24.10.79

Report Series
No 15

Management Services and the Microprocessor

September 1979

The Butler Cox Foundation

Abstract

Report Series No 15

Management Services and the Microprocessor

by Edward Goldblum
September 1979

The microprocessor or "silicon chip" has rapidly become the most topical subject in computing. Everyone knows – or thinks he knows – that the micro will have a great impact on the way we live. But few people have considered the impact it will have on data processing and information systems, the activities to which it is most closely related.

This report first describes the historical background of the modern integrated circuit, showing that it has been economics rather than technology that has determined the pattern and the pace of developments. Both the production and the use of microprocessors are analysed in terms of their key economic factors.

The report then describes the new opportunities in data processing that are presented by the introduction of very cheap microprocessors. These opportunities do not lie along the traditional lines of major developments in computing. Instead, they follow and enlarge upon the trends which were begun by the minicomputer.

The advantages and disadvantages of these new opportunities are discussed in detail, including the impact on various sectors of the computing industries, the threat to the centralised data processing function, and the concept of total system costs. The key role of microprocessor software is highlighted.

The report concludes with a practical guide for the management services department which wishes to make a constructive start in microprocessors without committing itself to a major expenditure of resources.

The Butler Cox Foundation is a research group which examines major developments in its field — computers, telecommunications, and office automation — on behalf of subscribing members. It provides a set of 'eyes and ears' on the world for the systems departments of some of Europe's largest concerns.

The Foundation collects its information in Europe and the US, where it has offices through its associated company. It transmits its findings to members in three main ways:

- As regular *written reports,* giving detailed findings and substantiating evidence.
- Through *management conferences,* stressing the policy implications of the subjects studied for management services directors and their senior colleagues.
- Through *professional and technical seminars,* where the members' own specialist managers and technicians can meet with the Foundation research teams to review their findings in depth.

The Foundation is controlled by a Management Board upon which the members are represented. Its responsibilities include the selection of topics for research, and approval of the Foundation's annual report and accounts, showing how the subscribed research funds have been employed.

Report Series No.15

MANAGEMENT SERVICES AND THE MICROPROCESSOR

by Edward Goldblum

September 1979

Butler Cox & Partners Limited Morley House 26 Holborn Viaduct London EC1A 2BP

ACKNOWLEDGEMENT

The author sincerely thanks those members of the Butler Cox Foundation who participated in the microprocessor survey which preceded the writing of this report, and who gave their time freely in the subsequent interviews.

Thanks are also due to a number of faculty members at the Massachusetts Institute of Technology and to executives of several semiconductor manufacturers in the USA. These individuals were candid and particularly helpful in shaping our thoughts, even though there was no prospect of their receiving this report. Among the latter group were employees of Data General, Digital Equipment Corporation, Intel, QANTEL, Western Digital and Zilog.

Finally, a special note of gratitude to those users of microprocessors who were interviewed during the study. Their experiences, not always pleasant but always pleasantly described, conveyed a vivid impression of the state of the art in applying microprocessors to everyday problems.

E.J.G.
September 1979.

CONTENTS

CHAPTER 1

INTRODUCTION

Beyond any doubt, the microprocessor – more widely but less accurately known as the silicon chip – has been the most eye-catching technological innovation of the 1970s. To a credulous and bemused public, the microprocessor combines the impenetrable complexity of the computer with the sinister and inexorable encroachment of a tiny, unseen enemy.

Many sources predict that these chips will sweep all before them and irreversibly alter the way we live and work. Most of these forecasts are profoundly pessimistic and even apocalyptic. Very few of them are firmly grounded in a thorough understanding of computers or of the true capabilities of microprocessors.

PURPOSE AND READERSHIP OF THIS REPORT

The purpose of this report is to present a neutral and unsensational description of the microprocessor and of its likely use and impact in just one sphere of application: data processing and information systems. While this use of the microprocessor is arguably neither the most far-reaching nor the most innovative one, it is nevertheless extremely important to the management of most types of organisation.

Although many individuals today express very strong views about the microprocessor, few people are well-prepared by virtue of either knowledge or training to understand its implications. To make matters worse, the media and more than a few "experts" frequently embellish their descriptions of microprocessor technology with considerable "gee-whizzery" and outright humbug.

This report attempts to avoid such exaggeration. It is intended for readers who understand computing in a functional sense and who are prepared to give the microprocessor a fair hearing. This group will include members of the data processing professions, management services staff, and systems-oriented managers in many kinds of organisations.

DEFINITIONS

There is probably no technical subject which contains more imprecise or fluid definitions of its principal implements than does computing. Changes in technology and in its applications render existing definitions useless or simply wrong every few years. Nevertheless, it is necessary to attempt to standardise the terminology which is used in this report. Therefore, while not expecting to achieve general agreement, we propose the following definitions:

1. The microprocessor (or "micro" as it is often called) is a complete central processing and arithmetic/logic unit (CPU) contained on a single large-scale integrated circuit. It may or may not include a certain amount of program memory. The internal word size of the micro may be 4, 8 or 16 bits. Micros are generally used only as components within other devices. They are, therefore, sold to end users by original equipment manufacturers (OEMs) and distributors, not by the manufacturers of the microprocessors themselves. A typical price of a popular 8-bit microprocessor chip in 1979 is about $10. An older 4-bit micro may cost less than $1.

2. A microcomputer consists of a microprocessor together with further features included on or interconnected with the CPU, to form a complete computer. These extra features typically include random-access or read-only memory, clock circuits, input/output drivers, analogue/digital converters, peripheral device interface logic and the like. The physical space required for all this circuitry may range from a single integrated circuit (the "computer on a chip" which has now become available) to an entire circuit board of roughly A4 paper size containing perhaps 50 chips. Like the microprocessor, microcomputers are used chiefly as components in larger devices (such as small computer systems) and so they also tend to be marketed to end users by OEMs rather than by manufacturers. A typical microcomputer on a board with 16K bytes of memory costs about $400.

3. Minicomputers span a very wide and expanding range of capability. At the low end, they are distinguishable from microcomputers only by physical size. We shall define the minicomputer to mean a computer with a word size of 12 to 32 bits, capable of operating as a stand-alone machine, and contained in a single box or cabinet. A memory size of 64K bytes is typical of today's minis, but is extravagant by the standards of a few years ago. These machines are used for a wide variety of purposes, ranging from dedicated applications, such as front-end communications processing, to stand-alone small-to-medium data processing systems. An average price for a minicomputer with 64K bytes of memory but no peripheral devices is about $15,000.

4. The mainframe computer is familiar to most computer professionals, even though its definition is as nebulous as the others. It may best be thought of as including everything left over at the high end of the preceding definitions. Most mainframes have an internal word length of 32 bits or more, and are used most often as medium-to-large general-purpose computers. They differ very significantly from all other computers in that they are marketed exclusively by their manufacturers directly to end users, and are almost always supplied with programming languages and comprehensive operating software systems. A typical memory size might be 512K bytes, and such a machine without its various peripherals (a most unlikely configuration) might cost anything from $100,000 to $2,000,000.

The foregoing definitions are rather arbitrary. This is true not only because the classifications depend upon usage as much as on intrinsic properties, but because there is considerable overlap at the boundaries of each of the three types of computer. There has correspondingly been some "leakage" of traditional customers of each type of computer into the adjacent types. More rather than less of this kind of blurring of definitions will occur in future. Much of it has been and will continue to be exacerbated by developments in microprocessor technology.

WHAT THIS REPORT IS NOT ABOUT

Some readers may be surprised or disappointed that this report addresses itself only to the use of microprocessors as components of information systems. This restriction is imposed for two reasons. First, the wider topic of microprocessors in general is so broad in its implications that any attempt to cover it in a report of this length would be uselessly superficial. Second, developments within data processing itself are likely to be directly relevant to management services managers and other "technically literate" persons.

This restriction does not mean that there are no other applications worthy of investigation. On the contrary, there are hundreds of current and potential applications of micros, some of which will affect society much more than others. A few of these will significantly affect industry and commerce and, therefore, deserve to be mentioned briefly here before we turn to the specific subject of this report.

1. Microprocessors in manufactured products

   If a "microprocessor revolution" occurs, this is where the first battles will be fought.

Existing industries that manufacture consumer products using traditional nut-and-bolt electro-mechanical components are today re-designing and re-tooling to incorporate microprocessors. Some products already on the market are cookers and washing machines that are crudely programmable, cars with micro-controlled ignition and electrical functions, and simple educational devices including some which can synthesise human speech.

Some obvious candidates for the near future are telephones, domestic heating and lighting, high-fidelity music systems (a particularly promising application), and leisure products, such as portable language translators and advanced computer games. This last category represents a class of completely new industries whose entire existence is due to the availability of cheap microprocessors. The personal calculator and the digital wristwatch industries are the forerunners of these. More will arise as microprocessors find new domestic applications (figure 1, below, illustrates the market penetration achieved by various domestic products).

