MULTI-MICROPROCESSOR SYSTEMS

Y. PAKER

APIC STUDIES IN DATA PROCESSING
NO. 18

Multi-microprocessor Systems

Y. Paker

The Polytechnic of Central London, UK

1983

Academic Press

A Subsidiary of Harcourt Brace Jovanovich, Publishers

LONDON · NEW YORK
PARIS · SAN DIEGO · SAN FRANCISCO · SÃO PAULO
SYDNEY · TOKYO · TORONTO

ACADEMIC PRESS INC. (LONDON) LTD., 24/28 Oval Road, London NW1 7DX

United States Edition published by ACADEMIC PRESS INC., 111 Fifth Avenue, New York, New York 10003

Second printing 1984

British Library Cataloguing in Publication Data
Paker, Y.
Multi-microprocessor systems.—(Automatic Programming Information Centre studies in data processing ISSN 066-4103;18)
1. Microprocessors
I. Title II. Series
001.64′04 QA76.5

ISBN 0-12-543980-6

LCCCN 82-72980

Typeset by CK Typesetters Ltd., 26 Mulgrave Road, Sutton, Surrey
Printed in Great Britain by St. Edmundsbury Press, Bury St. Edmunds, Suffolk

Preface

The quest for higher speed, higher performance and more reliable computers at lower costs (sometimes at any cost) goes on relentlessly. This is increasingly so today as the technologically advanced societies rely more and more on computers in nearly all aspects of life: central and local government, industry, commerce, defence and even leisure. Computers are not only replacing some traditional activities such as manual work on an assembly line, but are also giving rise to a whole new range of products and industries.

Even though current computing systems have progressed in a relatively short period of time over the last thirty years, in terms of tasks that they perform of data storage, retrieval, manipulation and input–output, their software and hardware concepts have not advanced very far from the early models. The next generation of computers are expected to handle far more complex systems, go beyond handling numbers and become capable of processing voice and pictures, run fully automated factories, fly inherently unstable aircraft, manage automated offices and so on... and communicate with humans interactively through speech and images. It is recognized that the next generation of computing systems, the so-called intelligent, knowledge-based systems, while taking advantage of the full power of circuit technology, will require new approaches to systems architecture.

Over the last decade, microprocessors have become key components feeding the progress of information technology. Extending the lower end of the computer spectrum, they have reduced computers to circuit components, thus opening up vast application areas. Their low cost has made computers available to practically everyone. Microprocessors also have helped the process of the convergence of computer and telecommunications technologies. Computers are no longer isolated point resources; their interconnection within a local or global framework is providing a whole new dimension, which is giving rise to structures better adapted to application requirements.

Microprocessors and ancillary circuits offer the designer a wide range of options, from a simple single-chip microcomputer up to a level of power and sophistication approaching conventional mainframe computers. The strides in the underlying circuit technologies continue and the rate of progress has yet to show signs of slowing down. The computing experience acquired over the last three decades, which has been the basis of the design of modern microprocessors, has made them very versatile components used for building a vast variety of systems as well as computers.

Microprocessors as powerful, and low-cost circuit components with relatively small size and power consumption, have brought a new design option whereby the aggregation of such devices can be used for a given application. Such structures are called multi-microprocessor systems.

There can be a variety of different motivations for building multi-microprocessor systems such as convenience, improved performance, high reliability and so on. For example, let us consider the performance factor. Fundamentally, progress in circuit technology could provide higher performance by faster clock rates (speed) and wider data paths (gate density). At a given stage of technological attainment, to improve performance, ingenuity could be applied to better arrange the given number of logic gates on a chip and/or a number of the same type of chips could be interconnected in some fashion to operate in parallel. The latter option assumes that an application could be defined as a combination of parallel processes. To improve reliability, the idea of using spare stand-by processors is an appealing possibility which favours a multi-microprocessor approach. Thus using a multiplicity of the same type, or a mix of microprocessors suitably interconnected is an important design option. This will become more important as we approach the technological limitations of circuit integration, hence the limited capabilities of a single microprocessor.

