iteration must be performed, again using the Floating-point Multiply Unit in a separate instruction. Thus a scalar quotient is normally computed in 29 clock periods, and forming an n-element vector quotient

requires approximately 3n clock periods.

### 6.4.2 Instruction Issue

As we saw in section 5.4, all instructions in the CRAY-1 are accessed from within the instruction buffers; if an instruction request cannot be satisfied by any of the four buffers, a 64-parcel block of instructions is transferred from memory into one of them. A new instruction is accessed whenever the P register (program counter) is updated. For sequential instructions this occurs as an instruction parcel enters the Next Instruction Parcel register (NIP in figure 6.11). From NIP the instruction parcel is copied into the Current Instruction Parcel register (CIP), where it waits to be issued. In the case of a 32-bit instruction the second parcel is contained in the Lower Instruction Parcel register (LIP) which is loaded in parallel with NIP.

As in the case of the CDC 7600, an instruction is only issued when the conditions in the functional units and operating registers are such that the instruction can be carried through to completion without conflicting with any previously issued, but as yet uncompleted instructions. Thus although any number of the V registers can, in principle, accept results in a single clock period, the A and S registers are similar to the X registers in the 7600; only one A register and one S register can accept a result in any one clock period. Issue of an instruction is therefore delayed if it would cause a result to arrive at either of these sets of registers at the same time as a result from a previously issued instruction.

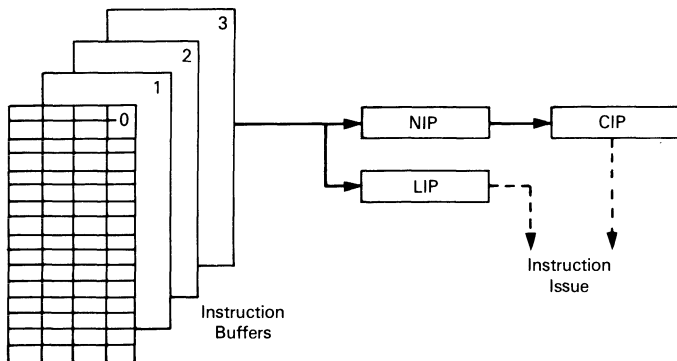


Figure 6.11 Instruction Issue Registers in the CRAY-1

The CRAY-1 also uses a reservation mechanism similar to

that in the CDC 7600. When an instruction is issued which will deliver a new result to an A, S or V register, a reservation is set for that register which prevents the issuing of any subsequent instruction requiring the use of that register until the result has been delivered. In the case of a V register the reservation is for the whole register, rather than individual elements, and furthermore, during the execution of a vector operation, reservations are placed on the operand V registers as well as on the result register. These reservations do not apply to an S register taking part in a vector operation, however, or to the VL register. The values in VL and a specified S register are copied into the unit carrying out the operation as the instruction is issued, and the registers themselves are then immediately free to participate in subsequent instructions.

The need to reserve the operand V registers arises from the nature of the integrated circuits used in the construction of the V registers. These each contain  $16 \times 4$  bits, representing 4 data bits in each of 16 vector register elements, and only one set of 4 bits can be accessed in any one clock period. If two vector instructions using the same operand V registers were in progress at the same time, they would require access to two different elements simultaneously. The same argument applies to the result register; it is impossible to read one element from within a vector register while a new value is being written into another element. The only exception to this rule occurs when an element value which is being delivered to a vector register can, in the same clock period, be routed back into another functional unit as an input operand. This arrangement allows 'chaining' of vector operations, and thus allows much more effective use to be made of the parallel functional units than was possible in the CDC 6600 or 7600.

### 6.4.3 Chaining

Chaining starts when a match occurs between one of the V register operand designators of an instruction awaiting issue in CIP, and the V register result designator of a previously issued instruction which has not yet returned its first result element. When this element becomes available for delivery to the result register, the instruction in CIP is issued (provided there are no other hold-ups) and the result element is forwarded with this instruction to the appropriate functional unit. Successive elements follow until the whole vector has been both written into its result register and forwarded to the second functional unit. The results of this second vector operation may themselves be chained into a third operation, and so on, as shown in the example in figure 6.12.

Figure 6.12 shows a schematic representation of the execution of the instruction sequence

$$\begin{aligned} V0 &\leftarrow \text{Memory} \\ V1 &\leftarrow V0 * S1 \\ V3 &\leftarrow V1 + V2 \end{aligned}$$

The first instruction causes 64 operands from a designated area in memory to be read out and copied in sequence into the 64 element positions in V0. Store requests are pipelined in such a way that the store appears to the processor as a pseudo functional unit. Thus after a start-up delay of seven clock periods, the first element of the vector from store becomes available for delivery to V0, and successive elements follow in successive clock periods.




Figure 6.12 Chained Vector Operations in CRAY-1

In the clock period following the issue of the first instruction, the second instruction in the sequence is copied into CIP, but the reservation on V0 prevents it from being issued immediately. This reservation is lifted, however, allowing the instruction to issue, during the clock period in which the first vector element arrives from store ready for delivery to V0. This clock period is known as 'chain slot time'. Chaining allows the vector elements being copied into V0 to flow directly from the memory read pipeline into the Floating-point Multiply Unit pipeline, where each element is multiplied by the value taken from S1 at the start of the

operation, to produce the vector V1.

The third instruction in the sequence becomes ready for issue in the clock period following issue of the second instruction, and it too is held up by a reservation on one of its input operands, this time V1. When the first element of V1 appears from the Floating-point Multiply Unit, the reservation on V1 is lifted, allowing this third instruction to issue. Now the elements emanating from the Floating-point Multiply Unit can flow directly into the Floating-point Add Unit pipeline as well as into the result register V1. Thus the memory read pipeline, and the Floating-point Multiply and Floating-point Add Unit pipelines are all chained together to produce the elements of V3, and the need for all pipelines to have the same segment time now becomes very apparent; pipelines such as those in the CDC 7600, which have different segment times, could not be chained together in this way.

An alternative representation of the sequence of events illustrated by figure 6.12 is shown in the timing diagrams of figure 6.13. Figure 6.13(a) shows in detail the activities involved in producing the first element of V3. The memory read instruction issues in clock period 0 and the first memory word arrives at V0 in clock period 8. In clock period 9 (chain slot time, marked \*) this element is transmitted to the Floating-point Multiply Unit (together with the value in S1 for this first element). In clock period 17 the first result element from the Multiply Unit is transmitted to V1, and in clock period 18 (chain slot time) this element is transmitted to the Floating-point Add Unit, together with the first element of V2. In clock period 25 the first result element of the composite operation ( $V3 = V0 * S1 + V2$ ) becomes available and is transmitted from the Floating-point Add Unit to V3. Figure 6.13(b) shows how successive elements of V3 are produced, each one staggered one clock period behind the one preceding it.

#### 6.4.4 Performance

Vector processing is only of any value firstly if users' numerical problems involve vectors to a sufficient extent, in terms of both number and length, and secondly if these vectors can be identified during the process of mapping the numerical problem on to the hardware. For CRAY-1 users the mapping problem is largely carried out by the FORTRAN compiler, which 'vectorises' DO loops and splits long vectors into 64-element segments. CRAY Research Inc. have produced results of performance studies [38] which show that the vector subroutines in their FORTRAN library of mathematical functions outperform repetitive use of the equivalent scalar subroutines

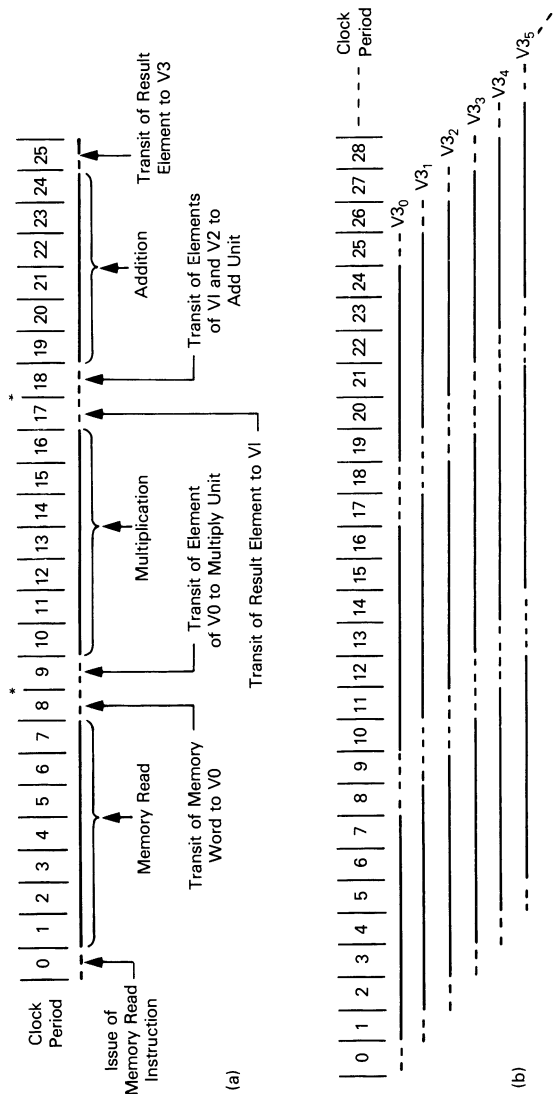


Figure 6.13 Timing Diagrams for Chained Vector Operations in CRAY-1



for vectors having as few as only two elements. These results are reproduced in figure 6.14, in which the cost, in clock periods per result, is plotted against vector length. For scalar subroutines the cost per result remains constant, since the subroutine must be called each time. For vector subroutines the cost per result drops very rapidly as the vector length increases from one and approaches a lower limit equal to about one seventh of the scalar cost.




Figure 6.14 Comparison of Vector and Scalar Performance

A second set of results, shown in figure 6.15, relates to matrix multiplication. This involves repeated scalar product operations, and as we saw in Chapter 4, this operation lends itself particularly well to vector processing. In the CRAY-1 a vector loop is used to multiply together up to 64 pairs of vector elements to form the same number of sub-products. These sub-products must then be added together in what is essentially a scalar loop to form the scalar product. However, by executing a vector add instruction in which one of the

operand registers is the result register of the multiplication, and the other is identical with the result register of the addition itself, a recursive effect is achieved in which 64 sub-products can be reduced to 8 in an operation chained with the multiplication. This recursive effect occurs because the element selection pointer of the result/operand vector does not move on from the first element until receipt of the first result. The value contained in this first element (initially set to zero) is therefore sent repeatedly until the first result element (with a value equal to the first multiplication result) is received in the eighth clock period of the operation. In the ninth clock period the value of this first multiplication result is returned as an operand to the Floating-point Add Unit to be added to the ninth multiplication result, with subsequent elements following on similarly. The result of this addition is then returned as an operand in the 17th clock period to be added to the 17th result, and so on, until at the end of the add operation elements 56 to 63 of the result vector from the Floating-point Add Unit contain eight partially summed sub-product results of the multiplication.




Figure 6.15 Matrix Multiplication Performance

In figure 6.15 the execution rate in MFLOPS is plotted against increasing matrix dimension for matrix multiplication operations on square matrices. Matrix sizes which are multiples of 64 give the best results, as would be expected, and execution rates of up to 140 MFLOPS are clearly shown to be possible. This is a best case example, of course, and execution rates such as this cannot be sustained over a more typical mix of user programs. Average rates in the range 20-50 MFLOPS have been reported by Hockney [39].

Matrix multiplication identifies another feature of vector processing. Vectors in user programs are themselves frequently elements of larger data structures. Thus when accessing a vector from store into a vector register in the CRAY-1, the user (or more usually the compiler) specifies a starting location and an increment. Since the increment can be non-unary, arrays, for example, can be accessed by row, by column or by diagonal. While this scheme offers complete flexibility, it leaves the entire management of all forms of data structure to the software. Vector processing is not really offered as a facility in the CRAY-1; it is a means for improving raw performance. Vector (and data structure) facilities are to be found, however, in machines such as the TI ASC (as we saw in Chapter 4), and in machines such as MU5 and its commercial derivatives at the top end of the ICL 2900 range, and in the CDC STAR-100 and its successors in the CYBER 200 Series. Chapter 7 will be concerned with these vector processing facilities.

