

# Architecture of High Performance Computers Volume I

R.N. Ibbett and  
N.P. Topham

# Architecture of High Performance Computers

## Volume I

Uniprocessors and vector processors

R. N. Ibbett and N. P. Topham

Department of Computer Science  
University of Edinburgh



**M**  
MACMILLAN

© Roland N. Ibbett and Nigel P. Topham 1989

All rights reserved. No reproduction, copy or transmission of this publication may be made without written permission.

No paragraph of this publication may be reproduced, copied or transmitted save with written permission or in accordance with the provisions of the Copyright Act 1956 (as amended), or under the terms of any licence permitting limited copying issued by the Copyright Licensing Agency, 33-4 Alfred Place, London WC1E 7DP.

Any person who does any unauthorised act in relation to this publication may be liable to criminal prosecution and civil claims for damages.

First published 1989

Published by  
MACMILLAN EDUCATION LTD  
Houndmills, Basingstoke, Hampshire RG21 2XS  
and London  
Companies and representatives  
throughout the world

Printed and bound in Great Britain at  
The Camelot Press Ltd, Southampton

British Library Cataloguing in Publication Data  
Ibbett, Roland N.

Architecture of high performance computers.

—(Macmillan computer science series).

1. High performance computer systems.

Design

I. Title      II. Topham, Nigel P.

001.64

ISBN 0-333-46362-5 V.1 Pbk

ISBN 0-333-48988-8 V.2 Pbk

## Macmillan Computer Science Series

*Consulting Editor*

Professor F. H. Sumner, University of Manchester

- S. T. Allworth and R. N. Zobel, *Introduction to Real-time Software Design, second edition*
- Ian O. Angell and Gareth Griffith, *High-resolution Computer Graphics Using FORTRAN 77*
- Ian O. Angell and Gareth Griffith, *High-resolution Computer Graphics Using Pascal*
- M. A. Azmoodeh, *Abstract Data Types and Algorithms*
- C. Bamford and P. Curran, *Data Structures, Files and Databases*
- Philip Barker, *Author Languages for CAL*
- A. N. Barrett and A. L. Mackay, *Spatial Structure and the Microcomputer*
- R. E. Berry, B. A. E. Meekings and M. D. Soren, *A Book on C, second edition*
- G. M. Birtwistle, *Discrete Event Modelling on Simula*
- B. G. Blundell, C. N. Daskalakis, N. A. E. Hayes and T. P. Hopkins, *An Introductory Guide to Silvar Lisco and HILO Simulators*
- B. G. Blundell and C. N. Daskalakis, *Using and Administering an Apollo Network*
- T. B. Boffey, *Graph Theory in Operations Research*
- Richard Bornat, *Understanding and Writing Compilers*
- Linda E. M. Brackenbury, *Design of VLSI Systems—A Practical Introduction*
- J. K. Buckle, *Software Configuration Management*
- W. D. Burnham and A. R. Hall, *Prolog Programming and Applications*
- P. C. Capon and P. J. Jinks, *Compiler Engineering Using Pascal*
- J. C. Cluley, *Interfacing to Microprocessors*
- J. C. Cluley, *Introduction to Low Level Programming for Microprocessors*
- Robert Cole, *Computer Communications, second edition*
- Derek Coleman, *A Structured Programming Approach to Data*
- Andrew J. T. Colin, *Fundamentals of Computer Science*
- Andrew J. T. Colin, *Programming and Problem-solving in Algol 68*
- S. M. Deen, *Fundamentals of Data Base Systems*
- S. M. Deen, *Principles and Practice of Database Systems*
- Tim Denvir, *Introduction to Discrete Mathematics for Software Engineering*
- P. M. Dew and K. R. James, *Introduction to Numerical Computation in Pascal*
- M. R. M. Dunsmuir and G. J. Davies, *Programming the UNIX System*
- D. England *et al.*, *A Sun User's Guide*
- K. C. E. Gee, *Introduction to Local Area Computer Networks*
- J. B. Gosling, *Design of Arithmetic Units for Digital Computers*
- M. G. Hartley, M. Healey and P. G. Depledge, *Mini and Microcomputer Systems*
- Roger Hutty, *Z80 Assembly Language Programming for Students*
- Roland N. Ibbett and Nigel P. Topham, *Architecture of High Performance Computers, Volume I*
- Roland N. Ibbett and Nigel P. Topham, *Architecture of High Performance Computers, Volume II*
- Patrick Jaulent, *The 68000—Hardware and Software*
- P. Jaulent, L. Baticle and P. Pillot, *68020–68030 Microprocessors and their Coprocessors*