At the outset, although building multi-microprocessor structures appears attractive, in practice one encounters some problems. The fundamental difficulty comes from the fact that the principle of microcomputers, like all stored program computers, is sequential operation. In other words, their working is based on the execution of one operation at a time. This means that, whatever the complexity, an application needs to be formulated as a list of sequential instructions which is then handled by an operating system which itself consists of the same type of instructions. The execution of these atomic steps one after the other at a very high speed gives the impression of apparent simultaneity of various activities. On the other hand, when a number of microprocessors operate at the same time to handle a given application, this introduces parallelism into the basic architecture, not a straightforward proposition in a sequential environment. By contrast, analogue computers, being inherently parallel machines, adapt relatively easily to the functional requirements of an application.

Multi-microprocessors represent an important opportunity for novel architectural approaches in building computing systems to better suit given application requirements initially and also to adapt to changes during the system's life cycle. The semiconductor manufacturers have recognized this. The newer types of microprocessors have features which enable them to be interconnected via a common shared bus. Recent work in local area networks is providing a framework to connect a number of microprocessors. Yet there is no coherent approach in building such systems to take advantage of modularity. Multi-microprocessor systems can still be considered in the experimental research phase.

This text is written to describe the main issues involved in the building of multi-microprocessor systems. This is done by taking into account the current state of the art and likely future developments. The commercially available microprocessor is taken as the basic building block. It is recognized that very large scale integration (VLSI) provides many opportunities for introducing parallelism within a chip. Unfortunately, this option is only available to a few organizations with access to VLSI design and production facilities. In the text, the subject is treated at the level of systems architecture, without entering too much into detailed circuit design aspects. The major architectural approaches, interconnection schemes and related software issues are considered at some length. To illustrate the concepts introduced, numerous machines are referred to as examples. Care is taken, however, to ensure a balanced presentation so that one particular example does not overshadow the whole text. One architectural approach, the variable topology multicomputer, developed by myself and my colleagues, is perhaps given slightly more coverage than the others.

In the text a new set of symbols are proposed to identify various basic components of a multi-microprocessor system such as an individual processor, memory, bus, etc. A full complement of these symbols can be seen in Fig. 1.9. It is hoped that this new representation will, at a glance, enable the reader to extract the major architectural features as well as identify the functions of each block, without the necessity of explicitly writing this on the diagram. Note that the symbols chosen are either those used for flow-charting or a combination of these shapes. The intention is that a standard flow-charting template, which can be obtained easily, can be used to draw these diagrams. It is expected that this convention will help the reader to understand better the diagrammatical representation of various architectures presented. Furthermore, I hope that this representation will have a much wider acceptance and be adopted by microprocessor designers.

The text commences with an overview of current microprocessors, their classification and a comparative presentation of the main architectural features of a number of well-known types. In Chapter 2 a historical background to the development of multi-microprocessors is presented with examples. Then a classification for multi-microprocessor structures is given. In Chapter 3 the main motivations for building multi-microprocessor systems and various possible architectural approaches are discussed in depth. Chapter 4 then deals with interconnection and communication aspects. Thus a common bus such as MULTIBUS, a contention bus such as Ethernet, a loop such as the Cambridge ring and network structures such as the variable topology multicomputer are treated in some detail. Chapter 5 considers the point of view of a systems designer and the main issues and options available for building a multi-microprocessor system. Chapter 6 covers the important problem of software for which there are a number of different approaches, yet few generally accepted rules. The software problem is covered by discussions of the operating system, communication and language aspects. A number of examples such as MICROS, CHORUS, MEDUSA and others are mentioned. Finally, the problem of building highly reliable computers taking advantage of the redundancy that is inherent in a multi-microprocessor system constitutes Chapter 7. Included are hardware and software redundancy issues and reconfigurable systems that can sustain some failures, yet are able to continue operation with some loss of performance.