# 7 *Vector Processes*

The need to provide facilities for processing sequences of vector elements was recognised in the very early days of digital computer design. The von Neumann concept, for example, included the notion of allowing instructions to be treated as data, which meant that the address part of an instruction accessing a vector element could be incremented during the execution of a program loop and thus produce the effect of processing a vector. In practice, however, this technique allows so much scope for program error that even the very first stored program computer (the Manchester Mark 1) used B-lines instead, and this latter technique has been used almost universally ever since. Thus virtually any digital computer can be used to process vectors. The differences between machines lie in the addressing facilities which they provide to support accesses to data structures, and whether or not they include instructions which implicitly process a sequence of vector elements. Computers with this latter facility have been described by Flynn [40] as Single Instruction Multiple Data (SIMD) arrangements, in contrast with the Single Instruction Single Data (SISD) arrangement of conventional computers. We have already met two examples of SIMD machines in the TI ASC and the CRAY-1 and in this chapter we shall be considering the CDC STAR-100 and CYBER 205 computers. First, however, we shall consider the SISD vector addressing facilities found in MU5 (and by implication machines at the top of the ICL 2900 Series), and also the SIMD string processing orders in MU5.

## **7.1 VECTOR FACILITIES IN MU5**

As we have seen in earlier chapters, the MU5 instruction set enables information on the nature of operands to be transmitted to the hardware to allow optimal operand buffering. In particular this allows scalar variables, of which only a small number are in frequent use at any one time, to be buffered separately from array elements, which are generally selected sequentially from a large group. Furthermore, hardware assistance is provided for accessing elements of arrays and other data structures, particularly with regard to bound checking. This is achieved through the

use of descriptors; an access to a data structure element involves a primary access to the data structure name (the result of which is the descriptor of that data structure), followed by a secondary access involving the descriptor and possibly a modifier. As we saw in section 2.5, descriptors are of two main types

Vector Descriptors	Tv	BOUND	ORIGIN
	8	24	32
String Descriptors	Ts	LENGTH	ORIGIN
	8	24	32

Vector elements may be of size 1, 4, 8, 16, 32 or 64 as defined by bits within the type field (Tv), whereas string elements are constrained to be of size 8 bits, and all are close-packed in store. A typical access to a vector element would involve the following instructions

```
B = i
ACC = X[B]
```

The first of these instructions loads the variable *i* into the B register for use by the second instruction as a modifier. This instruction loads the named variable X into a descriptor register (DR) for interpretation. The hardware associated with this register first forms the address of the element, by scaling the modifier according to the element size and adding the result to the origin address contained within the descriptor. It also checks that the modifier is greater than or equal to zero, and less than the bound; if not an interrupt is generated. This check incurs virtually no time penalty and removes the need for the costly software bound check which is often required on more conventional machines, particularly during program development. The address is then used to access the store, and the required vector element is selected from within the corresponding store word, right justified, and zero-filled for presentation to the designated arithmetic unit. This mechanism is equally applicable to dynamic arrays whose starting address and bounds are known only at run time, as in Algol-like languages, or to arrays where this information is available at load time, as in FORTRAN.

Operations on whole vectors are organised by program loops such as that in the example already given in section 4.2, which forms the scalar product of two vectors X and Y each of length LIMIT

```

      B = 0
      ACC = 0
L1: ACC *= X[B]   :: Stack ACC and load with element of X
      ACC * Y[B]  :: Multiply by element of Y
      B CINC LIMIT :: Compare B against LIMIT and increment B
      ACC + STACK :: Add Stack content to ACC
      IF /=, -> L1 :: Continue round loop until B = LIMIT

```

The arguments for processing vectors in this way, rather than by specific vector orders, are that in MU5 the hardware can distinguish control and address calculation instructions from those operating on the vector elements, and can execute the former at the peak rate of the Primary Operand Unit pipeline (section 4.1), while the latter are queued in the Secondary Operand Unit pipeline awaiting operand fetches (figure 7.1). Furthermore, the compilers are only required to generate conventional code in a straightforward way without using intricate algorithms to decompile back to vector operations. There are, however, other features of vector processing which would make vector orders attractive in an MU5-like machine. The scalar product loop is a particular example; as we saw in section 4.2, the running total in MU5 has to be maintained either in the top-of-stack location (as in the coding given above), or in some named variable location. In either case this requires that a new value be written into a buffer store (the Name Store) each time the loop is obeyed, and then read out again by the next ACC order but one. In the absence of complex data forwarding techniques such as those used in the IBM System/360 Model 91, this situation creates problems in a pipeline environment. A vector processing mechanism such as that in the TI ASC, which takes pairs of operands, multiplies them together and adds the result to a running total held within the arithmetic unit offers obvious performance advantages.

By contrast, operations on strings in MU5 are carried out by string processing (store-to-store) orders which provide facilities for the manipulation of data records in languages such as COBOL, PL/1 and Algol 68 [41], and which operate on complete strings of byte elements defined by String Descriptors. We shall deal with these in some detail in section 7.2. Before doing so, however, we shall first examine the hardware used in the MU5 Secondary Operand Unit, and some of the more complex descriptor mechanisms which it implements.

### 7.1.1 The MU5 Secondary Instruction Pipeline

The MU5 Secondary Instruction Pipeline (figure 7.1) is mainly concerned with accessing and processing the data structure elements specified as secondary operands by means of

descriptors. It consists of three major sections, the D-unit, the Operand Buffer System and the A-unit (containing the ACC register). The D-unit is itself made up of two units, the Descriptor Addressing Unit (Dr) and the Descriptor Operand Processing Unit (Dop). The D-unit and the Operand Buffer System together constitute the Secondary Operand Unit (SEOP) of the processor (c.f. figure 3.5).

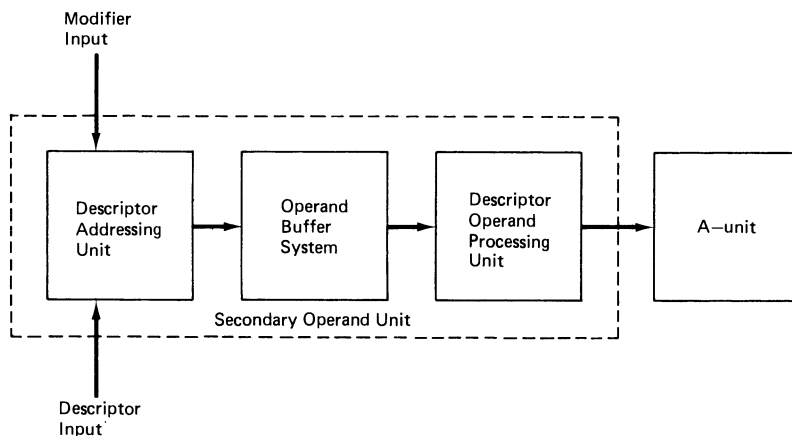


Figure 7.1 The MU5 Secondary Instruction Pipeline

Instructions enter the Secondary Operand Unit from the Primary Operand Unit (PROP). For an instruction specifying an array element as its operand, the descriptor supplied by PROP is loaded into the DR Register, and Dr proceeds to generate the required address using the content of the B-Register as a modifier. This address is then sent to OBS which, like the Name Store in PROP, is invisible to the programmer and exists simply to improve the instruction execution rate. OBS makes the necessary store request and sends the required 64-bit word to Dop, which contains masking and shifting circuitry to select the array element from the appropriate position in the word. Dop is controlled by a combination of a control field generated by Dr (at the same time as the address) and the original function code, which accompanies the instruction as it passes through the pipeline. Instructions leaving Dop are normally sent direct to the A-unit, but may also be sent to PROP, the B-unit or to Dr. In the case of the store-to-store functions, the operand processing takes place in Dop.

### 7.1.2 Descriptor Processing Hardware

The basic tasks of the D-unit are comparatively simple. They involve the formation of the address of an array element, by

the addition of the modifier to the origin address contained within the descriptor, and the subsequent selection of that element from within the corresponding store word. These tasks are complicated, however, by a number of factors, such as the existence of different descriptor types and different operand sizes, the possibility of an operand straddling across two store words, and the need for bound checking. In addition, the two 64-bit descriptor registers, DR and XDR, can be manipulated in whole or in part by various functions. Also the origin and length fields of one or both of these registers must be incremented and decremented, respectively, during each cycle of the execution of the store-to-store orders (section 7.2.1). For both this task and the bound checking operation, it would clearly be desirable to use a subtractor, separate from the adder used either to increment the origin or to add the modifier and origin, since the appropriate additions and subtractions could then proceed in parallel. However, the MU5 SEOP actually uses one adder to carry out these tasks sequentially, since this arrangement involves less hardware.

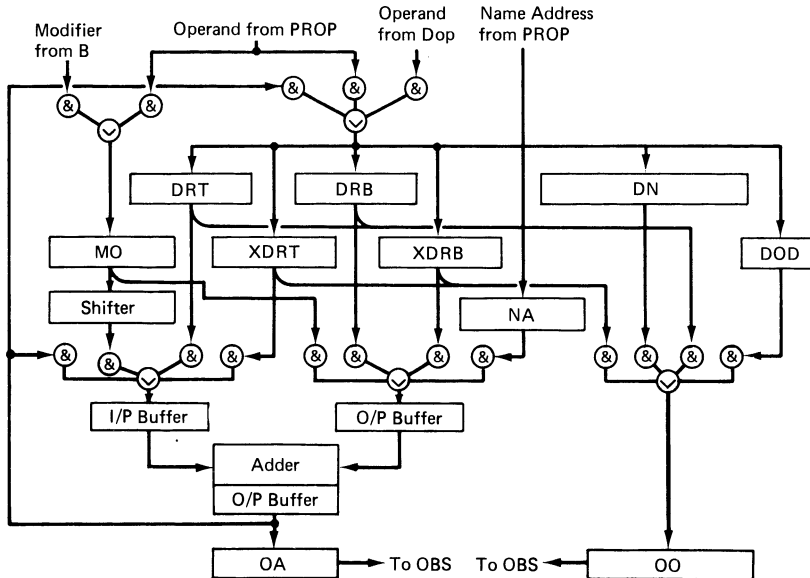


Figure 7.2 The MU5 Descriptor Addressing Unit

The two descriptor registers each consist of a top and bottom half (DRT and DRB, XDRB and XDRB). DR is the main descriptor register, used for array accesses and the destination address in store-to-store orders. DRB contains the origin and DRT the type and bound or length fields. XDR is similarly divided and is used to hold the source address



required by some of the store-to-store orders. DR and XDR are normally loaded from PROP, but may in some circumstances be loaded from Dop.

The 32-bit modifier is held in register MO, which is loaded from the B-unit during an array access or from PROP during the execution of functions which manipulate DR or XDR. The true output of MO is connected to one side of the 32-bit adder via a shift network, and the complementary output directly to the other side. The modifier normally refers to the position of an element within an array, regardless of its size, and must therefore be shifted relative to the origin, which always refers to a byte address. The complementary output is used in the bound checking operation, where the subtraction of the modifier from the bound (equal to the number of elements in the array) is performed by adding the complement of the modifier to the bound and forcing a carry into the least significant digit position of the adder. Following the formation of the address, the output of the adder is passed on to OBS, via the OBS Address buffer (OA). This address may also be used internally for the updating action of the store-to-store orders and for array accesses which involve elements straddling store-word boundaries.




Figure 7.3 The MU5 Descriptor Operand Processing Unit

Register DN is a 64-bit temporary buffer used for holding primary operands (names, literals or internal register values) associated with those ACC functions which do not require a secondary operand access. These functions pass through the Secondary Instruction Pipeline en route to the A-unit in order

that the correct program sequence be preserved, but must not disturb the contents of the descriptor registers in Dr. The contents of the descriptor registers themselves also pass through the primary operand route to OBS (via the OBS Operand Buffer, 00) when required for orders such as 'DR =>'. A route through Dr is also required for addresses of names used with ACC functions and held in the OBS Name Store (section 3.4). This is provided via register NA, which is loaded from PROP and can be gated as an input to the Dr address adder. Register DOD is concerned with D-unit interrupts.

Orders leaving Dr are forwarded to OBS, which obtains the store word containing the operand and sends it to Dop. Control information generated by Dr during the address formation is buffered with the function in OBS, and is supplied to Dop with the store word. Using this information, Dop selects the required array element from within the store word, and in the case of a load order, routes it to the least significant end of its output highway. In the case of a store order, Dop updates the appropriate part of the store word and returns the updated version to OBS.