Figure 1 Electrical products in the US market

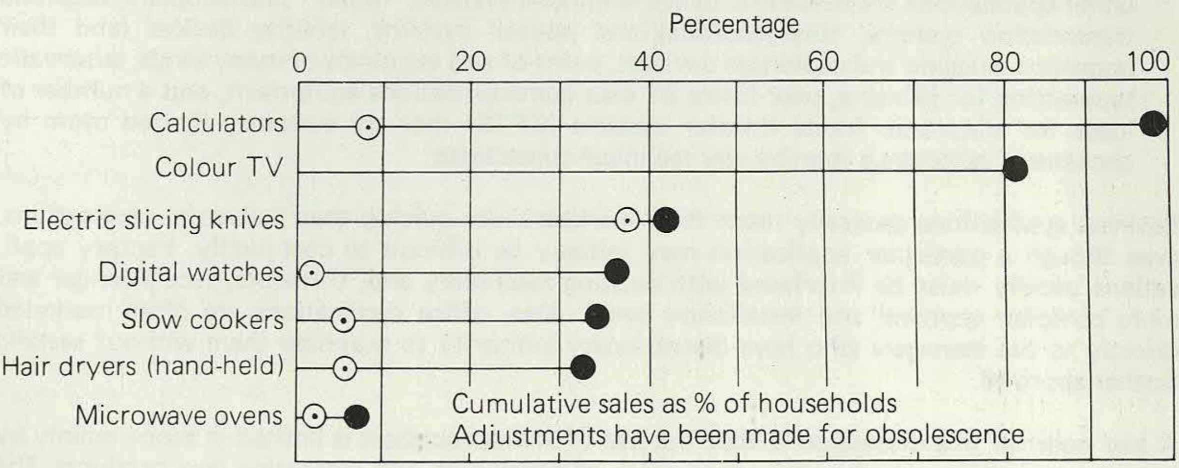

Each line shows the cumulative penetration of a product as a percentage of all American households in 1972, shown as ⊙, and in 1977, shown as ●.

The pattern for calculators shows how quickly an enormous market can be achieved by a cheap product that has universal appeal.

(Source: *Merchandising*)

2. **Microprocessors in industry**
The potential contribution to manufacturing industry is huge, and is concentrated in two key aspects of manufacturing in which British industry is traditionally weak: quality control and plant economy.

Many production processes can be vastly improved by carefully introducing microprocessors to assist in measuring, monitoring, and controlling those processes. Typical applications include controlling temperature and pressure, monitoring power consumption, fatigue testing, operating valves for fluid or gas flows, and optimising production schedules.

Away from the production line itself but still within the factory, applications will include automating warehousing and stock control, counting and sorting finished products, testing and quality control, and integrating factory systems with central computer systems.

In order to achieve these goals, a new generation of microprocessor-controlled precision instruments and components is rapidly appearing in industrial markets. These components are rendering obsolete the crude electro-mechanical equivalents which served industry for generations. Analogue-to-digital converters are experiencing something of a renaissance. The motto of engineers in this field seems to be, "If you can measure it, you can control it with a microprocessor".

3. Microprocessors in commerce

   There are many potential uses of microprocessors in offices, administrative centres, and non-industrial commerce. Some applications have been present in an under-developed form for years, and are only now achieving general acceptance. The word processor is the best example of this type of application.

   Other applications are automatic office telephone systems, "smart" photocopiers, facsimile transmission systems, time-recording and payroll systems, security devices (and their opposites, bugging and espionage devices), point-of-sale terminals of many kinds, automatic typesetting for printing, new forms of data communications equipment, and a number of ideas for electronic funds transfer systems (EFTS) that are currently limited more by consumers' reluctance than by any technical constraints.

Business applications generally reach their markets more quickly than industrial applications, even though a particular application may initially be difficult to cost-justify. Factory applications usually must be interfaced with existing machinery and, therefore, face a longer and more complex approval and installation cycle. Also, office applications are often marketed directly to the managers who have discretionary authority to purchase them without seeking further approval.

A key point in this discussion is that the use of microprocessors is limited in scope mainly by human imagination and the sheer hard work of developing and marketing new products. The technology and tools for the job have for the most part been available for some years. Thus, it is not technology *per se* which brings new applications into the market (although it does tend to force the pace); rather it is innovation, entrepreneurship and economics.

THE WIDER CONCERN

The remaining important topic that this report does not address is the wider impact of microprocessors on society as a whole. It is difficult to avoid this subject entirely because most of the public discussion surrounding the micro has been centred on it. Interestingly and unusually, this discussion has occurred *before* most of the applications have materialised. The predicted effects have mainly related to work habits, unemployment patterns, and the nation's general prosperity, most of which have been forecast to be drastically and unfavourably affected.

Despite this publicity, it is too early to predict with any confidence the degree of general change that microprocessor technology may bring about. Only a tiny amount of credible data about the micro's impact is available from any source, and even less from within western Europe. The little data that has emerged from the USA so far does not, however, support the pessimistic forecasts made for the United Kingdom.

In any event, there is nothing to be gained in condemning or praising out of hand a device which in most cases will be used as a tool to implement new ideas. It is preferable to judge each implementation on its own merits. No doubt some uses will be badly conceived and possibly harmful, but the tool itself is neutral.

A good analogy to the role of the microprocessor is that of the electric motor. Most people live out their entire lives without ever buying an electric motor as such. Nevertheless, motors are found in countless places around the home and office, performing commonplace tasks in which they are taken completely for granted. (It is not unusual to count ten or more just in a typical modern kitchen.) No one would think of banning the electric motor because in the form of an electric drill it sometimes assists the safe-cracking criminal as well as the do-it-yourself enthusiast.

This objective view should be applied to microprocessors as well. The micro itself, like the electric motor, has no real existence of its own outside the device of which it is a component. The desirability or otherwise of the microprocessor in our lives and society is not questioned in this report. But the morality, and even the legality, of every *use* to which the micro is put should always be open to question.

CHAPTER 2

A GUIDE TO THE WORLD OF SILICON

The microprocessor is not really new. The device itself has been available for eight years, the technology that makes it possible has been in production for nearly twenty years, and the underlying physical principles were discovered over thirty years ago.

To understand the direction which developments are taking today, it is first useful to understand the basic elements of the technology that underpins the microprocessor. This chapter presents for the interested layman an outline of the history and the technology surrounding this remarkable marriage between solid-state physics and digital electronics.

THE ORIGINS OF A TECHNOLOGY

Quite unlike its glamorous and pioneering successor of today, the electronics industry in 1947 was firmly grounded in a technology which had not changed appreciably in twenty years.

The war gave the electronics industry a sudden huge investment in the form of soaring demand for military communications and radar equipment, and television was poised for a mass public launching. Nevertheless, electronic equipment remained expensive and unreliable, principally because of the physical nature of the circuit components that designers had to incorporate into their devices.

These components included the "passive" circuit elements (the resistor, capacitor and inductor), and the "active" elements based on the vacuum-tube valve. Only active elements had the capability of gain or amplification, i.e. of transforming a low-power signal into a high-power one by drawing on an external power source. But the valve was expensive to produce, fragile in its glass shell, and extremely hungry for electric power to keep its filament glowing and to maintain the high voltages at which it functioned.

In 1947, John Bardeen, Walter Brattain and William Shockley of the Bell Telephone Laboratories in New Jersey made a fundamentally important invention: the bipolar or junction transistor based on semiconductors.

A semiconductor is a piece of pure crystalline material (usually silicon, obtained from sand) into which tiny amounts of other materials have been added. Typically, these materials are phosphorus and boron, and are known as impurities or dopants. They alter the electrical properties of the silicon in a special way.

Pure crystalline silicon is a good electrical conductor, but when "doped" with such impurities, the atoms of the substance exhibit either a deficiency or an excess of electrons when a small voltage is applied. The deficiency, also called a "hole", results from a *p*-type (for *p*ositive) impurity; an excess results from an *n*-type (for *n*egative) impurity. Placing a *p*-type and an *n*-type crystal together results in a device that will conduct electrical current only in one direction. This is the simplest semiconductor device of all: a diode.

A diode is not capable of amplification, and hence it cannot replace the vacuum-tube valve. However, a much more useful device, the *transistor,* may be made by introducing a third element and forming a "sandwich" of *p*- and *n*-type materials, as either *pnp* or *npn*. This configuration permits the electrical current flow between two of the elements to be very effectively regulated by a small voltage applied to the third, and thus achieves gain. It is called bipolar because materials of both polarities (*p* and *n*) are involved in its operation.

The transistor was clearly superior to the vacuum-tube valve in almost every way. It was smaller, faster, operated at lower voltages, consumed less power, was much more rugged, lasted longer, and was cheaper to produce. However, for the three reasons described below, the electronics industry did not immediately embrace the new device and abandon its less satisfactory predecessor.

First, the existing manufacturers of valves had no experience of producing such a new device. Virtually all of their existing production lines were quite useless for manufacturing a product which was based on a completely different technology.

Second, the transistor could not simply be substituted one-for-one for the valve. Many circuits had first to be re-designed in order to incorporate transistors into existing electronic products. But circuit designers were unfamiliar with the electrical operating characteristics of the transistor and preferred valves, which they already knew well.

Third, the manufacturers and distributors of electronic products were somewhat reluctant to risk introducing completely new product lines which might compete with their existing products. Introducing solid-state electronics (as transistor technology was called), which was so obviously superior, would quickly make obsolete the older products based on valves. It might well undermine the manufacturers' revenue base because each new product would be cheaper than the one it replaced.