(continued overleaf)

J. M. King and J. P. Pardoe, *Program Design Using JSP—A Practical Introduction*  
 E. V. Krishnamurthy, *Introductory Theory of Computer Science*  
 V. P. Lane, *Security of Computer Based Information Systems*  
 Graham Lee, *From Hardware to Software—an introduction to computers*  
 A. M. Lister and R. D. Eager, *Fundamentals of Operating Systems, fourth edition*  
 Tom Manns and Michael Coleman, *Software Quality Assurance*  
 G. P. McKeown and V. J. Rayward-Smith, *Mathematics for Computing*  
 Brian Meek, *Fortran, PL/I and the Algols*  
 A. Mével and T. Guéguen, *Smalltalk-80*  
 Barry Morrell and Peter Whittle, *CP/M 80 Programmer's Guide*  
 Derrick Morris, *System Programming Based on the PDP11*  
 Y. Nishinuma and R. Espesser, *Unix – First contact*  
 Pim Oets, *MS-DOS and PC-DOS—A Practical Guide, second edition*  
 Christian Queinnec, *LISP*  
 E. R. Redfern, *Introduction to Pascal for Computational Mathematics*  
 Gordon Reece, *Microcomputer Modelling by Finite Differences*  
 W. P. Salman, O. Tisserand and B. Toulout, *FORTH*  
 L. E. Scales, *Introduction to Non-linear Optimization*  
 Peter S. Sell, *Expert Systems—A Practical Introduction*  
 A. G. Sutcliffe, *Human-Computer Interface Design*  
 Colin J. Theaker and Graham R. Brookes, *A Practical Course on Operating Systems*  
 M. R. Tolhurst *et al.*, *Open Systems Interconnection*  
 J-M. Trio, *8086-8088 Architecture and Programming*  
 M. J. Usher, *Information Theory for Information Technologists*  
 B. S. Walker, *Understanding Microprocessors*  
 Peter J. L. Wallis, *Portable Programming*  
 Colin Walls, *Programming Dedicated Microprocessors*  
 I. R. Wilson and A. M. Addyman, *A Practical Introduction to Pascal—with BS6192, second edition*

#### **Non-series**

Roy Anderson, *Management, Information Systems and Computers*  
 I. O. Angell, *Advanced Graphics with the IBM Personal Computer*  
 J. E. Bingham and G. W. P. Davies, *A Handbook of Systems Analysis, second edition*  
 J. E. Bingham and G. W. P. Davies, *Planning for Data Communications*  
 B. V. Cordingley and D. Chamund, *Advanced BASIC Scientific Subroutines*  
 N. Frude, *A Guide to SPSS/PC+*  
 Barry Thomas, *A PostScript Cookbook*

# Preface

This second edition of *The Architecture of High Performance Computers* has been produced as a two volume set. Volume I is a revised, updated and slightly expanded version of the first edition, dealing mainly with techniques used in uniprocessor architectures to attain high performance. Many of these techniques involve some form of parallelism, but this parallelism is largely hidden from the user. Machines which are explicitly parallel in nature are dealt with in Volume II, which concentrates on the architecture of systems in which a number of processors operate in concert to achieve high performance. The high performance structures described in Volume I are naturally applicable to the design of the elements within parallel processors, and therefore Volume II also represents a historical progression from Volume I.