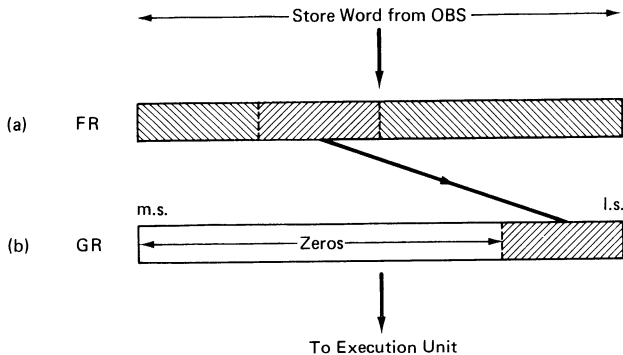


Figure 7.4 Dop Actions for a Load Order

Dop consists of two sections (figure 7.3), one with a 64-bit data path used for operands associated with computational orders, and one with an 8-bit data path used for operands associated with the store-to-store orders. The main input register (FR) has masking facilities which permit the selection of left or right hand masks to any bit position over the full 64-bit width, while the shift mechanism allows any circular shift from 0 to 63 bits in single bit increments. This is achieved by three levels of logic, the first allowing shifts of 0, 16, 32 or 48 bits, the second shifts of 0, 4, 8, or 12 bits and the third shifts of 0, 1, 2 or 3 bits. The output of the shifter is copied into the main output register

(GR) if the order is destined for the A-unit or the Dr unit, or gated on to the processor Highway if the order is destined for PROP or the B-unit. The route from GR back into FR is used during the execution of store orders.

Figure 7.4 illustrates the actions taken for a load order requiring a 16-bit operand. The store word containing the operand is loaded into FR (a), and the shifter routes the required 16-bit operand to the least significant end of the input to GR. At the same time, the most significant 48 bits of the input to GR are set to zero by inhibiting the appropriate section of the data path in the third level of the shifter (b). GR is strobed when sufficient time has elapsed for the shifter outputs to settle, following the receipt of a store word in FR, and when the previous order in GR has been accepted by another unit.

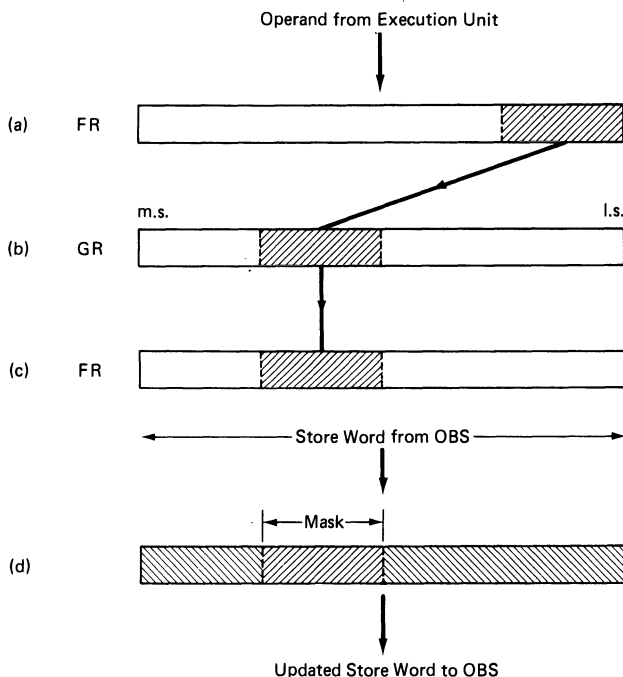


Figure 7.5 Dop Actions for a Store Order

The actions for a store order are more complex. The order first passes through Dop as if it were a load order, en route to the appropriate execution unit, and is only processed by Dop when the execution unit indicates that the required operand is available. The operand is received from the

processor Highway (figure 7.3) and is copied into the least significant end of register FR (figure 7.5(a)). It is then shifted to its eventual alignment in the store word, copied into GR (b) and copied back from GR into FR (c). The store word of which it is to form part is then sent again from OBS and selectively copied into FR around the operand, using the masking facility (d). The updated version of the store word is then returned from FR back to OBS.

### 7.1.3 Complex Descriptor Mechanisms

The MU5 instruction set has provision for a number of descriptor mechanisms more complex than those described so far. Thus one of the special descriptor types (Indirect) may be used to describe a structure containing other descriptors. Hence

```

B = i
DR = X[B]
B = j
ACC = D[j]    :: Use the current descriptor in DR

```

could be used to access an element  $X[i,j]$  of a two dimensional array using an addressing vector technique. This mechanism has the disadvantage in MU5 of requiring a descriptor to be accessed through the Secondary Operand Unit and returned to the front of it before the element access can begin. This destroys the overlap in the secondary pipeline and for simple languages an alternative technique involving the use of the multiplication facilities in the B-unit has the advantage that the subscript calculations take place independently of operations queued in the secondary pipeline and can be substantially overlapped with them. The previous sequence becomes

```

B = i
B * n
B + j
ACC = X[B]

```

with a bound check on the overall limits of the array, rather than on each dimension. For languages where the lower bound of arrays is 1, a load and decrement B order is provided. A further, more elaborate, dope vector mechanism is convenient for arrays with dynamic lower bounds or for slices of arrays. For each dimension the dope vector contains three 32-bit elements, a lower bound which is subtracted from the subscript, an upper bound against which the subscript is checked, and a stride by which it is multiplied. The SUB1 and SUB2 instructions implement this facility, so  $X[i,j]$  becomes

```
B = i
SUB1 X1      :: process first subscript, using dope vector X1
B = j
SUB2        :: process next subscript
B = DO      :: move composite subscript to B (DO = Origin
ACC = X[B]  :: field of DR)
```

From a performance point of view these orders are worse than the Indirect descriptor mechanism, since several of the actions implied by SUB1 and SUB2 require accesses to quantities which are to be returned to the front of the secondary pipeline. Some quite different structure from that of MU5 would be needed to obtain good performance using these orders.

#### 7.1.4 The Operand Buffer System

The need for an Operand Buffer System (OBS) in the MU5 Processor arises from the disparity between the data rate and access time of the Local Store. In principle the MU5 Local Store can supply 128-bit words at intervals of 65 ns, but the time gap between the generation of a secondary operand address, and the receipt of the corresponding store word, is over 600 ns. When MU5 was designed it was expected that floating-point addition and subtraction would take only around 100 ns, and so the store access time was unacceptably high. The difference between the access patterns for named variables (of which a small group is generally used repeatedly) and data structure elements (which are generally selected sequentially from among a large group) precluded the use of a system corresponding directly to the Name Store for secondary operands. Instead a first-in/first-out 'Function Queue' is used; as each address is generated by Dr, the corresponding function, together with control information, is entered into the Queue. A function leaves the Queue when the word containing its operand is received from store. Since the store accessing system is itself a form of pipeline, the effective access time is reduced by a factor corresponding to the number of positions in the Queue.

In a synchronous system, no extra operand buffering would be required, but because the MU5 Processor operates asynchronously, it is essential that as many buffers as Queue positions be available to receive the returning store words. Thus the Operand Buffer System contains eight 128-bit buffers, for which the corresponding virtual addresses are held in an associative store. Fetching a 128-bit store word whenever an operand is accessed, and retaining each of these words in associatively addressed buffers means that many operands are automatically pre-fetched, and the corresponding store

requests avoided.

An additional problem arises as a result of page fault interrupts; a significant fraction of the access time for secondary operands is taken up with generating a virtual address and obtaining either a real address or a page fault interrupt. When such an interrupt arises, the Control Register at the end of the Primary Operand Unit will have been incremented on beyond the address of the corresponding instruction and that instruction will therefore have become irrecoverable as far as the program is concerned. Thus in order to be able to preserve this and subsequent instructions and re-execute them after the interrupt has been serviced, the Queue is designed to contain all functions for which store requests are outstanding. In the event of a page fault, the Page Registers can then be manipulated by the Operating System using orders which bypass the Queue or, if a process change is required, the whole of the Queue and its associated buffer registers can be preserved (for subsequent restoration) in the store, thereby unblocking the Queue and allowing other processes to be run.

## 7.2 STRING OPERATIONS IN MU5

The orders provided in MU5 for the string processing functions which occur in languages such as COBOL and PL/1 may be divided into two groups, string-to-string and byte-to-string. The string-to-string orders operate on two fields, or strings, each described by a descriptor. The descriptor of the destination string is held in the DR register while that for the source string is held in XDR. As the operation of the instruction proceeds, the descriptors in DR and XDR move along the strings. No visible register is used by the strings themselves. The operand of the order is an 8-bit mask, that determines which bits within each byte are to be operated on, together with an 8-bit filler and in some cases four 'function' digits used as described below. Provision is made in the hardware for these orders to be interrupted (section 7.2.3). Examples of the orders are

- BMVE    Use the filler byte as source and copy to all bytes of the destination string.
  
- BMVB    Copy the filler byte to the first byte of the destination string.
  
- SMVB    Move one byte from the source to the destination string, or use the filler byte if the source is exhausted.

- SMVF Move the whole source string to the destination string followed by filler bytes if the source is shorter than the destination.
- SCMP Compare the source and destination strings byte by byte ending when inequality is found, or the destination string is exhausted.
- SLGC Logically combine the source and destination strings into the destination. The form of combination (logical OR, for example) is selected by the 'function' bits in the operand.

An example of the use of these orders is in the implementation of the MOVE verb in COBOL. Suppose that two fields C and D are specified

```

02   C   PIC   X(7)
02   D   PIC   X(7)

```

In MU5 descriptors would be created at compile time for C and D, each describing a 7-byte field starting at the required byte address. The COBOL sentence

```
MOVE C TO D
```

would then become in MU5 instructions

```

XDR = C      set source descriptor for C
DR  = D      set destination descriptor for D
SMVF        move the field described by XDR to that
            described by DR

```

If D is specified as

```
02   D   PIC   X(9)
```

then the final two bytes of D must be spaces. The filler option of SMVF allows this to be carried out automatically. The sequence becomes

```

XDR = C
DR  = D
SMVF 'space'

```

If the source field is too long, then the SMVF order terminates when the destination field is full, and an optional interrupt enables this condition to be monitored if required.

An extension of the above technique to vector operations of

a mathematical form would be fairly straightforward. For example, a vector add of the form

$$F := F + E$$

would become

```
DR = F
XDR = E
VECTOR ADD
```

However, as we have already seen, operations of this type are programmed out into loops in MU5. Thus the idea of including vector orders in the MU5 instruction set was dropped in favour of a pipeline approach which would lead to execution rates for such loops approaching the peak rate at which the store could deliver the vector operands.

### 7.2.1 Operation of the Store-to-Store Orders

The store-to-store orders are controlled by logic in the Dr unit which, by incrementing the origin field(s) of the descriptor(s) at each access, generates the appropriate series of addresses, and by decrementing the length field(s), until zero is reached, determines the point at which to terminate the order. The descriptors themselves are loaded into DR and XDR by preceding orders, while the filler/mask, which is the operand associated with the actual store-to-store order, is automatically transferred to Dop by a preliminary access during the execution of the store-to-store order itself. A 'Dop bit' accompanying this access indicates its special nature and causes the operand to be loaded into the filler/mask register within the store-to-store processing logic in Dop (figure 7.3). This logic allows source bytes to be copied to the destination, source and destination bytes to be logically combined or compared, and so on, in accordance with the requirements of the various store-to-store orders.

During the execution of these orders, the main data path is used to extract bytes from the source and destination store words, and to return them to the destination store words as appropriate. Thus, for the SMVB and SMVF orders, for example, Dop performs 'load' operations to extract bytes from the source string, and for the SMVB, SMVF, BMVB and BMVE orders, performs 'store' operations to return bytes to the destination. For the logic orders, SLGC and BLGC, bytes are also extracted from the destination string before being logically manipulated and returned to the destination string. For the comparison orders, only 'load' operations are required since the destination string is not updated.



The sequence of operations in a store-to-store order is best illustrated by reference to the patterns of bytes in the strings as the order proceeds. Figure 7.6(a) shows the pattern of bytes in the destination string at the start of a BMVE order. The origin of the descriptor initially points to the first byte of the string, and the length is set to the total number of bytes  $L$ . At the end of the first cycle, the position is as shown in (b). The (masked) filler has been moved into the first byte, but the origin and length are unchanged; the only route to the OBS address highway is through the Dr adder, and since the first address to be sent must be that contained in the descriptor origin, the updated value is not available for loading into the descriptor. For the second and subsequent cycles, the address sent to OBS is the origin plus 1. The origin is then updated with this new value, and finally the length field is decremented by 1. The string patterns are as shown in (c) and (d).