Thus, the electronics market, which would have benefited in many ways from the rapid introduction of the transistor in electronic products, was thwarted for a time by the internal economics of the manufacturers. It was not until the impetus of the American military and space programmes of the late 1950s that this attitude changed and solid-state devices became commonplace.

The consequences of this delay are interesting to note in retrospect:

1. A few of the valve manufacturers managed to make the transition to producing semiconductors, but most of them eventually got out of the business altogether.

2. A completely new industry was formed around semiconductor production.

3. The Japanese were afforded a decade's head start in producing transistorised consumer products, a lead which they exploited skilfully and largely still enjoy today.

4. As a result of this lead, Western electronics manufacturers lost much of their domestic markets to Japanese imports.

5. The market eventually got the products it wanted, though rather late and at the expense of the domestic manufacturers.

There is a considerable similarity between this experience and the factors influencing the market for microelectronics today.

By the 1950s, the semiconductor industry had become highly competitive, and engineering work was concentrating on reducing the cost of manufacturing transistors in order to compete more keenly. The process of fabricating individual transistors gave way in the mid-1950s to a new method of batch processing many small transistors on a single thin slice from a large silicon crystal. The electrical properties of the transistor did not depend on its physical size.

However, each transistor was still only a single electronic component. Even though several hundred transistors could be fabricated on one silicon slice or wafer, they still had then to be physically separated, assembled individually with tiny wires, and placed into protective housings to take their place with resistors and capacitors in electronic circuits. These steps were very labour-intensive, and prevented the prices of complex electronic equipment from falling very far.

Progress took a giant step forward in 1959 with the development of the integrated circuit (or IC) at Fairchild Semiconductor. This device fabricated, separated, and interconnected the transistors and all the other circuit elements *electrically* instead of physically.

Integral resistors were constructed in the silicon wafer by utilising the body resistance of the semiconductor itself, while the inherent capacitance of the junctions between the *p* and *n* regions of the same material provided small but adequate capacitors. Thus, all three major circuit elements — resistors, capacitors, and transistors — could be realised electrically *in the same material* without external interconnections to one another. (Inductors have never been produced successfully on integrated circuits.) This meant that *an entire electronic circuit* — not merely one element — could be produced at one time, including all the interconnections amongst the elements.

On this occasion, industry did not ignore the opportunity presented by a new electronic device. There were two strong market forces compelling the early use of integrated circuits: the accelerating American strategic arms and space programmes, which placed a high premium on miniaturising components, and the newly-emerging digital computer industry.

The application of integrated circuits to digital computers was an ideal one. Digital logic required very large numbers of active circuit elements compared with devices that employed analogue amplification, such as radios. Computers could make use of large quantities of identical components, because the basic circuits were replicated many times in the architecture of a computer. And, finally, the computer's use of transistors as binary switches, which had only two operating states, overcame the practical difficulty of producing integrated circuits with small tolerances and precise electrical characteristics.

The computer industry and microelectronics have advanced hand in hand since the mid-1960s. The achievements of each have led to greatly increased demand for the other. The cost reductions in computers were due in large part to the better and better techniques for producing integrated circuits. These techniques in turn were encouraged because of the growing market for ICs provided by the increasing number of ever-cheaper computers being sold. Figure 2, on the opposite page, shows how the cost of making and using an integrated circuit is built up.

Integrated circuits continued to grow in complexity (but not in size) as more and more functions were squeezed onto a single chip by photo-reduction techniques. But until 1971, all ICs were merely used as components within more complex devices. In that year, the Intel Corporation produced the first microprocessor for use in a new, small desk calculator to be made by a Japanese company.

Although it was possible to build the calculator from custom ICs, designing several complex custom circuits would be expensive, and the resulting ICs would serve only the one purpose for which they had been designed. Being complex digital devices, they also might very well contain errors in their design that might not be detected until many ICs had already been made.

Figure 2 The cost structure of an integrated circuit

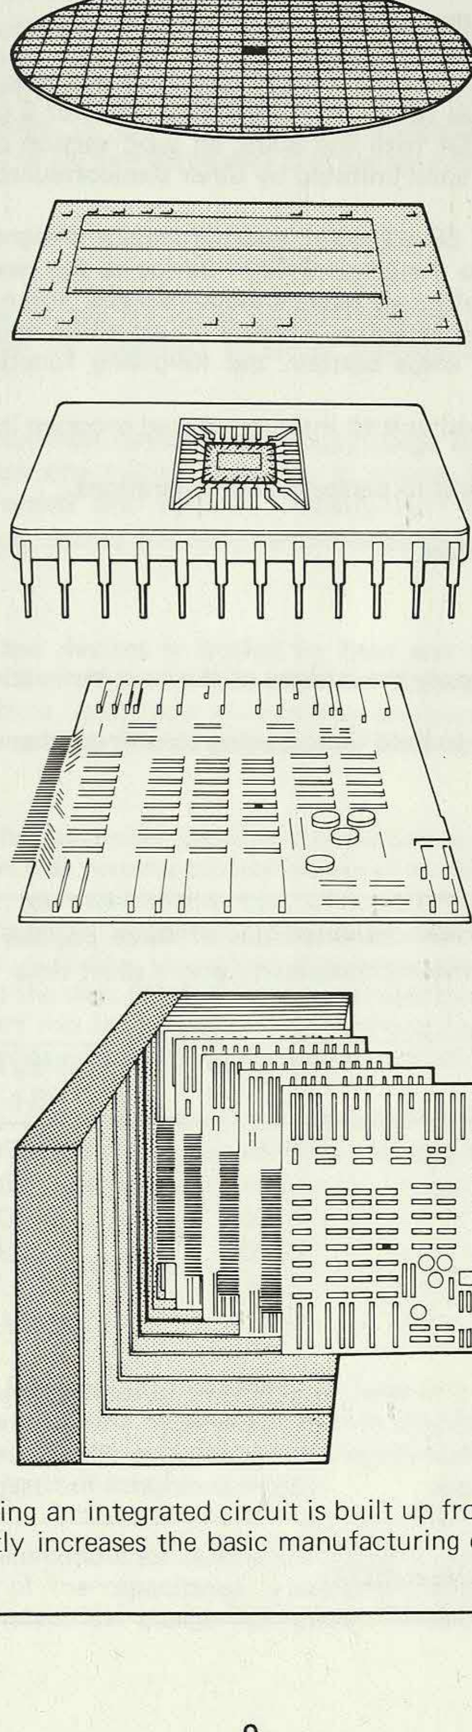

	Average cost $	Cumulative cost $
CHIP IN WAFER, UNTESTED	0.10	0.10
TESTING AND YIELD PER GOOD CHIP	1.00	1.10
PACKAGE, PACKAGING AND TESTING	0.50	1.60
SPACE ON PRINTED-CIRCUIT BOARD	1.00	2.60
SHARE OF BACK PANEL, WIRING, CABINET AND POWER SUPPLY	0.35	2.95

The cost of making and using an integrated circuit is built up from the costs of its constituent components, and this greatly increases the basic manufacturing cost which is about 10 cents.

Intel's novel idea — which seems obvious today — was to make the circuits simpler but more general, and to allow them to be *programmed* in an easy way. Such an approach would circumvent the custom-design errors, and at the same time would permit the circuits to be employed elsewhere in a different role. This idea resulted in a tiny general-purpose programmable processor, the Intel 4004, which could become a special-purpose device through specific programming rather than through custom manufacturing.

The 4004 had an internal word size of only 4 bits, which was adequate for representing decimal digits in a calculator but was hardly sufficient as a basis for a general-purpose computer. In 1972, Intel followed the 4004 with the 8008, an 8-bit version of the microprocessor, which became widely used and was soon imitated by other semiconductor manufacturers.

There are currently at least 50 different microprocessor designs on the market. They range from 4-bit models similar to the original 4004 device to the new 16-bit micros which incorporate the instruction sets of much larger minicomputers. Some of the characteristics of old and new micros are tabulated in figure 3 below. These chips range in price from $1 to about $200. Most microprocessor chips contain the following functional units on a single chip:

- A decoding and control unit to interpret stored program instructions.
- An arithmetic/logic unit to perform basic operations.
- General-purpose registers.
- An accumulator.
- Address buffers to supply the address of the next instruction.
- Input/output buffers to hold data flowing to or from the microprocessor.

Figure 3 A comparison of some characteristics of three popular microprocessor chips, showing the trend toward increasing complexity over a short time

	Intel 8080 (8-bit)	Zilog Z80 (8-bit)	Zilog Z8000 (16-bit)
Year introduced	1974	1976	1978
Power consumption (W)	1.2	1.0	1.5
Number of transistors	4,800	8,200	17,500
Number of gates	1,600	2,733	5,833
Chip size (mm^2)	22.3	27.1	39.3
Density (gates/mm^2)	72	101	148
Number of distinct instructions in set	34	52	81
Combination of number of distinct instructions, data types and addressing modes	65	128	414

(Source: *Electronics*, 21 December 1978)

Usually the clock circuits and the main memory are supplied on separate chips, but newer models of some microprocessors contain both this circuitry and the central processor functions on a single chip, thus comprising a complete programmable computer only one quarter-inch square.

Developments in microelectronics continue today with even more impressive results than ever before. The technologies involved have not yet approached the fundamental limits of size or speed imposed by the laws of physics. However, any attempt either to catalogue all the devices available on the market today, or to predict with confidence what the market will contain in two years' time, is destined to be overtaken very quickly by events.