It is hoped that this book will serve researchers and students of computer architecture as well as managers and design engineers in industry to better understand the underlying principles, opportunities and difficulties in building multi-microprocessor systems. This is a very rapidly changing field, yet with every new generation of microprocessors the motivation for their use in large numbers is becoming even greater. To cope with the requirements of the next generation of computers, multi-microprocessors are bound to play an important part.

This book is the culmination of my work in real-time, parallel and distributed computing over a number of years, supported by the Science and Engineering Research Council, UK, United States European Research Office, the Polytechnic of Central London and Rennes University, France. I gratefully acknowledge the support that I have received. An international seminar in multi-microprocessor systems that I organized at PCL with INRIA in 1978, many courses and meetings, visits and discussions with various research groups in Europe and USA contributed much to this book. I am, therefore, grateful to the many people who directly or indirectly helped in the writing of this text. In particular I should like to mention the contribution of Dr M. Bozyigit who worked with me on the variable

topology multi-computer project. Finally the high quality of drawings owe much to the meticulous work of Carol Hughes, of the Audio Visual Aids department of the Polytechnic of Central London.

February 1983

Y. Paker
London

Contents

Chapter 1

Introduction to Microprocessors

1.1 Overview of Microprocessor Technologies and Trends

The availability of low cost yet increasingly more powerful microprocessors and related high-density integrated circuits is opening up new application areas which were inconceivable only a few years ago. The rapid technological progress that we have been experiencing in microelectronics, which has been a major driving force in the computer industry, does not appear to be levelling off and is expected to continue at the same rate for some years to come, till the mid-eighties and possibly beyond [1]. The usual method of indicating this progress is to draw a graph showing the number of components (or transistors) that can be built on a single chip. The year 1959 is taken as the origin, where the first planar transistor was built. The simple straight line plot, the so-called Moore's law [2], is remarkable in that it shows that the number of components that can be placed on a chip doubles every year; a unique human experience in technological progress (Fig. 1.1).

It is customary to differentiate the various stages of development of integrated circuits as shown in Table 1.1. The figures given in this table should be considered as order of magnitude quantities. The implications of each stage, however, in digital electronics and computer design have been crucial. For example, small scale integration (SSI) provided elementary logic functions which previously had to be built using discrete components (transistors, resistors). Medium scale integration (MSI) then provided logical building blocks, each incorporating several logic functions. At each stage the reliability improved, and the cost and power consumption per function decreased. This made it possible to design and build more and more complex circuits using functional blocks.

This progression, however, could not follow the same pattern when transition had to take place to large scale integration (LSI). MSI technology

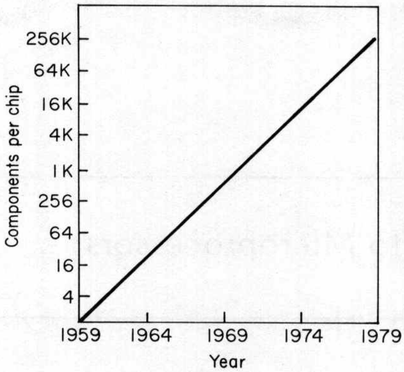

Fig. 1.1 The number of components per chip as a function of time.

has been able to produce a consistent set of circuits (building blocks) to satisfy most requirements for manufacturing a computer or complex digital systems. In fact a number of technologies have been developed like TTL (transistor transistor logic), ECL (emitter coupled logic) and others to satisfy performance, speed, and other specific requirements.

Table 1.1 Progress in integrated circuits.