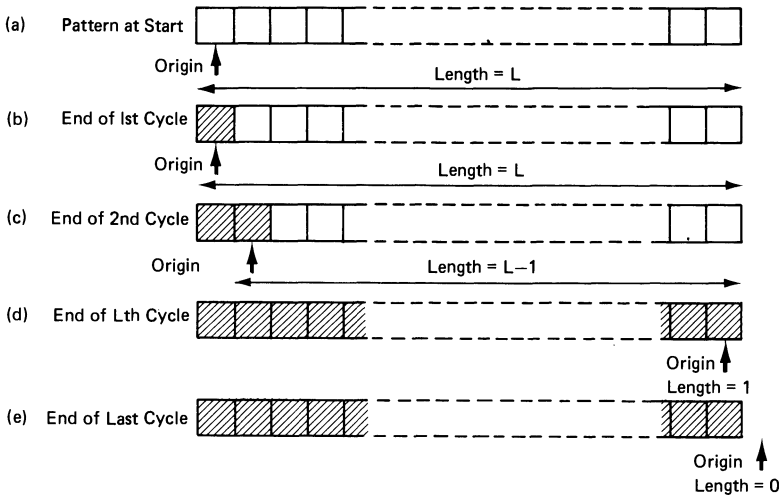


Figure 7.6 String Patterns in a Byte-string Order

The end of the move sequence is reached when the length becomes equal to 1. This indicates that the last byte has been filled, as shown in (d). Since the descriptor must point after the end of the string when the order terminates, however, an extra cycle is needed (without a store access) to increment the origin by 1 and decrement the length by 1. The final string pattern is shown in (e) and figure 7.7 shows the event sequences in the various cycles.

The event sequence for a 'normal' cycle of a string-string order is shown in figure 7.8. The source byte address is

generated first and sent to OBS, and then the source origin is updated. The destination byte is then accessed in the same way and the destination origin is updated. Finally, the two length fields are decremented by 1. The order is normally terminated when the destination length reaches zero, or when equality is found in SCMP. If the source length reaches zero before the destination length, the action taken depends on the order involved. SLGC is terminated immediately, and an interrupt is generated, whereas SMVF and SCMP are converted to the corresponding byte orders. This is illustrated in the complete event sequences for SMVF (figure 7.9).

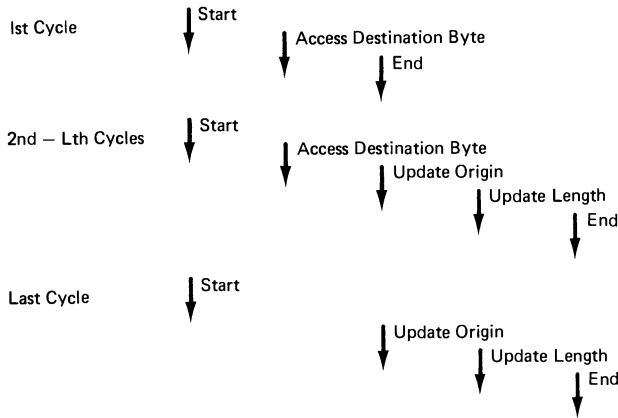


Figure 7.7 Event Sequence in Byte-string Order Cycles

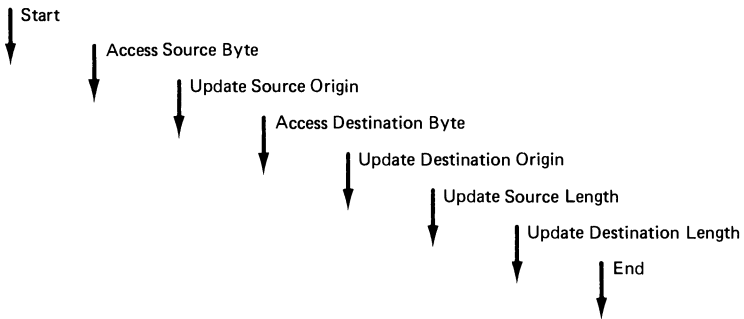


Figure 7.8 Event Sequence in a String-string Normal Cycle

In the first cycle, the origin fields are sent out directly as addresses and the two descriptors are left unchanged at the

end of the cycle. 'Normal' cycles are then executed until the source string runs out. For subsequent cycles only the destination address is generated, and the function code associated with these accesses is changed to BMVE, so that the rest of the destination string is filled with the filler byte held in Dop. The first of the BMVE cycles is also a 'last' cycle for the source descriptor, in which the origin and length fields are updated so that the descriptor finally points after the end of the string and has zero length.




Figure 7.9 Complete Event Sequence for a String-string Order

### 7.2.2 Special Store Management Problems

Since the addresses generated by the store-to-store orders are virtual addresses, the execution of a store-to-store order will generally involve the crossing of one or more page boundaries, and hence require access to a page of data not currently available. A system which attempted to overcome this problem by ensuring the availability of all the necessary pages of data before the start of the order would involve different hardware constraints on the length of strings which could be handled in different implementations of the instruction set, and was therefore not considered. The alternative approach adopted in MU5 is to allow the store-to-store orders to be interruptable, so that only parts of each string need be resident in the Local Store at any one time. Thus Dr samples the interrupt signal at the start of each cycle, and when an interrupt occurs it converts the new cycle to a 'last' cycle in order to terminate the order normally with the origin(s) pointing to the byte(s) immediately following those for which requests have already been sent to OBS. Requests trapped in the OBS Queue as a result of the interrupt are dealt with as for any page fault interrupt (section 7.1.4) and are automatically re-issued after the interrupt has been serviced. When an interrupted

store-to-store order is re-started, the first cycle generates the address(es) of the next byte(s) in sequence. Thus it is possible to interrupt execution of a store-to-store order at a partially completed stage and re-start it from the same point when a new page of data has been made available.

### 7.3 THE CDC STAR-100

In chapter 6 we followed the development of the CDC 6600 design as it evolved first into the 7600 and then into the CRAY-1. In 1965, however, CDC started a separate line of development in response to a requirement of the Lawrence Livermore Laboratory for a vector processor capable of executing 100 MFLOPS. The machine which resulted was the STAR-100 [42]. A great deal of controversy raged about this machine in its early years, and many of the essential design issues and performance goals have been obscured. An interesting analysis of these issues was made by Lincoln [43] in a paper presented at the AFIPS National Computer Conference in 1978. Despite the many difficulties which arose in the course of the STAR-100 programme, CDC remained convinced that the underlying architectural concepts of the STAR-100 were sound, and went on to produce a second version, the STAR-100A, which appeared commercially as the CYBER 203, and a completely re-engineered version, the STAR-100C, which is now produced commercially as the CYBER 205. In this section we shall consider the design philosophy of the original STAR-100, and deal in more detail with the instruction set and operation of the current version, the CYBER 205, in section 7.4.

The fundamental premise on which the STAR-100 project was based was the provision of as much immediate access storage as was possible within engineering constraints. In the mid-1960s this meant using ferrite core technology; semiconductor storage technology, while holding out a great deal of promise, was not sufficiently mature for its use in large quantities to be a realistic venture. The limitations on the amount of storage which could be provided using cores were physical space requirements, the physical interconnection problem, and overall reliability. These limitations led to a maximum memory size of one million 64-bit words using the 2D core memory developed for use as Extended Core Storage in the CDC 6600. These stores had a 280 ns access time and a 1280 ns cycle time. (The limitations on store size have not gone away with the improvements in storage density and access time provided by semiconductor technology, because logic densities and processing speeds have also improved and have led to higher user expectations; transmission delays still account for roughly the same percentage of the total STAR system delay in the CYBER 205 as they did in the original STAR-100.)

Coupled with the provision of large amounts of storage in the STAR-100 was a high bandwidth for the transfer of data between the store and the central processing unit. Because the access time to an individual store word was relatively long, however, it was clear from the outset that it would only be possible to use this bandwidth effectively if operand pre-fetching could be used in the same way as instruction pre-fetching is used in conventional machines. Instruction pre-fetching relies on the sequential nature of instruction execution; operand pre-fetching is similarly only effective if the processor can generate long sequences of addresses, a situation which arises quite naturally in vector processing. Thus the STAR-100 project was naturally geared to the provision of computing power for highly vectorisable problems, and vectors with up to 65,536 elements may be processed by a single vector instruction.

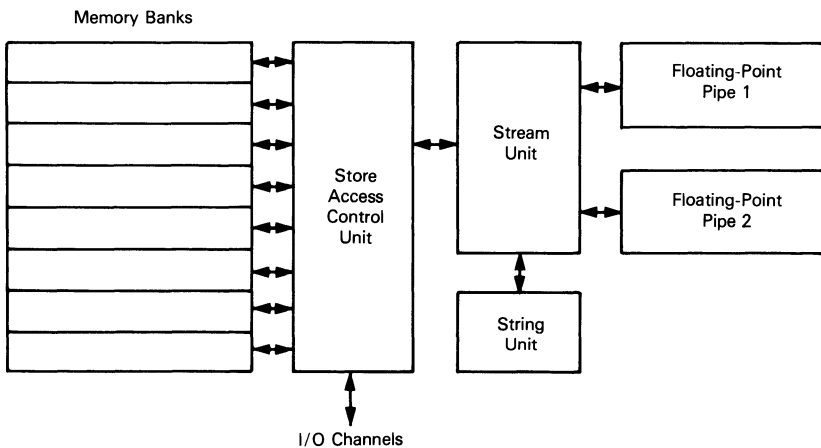


Figure 7.10 STAR-100 Processor Organisation

Memory bandwidth is obtained in the STAR-100 by a combination of eight-way interleaving of memory banks with wide store words (512 bits accessed per store cycle) and the pipelined transmission of each of these 'superwords' to the processor in four sequential groups of 128 bits. In order to handle this data rate in the processor, however, the designers were faced with choices ranging, in the extreme, between using a multiplicity of 6600-like arithmetic units and a single pipelined arithmetic unit capable of executing 32-bit operations at a rate of one every 10 ns. In practice the solution chosen for the STAR-100 uses two pipelined arithmetic units as shown in figure 7.10, each of which can act as a single 64-bit or a twin 32-bit unit. Floating-point Pipe 1

performs floating-point operations such as add, subtract, multiply and compare, and also fixed-point address operations. Floating-point Pipe 2 is similar, except that it includes a divide unit and a multi-purpose unit which is used for special operations such as square root. For operations such as floating-point addition the two pipes act in parallel and can therefore accept the full 128 bits available from store in each clock period. For 32-bit operations this leads to an execution rate of four operations per 40 ns clock period, equivalent to 100 MFLOPS.

The store is made up of eight sections, each containing four banks of 16K 64-bit words. 512 bits are read out from a bank in one store cycle, and the outputs of the four banks in any one section are multiplexed at 40 ns intervals on to a 128-bit data trunk. The data trunks from each section are then multiplexed together in the Store Access Control Unit which provides a stream of operand store words to the processor. The outputs from the Store Access Control Unit cannot be fed directly into the arithmetic units, however, since in practice vectors cannot necessarily be expected to start on 512-bit word boundaries and the two operands making up a pair of source operands are normally taken from separate superwords. Thus the Stream Unit contains input shift networks which align the input operands correctly and a 'front-end pre-count' mechanism which discards any unwanted elements from the first superword received at the start of a vector order. (The parallels between operand accessing in the STAR-100 and techniques used elsewhere in instruction accessing can also be seen in this mechanism. A similar system is used in the MU5 Instruction Buffer Unit to allow control transfers to jump to any 16-bit parcel within a 128-bit store word.)

The Stream Unit must also align the result vector correctly, using an output shift network, and at the start and end of a result vector must ensure that only the appropriate parts of the store word are overwritten. This problem also occurs in exactly the same way in MU5 and is dealt with by the Descriptor Operand Processing Unit (section 7.1.2) which has facilities for updating store words down to the bit level. In the STAR-100 such facilities exist in the String Unit, but for arithmetic operations the smallest unit is a 32-bit word, and selective updating of the 512-bit superword in a given memory bank is carried out, on a 32-bit word basis, at the store interface. Thus the Stream Unit sends a four-bit 'write enable' control field with each 128-bit result, as illustrated in figure 7.11 which shows the effect for eight 32-bit results positioned arbitrarily within a 512-bit superword.

Any read or write operation on a core store involves a

read-write cycle and in the STAR-100 a word read out from a complete memory bank is held in a 512-bit holding register before being written back. The write enables control the updating of the contents of this register by the results sent from the processor, giving the effect of a read-modify-write operation on the selected superword in the memory bank. Once the use of these write enables had been established as part of the design of the STAR-100, the concept of a control vector was an obvious extension. Control vectors became an important feature of the STAR-100 design and are used extensively, not only to define write enables, but also to control vector restructuring operations such as MASK, COMPRESS, and MERGE, which derive directly from the programming language APL. (Iverson's book on APL [44], which describes these operations and the notion of a control vector, was a significant influence on the STAR-100 architects.) These operations will be discussed in more detail in section 7.4; there is still another problem to be discussed in connection with the writing of results back to store.