CLASSES OF MEMORY

All general-purpose digital computers make use of *memory* for storing programs and data. These digital memories can be classified as either *moving-surface devices* or *all-electronic devices.*

In the former category, common devices used today range from the simplest magnetic tape cassette (containing perhaps one million data bits in a serial format at a cost of about 10^{-5} cents/bit) to large block-access disc systems (holding 10^{10} bits for about 10^{-4} cents/bit). Figure 4, on the next page, shows the cost/performance characteristics of various types of memory.

The speed of magnetic tape devices is limited by heat and frictional wear because of the physical contact between the magnetic tape and the read/write heads. The speed of discs is not so limited, in principle, but there are complex engineering problems arising from the need for extremely accurate head positioning and for uniform magnetic flux properties over a large disc surface.

Despite these problems, the computer industry's requirement for fast random-access storage has enabled the capacity of disc systems to double approximately every 30 months for the last 12 years. The disc, in fact, has outlived a number of competitors that appeared to offer better price/performance characteristics at the time they were introduced. Many of these devices were launched with great publicity by well-known computer manufacturers. Every one of them has failed to displace the disc. Most of them have been quietly dropped or ignominiously interred after a fairly short run in the market. Still others, based on lasers and optical techniques, are in the pipeline today awaiting commercial exploitation.

Moving-surface devices do not, however, provide a satisfactory medium for a computer's program storage. This is supplied by all-electronic random-access memories (RAMs) that contain no moving parts and have access times comparable to the internal speed of the computer itself.

RAM, ROM, PROM and EAROM

Before 1970, most computer memories consisted of large arrays of ferrite cores. These were tiny rings of magnetic material one millimetre or so in diameter, strung by the hundreds of thousands on grids of wires. Ferrite cores have been largely superseded by new semiconductor memories, thanks to the advent of integrated circuits.

In its simplest form, the semiconductor read/write memory cell consists of one transistor and one capacitor. The value of the capacitance is extremely small (about 10^{-14} farad) but is adequate to store a small electrical charge indicating a binary 1, while the absence of this

Figure 4 Comparison of memory costs and access times for various memory technologies

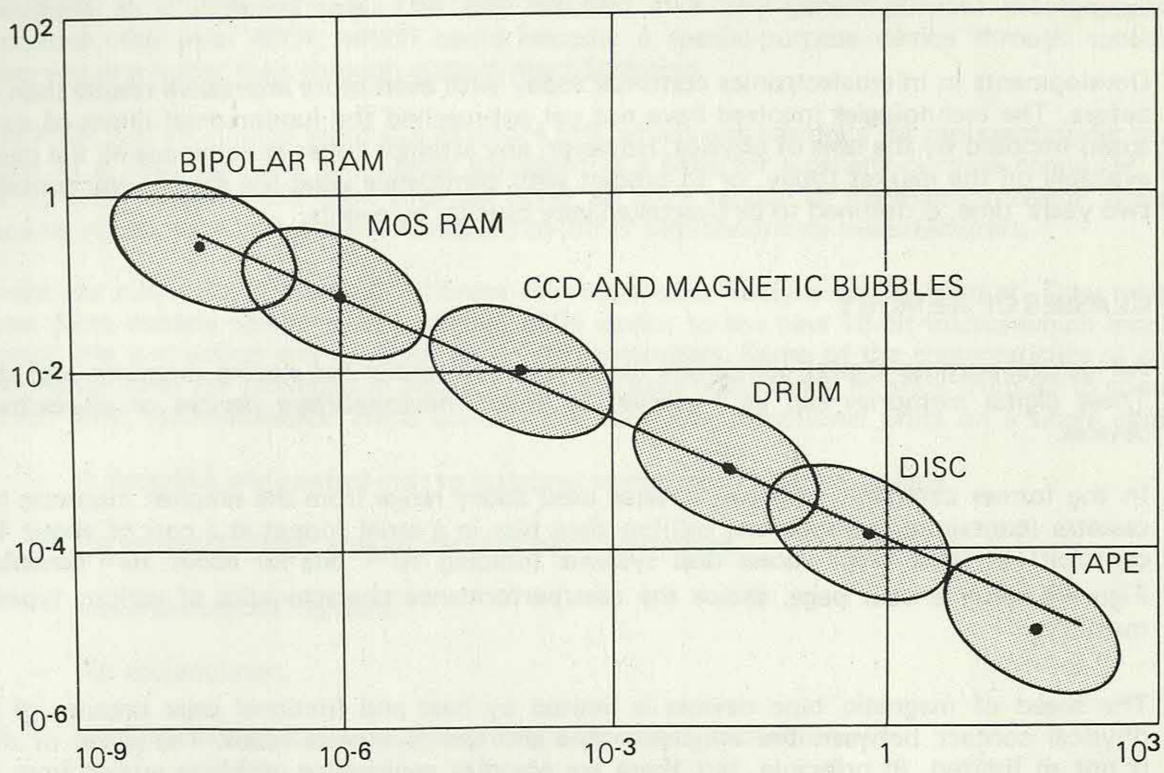

The values indicated were current in 1977, but the cost of all types of memories continues to fall.

(Source: *Scientific American*)

charge represents a binary 0. The transistor is used as a switch to connect the storage capacitor to a data line when the cell is selected for reading or writing.

However, this type of storage cell loses its stored information each time it is read, and also loses it by leakage of the capacitor's charge. Leakage can take place in as little as a few milliseconds. This kind of memory cell, known as a "dynamic RAM", therefore requires its stored charge to be refreshed about once every two milliseconds, as well as after every read operation. Other designs, called "static RAMs", do not require refreshing, but they incur the penalty of requiring additional transistors, which correspondingly take up more chip area and result in a higher cost per bit.

Today, single-chip random-access memories of 16K bits are commonplace, and 64K bit devices are available. Some unforeseen problems caused by radiation-induced memory failures may slow the development of even denser RAMs.

Some applications require random-access memories containing permanently-stored information, such as control program instructions or constant data values. A practical example is a pocket calculator's control program, which is never changed. This type of storage is provided by the read-only memory (or ROM). In its simplest form, the storage capacitor of the RAM may be replaced either by an open circuit or by a direct connection to earth, representing a binary 0 or 1 respectively. The desired data pattern is fabricated on the chip itself, and incurs a high initial production cost.

A popular alternative is to manufacture a "blank" memory which contains tiny fusible wires in all bit positions. By applying a suitably high voltage in an appropriate pattern, the undesired links may be "blown", leaving an arrangement of intact wires representing the required data bit pattern. This type of memory is also known as a programmable read-only memory (or PROM). Obviously, the program is a permanent one, because the fusible links cannot be reformed. It has the compensating advantage that its information is always retained when external power is removed, i.e. it is non-volatile. This is not generally true of semiconductor RAMs.

Still other applications require a memory that can be altered, but for which read operations are much more frequent than write operations. A solution provided by microelectronics is the so-called "read-mostly" memory, of which there are two principal variants. In the commonest form, a store of charges on pairs of electrodes may be built up to form the desired pattern of information. This charge will remain reliably in place for years. To alter the pattern, however, the contents of the entire memory must be erased by exposing the chip to ultraviolet radiation, which allows the charges to leak away rapidly. After this erasure, a new pattern of information may be imposed. These devices are called UV-erasable PROMs.

The alternative device is called the electrically-alterable read-only memory (or EAROM), which can be altered selectively without the need to erase the entire array of data. This device employs special materials during its fabrication to build small circuits that may be selectively charged and discharged. Each cell will hold a charge until it is erased by a strong pulse of current, after which a new charge may be electrically imposed on the cell while leaving all other cells unchanged. This type of device has not yet been developed to an acceptable standard for widespread commercial use, but it is clearly a potentially useful component.

BUBBLES AND TUNNELS

All the devices described above are random-access memories, in which a desired item of information can be addressed directly by its location in an array of other items. A totally different kind of device is the *magnetic bubble memory.* It has so little in common with semiconductor technology that logically it does not really belong in this discussion. However, it promises to become an important component of microcomputer systems, and, therefore, merits a closer look.

The magnetic bubble memory is a *serial* device which cannot be randomly accessed. It exploits a rather obscure physical phenomenon, the local variations created in uniform magnetic fields in thin films of certain magnetic materials, such as garnet. Both the materials and the physical principles employed are thus very different from those of transistor-based semiconductor ICs. However, the device can be fabricated on chips of a sort, and it is, therefore, "microelectronic" in the broad sense of the word.

Data in a bubble memory may be thought of as circulating round and round a closed loop within the device. The data passes a single "window" at regular intervals, at which times it may be read or written. It is necessary to wait for the desired item to pass by the window before accessing it. The device is thus inherently serial in operation, and so is very slow compared with semiconductor memory, having an access time of tens of milliseconds instead of a fraction of a microsecond. However, it has the great advantage that it is non-volatile, requires no external power to maintain its data pattern, and is completely electronic.

The most attractive potential application of the bubble memory is to replace small disc and tape memories with a capacity of up to 10^7 bits or so. When produced in quantity, it should have a price advantage over small discs of comparable capacity. It should also have similar performance characteristics, as well as being much smaller and more reliable. Already several semiconductor manufacturers have announced million-bit bubble memories for the commercial market at prices around $2,000.