Computer architecture is an extensive subject, with a large body of mostly descriptive literature, and any treatment of the subject is necessarily incomplete. There are many high performance architectures, both on the market and within research environments, far too many to cover in a student text. We have therefore attempted to extract the fundamental principles of high performance architectures and set them in perspective with case studies. Where possible we have used commercially available machines as our examples. The two volumes of this book are designed to accompany undergraduate courses in computer architecture, and constitute a core of material presented in third and fourth year courses in the Computer Science Department at Edinburgh University.

Many people gave advice and assistance in the preparation of the first edition, particularly former colleagues at the University of Manchester, and much of this has carried over into the second edition. 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, particularly Chuck Purcell, vetted parts of the manuscript and provided much useful information. In preparing this first volume of the second edition the authors are indebted to William White of Cray Research for his comments on the content of the manuscript. Preparation of the manuscript involved many hours at computer terminals and the authors would like to thank Alison Fleming for her assistance and expert advice on the use of L<sup>A</sup>T<sub>E</sub>X.

Roland Ibbett  
Nigel Topham

# Contents

<b>Preface</b>	<b>viii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Historical developments . . . . .	1
1.2 Techniques for improving performance . . . . .	2
1.3 An architectural design example . . . . .	3
<b>2 Instructions and Addresses</b>	<b>7</b>
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
2.7 Reduced instruction sets . . . . .	25
<b>3 Storage Hierarchies</b>	<b>26</b>
3.1 Store interleaving . . . . .	26
3.2 The Atlas paging system . . . . .	29
3.2.1 Real and virtual addresses . . . . .	30
3.2.2 Page address registers . . . . .	31
3.3 IBM cache stores . . . . .	33
3.3.1 The System/360 Model 85 cache . . . . .	33
3.3.2 The System/370 Model 165 cache . . . . .	35
3.3.3 The 3090 cache . . . . .	37
3.4 The MU5 Name Store . . . . .	37
3.4.1 Normal operation . . . . .	38
3.4.2 Non-equivalence actions . . . . .	40
3.4.3 Actions for ACC orders . . . . .	43
3.4.4 Name Store performance . . . . .	43
3.5 Data transfers in the MU5 storage hierarchy . . . . .	44
3.5.1 The Exchange . . . . .	44
3.5.2 The Exchange priority system . . . . .	48

<b>4</b>	<b>Pipelines</b>	<b>49</b>
4.1	Principles of pipeline concurrency . . . . .	49
4.2	The MU5 Primary Operand Unit pipeline . . . . .	52
4.2.1	Synchronous and asynchronous timing . . . . .	57
4.2.2	Variations among instruction types . . . . .	59
4.2.3	Control transfers . . . . .	62
4.2.4	Performance measurements . . . . .	63
4.3	Arithmetic pipelines — the TI ASC . . . . .	64
4.4	The IBM System/360 Model 91 Common Data Bus . . . . .	70
<b>5</b>	<b>Instruction Buffers</b>	<b>74</b>
5.1	The IBM System/360 Model 195 instruction processor . . . . .	74
5.1.1	Sequential instruction fetching . . . . .	75
5.1.2	Conditional mode . . . . .	77
5.1.3	Loop mode . . . . .	78
5.2	Instruction buffering in CDC computers . . . . .	79
5.2.1	The CDC 6600 instruction stack . . . . .	79
5.2.2	The CDC 7600 instruction stack . . . . .	82
5.3	The MU5 Instruction Buffer Unit . . . . .	83
5.3.1	Sequential instruction fetching . . . . .	86
5.3.2	Instruction effects in the Jump Trace . . . . .	87
5.4	The CRAY-1 instruction buffers . . . . .	88
5.5	Position of the Control Point . . . . .	90
<b>6</b>	<b>Parallel Functional Units</b>	<b>96</b>
6.1	The CDC 6600 central processor . . . . .	96
6.1.1	Functional units in the CDC 6600 . . . . .	97
6.1.2	Instruction dependencies . . . . .	99
6.1.3	Data highways . . . . .	104
6.2	The CDC 7600 central processor . . . . .	105
6.2.1	Functional units in the CDC 7600 . . . . .	107
6.2.2	Instruction issue . . . . .	108
6.3	Performance . . . . .	111
<b>7</b>	<b>The CRAY Series</b>	<b>113</b>
7.1	The CRAY-1 . . . . .	113
7.1.1	Functional units in the CRAY-1 . . . . .	116
7.1.2	Instruction issue . . . . .	119
7.1.3	Chaining . . . . .	121
7.1.4	Performance . . . . .	124
7.2	The CRAY X-MP . . . . .	127
7.2.1	Central memory organisation . . . . .	128
7.2.2	Chaining . . . . .	131