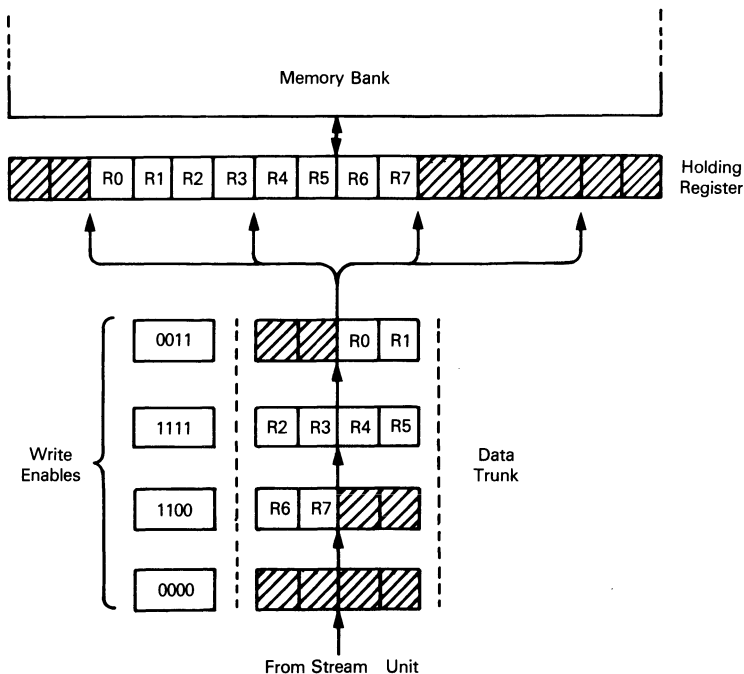


Figure 7.11 Result Operand Alignment in a STAR-100 Superword

The STAR-100 uses a virtual memory organisation based on that used in the Atlas computer, and in executing a long

vector operation inevitably encounters page faults in exactly the same way as the store-to-store orders in MU5. The solution to this problem is relatively simple in the MU5 case because destination addresses are generated before any action is taken to produce the corresponding results (section 7.2.2). Updating of the store words takes place within the processor on a byte-by-byte basis, so that whole words are fetched from store before any new values can be written into them. In the STAR-100 vector orders are normally three-address, and destination addresses are only generated after the results have been produced. Thus in order to achieve the full vector rate, while still responding to page fault interrupts, a one-page look ahead is required on the output stream to ensure that operands progressing through the long processor pipeline can be returned to store before the instruction is interrupted.

#### 7.4 THE CDC CYBER 205

The STAR-100 was criticised on a number of grounds by users who wished to apply it to more general computing problems than those for which it was designed. The grounds for criticism were mainly the long vector start up time and poor performance on scalar arithmetic, both of which were inevitable consequences of the design. These problems are largely overcome in the CYBER 205 by the use of a very much faster (80 ns access time) semiconductor memory and by the inclusion of a high performance scalar unit. The overall performance of the CYBER 205 is further enhanced by its implementation in specially developed ECL LSI Uncommitted Logic Array technology, allowing a reduction of the clock period from the 40 ns used in the STAR-100 to 20 ns.

Figure 7.12 shows the overall design of the CYBER 205 computer system [45], the major components of which are the central processing unit (containing the vector processor, scalar processor and eight or 16 I/O ports), the central memory, and the maintenance control unit (linked to a standard I/O port via a system channel adaptor). A basic system has one million words of central memory, expandable to two or four million words (and with a capability to expand to eight million words as memory technology progresses), with each million words being connected to the memory interface via a 512-bit data highway. Within the vector processor there can be one, two or four floating-point pipelines.

Instruction execution is controlled by a number of independent high-speed microcode memories, each containing 256, 512 or 1024 words of microcode with between 48 and 120 bits per word. During instruction execution each memory



operates as a read-only memory, having been loaded at start-up time by the maintenance control unit.

The scalar processor receives instructions from central memory, decodes them, and (according to type) either executes them itself or sends them to the vector processor. A typical vector instruction has the format shown in figure 7.13(a), and causes the two source vectors defined by A and B to be combined and copied to a result vector defined by C. Each address field within the instruction is an eight-bit designator referring to one of 256 64-bit locations within a Register File in the scalar processor.

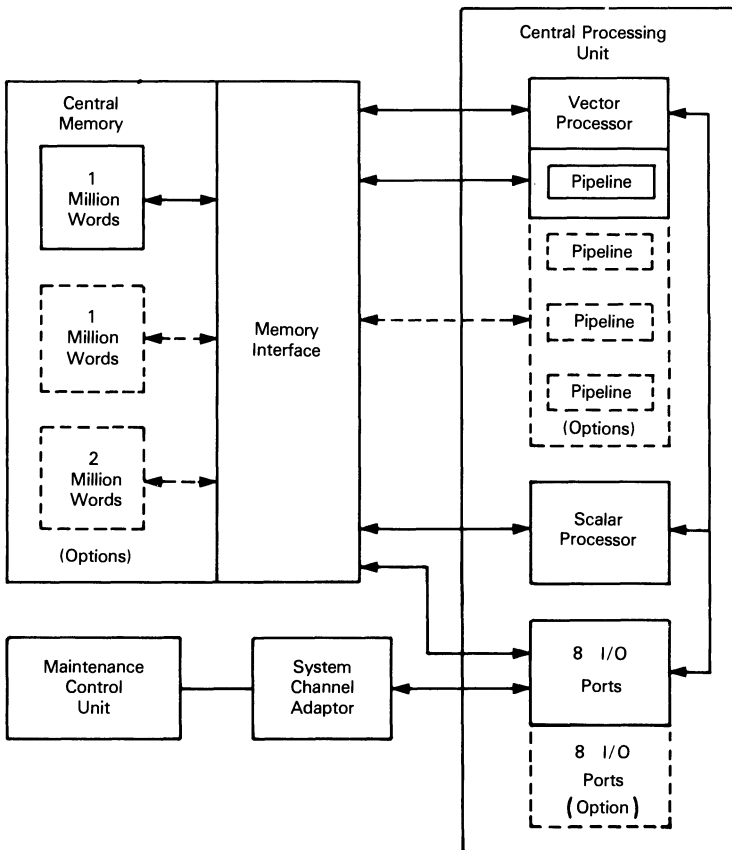


Figure 7.12 The CYBER 205 Computer System

Whenever the scalar processor issues an instruction to the vector processor it also sends, via two 64-bit highways, the contents of the required Register File locations. This



Operands are taken from locations in the Register File defined by R and S, and the result is returned to a location in the Register File defined by T. Transfers between the Register File and central memory are controlled by instructions executed in the Load/Store Unit. In these instructions R and S typically specify an indexed central memory address and T the location of the operand in the Register File.

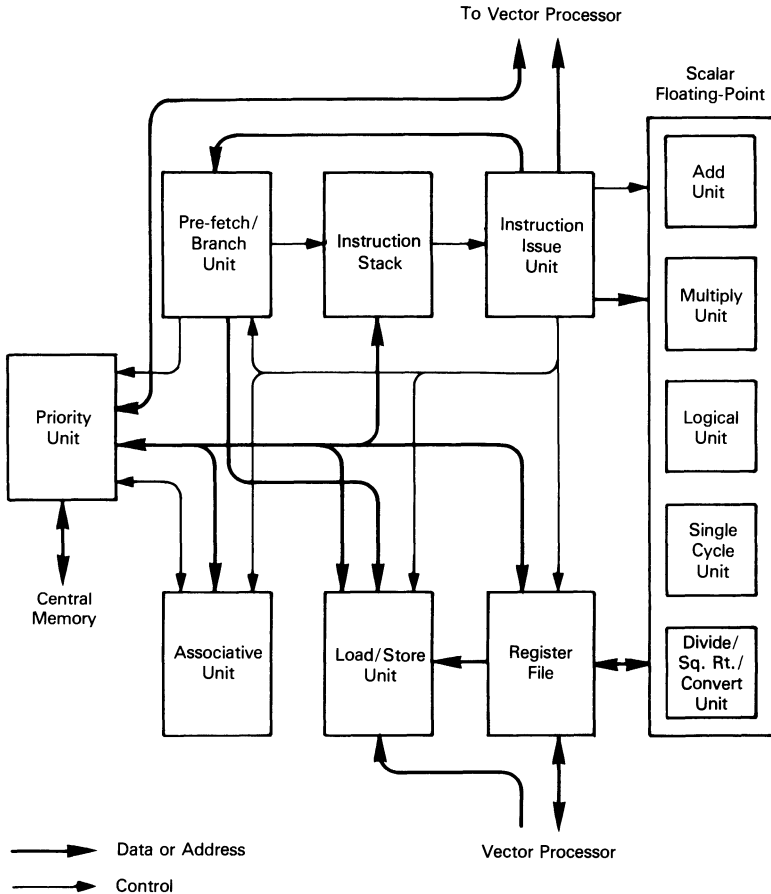


Figure 7.14 The CYBER 205 Scalar Processor

The Add/Subtract and Multiply Units each take five clock periods to produce a result and return it to the input of another unit. The Logical Unit takes three, and the Single Cycle Unit, which is used, for example, to load a literal operand into a Register File location, takes one. The Divide/Square Root/Convert Unit is not pipelined and takes

between 21 and 54 clock periods to produce its result, depending on the function. These times are all 'shortstop' times. Shortstopping is a technique which allows a result required as a source operand by a succeeding instruction to be returned directly to the input of an arithmetic unit. This is the scalar equivalent of chaining in the CRAY-1, and although similar in some respects to data forwarding in the IBM System/360 Model 91, it is different in that a shortstopped result in the CYBER 205 is sent as an input operand to an arithmetic unit at the same time as being written into the Register File. In the Model 91 a result is not written into a register at all if it is forwarded.

Shortstopping is controlled by the Instruction Issue Unit, which is implemented as a three-stage pipeline and issues one instruction per clock period unless held up by an instruction dependency. These dependencies are similar to those found in the 7600 (section 6.2.2). Thus an instruction may not issue if it specifies as its result operand or as one of its source operands a location in the Register File which is awaiting the result of some other previously issued but as yet uncompleted instruction. Source operand conflicts are resolved in the CYBER 205 by the use of 16 Result Address Registers, which hold the Register File addresses of the result operands of uncompleted instructions. Before an instruction is issued its source operand addresses are checked, simultaneously, against all 16 Result Address Registers, and if a match is found with a valid address, issuing of the instruction is held up until the conflict is resolved.

Shortstopping allows these conflicts to be resolved more quickly in the CYBER 205 than in the 7600, since the shortstop route allows a result to be returned as a source operand before it has been written into the Register File. Only one shortstopped result can appear in any one clock period, since as in the case of the 7600, only one result can be written into the Register File in any one clock period. (The Register File can, however, deliver two source operands in one clock period, as well as receiving one result.) Thus the issuing of an instruction must be delayed if its result would arrive at the Register File at the same time as the result from a previously issued, but slower, instruction. This type of conflict, and result operand conflicts, are resolved by means of a result position timing chain, similar in principle to the X Register Access Control Logic in the 7600 (section 6.2.2). Before an instruction is issued its result operand address is checked against result addresses in the timing chain, and if a match is found, or if the required Register File access reservation register is occupied, issuing of the instruction is held up until the conflict is resolved.

Vector and scalar operations can be executed in parallel, provided there are no Register File reference conflicts, and provided that no central memory references are required by scalar instructions during the execution of a previously issued vector instruction. When the Instruction Issue Unit decodes a vector instruction, and the vector processor is not busy, it supplies the vector processor with the function and with the contents of all the Register File locations specified in the instruction. The Instruction Issue Unit reserves any Register File locations which the vector instruction may update during the course of its operation, and once the vector processor has started execution of the vector instruction, the Instruction Issue Unit is free to issue subsequent (scalar) instructions.