Magnetic bubble devices are competing with *charge-coupled devices* (or CCDs), which operate on similar principles but which use capacitor cells that are refreshed during the cycling process. Unlike bubbles, the CCD's contents are volatile.

The type of microelectronic memory most recently developed is the so-called tunnel junction or Josephson memory cell. Once again, a different technology has been exploited, in this case the superconductivity of materials operating at extremely low temperatures. This phenomenon enables circuits to perform switching operations 10 to 100 times faster than ordinary semiconductor circuits, while simultaneously consuming 100 to 1000 times less power.

The very low power consumption means that such circuits can be packed extremely densely without creating problems of excessive heat dissipation in a small chip area. Interconnecting lines between elements may be kept shorter in denser circuits, thus reducing signal-propagation times. These devices potentially offer memories of very small physical size operating at speeds 100 times faster than today's fastest memories. Such a development could have a major impact on computer system design in the future.

However, the one outstanding (and to date the conclusive) drawback of the tunnel junction is that it must operate at temperatures close to absolute zero (-273°C), which are usually achieved by immersion in liquid helium. Mechanical stresses resulting from extreme temperature cycling must, therefore, be overcome. New packaging techniques will have to be developed to achieve the desired high packing densities. These requirements suggest that it will be some time yet before this device becomes commercially viable, although several major semiconductor manufacturers are known to be developing it.

MAKING A MICRO

A description of basic microelectronic technology must include at least a brief review of the process of manufacturing integrated circuits, partly because it is exotic and interesting, but also because it is an important factor in the economics of producing (and thus of using) the microprocessor.

The structure of an integrated circuit is complex. Each device has an intricate three-dimensional architecture which must be reproduced precisely in each circuit. The two-dimensional layout of an IC may be glimpsed from a photographic enlargement of a completed chip, but the arrangement of the several vertical layers is not visible.

In layman's terms, producing an integrated circuit requires the following steps, which are also illustrated in figure 5 on the opposite page:

1. Specification and design

   The designers who have conceived the new circuit specify its functional characteristics and select the processing steps that will be required to produce it. The physical size and location of each element in the circuit are then estimated. Computer-aided design facilities are heavily relied upon during this stage. A computer may also be used to simulate the operation of the circuit that is being designed. Corrections and improvements may be made as a result of these simulations.

Figure 5 Simplified schematic of the manufacturing of integrated circuits

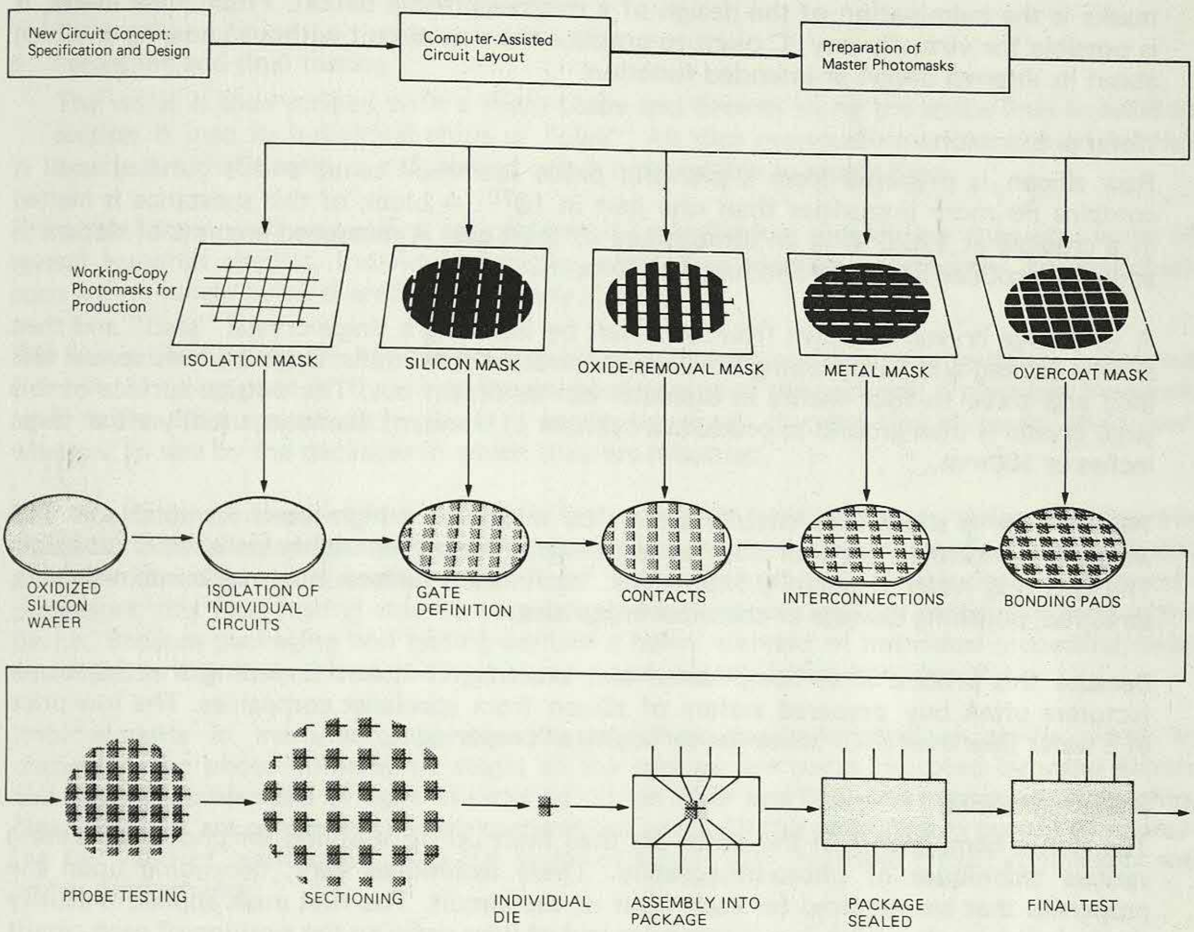

A number of photomasks are required in this complex process, but most of the manufacturing costs are incurred after the wafer has been sectioned into individual dice, when handling costs begin to grow.

(Source: *Scientific American*)

The final layout of all circuit elements is then performed. The goal of this layout is to achieve the desired function of the device within the smallest possible space. During this design phase, engineers are working with "human-sized" models of a circuit which will eventually be about 500 times smaller.

The time required to complete the design and layout makes it the longest step in the fabrication process. For a new microprocessor (the most complex form of IC), this process may take two years or more. For a semiconductor memory with a few basic repetitive circuit patterns it may require only a few months.

2. Production of photomasks

From the detailed circuit layout, a set of photomasks is prepared. Each of these is a glass

plate about five inches square containing the pattern for a single layer of the circuit. Because many individual circuits will be fabricated together, the photomask contains this single pattern repeated many times over its surface. A complete and correct set of masks is the culmination of the design of a microelectronic circuit. From these masks, it is possible for virtually any IC plant to produce the new circuit without knowing anything about its internal design or intended function.

3. Wafer preparation

   Raw silicon is prepared from a plentiful oxide (common sand) and is purified until it contains no more impurities than one part in 10^{10}. A block of this substance is melted in a crucible at 1,420°C in an atmosphere of inert gas. A measured amount of dopant is added to produce a specific conductivity (*p* or *n*).

   A very large crystal is grown from this melt by inserting a single-crystal "seed" and then slowly turning and withdrawing it. Under controlled conditions, single crystals several feet long and three to four inches in diameter can be drawn out. The outside surface of this large crystal is then ground to produce a cylinder of standard diameter, usually either three inches or 100mm.

   The cylinder is sliced into circular wafers by means of a high-speed diamond saw. The wafers are smoothed on both sides by grinding, and are then highly polished on one side. The resulting wafer is about 0.5mm thick. Its finished surface must not contain defects, scratches, polishing damage or chemical impurities.

   Because this process is so complicated and exacting, it is not surprising that IC manufacturers often buy prepared wafers of silicon from specialist companies. The low price of a wafer (less than $10) belies the difficulty of preparing it.

4. Wafer fabrication

   The circuit components of the wafer are then built up by applying the photomasks using various techniques of photo-lithography. These techniques vary, depending upon the properties that are required for each layer of the circuit. The first mask applied is usually the isolation mask, which is a rectangular grid of lines defining the position of each circuit within the wafer.

   One of the reasons why silicon plays such a dominant role in microelectronics is that its main oxide, SiO_2, has very valuable properties in circuit manufacturing. If a wafer of silicon, a good electrical conductor, is heated in an oxygen atmosphere, a layer of silicon dioxide soon forms on its surface. This layer is an excellent insulator and, therefore, permits another conducting layer to be deposited on top of it, provided that contact "windows" have been placed at desired points in the oxide layer to allow the conductor to make electrical contact with the silicon substrate below. A number of layers can be built up in this fashion, the complexity depending upon the design of the device and the materials chosen for its fabrication.

   The result is a single wafer with many identical tiny integrated circuits fabricated upon it. Typically, several hundred wafers are batch-processed at one time through this phase, each wafer containing about 250 individual circuits.