7.2.3	Interprocessor communication . . . . .	132
7.2.4	The CRAY Y-MP . . . . .	133
7.3	The CRAY-2 . . . . .	133
7.3.1	Background Processor architecture . . . . .	136
7.3.2	Interprocessor communication . . . . .	137
7.3.3	Technology . . . . .	138
7.4	Beyond the CRAY-2 . . . . .	139
<b>8</b>	<b>Vector Facilities in MU5</b>	<b>140</b>
8.1	Introduction . . . . .	140
8.1.1	The MU5 secondary instruction pipeline . . . . .	142
8.1.2	Descriptor processing hardware . . . . .	143
8.1.3	Complex descriptor mechanisms . . . . .	147
8.1.4	The Operand Buffer System . . . . .	149
8.2	String operations in MU5 . . . . .	150
8.2.1	Operation of the store-to-store orders . . . . .	152
8.2.2	Special store management problems . . . . .	155
<b>9</b>	<b>The CDC Series</b>	<b>156</b>
9.1	The CDC STAR-100 . . . . .	156
9.2	The CDC CYBER 205 . . . . .	160
9.2.1	The scalar processor . . . . .	162
9.2.2	Virtual address translation . . . . .	165
9.2.3	The vector processor . . . . .	166
9.2.4	The vector floating-point pipeline . . . . .	169
9.2.5	Sparse vector operations . . . . .	172
9.2.6	Indexed list operations . . . . .	174
9.2.7	Performance . . . . .	176
9.3	The ETA <sup>10</sup> . . . . .	177
9.3.1	Technology . . . . .	179
<b>10</b>	<b>Performance of Vector Machines</b>	<b>180</b>
10.1	Hardware models . . . . .	180
10.1.1	Efficiency of vector operations . . . . .	181
10.2	Measuring vector efficiency . . . . .	182
10.2.1	The effect of technology . . . . .	183
10.2.2	Optimal pipelining . . . . .	183
10.2.3	Mixed-mode performance . . . . .	185
10.3	Comparing vector machines . . . . .	189
	<b>Bibliography</b>	<b>193</b>
	<b>Index</b>	<b>198</b>

# 1 *Introduction*

Computer architecture has been defined in a number of ways by different authors. Amdahl, Blaauw and Brooks [ABB64], for example, the designers of the IBM/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 [Sto75], 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 [Lav75]), 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 the 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 eased dramatically in the late

1950s and early 1960s by the commercial availability of transistors, and the first generation of ‘supercomputers’ such as Atlas [KELS62], Stretch [Buc62], MULTICS [Org72] and the CDC 6600 [Tho70] appeared. These machines incorporated features which have influenced the design of subsequent 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 principal 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 speeds 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. Various authors have attempted to classify computers on the basis of the type of concurrency they exhibit. The most well known of these taxonomies is that proposed by Flynn [Fly72], who classified processors as Single Instruction Single Data (SISD) machines, Single Instruction Multiple Data (SIMD) vector machines, Single Instruction Multiple Data array processors, and Multiple Instruction Multiple Data (MIMD) machines. Volume I of this book deals mainly with techniques which are applicable within the design of SISD computers and with SIMD vector machines. SIMD array architectures and MIMD (multiprocessor) machines are dealt with in Volume II.