#### **7.4.2 Virtual Address Translation**

As in the case of the STAR-100, the CYBER 205 uses a virtual addressing scheme derived from that used on Atlas. Thus addresses generated by the scalar and vector processors are normally logical (virtual) addresses which must be translated into absolute (real) storage addresses before being sent to central memory. This translation is effected by the Associative Unit (figure 7.14) which contains 16 associative address registers. In the Atlas scheme (section 3.2) each associative page register bore a one-to-one correspondence with a real page in the core store. The CYBER 205 Associative Unit is more closely related to the MU5 Store Access Control Unit (section 3.5), however, in that the number of associative registers is considerably smaller than the number of pages in real store, and the associative registers are therefore current page registers containing the most recently used virtual pages. The remainder of the page address translation table (the 'space table') is held in fixed locations in central memory in the CYBER 205 and is searched sequentially under microprogram control if a particular page address is not found in the associative registers (the 'page table').

Entries in the page table registers (AR00 - 15) are held as an ordered list of most recent use. Whenever a virtual address is presented for association and a match occurs, the content of the register containing the matching address is moved to the top of the list (register AR00) and the contents of all those above the matching register move down one position. If no match is found, the content of AR15 is copied into a buffer register and the contents of registers AR00 - 14 move down one position leaving a null entry in AR00. The contents of all the associative registers are then copied into central memory and the contents of the space table are examined in sequence using the associative facilities of the page table registers. Thus

the first entry in the space table is read and its location in central memory replaced by the previous content of AR15, held in the buffer register. After this first word has been examined it is itself copied into the buffer register. If no match occurs this cycle of events is repeated, with the second space table entry being read and replaced in central memory by the first word (now held in the buffer register), and so on. When a match occurs the page table registers are re-loaded from central memory and the content of the matching address, held in the buffer register, is copied into AROO. Thus the ordering of entries in the page table also extends into the space table.

Since different programs running on the CYBER 205 require differing numbers of pages, a mechanism is required to prevent the hardware searching invalid page or space table entries. An 'end-of-table' marker is used for this purpose. If this marker is found in a page table register when there is a page fault, an interrupt is generated immediately and the space table is not searched. If the end-of-table marker is in the space table, however, then it will only be read, and an interrupt generated, after all valid entries in the space table have been examined.

The real addresses generated by the Associative Unit are passed to the Priority Unit (figure 7.14), which co-ordinates central memory requests from different sources. In the event of two requests arriving simultaneously, the one with higher priority is serviced first and the other is delayed. I/O requests, for example, have the highest priority and are always accepted immediately. Requests may also be delayed because of bank busy conflicts in central memory. In the case of vector requests being delayed, the inertia of the address generation pipeline is such that a Re-try Unit is required to buffer unsatisfied requests. When the Re-try Unit is activated, address generation in the vector processor is halted, but as many as three additional requests may arrive before the flow ceases, and the Re-try Unit therefore has buffering capacity for four requests. During the fourth clock period after activation the Re-try Unit re-transmits the buffered requests to the Priority Unit and, if the initial request is accepted, releases the address generation hold-up in the vector processor. The first new request then arrives at the Priority Unit just after the last buffered request is transmitted from the Re-try Unit.

### 7.4.3 The Vector Processor

Figure 7.15 shows the overall structure of the CYBER 205 vector processor. The Vector Control Unit receives vector

functions sent from the scalar processor, together with the appropriate Register File contents, and controls the setting up and execution of each instruction within the vector processor. During the execution of a typical vector instruction, addresses are generated by the Stream Addressing Pipeline and sent via the Associative Unit in the scalar processor to central memory. The corresponding store words accessed from central memory are received by the Input Stream Unit which aligns the various input operand streams correctly for transmission to the Floating-point Pipeline or String Unit. Results from the Floating-point Pipeline or String Unit flow into the Output Stream Unit, and thence back to central memory. The Stream Units act in a similar way to the Stream Unit in the STAR-100 (section 7.3), not only in aligning the input streams correctly, but also in organising proper starting and finishing points for vectors which are not aligned on 512-bit store word boundaries.

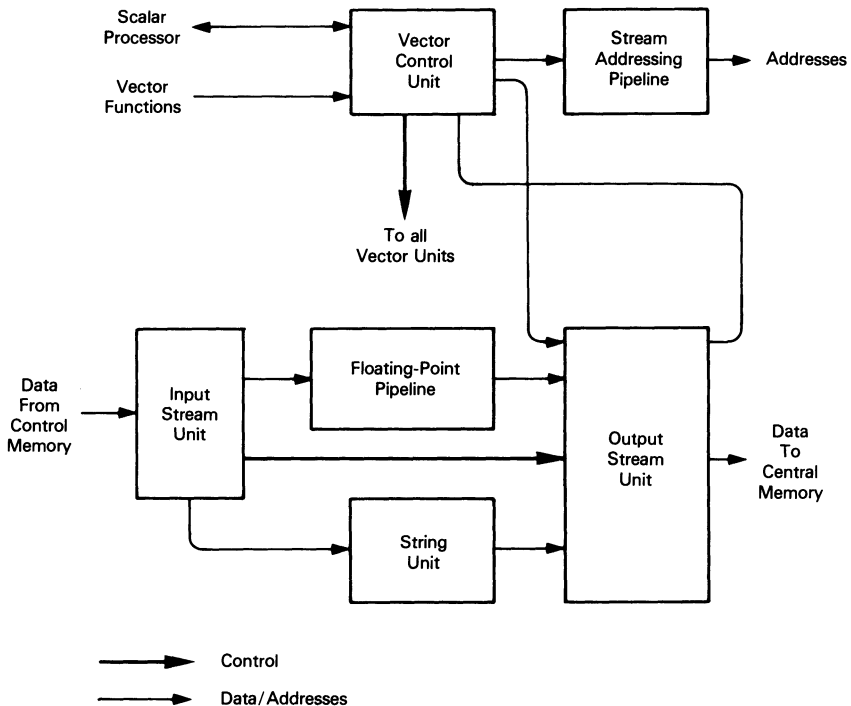


Figure 7.15 The CYBER 205 Vector Processor

Most vector instructions process data using two data input streams, one control vector input stream and one data output stream. These streams are specified by the origin addresses,

field lengths and off-sets or indices selected from locations in the Register File by the instruction before it leaves the scalar processor, and sent to the Vector Control Unit when the instruction is issued. During the setting up of a vector instruction within the vector processor, off-set values are added to the corresponding origin addresses and subtracted from the field lengths before being sent to the Stream Addressing Unit. For each data stream the Stream Addressing Unit generates sequential store addresses (normally to 512-bit words) by incrementing the origin address until the field length, which is decremented correspondingly, reaches zero.

The generation of addresses for each input and output stream is treated as an independent operation within the Stream Addressing Unit. This allows each operation to be stopped and started independently, and allows the Stream Addressing Unit to avoid bank busy conflicts in central memory. Since the banks have an 80 ns cycle time, a request to any one bank causes a four clock-period busy condition, and a subsequent request to the same bank is not allowed to proceed until a further three clock periods have elapsed. In addition the priority of requests must be taken into account; within the vector processor input operand requests have highest priority, vector write requests have second priority, string write requests third and control vector read requests fourth. A request of low priority must therefore be checked against higher priority requests to the same bank up to three clock periods ahead and up to three clock periods behind the clock period in which it will itself access central memory. A timing chain mechanism similar to that used to control entry of results into the Register File (section 7.4.1) is used for this purpose.

The Stream Addressing Unit also controls the flow of data through the Input and Output Stream Units. The Input Stream Unit contains a buffer for each input stream large enough to hold all data requested but not used. Associated with each buffer is a counter in the Stream Addressing Unit which is incremented each time a 128-bit word is requested from central memory, and decremented each time a 128-bit word is removed from the buffer for use. If a counter indicates that the corresponding buffer is full, then further requests for that input stream are held up. For the Output Stream Unit a counter in the Stream Addressing Unit is incremented each time a result is entered into the output buffer, and decremented each time a write request is sent to central memory.

A vector order terminates once all the input data has passed through the Input Stream Unit. For each of the two data input streams a field length register in the Input Stream Unit



is loaded at start-up time. These registers are decremented whenever data is sent to a functional unit and, when they reach zero, an 'empty' signal is sent to the Vector Control Unit. Processing of the next instruction can then begin, even before the last few operands of the current instruction have been processed and the results written back to central memory. If an interrupt occurs before an instruction terminates (as a result of a page fault, for example) the current values of operand addresses and field lengths are preserved in central memory. Once the interrupt has been serviced, these values are copied back into the vector processor, and the interrupted instruction is re-started.

#### 7.4.4 The Vector Floating-point Pipeline

The vector Floating-point Pipeline provides logical and arithmetic operand processing for all vector instructions. It is made up of five pipelined operand processing units, a Data Interchange and associated control logic (figure 7.16). The typical CYBER 205 system has a two-pipeline configuration in which two 64-bit or four 32-bit operands presented on each of the A and B inputs can be processed simultaneously. Except in the case of divide and square root operations, these pipelines can accept a new pair of inputs in every clock period. A one-pipeline configuration processes the two halves of the 128-bit words supplied to it sequentially and therefore runs at half the speed of the two-pipeline version. The four-pipeline version is essentially two two-pipeline processors, giving a total data path width of 256 bits. Floating-point numbers use a binary exponent and have the formats shown in figure 7.17, where all exponents and mantissae are represented by 2's complement integers.

In a normal two-pipeline configuration the Add Unit receives operands from the Data Interchange over two 128-bit data highways and returns results to the Data Interchange over a single 128-bit highway. In a straightforward add or subtract operation successive elements of vector B are added to or subtracted from successive elements of vector A, and the results written to successive elements of vector C. Feedback paths within the Add Unit allow other types of add operation to be implemented, however, so that a single result element C may be obtained by summing all the elements of vector A, for example, or a vector C may be produced by repeated addition of a scalar element B to the initial value of scalar element A.

The Multiply Unit is similarly connected to the Data Interchange by two 128-bit input operand highways, and one 128-bit result highway. In a typical multiply operation successive elements of vector A are multiplied by successive

elements of vector B and the results written to successive elements of vector C. An internal feedback path allows a single result to be formed by multiplying together all the elements of vector A, and the Multiply Unit also contains the logic for divide and square root operations.

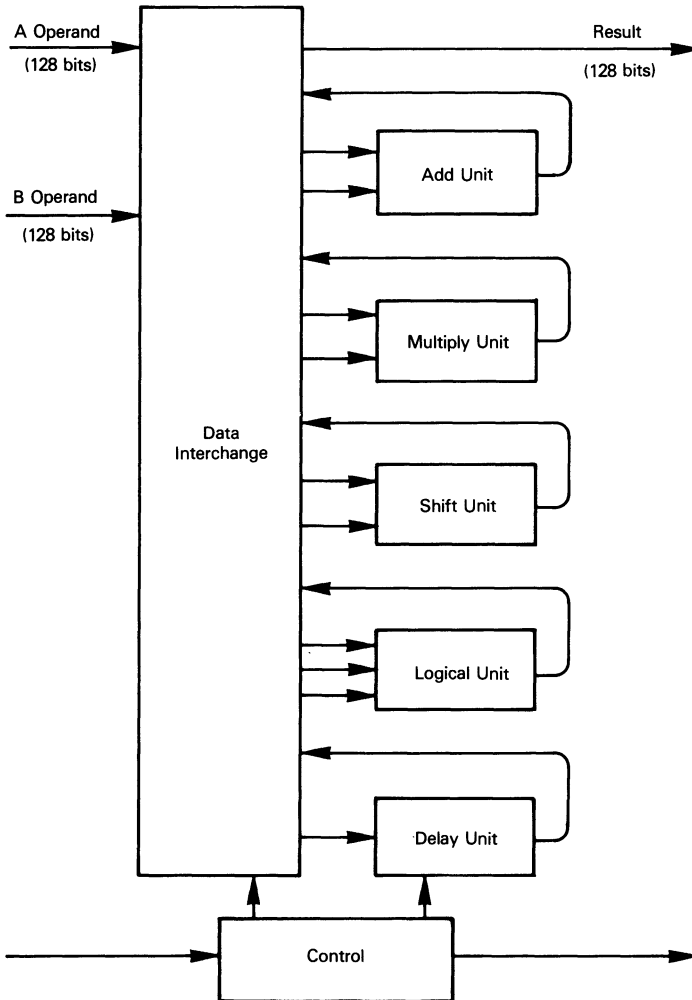


Figure 7.16 The CYBER 205 Vector Floating-point Pipeline

The Shift Unit has one 128-bit data input highway (A) and one 14-bit input highway (B) which supplies a 7-bit shift count for each half of the pipeline. Each 64-bit element of vector A is shifted left or right according to the most

significant bit of the corresponding 7-bit element of vector B, with the number of bit position shifts being determined by the six least significant bits of B.