5. Probe testing and selection

   While still part of the wafer, each circuit is rapidly tested by an automatic device which uses small pointed probes to touch each of the IC's contact pads and perform a set of electrical measurements. These tests are fairly crude, and are intended merely to establish whether the fabrication process has permitted any basic errors to occur that render the IC worthless. Each circuit found to be defective is marked with a small dot of ink. The machine then steps to the next circuit on the wafer and repeats the cycle. This process is performed without human intervention, usually under computer control. The computer keeps extensive statistics on the relative incidence of failures, their locations on the wafer, etc.

6. Packaging and final testing

   The wafer is then scribed with a sharp blade and broken along the scribe lines in order to section it into its individual chips or "dice". All dice previously marked with an ink dot by the automatic tester are discarded – it is not possible to repair them.

From this step on, all operations are performed on individual dice, rather than on a wafer of several hundred circuits. Individual handling costs then increase enormously, because these costs are no longer being shared among many circuits.

The unmarked circuits are bonded into standard packages, and fine wire leads are connected from the bonding pads of the die to the connecting pins of the package. A plastic or a ceramic cover is moulded around each die. It is somewhat ironic that the tiny chips are often overwhelmed in size by the packages in which they are mounted.

The fabrication process is then complete, but each unit must undergo a series of exhaustive electrical tests to ensure that it actually performs its required functions perfectly. More faulty units are identified and discarded at this point. For circuits of great complexity, such as microprocessors, this final testing may be very lengthy and yet may still not reveal all defects in the device. Because packaging and testing contain a heavy element of individual processing, these costs tend to dominate the overall cost of manufacturing an integrated circuit.

Improvements in manufacturing microelectronic devices still continue to be made. In particular, the photo-lithographic stages of the process are being improved by using shorter and shorter wavelengths of light in order to obtain finer and finer line resolution on the chip. New materials are constantly being experimented with. All this activity is in pursuit of improving the yield of marketable chips per wafer of input, with increased circuit complexity and higher performance.

SUMMARY

Today's integrated circuits are the lineal descendants of the transistor, which was invented in 1947. The microprocessor itself is merely a particularly complex kind of integrated circuit with the capability of being programmed. This capability means that it can take the place of many special-purpose circuits in a wide variety of applications, by being individually programmed for each application.

Microprocessors on the market today range from the original 4-bit design to the newer 16-bit designs, and they come in a variety of architectures. Some have extensive instruction sets, a memory, and on-chip circuitry for interconnecting with external devices.

Electronic memories have kept pace with other microelectronic devices in improving capacity and speed. The most common type today is the semiconductor RAM, which has completely displaced the magnetic ferrite core memory. Other types of memory based on exotic technologies offer non-volatile storage of data and extremely fast operation. These are only now beginning to reach the commercial market.

Fabrication processes are constantly being improved, with two goals always uppermost: increased complexity (integration) on a single chip, and lower unit costs.

CHAPTER 3

THE ECONOMICS OF THE MICROPROCESSOR

In Chapter 2 we observed that the impetus for using new electronic devices as they became available has been principally an *economic* one. New technology has been accepted only when it appeared to offer either a significant reduction in cost or an improvement in performance for the same cost. Conversely, new devices have been perfectly capable of being ignored or avoided if the economics they represented appeared to be unfavourable.

FIRST PRINCIPLES

In thirty years, the only part of the innovative process that has changed is the state of the technology itself. In every other respect, organisations react today very much as they have in the past to the introduction of a new technological product. Their analysis is still based on the perceived economics of use. Their conclusions still reflect a judgment of whether their own micro-economics appears to benefit from an involvement with the new product.

Stated in this way, the idea that economics determines how and whether a new product is used appears to be almost a truism. Nevertheless, the point is well worth keeping firmly in mind when considering the future of the microprocessor. There are two reasons for this:

1. There is a temptation today to believe that all the products that technology provides must inevitably be incorporated into our economy and culture. This is untrue. Those new products that are finally adopted are the visible tip of a much larger iceberg representing a huge investment in research and development of new products, many of which either fail utterly or are never introduced.

In fact, the market is very conservative in accepting new entrants. History suggests that industrial experience favours an *existing* technology, and creates a cost barrier to the introduction of any new technology. For example, an enormous effort has been channelled into innovation in computer memories over the past twenty years or so. Yet only one established type of memory has been displaced by a new one (the ferrite core by the semiconductor RAM). For very large memories, the disc still reigns supreme, and is not seriously challenged even today.

2. A related idea is that anything that is technologically possible will inevitably be developed and made available. This too is false. Products that would be extremely useful and economical for certain applications may never be developed at all, or may even be kept back from the market intentionally. A prosaic example was the very late introduction of the stainless-steel razor blade by certain companies that reckoned that they were better off if shavers replaced their worn blades every day or two.

The same forces operate in the computer industry. They are economic forces, not technological ones. In every case, the economics of two groups (usually manufacturers and consumers) do not coincide. When this happens, the existing market tends to be perpetuated and the *status quo* prevails, if only temporarily.

This chapter describes the basic economic factors that control the production and use of the microprocessor from factory to end user.

A MOST REMARKABLE INDUSTRY

The microelectronics industry can best be understood in terms of the classical industrial learning curve. Most industries reduce the unit costs of their products by 20 to 30 per cent in constant currency each time their cumulative output doubles. Integrated circuit costs have declined by about 28 per cent with each doubling of the industry's cumulative output, a result which is in agreement with this theory. However, *the industry has doubled its output nearly every year since 1965.* This means that the cost of a given electronic component has *decreased by a factor of 100* in real terms in fourteen years.

Another way of viewing this astounding record is to observe that an individual IC produced today incorporates more electronic elements than the most complex piece of electronic equipment that could be built in 1950. A 1979 microcomputer, costing perhaps $200, has more computing capacity than the first large all-electronic computer, ENIAC, which was developed in 1945. The microprocessor is about twenty times faster, has a larger memory, occupies 1/30,000th the volume and costs 1/10,000th as much. ENIAC had 18,000 vacuum-tube valves, which consumed many kilowatts of power, and some of which almost constantly required replacing. By contrast, the microcomputer is almost totally reliable, requires no maintenance, and consumes only the power of a small light bulb.

The earliest integrated circuits contained only a few dozen circuit elements. In 1964, Gordon Moore of Fairchild Semiconductor predicted that circuit complexity would continue to double every year. Today, there are circuits containing 2^{20} circuit elements in experimental production, and there has not yet been any significant departure from "Moore's Law". Some deviation from an exponential growth curve is, however, eventually inevitable.

In view of this progress, it is significant to observe that there have been no fundamental breakthroughs in producing integrated circuits since they were first developed. The processes used then are recognisably the same as those used today. Instead, progress has been achieved in three ways:

1. By developing increasingly complex circuits, thus lowering the cost per function.

2. By improving production methods and controls in the factory, thus increasing the yield of good ICs per wafer of input to the process. This is achieved mainly through scrupulous plant cleanliness and by controlling processing temperatures to within 0.5°C.

3. By reducing the size of the basic circuit elements on each IC, thus enabling more complex circuits to be crowded into a given chip area. This has been achieved by improving the photolithographic processes that reduce the original circuit designs down to a quarter-inch square on the photomasks.

There are important manufacturing trade-offs between chip size, circuit complexity, and yield, which influence the types of integrated circuits that can be made economically.

1. The size of a chip can be made an arbitrary fraction of the total wafer size. In practice, the incidence of random manufacturing defects is proportional to the surface area involved, and all wafers result in some defects. Therefore, if a chip were made to occupy the entire surface of the wafer, every one would contain defects, and the yield would be zero. In general, the smaller the chip size, the less the probability of its having a manufacturing defect.

2. Complexity of circuit design means that, for a given size of chip, more components must

be squeezed onto the same area of silicon. This requirement leads to very expensive research by the semiconductor companies into new methods of fabrication. The costs of designing a chip increase rapidly with circuit complexity, sometimes even increasing as the square of the number of components on the chip. It is also generally true that a more complex device is more specialised in its function than a simpler one, and is, therefore, applicable to a smaller variety of applications. Consequently, the manufacturer may not be able to sell very many of them. Complex circuits are more difficult to test, and may require larger and more costly packaging that involves more external connections.

3. The yield of an IC is the percentage of the theoretical input number that is ultimately usable. The manufacturer's profit on a particular IC is almost directly proportional to this figure. Below certain levels of yield, the benefits of mass production will be outweighed by the high fixed costs associated with research and development, factory operation, and depreciation of expensive, short-lived capital equipment. Actual yield figures are jealously guarded by manufacturers, but a yield of 15 to 20 per cent is considered good today. A comparable figure of ten years ago was 5 per cent. Clearly, there is still scope for greatly improving the yield of the manufacturing processes.

MAKING A PROFIT ON MICROPROCESSORS

The semiconductor manufacturer's profit is made almost entirely by producing and selling integrated circuits and boards (circuit boards containing a number of interconnected ICs). Not all the ICs produced by a given manufacturer are of his own proprietary design. There is a complex network of cross-licensing agreements in the industry which permits other manufacturers to produce circuits from an original design. At the other end of the scale, there is a good deal of outright stealing of other companies' designs, and something of a grey area in between.

Competition in this industry is ferocious. Many of the key employees of the principal manufacturers know one another from previous employment, especially in the "Silicon Valley" area south of San Francisco. As a result, information and ideas tend to spread rapidly (and often unofficially) through the industry, thereby forcing the pace of competition even further.