	Component per chip	Year
Small scale integration (SSI)	1-10	1960-65
Medium scale integration (MSI)	10-1000	1965-70
Large scale integration (LSI)	1000-100 000	1970-80
Very large scale integration (VLSI)	100 000-(1 000 000)	1980-

LSI, and later on very large scale integration (VLSI), on the other hand, made it possible to build a complete digital system on a single chip. To ensure the widest possible utilization of a given circuit type, which is the main objective of the semiconductor manufacturers geared for volume production, it had to be as versatile as possible. This principle demonstrates the success of microprocessors, where a single chip can be used for many different purposes by simply changing the program stored in its memory. Thus we observe that while the custom-specified circuits are technically feasible to manufacture, because of the high costs involved this is possible only for big consumers like the automobile and communication industries. The main trend, however, is to standardize and reduce the types in the VLSI industry, in particular microprocessors, memory, and auxiliary circuits. The vast market that exists for a few standard types of microprocessors is ensuring the increased volume of production, hence reduction of costs.

Thus a microprocessor has now become a standard VLSI circuit element for the design and construction of a majority of digital systems, opening up new and unexplored application areas.

Currently the most successful technology used for building a microprocessor is the MOS (metal oxide silicon) process, in which a silicon wafer is used as the basic substrate to build transistors, using photo-etching, doping, and metal evaporation techniques. A measure of power of a VLSI chip is the number of gates per chip multiplied by clock speed, per gate. The current figure for this is 10^{11} gate–hertz. A US Department of Defence Project is currently underway to achieve a figure of 10^{13} gate–hertz [3]. Advances in silicon fabrication techniques alone contributed to the achievement of the current figure by increasing chip area (better yield), by introducing innovative circuits, and by scaling down the dimensions of gates and interconnections. The above project has the objectives of producing, in 1986, a processor containing 250 000 gates, operating at clock speeds of at least 25 MHz and performing between several million and several billion operations per second. The ultimate limit of MOS technology is seen as a circuit (CMOS-Complementary MOS) that will operate with a supply voltage of 400 mV, a minimum line width of 0·25 μm, dissipate 1 W at an operating frequency of 100 MHz with a chip size of 50×50 mm^2 and complexity of 100 million gates [4]. Currently, a typical complex microprocessor available is the Motorola 68000 which contains 70 000 transistors and operates at 8 MHz.

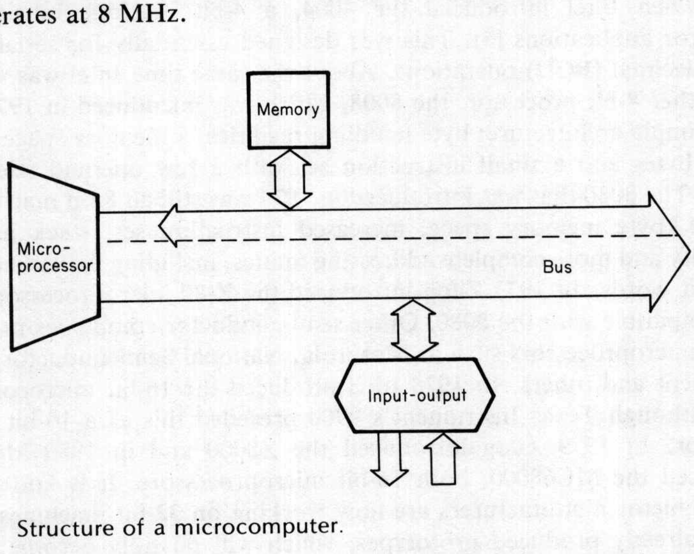

Fig. 1.2 Structure of a microcomputer.

At this stage it is perhaps necessary to give a definition of a microprocessor. It is a single integrated circuit containing all the functions of a central processing unit of a stored program digital computer. For a microprocessor to become an operational unit the minimum requirements are memory and some input–output devices. A microprocessor with memory and input–output units constitutes a digital computer which can be called a microcomputer, as shown in Fig. 1.2. Note the different shapes used for the functional components. If a single integrated circuit contains all the three elements (processor, memory, and input–output) then this is called a single-chip microcomputer. There is another approach, so-called bit-slice machines, where a central processor is built using a set of LSI circuits and memory. Normally these are microprogrammed processors where the user defines the microcode which is then stored in microprogram memory. In this case such a processor does not qualify as a microprocessor according to the above definition. Unlike a bit-slice processor, in a microprocessor there is no possibility of having access to the internal control signals nor of making changes, such as adding more instructions. The only means of interaction between a microprocessor and the outside world is the external bus.