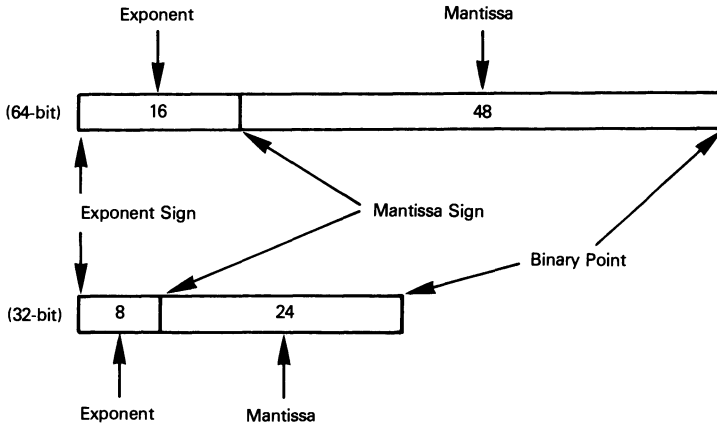


Figure 7.17 CYBER 205 Floating-point Formats

The Logical Unit carries out bit-by-bit logical operations between pairs of A and B elements supplied via 128-bit input highways and returns its results to the Data Interchange via a 128-bit result highway. It also carries out pack and unpack operations on floating-point numbers (similar to those implemented in the CDC 6600 and 7600) and a masked compare instruction for which a third input is required, containing the 128-bit mask. This instruction searches elements of vector A in sequence for a bit-by-bit match with the single element B; bit positions for which the bit in the mask is zero are assumed to match. As each word is examined, the Register File location containing the index of A is updated, so that, if a match is found, the index provides a means of locating the position of the matching element. If no match is found the index is left pointing to the end of the vector. When the instruction terminates a condition code is set to indicate the result.

For simple vector instructions the Data Interchange is configured to connect the input and output highways to the appropriate processing unit. The 'Select Link' instruction, however, causes the Data Interchange to be configured such that the succeeding two instructions in the code sequence become chained together. In this case the output of the unit used by the first instruction of the pair is routed to the input of the unit used by the second. Only two vector streams may be used in total, but this does allow commonly occurring triadic operations such as

$$\text{Vector C} = \text{Vector A} * \text{Constant} - \text{Vector B}$$

to be implemented in this manner. Not only does this allow the Multiply and Add Units to operate in parallel, but it also avoids the need to write the intermediate result vector into central memory and then read it out again.

Chaining of this sort occurs automatically during the execution of vector macro instructions such as the scalar product instruction, for example, in which pairs of input operands are multiplied together in the Multiply Unit and their results then summed in the Add Unit. This summation also involves the use of the Delay Unit, which contains a 16-word temporary buffer store. Successive 128-bit words sent to the Delay Unit are written into successive locations selected on a cyclic basis by a write counter. These same words are then read out again, and returned to the Data Interchange, under the control of a read counter. The delay function is implemented by offsetting the read and write counters by the required number of clock cycles of delay.

In the summation of results from the Multiply Unit in the scalar product operation (and similarly in the summation of all the elements of a single vector), the addition is performed by feeding values into input A of the Add Unit, and routing the output of the Add Unit back into input B. This produces a recursive effect similar to that found in the CRAY-1 (section 6.4.4). Because of the delay through the Add Unit pipeline, the zeroth input value, having passed through the Add Unit, is returned to input B in time to be added to the eighth value, the first in time to be added to the ninth, and so on. After a further pass through the Add Unit the sum of the zeroth and eighth values is ready to be added to the sixteenth value, and so on, until the input stream is exhausted and eight partial sums are left circulating in the Add Unit. Adding these partial sums together involves the use of the Delay Unit. The first four partial sums are delayed by four clock periods so that they become aligned with the last four at the input to the Add Unit. The four results obtained from this operation are then further added in pairs using a two-clock period delay, and a final add involving a single clock period delay produces the desired result ready to be written into the Register File.

#### **7.4.5 Sparse Vector Operations**

One of the distinctive features of the CYBER 205 architecture is the provision in hardware of facilities for handling sparse vectors. In many computational problems vectors with large numbers of zero or near-zero elements occur, and the ability

to store and process these vectors in sparse form allows considerable savings to be made in both storage space and processing time. The key to the handling of sparsity in the CYBER 205 is the control or 'order' vector which, as we saw in section 7.3, was introduced in the STAR-100 to overcome the problem of writing data to only part of a word in store. Thus a sparse vector is made up of two parts, a data vector and an order vector. The order vector contains one bit for each element of the full vector indicating whether or not a (non-zero) value exists for that element. The data vector is a sequence of 32-bit or 64-bit floating-point numbers constituting the non-zero values.

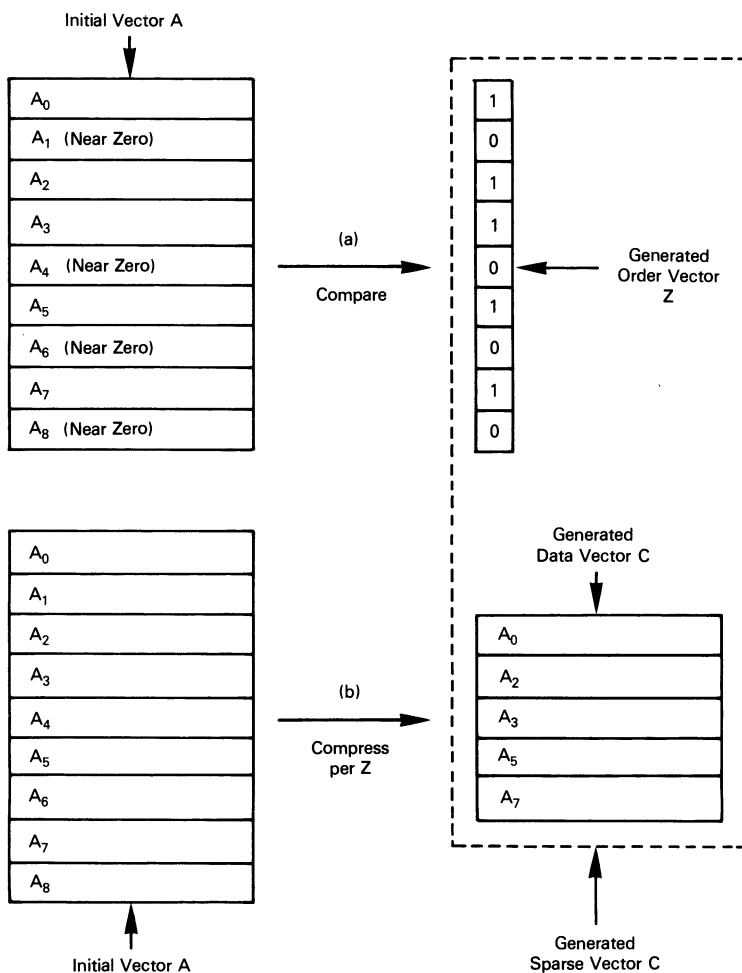


Figure 7.18 Formation of a Sparse Vector

Figure 7.18 shows how a sparse vector might typically be formed. A COMPARE instruction is first used in which each element of the initial vector A is compared against the limiting small value below which elements are to be treated as zero. For each element below this limit a zero is generated in the result control vector Z, while for each element above the limit a one is generated in Z. A COMPRESS instruction is then used in which the order vector Z generated by the COMPARE instruction is used to determine whether or not values in the initial vector A are to be copied into the data vector C. Only values for which the corresponding bit in the order vector is a one are copied to vector C, and these are written into sequential store locations. Thus there is no longer any direct correspondence between the position of an element in the vector and its position in store relative to the vector origin, and in sparse vector operations the order vector must be used to identify the positional significance of each element in the data vector. These order vectors are specified in sparse vector instructions by the X, Y and Z fields. For the source vectors this involves interpreting the contents of the locations in the Register File specified by the X and Y designators (figure 7.13(a)) as origin + field length pairs rather than as offsets.

When two sparse vectors are combined in arithmetic operations, their order vectors are used to align corresponding elements correctly and to determine the outcome of the operation. When two sparse vectors are combined by multiplication or division, for example, a result element is normally produced, and a one written into the result order vector, only when corresponding elements in the input vector are both non-zero, as indicated by the presence of a one in each of the source order vectors. When two sparse vectors are added or subtracted, a result element is normally produced when there is a non-zero element in corresponding positions in either source vector.

The actual logical combination of source order bit pairs used to produce the bit to be written into the result order vector, and hence to determine whether or not a value is to be written into the result data vector, is determined by bits in the sub-function field of the instruction (G in figure 7.13(a)). For multiplication or division the AND operation is normally used, while for addition or subtraction the OR operation is normal. However, the OR operation may be specified for the former, the AND operation for the latter, and Exclusive OR or Implication for either. In a multiplication instruction where Exclusive OR is specified, for example, a non-zero element in one source vector is implicitly multiplied by 1 when no corresponding non-zero

element exists in the other source vector, and no result is produced if neither or both source elements exist.

#### 7.4.6 Indexed List Operations

The memory of the CYBER 205 is organised such that within each one million-word unit there are 16 interleaved 32-bit stacks allowing eight sequentially addressed 64-bit words (512 bits) to be accessed in parallel. Each stack is divided into eight independent banks, each holding 16K words. A new bank read access can be initiated in every clock period, provided it is to a bank which is not already busy, and the banks can share a single stack data path since this path is only required during the last clock period of the four clock-period bank cycle. One of the criticisms levelled at the CYBER 205 is that vectors may only be accessed on the basis of a unit increment between successive elements, so that matrices stored by row, for example, cannot easily be accessed by column. The high memory bandwidth of the CYBER 205 is only obtainable because many adjacent elements can be read out in parallel in each memory cycle, however, and to allow different accessing patterns in standard vector orders would severely degrade performance. Provision of increments different from unity would also require changes to the instruction format, and alternative general mechanisms for the re-ordering of operands are provided instead. These mechanisms are invoked by Indexed List Operations which allow for the gathering or scattering of periodic or random data within a data structure.

During execution of a scatter operation (Transmit List -> Indexed C), for example, two source vectors (A and B) are normally read from memory and a result vector (C) is written to memory. Source vector A is a list of indices referring to positions in the result vector, while source vector B is the sequential list of 32-bit or 64-bit operands to be written. During each cycle of operation, the index taken from vector A is shifted appropriately and added to the result vector origin to form a memory address. This address is then sent to memory, accompanied by appropriate write control bits and the corresponding operand taken from vector B. The index list may be formed either by geometrical considerations or by data dependent considerations. Sparse vectors with a population density of the order of 0.1 per cent, for example, can be dealt with more efficiently by this mechanism than by the order vector technique described in section 7.4.5, which is suitable for population densities in the range 1 - 10 per cent.

In a gather operation (Transmit Indexed List -> C), source vector A is again a list of indices, but these indices refer

to positions in source vector B rather than the result vector. Addresses formed by adding each index from vector A in turn to the origin of vector B are used to access operands from vector B, and these operands are then written to sequential result vector locations.

A number of variations on the general form of the gather and scatter operations may be invoked by appropriate settings of bits in the G field of the instruction (figure 7.13(a)). In particular, one option allows source vector A to be replaced by a fixed increment (a) taken from the Register File location specified by A. In this case source vector addresses in a gather operation are of the form  $b$ ,  $b+a$ ,  $b+2a$ ,  $b+3a$ , etc. (where  $b$  is the origin address of B), allowing access by column to a matrix stored by row, for example. An alternative variation allows groups of sequential B operands to be selected for each element of vector A.

As with ordinary vector operations (section 7.4.3), the Stream Units must avoid bank busy conflicts in memory during the execution of gather and scatter operations. A request to any one bank causes a four-clock busy condition, and a subsequent request to the same bank is not allowed to proceed until a further three clock periods have elapsed. In scatter operations these conflicts can occur on the write accesses, while in gather operations they can occur on read accesses to source vector B. These conflicts might be expected to affect performance quite severely in the case of non-sequential accesses. In practice, however, the number of memory banks provided is so large that, except when a specific sequence is presented in which each address has the same value modulo 64, few are found to be busy.