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 [MI79].

### 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 interconnec-

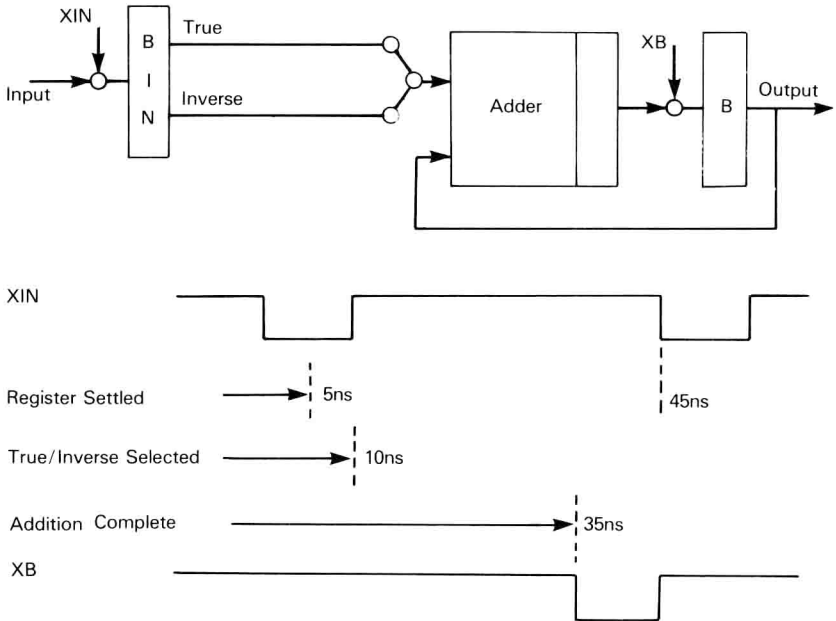



Figure 1.1 An index arithmetic unit

tions) 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, thus totalling 6" 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.

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



*Figure 1.2 The MU5 name store*

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. Thus, in the absence of some additional techniques, 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 were in use at any one time. Thus, a system which kept these operands in fast programmable registers would be able to achieve high performance. Such a technique is now in widespread use. However, this is exactly the sort of hardware features which causes compiler complexity, and which the designers of MU5 sought to avoid.

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, known as 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 a buffer register, 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 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 [Tho70]. An understanding of its organisation and operation provides a proper background not only for an appreciation of the design of the 7600 system, but



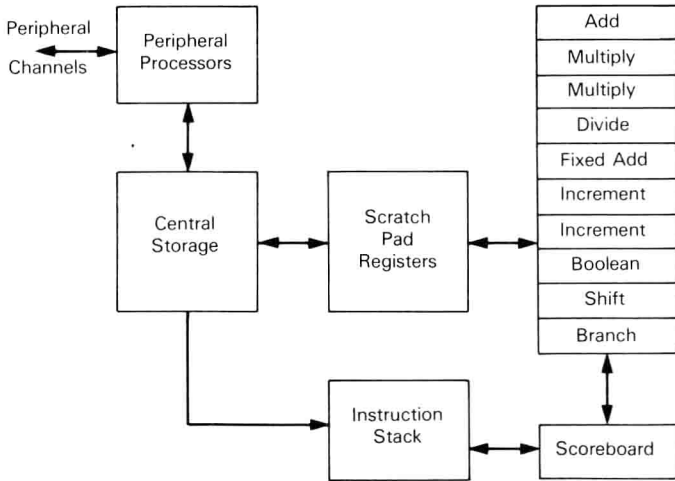


Figure 2.1 CDC 6600 processor organisation

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 chapters 6 and 7.

The CDC 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. This in turn required an instruction set and processor organisation which could exploit this parallelism, while at the same time maintaining at least the illusion of strictly sequential execution of instructions. 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. Despite the potential for instruction overlap, dependencies between instructions can still occur in a three-address system. For example, where one instruction requires as its input the result of an immediately preceding instruction, 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.