The development of a microprocessor by a manufacturer is a highly-speculative venture that carries no assurance of commercial success. A microprocessor is, of course, just a particularly complex integrated circuit, but it has little other similarity to an integrated circuit.

The manufacturer designs a new microprocessor to satisfy a market which does not yet exist but which he hopes will be present at the right time. There is no way of guaranteeing this market. The design costs are very high indeed, and the process may require two years or more to develop before pilot production can commence. Initial sales are very small because the market as yet has very few applications for the new device. No one in the field has any experience of programming it. The price at this stage is accordingly very high because the design and tooling costs have not yet been amortised by mass production.

After a considerable time, perhaps eighteen months, the microprocessor will have been designed into new OEM products, which themselves are about to enter mass production. At that time, the manufacturer can begin to produce and sell large quantities of his "new" microprocessor, the design of which may by then be three to four years old. He will have streamlined the manufacturing process and improved its yield so that he can meet this increased demand very quickly. The usual terms of sale are cash-with-order.

If the manufacturer has been exceptionally fortunate, the market will have chosen his micro as one of the three or four it favours most. This is by no means purely a measure of the device's inherent power and speed. Strong irrational forces also influence the market's choice.

Prices then begin to tumble. Discounts for quantity purchases are substantial. If the micro has proved to be popular, it is likely that other manufacturers will have copied the design and are now striving to capture a piece of the market for themselves. Speed and volume of production must be sustained in order to compete on price. Given sufficient volume, the cost of producing *any* microprocessor begins to approach a constant lower limit.

At some point during this part of the cycle the large fixed costs of designing and developing the micro will have been fully recovered, and only the incremental costs of production will remain. Provided that demand for the device is sustained, and that prices have not been forced down too far by competition, the manufacturer then stands to make a profit on his microprocessor. Nearly all the profit is made on this "down side" of the product cycle.

A common way of increasing unit profit on microprocessor sales is to market complete boards, rather than individual chips. The margin on the associated support circuitry and additional memory chips is often much higher than that of the microprocessor chip itself. Most manufacturers are now pushing strongly to sell their products in this way, having realised that they cannot become rich selling microprocessor chips for $1. Semiconductor memory chips, in particular, represent a potential gold mine for the manufacturers. It is no accident that new microprocessor designs can address a very much larger memory space than could the early devices — manufacturers are shaping a large market for their more lucrative memory chips. The estimated growth of this industry is shown in figure 6 on the next page.

The microprocessor industry is obviously not well suited to the faint-hearted or risk-averse entrepreneur. Its fast-paced technology and intense competitiveness make it a casebook study in new-venture enterprise. However, there are some factors which make the industry a less hostile and less insecure place than the description above suggests.

First, microprocessors still do not account for a major part of the revenue of most semiconductor manufacturers. These companies depend instead on bread-and-butter contracts for producing other kinds of ICs at a known price over a known period of time. This constant production assures steady income and helps the companies to amortise their expensive production facilities. Figure 7, on page 23, shows the leading manufacturers of microprocessors in 1978.

Second, many of the newer semiconductor companies are backed by the financial resources of very large corporations. This support makes it possible for a firm to survive either a lean period or the disastrous launching of a costly new product that fails.

Third, not all firms have to sell what they produce. IBM is reputed to have the largest production of microelectronics in the world, but it produces only for its own needs. Similarly, Texas Instruments, an innovator in semiconductor technology, is thought to be its own best customer for the microprocessors it manufactures.

Nevertheless, the economics of success in this industry compels each manufacturer constantly to juggle the three variables of size, complexity and yield. Not every company will be able to strike a profitable balance.

THE SHADOWY WORLD OF THE OEM

The category of "original equipment manufacturer" includes all those companies which incorporate microprocessors into products which are then sold to other customers. It is a broad category which includes the makers of digital wristwatches as well as manufacturers of small computer systems.

The world of OEMs appears shadowy and vague from the outside because it contains such a

Figure 6 The estimated revenue growth of the microprocessor industry

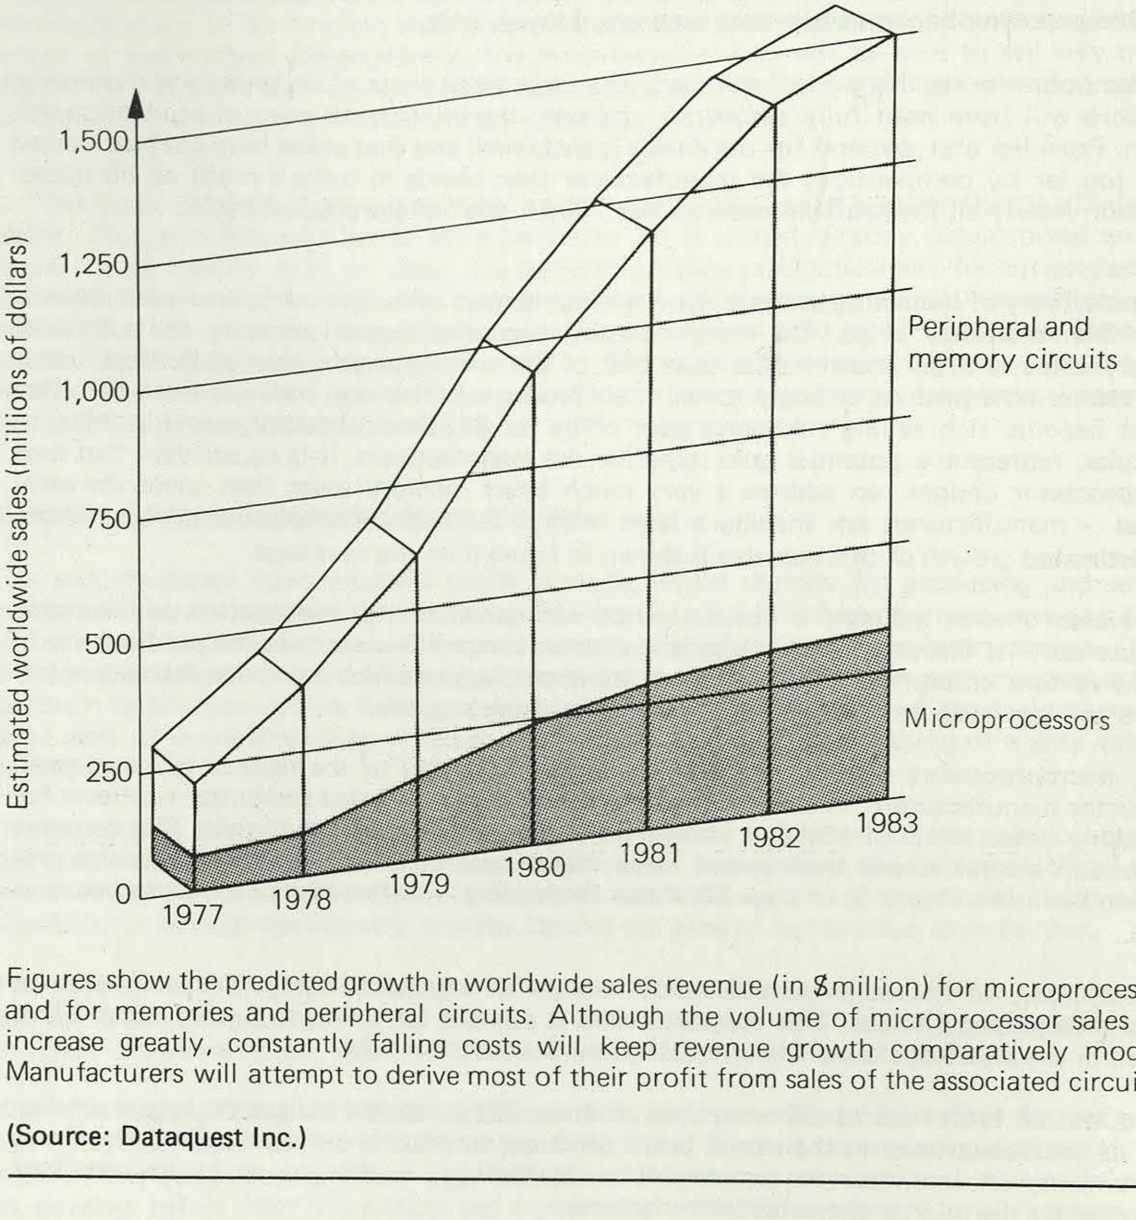

Figures show the predicted growth in worldwide sales revenue (in $million) for microprocessors and for memories and peripheral circuits. Although the volume of microprocessor sales will increase greatly, constantly falling costs will keep revenue growth comparatively modest. Manufacturers will attempt to derive most of their profit from sales of the associated circuitry.

(Source: Dataquest Inc.)

heterogeneous collection of inhabitants. They do not speak with a single voice, and do not constitute a single industry as such. Yet they play a crucial economic role in the use of the microprocessor.

An OEM makes a profit by selling a product made from a number of bought-in components, together with software which the OEM usually adds. This software may range from the trivial (the control program of a four-function calculator) to the complex (an operating system and compiler for a small computer system). The cost of the microprocessor chips in these two devices may be about the same.