1.2 Microprocessor Classification

To understand the various types of microprocessors that are available today it is worthwhile considering the historical progress that has taken place since 1971, when Intel introduced the 4004, a 4-bit machine intended for calculator applications [5]. This was designed essentially for serial binary coded decimal (BCD) operations. About the same time Intel was working on another 8-bit processor, the 8008, which was introduced in 1972. This had a simple architecture: byte handling facilities, a memory space limited to 16 kbytes and a small instruction set with a few operand addressing modes. The 8080 that was introduced in 1974 was still an 8-bit machine but with 64 kbyte memory space, increased instruction set, stack handling capability and more complete addressing modes, including limited handling of 16-bit words. In 1975 Zilog introduced the Z-80 microprocessor which was compatible with the 8080. Other semi-conductor companies produced similar microprocessors such as Motorola, National Semiconductor, Texas Instrument and others. In 1978 Intel produced the 16-bit microcomputer 8086, although Texas Instrument's 9900 preceded this as a 16-bit microprocessor. In 1979 Zilog introduced the Z8000 and in 1980 Motorola introduced the MC68000, both 16-bit microprocessors. It is known that semiconductor manufacturers are now working on 32-bit machines, some having already produced prototypes, which will possibly become widely available during the mid eighties.

We can differentiate essentially three types of microprocessors.

Single-chip Microcomputers

Single-chip microcomputers are complete microcomputers where processor, memory, and input–output, as well as auxiliary electronics such as clock oscillators, are all integrated on a single chip. They follow the tradition of the first 4-bit microprocessors. The silicon area that the advanced technology makes it available is used to integrate more functions such as memory, input–output, rather than increase the power and architectural complexity of the processor. Such a circuit computer contains, typically, a small data memory of 64–256 bytes and a program memory (ROM—read only memory) of about 2 kbytes. The program memory is fixed during the manufacturing process and therefore it needs to be defined during the masking stage. This is an expensive option available only for volume customers. Intel produces a relatively more expensive microcomputer model with EPROM memory for low volume users. The first such single-chip microcomputer was introduced by Intel in late 1976 (MCS-48) [6]. More recently Zilog produced a similar microcomputer, the Z-8, and Motorola the MC6081.

Single-chip microcomputers are meant for low cost, high volume embedded applications such as in home appliances, video games, cars, etc.

In a multi-processor environment, in spite of the low cost advantage (which is important when building machines with large numbers of such components), the restricted input–output facilities, memory size, and the computing power of single-chip microcomputer makes them unattractive for such applications. They are ideally suited to perform a well-defined function where requirements can be met with the limited processing power and memory available with such circuits. In this text such single-chip microcomputers will not be considered.

General Purpose Microprocessors

General purpose microprocessors are the 8-bit microprocessors with proven architecture which have attained today a certain level of maturity. They are low cost, yet have considerable computing power, intended for wide-ranging general purpose markets. Normally, a well-integrated family of circuits exist such as various memory types, serial and parallel input–output circuits, clock and DMA (direct memory access) circuits so that a designer has flexibility to assemble those circuits which are required by the specifications. Intel 8085, Zilog Z-80 and Motorola MC6800 are typical examples of such machines. Vast amounts of software have already accumulated for these machines.