#### 7.4.7 Performance

In the typical two-pipeline CYBER 205 system two 64-bit operands can be transferred into the Floating-point Pipeline on each of its two input highways, and two 64-bit results transferred out, in one 20 ns clock period. This corresponds to an execution rate for 64-bit floating-point numbers of 100 MFLOPS, or 200 MFLOPS for 32-bit numbers. For linked triadic operations, or during the execution of the scalar product operation, this rate is doubled to 200 MFLOPS for 64-bit numbers or 400 MFLOPS for 32-bit numbers. The memory bandwidth required to support this execution rate is six 64-bit words per 20 ns interval, or 300 Mwords/s. Since the data highway between each one million words of central memory and the memory interface is 512 bits wide, and a read or write transfer can occur every 20 ns, memory bandwidth is actually 400 Mwords/s, and a one million word memory configuration can



therefore support a two-pipeline system without difficulty. For the four-pipeline system at least two million words of central memory are required, with each one million words being connected to the memory interface via its own data highway. In this case the maximum execution rate is 800 MFLOPS for 32-bit linked triadic or scalar product operations.

Where accessing patterns involving non-unary increments are required, performance is reduced because the vector arithmetic instruction must be preceded by a gather instruction or succeeded by a scatter instruction. These instructions themselves run at less than the full arithmetic rate and have been measured to proceed at an average rate of 40 million operations per second. This performance is an average of many operations running at a rate of 50 million per second with occasional periods of 12.5 million per second.

Perhaps a more serious criticism of the CYBER 205 from the point of view of performance on general problems is the long vector start-up time. Thus whereas an execution rate of 100 MFLOPS for 64-bit numbers can in principle be obtained with a two-pipeline configuration, this can only be achieved in practice when very long vectors are used. As the length of the vectors being used becomes shorter, the start-up time has a progressively more serious effect, and Hockney [38] has used the vector length at which performance is halved as a measure of vector efficiency. For the CYBER 205 the nominal start-up time is 1  $\mu$ s, and for a nominal result rate of 100 MFLOPS this length is clearly 100. For the CRAY-1 Hockney quotes values in the range 10 - 20, though even lower values are possible. However, the CYBER 205 was intended to be used for problems involving long vectors, and as we observed in section 7.4.3, part of the start-up time can be eliminated for successive vector operations since the processing of one instruction can often begin before the last few elements of the previous instruction have been fully processed. Any attempt at a generalised comparison between the CYBER 205 and the CRAY-1 is largely irrelevant since the performance of each is critically dependent on the problem being solved and the way it is mapped on to the hardware. Both machines represent significant improvements over previously available commercial supercomputers, and there appears to be room in the marketplace for both.

## 7.5 FUTURE DEVELOPMENTS

Almost any attempt to predict future developments in the world of computers is doomed to failure, since some technological developments which appear promising fail to materialise, while others perform beyond all reasonable expectations. What is

clear is that in the past architectural techniques pioneered on large machines have eventually found their way into small machines, and there is no sign that this trend is abating as microprocessors become more complex and more sophisticated. The manufacturers of high performance systems constantly strive to produce better products, and for students of computer architecture the design of the latest supercomputer will always be important.

# References

1. G. M. Amdahl, G. A. Blaauw and F. P. Brooks, 'Architecture of the IBM System/360', IBM Journal of R & D, 8 (1964) 87-101.
2. H. S. Stone, Introduction to Computer Architecture, (Science Research Associates, Chicago, 1975).
3. S. H. Lavington, History of Manchester Computers, (NCC, 1975).
4. T. Kilburn, D. B. G. Edwards, M. J. Lanigan and F. H. Sumner, 'One-level Storage System', IRE Trans. EC-11 (1962) 223-234.
5. W. Buchholz, ed., Planning a Computer System, (McGraw-Hill, 1962).
6. E. I. Organick, The Multics System, (MIT Press, Cambridge, Mass., 1972).
7. J. E. Thornton, Design of a Computer: The Control Data 6600, (Scott Foresman & Co., 1970).
8. D. Morris and R. N. Ibbett, The MU5 Computer System, (The Macmillan Press, 1979).
9. 'IBM System/360 Principles of Operation', Form A22-6821-0, IBM Corporation.
10. R. P. Case and A. Padege, 'Architecture of the IBM System/370', Communications of the ACM, 21 (1978) 73-96.
11. C. G. Bell and W. D. Strecker, 'Computer Structures: What we have learned from the PDP-11', Proc. 3rd Annual Symposium on Computer Architecture, (1976) 1-14.
12. C. G. Bell, A. Kotek, T. N. Hastings and R. Hill, 'The Evolution of the DECsystem 10', Communications of the ACM, 21 (1978) 44-62.

- 166        The Architecture of High Performance Computers
13. J. K. Buckle, The ICL 2900 Series, (The Macmillan Press, 1978).
  14. E. I. Organick, Computer System Organisation - The B5700/B6700 Series, (Academic Press, 1973).
  15. M. R. Barbacci, D. Siewiorek, R. Gordon, R. Howbrigg and S. Zuckerman, 'An Architectural Research Facility - ISP Descriptions, Simulation, Data Collection', AFIPS NCC Conf. Proc., 46 (1977) 161-173.
  16. C. G. Bell and A. Newell, Computer Structures: Readings and Examples, (McGraw-Hill, 1971).
  17. J. Djordjevic, R. N. Ibbett and F. H. Sumner, 'Evaluation of Some Proposed Name-space Architectures Using ISPS', IEE Proc., 127E (1980) 120-125.
  18. H. J. Curnow and B. A. Wichmann, 'A Synthetic Benchmark', Computer Journal, 19 (1976) 43-49.
  19. S. H. Fuller, W. E. Burr, P. Shaman and D. A. Lamb, 'Evaluation of Computer Architectures via Test Programs', AFIPS NCC Conf. Proc., 46 (1977) 147-160.
  20. A. M. Lister, Fundamentals of Operating Systems, (The Macmillan Press, Second Edition 1979).
  21. P. J. Denning, 'Virtual Memory', Computing Surveys, 2, (1970) 153-189.
  22. J. S. Liptay, 'Structural Aspects of the System/360 Model 85 - The Cache', IBM Systems Journal, 7 (1968) 15-21.
  23. J. O. Murphy and R. M. Wade, 'The IBM 360/195', Datamation, (April 1970) 72-79.
  24. R. N. Ibbett and M. A. Husband, 'The MU5 Name Store', Computer Journal, 20 (1977) 227-231.
  25. S. H. Lavington, G. Thomas and D. B. G. Edwards, 'The MU5 Multicomputer Communication System', Trans. IEEE, C-26 (1977) 19-28.
  26. D. J. Kinniment and J. V. Woods, 'Synchronisation and Arbitration Circuits in Digital Systems', Proc. IEE, 123 (1976) 961-966.
  27. C. V. Ramamoorthy and H. I. Li, 'Pipeline Architecture', Computing Surveys, 9 (1977) 61-102.

28. N. A. Yannacopoulos, R. N. Ibbett and R. W. Holgate, 'Performance Measurements of the MU5 Primary Instruction Pipeline', in Information Processing 77 (North Holland, Amsterdam, 1977).
29. W. J. Watson, 'The TI ASC - A Highly Modular and Flexible Super Computer Architecture', AFIPS FJCC Conf. Proc., 41 (1972) 221-228.
30. M. J. Flynn and P. R. Low, 'The IBM System/360 Model 91: Some Remarks on System Development', IBM Journal of R and D, 11 (1971) 2-7.
31. R. M. Tomasulo, 'An Efficient Algorithm for Exploiting Multiple Arithmetic Units', IBM Journal of R and D, 11 (1971) 25-33.
32. D. W. Anderson, F. J. Sparacio and R. M. Tomasulo, 'The IBM System/360 Model 91: Machine Philosophy and Instruction Handling', IBM Journal of R and D, 11 (1971) 8-24.
33. 'Control Data 7600 / CYBER 70 Model 76 Computer Systems - Hardware Reference Manual' (Control Data Corporation, St. Paul, Minnesota, 1977).
34. R. M. Russell, 'The CRAY-1 Computer System', Communications of the ACM, 21 (1978) 63-72.
35. M. J. Flynn, 'Trends and Problems in Computer Organisation', in Information Processing 74 (North Holland, Amsterdam, 1974).
36. R. W. Holgate and R. N. Ibbett, 'An Analysis of Instruction Fetching Strategies in Pipelined Computers', Trans. IEEE, C-29 (1980) 325-329.
37. J. B. Gosling, Design of Arithmetic Units for Digital Computers, (The Macmillan Press, 1980).
38. P. M. Johnson, 'An Introduction to Vector Processing', Computer Design (Feb. 1978) 89-97.
39. R. W. Hockney and C. R. Jesshope, Parallel Computers, (Adam Hilger, Bristol, 1981).
40. M. J. Flynn, 'Some Computer Organisations and their Effectiveness', Trans. IEEE, C-21 (1972) 948-960.

41. P. C. Capon, R. N. Ibbett and C. R. C. B. Parker, 'The Implementation of Record Processing in MU5', IEE Conference Proceedings No. 121 (London, 1974).
42. R. G. Hintz and D. P. Tate. 'Control Data STAR-100 Processor Design', COMPCON '72 Digest (1972) 1-4.
43. N. Lincoln, 'A Safari Through the Control Data STAR-100 with Gun and Camera', AFIPS, National Computer Conference (1978).
44. K. E. Iverson, A Programming Language, (Wiley, New York, 1962).
45. 'Control Data CYBER 200 Model 205 Computer System - Hardware Reference Manual', (Control Data Corporation, St. Paul, Minnesota, 1981).

# *Index*

Addition 4, 65, 96, 106  
APL 145  
Asynchronous timing 54  
Atlas 2, 13, 26 29

B-lines 1, 13  
Block transfer unit 46  
B-store 13  
Burroughs B5500 16

Cache store 33  
CDC 6600 2, 7, 28, 77, 95  
CDC 7600 7, 28, 80, 104  
CFA programs 23  
Chaining 119  
Common data bus 67  
Concurrency 3 95  
Control point 60 89  
Control transfers 59, 75 79, 83, 89, 98  
Control vector 145  
CRAY-1 87, 112, 161  
Current page registers 44, 151  
CYBER 205 29, 80, 146

DEC PDP-10 12, 23, 33  
Descriptor 21, 127  
Division 98, 107, 117, 149  
Dope vectors 134

ECL 3, 56, 146  
ECLIPSE 15  
Extracodes 14  
Extra Name Base 21

Fixed-head disc 46  
Floating-point arithmetic 26, 62, 68, 97, 106, 116  
Floating-point formats 64, 97, 117, 157  
Function queue 135

Gather operations 161

IBM System/360 10, 33, 67, 72  
IBM System/370 12, 35  
ICL 2900 12, 126  
Indexed list operations 161  
Instruction dependencies 99, 108, 150  
Instruction hold-ups 58  
Instruction pre-fetching 72, 80, 84  
Interleaving 26  
ISPS 23

Jump trace 85, 92

Loop catching 72, 76, 79, 80, 85

Mark 1 1  
Mercury 1, 26  
MU5 3, 17, 37, 49, 82, 126  
MU5/2 23  
MU5/3 23  
MU5 Exchange 44  
MU5 Instruction Buffer Unit 50, 60, 82  
MULTICS 2, 14  
Multiplication 66, 98, 107, 117, 149, 155

Name Base 21  
Name Store 5, 18, 37, 50

One-address systems 12, 22  
One-level store 14, 30, 44  
Operand Buffer System 40, 129, 135

Page address registers 31  
Paging 14, 29, 136, 145, 155  
PDP-11 15, 44  
Performance 43, 61, 91, 110, 121, 161  
Peripheral processors 10  
Pipeline 6, 49  
Pipeline delay chain 53  
Platters 3, 37  
Population count 107, 116  
Primary Operand Unit 37, 49

Read-after-write 58, 67  
Read flags 101  
Real address 14, 30, 51  
Register file 147  
Reverse operations 19  
Reverse Polish 16

Scalar product loop 67, 123, 127, 158



Scatter operations 161  
Scoreboard 9, 95, 99  
Scratch pad 9, 96, 105  
Secondary Operand Unit 40, 128  
Segmentation 14, 21, 37  
Shortstopping 149  
Sparse vectors 158  
Stack Front 21, 56  
Stacking machine 16, 18  
STAR-100 33, 80, 142  
Store-to-store orders 158  
Stretch 2  
String operations 136  
Synchronous timing 54

Texas Instruments ASC 62  
Three-address systems 7, 22  
Triadic operations 157  
Two-address systems 10, 22

Vector processing 62, 112, 126, 142, 152  
Vector Parameter File 63  
Virtual address 12, 14, 30, 51, 145, 151  
V-store 15

Wichmann benchmark 23

X register access control 108, 150

Z8000 15  
Zero-address systems 51, 22