High Performance Microprocessors

High performance microprocessors are 16-bit machines at the fore-front of technology where the aim is to achieve minicomputer performance. Moving from 8- to 16-bit processors represents a quantum jump for the manufacturers where some preferred to have a fresh start rather than be restricted in design options due to compatibility considerations. One common point about the 16-bit microprocessor is the ability to handle large memory space and the introduction of memory management techniques by an additional circuit (Fig. 1.3). For example, the Z8000 can directly access 8 Mbytes of address space and there are six such distinct address spaces allocated for distinct purposes such as instruction, data, stack, etc. In addition to more powerful instruction set and data manipulation capabilities, such microprocessors separate system and user modes, making it easier to implement multi-task applications. Clearly, the instruction set is also aimed at providing a suitable environment for high-level language implementation.

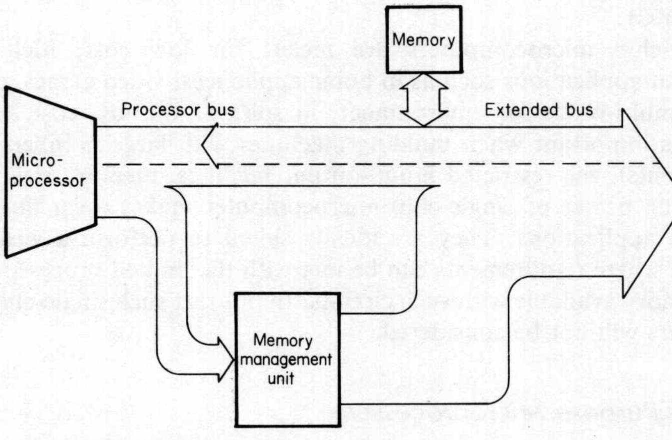

Fig. 1.3 High performance microprocessor.

Finally, these microprocessors include hardware features aimed at connecting a number of them in a multi-microprocessor configuration. This will be considered in more detail.

1.3 Microprocessor Architectures

Different forces have been at play in the design of the microprocessor, depending on the technological constraints, market demands, and the types of microprocessors, as explained above. For the single-chip microcomputer market the emphasis has been to incorporate more functions, hence a reduction of external circuits, in some cases even analogue functions, rather than improve computational power or increase memory. Hence the main objective has been to lower the overall hardware costs due to the use of a microcomputer, which results in the reduction of external electronics. In some applications this requirement may become crucial, say, due to weight or size limitations. The cost of software due to restricted memory, is not considered an important factor, more so since it is shared among many thousands to millions of units manufactured.

General purpose microprocessors have been designed to provide a decent processing capability, to satisfy a vast variety of applications, embedded or otherwise. Thus the silicon space available has been utilized to improve processor capabilities in terms of instruction set, addressing, and internal registers. In the mid seventies limited technology imposed the 8-bit word size, but experience with such microprocessors shows that they permitted the construction of amazingly wide-ranging embedded systems, as well as general purpose computers, albeit of restricted performance. Yet increasingly software has become an important factor, constituting from 50 to 90% of the cost of new designs [7].

It is not surprising, then, that high performance microprocessors are designed with a view to providing as many features as possible to support high-level languages and systems programs. In that sense they constitute a departure from the 8-bit architecture. The technological advances are exploited boldly to produce microprocessors compatible with the power and performance of current minicomputers, and to introduce new architectures more with a view of being compatible upward (by the addition of more facilities and perhaps compatibility with the future 32-bit machines) rather than maintaining compatibility with the current 8-bit machines. In the computer industry the introduction of a brand new machine line is costly. Thus the advent of the 16-bit machines indicates the confidence of the manufacturers in the future, and this belief that high-level languages will become a major software development tool for microcomputers.

For a systems designer, using a microprocessor presents certain opportunities since there is full access to the external bus of the processor. This is not true of a minicomputer, where access occurs via an input–output device. On the other hand, no internal processor signals are available to the systems designer. Therefore the understanding of the functioning of a microprocessor's external bus is essential. The external bus also provides a means of standardization where devices can be built to be interfaced to a given bus. Currently, each manufacturer maintains its own bus standard and provides compatible circuits. Memory, however, is a notable exception where de facto accepted standards exist, clearly to everyone's benefit.

In this section, without attempting to be comprehensive, brief outlines of