To an OEM, the launching of a product based on a microprocessor represents a major invest-

Figure 7 The leading manufacturers of microprocessors

Figures show the estimated shipments of microprocessors by size for the year 1978.

Company	1978 shipments (thousand units)
4-bit micros	
Texas Instruments	9,400
National Semiconductor	2,535
Rockwell International	2,475
Nippon Electric	1,600
Total	16,010
8-bit micros	
Intel	1,661
Fairchild	951
Motorola	950
National Semiconductor	750
Synertek	710
Mostek	670
Rockwell International	660
Zilog	540
Advanced Micro Devices	445
General Instrument	425
Total	7,762
16-bit micros	
Texas Instruments	182
National Semiconductor	91
General Instrument	60
Intel	31
Total	364

The ratio of shipments of 4-bit:8-bit:16-bit micros was 44:21:1.
Total units shipped were 24,136,000.

(Source: Dataquest Inc.)

ment. He, therefore, chooses with care the microprocessor he will use. This does not automatically mean that he will choose the newest, fastest, cheapest or most powerful micro. Obviously, the micro must have the capability which is required for the application, but the OEM's final selection will depend upon:

— Assured supplies, which usually require a second source of the chip.

- The general popularity of the micro, which should imply a long product life and continued availability.
- Its suitability for a particular application (a PDP-11 is not required to control a wristwatch).
- The product designer's familiarity with a particular micro.
- The range of support chips available.
- An existing relationship with a micro manufacturer.
- A wholly subjective "feel" for a particular micro.

The price of the chip itself is not usually a determining factor, since most microprocessors of comparable power are very competitively priced.

OEMs collectively represent virtually the entire primary market for microprocessors. Nearly every micro is sold in the first instance to an OEM, who adds value (and cost) by designing it into some product, which is then sold to an end user of that product. (Some products, such as VDUs, may go through another intermediate stage before reaching their ultimate market.)

It is, therefore, the OEM to whom the manufacturer of microprocessors listens when he is making decisions affecting new product designs. The OEM is virtually his only customer. The manufacturer has very little contact with the end users of his micros, and he may well not even know what products on the market contain micros that he has produced.

The most important consequence of this triangular relationship is that end users of microprocessors have very little direct influence on the design of new micros. Their reactions must be channelled through the OEMs who deal directly with the manufacturers. However, this indirect link is not a dependable one, because the respective interests of the OEMs and of the end users do not necessarily coincide.

This lack of direct communication between manufacturer and user becomes even more significant when the OEM's product is a micro-based small computer system. Considerations of software quality and programmability should be high among the manufacturer's priorities. But unless the feedback from his OEMs is particularly good, the chip manufacturer will remain largely uninformed about the needs of end users, and will continue to design his microprocessors as he thinks best.

From the discussion above, the following conclusions may be drawn:

1. OEMs strongly influence the supply of microprocessors by providing the primary market for them. This market determines the quantity of chips sold, but does not necessarily reflect the eventual uses to which the micros are put.

2. Chip manufacturers, who are interested in mass production but little else, will in effect favour a high-volume application (such as digital wristwatches) over a more sophisticated but lower-volume application (such as small computer systems). The range of micros available in the marketplace will reflect this preference.

3. The role of the OEM as the provider of most of the software means that the manufacturers continue to be ignorant of the software requirements of end users. A particular problem is that micros tend to be designed by electronics engineers, rather than by experts in software or systems.

Because demand for microprocessor-based products has continually exceeded the supply, hundreds of new OEMs have come into existence over a short time. With hardware costs falling, the capital required to start a business as an OEM is now quite small. Thus OEMs range from very large and reputable equipment manufacturers to two-man companies working out of garages and spare rooms.

Accordingly, there is a wide range of quality and price in the OEM market for small computers. These machines are very difficult to compare and evaluate, particularly for the first-time user who has no computer experience to guide him. Some of the small computers currently on offer can only be described as useless, as the reaction of some of their users confirms. While many small computers may be based on the same microprocessor, some include idiosyncratic software, poor or non-existent documentation, and unreliable peripherals, such as cheap audio cassette recorders that are used to store data. Although the hobbyist or student may well be delighted with such a system, the businessman would do better to look elsewhere for a reliable stand-alone system on which he can process, say, his company's accounts and VAT receipts.

SUMMARY

Unlike many products, each new microprocessor is a tool looking for an application. Manufacturers speculatively launch micros into a very competitive market, which is unpredictably fickle in its taste for new products. For the few micros that are widely adopted (not for altogether rational reasons) there is the prospect of large profits for the manufacturers, provided that they have struck an economic balance among the key factors of chip size, complexity and yield. Sales of microprocessors do not form the basis of most manufacturers' financial success.

OEMs collectively form the primary market for microprocessors. They comprise a diverse assortment of companies, and their number has grown swiftly and randomly over the last few years. This growth has resulted in an undesirably wide spread of quality in the products they offer to end users. Many of these companies are likely to disappear rather quickly as competition intensifies.

Because they are remote from end users, the microprocessor manufacturers have not fully understood the needs of those users, particularly with regard to software. The designs of current microprocessors reflect this ignorance. More recently, the manufacturers have invested heavily in operating software for their newer micros in an attempt to integrate upwards and become "systems suppliers". Their history does not equip them very well for such a role, but it is too early to say whether the move will be successful.

CHAPTER 4

NEW OPPORTUNITIES, NEW PROBLEMS

Although microprocessors were not developed with a view to their being used in large information systems, nevertheless this has begun to happen. The new economics of computing power brought about by cheap microprocessors has spurred system designers to utilise these devices in novel ways. Some of these applications will fall by the wayside as being uneconomical or functionally undesirable, but others may alter the appearance of the computing industry very considerably in future.

This chapter presents a summary of the most promising potential developments of the microprocessor in data processing, and some of the problems which will confront those who attempt to exploit that potential.

THE TWO PRINCIPAL OPPORTUNITIES

The goals and information needs of organisations have not changed at all as a result of the microprocessor. *What has changed is the way organisations can apply data processing to their needs.* This change offers two major new opportunities in computing:

1. Small, low-volume computer applications are now economically feasible. The cost of providing computer hardware for small applications has fallen so low that, in many cases, a formal cost-justification may be unnecessary. This means that a very large number of completely new computer applications will become candidates for serious consideration.

Using a dedicated computer to attack a small, one-off application is now perfectly acceptable, and it may be a more economical approach than any other. Until recently, this idea was not worth a second thought.

2. Applications no longer need to be machine-efficient. A preoccupation with hardware efficiency has haunted the computer industry since its earliest days because of the presumed need to utilise efficiently the scarcest resource, the processor. Now that processor costs can be treated as a minor component of total costs, the structure of a computer application can be made to fit the problem rather than to fit the computer. Under-utilisation of the computer can be ignored or taken for granted. The human resource will in future be treated as the most valuable resource.

As a result of these changes, computer applications will become almost entirely concerned with *economics* and *functionality of use.* This re-orientation will require a profound shift in the outlook of those who develop computer systems. The emphasis will change rapidly from producing crude computer analogues of user problems to helping users cope with and solve their problems *in situ* with the assistance of computers.

Both of these trends in data processing began with the introduction of the minicomputer. They are now accelerating because of the much cheaper processing power of the microprocessor. There is an important caveat, however. *Cheap processors do not guarantee cheap systems.* To achieve an overall low-cost system means controlling many other variables as well. This topic is discussed later in this chapter.

SIMPLICITY AND MODULARITY

Today's highly-complex mainframe computers, with their huge, cumbersome operating systems, were developed for two reasons:

1. To achieve certain economies of usage by "efficiently" sharing the most expensive resource — the central processor — amongst a number of competing applications.

2. To lock the user into a particular computer manufacturer by making it very difficult for the user to change systems.

The growth of these operating systems has greatly altered the application programming environment. Most computers in the early 1960s ran single-stream batch jobs, supported by straightforward operating system facilities which were reasonably familiar to most programmers. By the late 1970s, most mainframes were multi-programming several jobs concurrently, supporting complex schedulers, partitioning very large memories to fit more jobs in, swapping or paging out the jobs that would not fit, and offering virtual storage for still more. When several layers of telecommunications software were superimposed on this structure, a nightmare sometimes resulted.

During this evolution, simplicity was lost. Nothing about a modern mainframe is simple or straightforward. There is nothing either to encourage designers to develop simple systems or to make it easy for them to do so. Every facility seems biased towards complex patterns of use.

The thrust towards integrated applications has added yet another dimension to the existing complexity. Often, several large, monolithic suites of programs struggle to combine their results, while running jointly under the supervision of a large, flawed operating system which no one fully understands. The facility for allocating the computer's resources has itself become a major consumer of those resources.

In such an environment, it is not surprising that many application programs are large, unstructured and complex, because that approach represents the path of least resistance for the designer. Nor is it surprising that such systems take a very long time to be developed and are frequently unreliable in use.

A *modular approach to system design* has long been advocated by many authorities in order to reduce overall system complexity and shorten the time required to develop new systems. Some good work has followed these lines, but, in general, the benefits have not been great. Structured programming was one such spin-off activity.

Design teams have often encountered practical difficulties in applying a modular approach using only a single computer. The microprocessor will enable system designers to incorporate true modularity into new systems. If a problem can be logically decomposed into a number of quasi-independent, simpler sub-systems, then each one can be developed separately with a microprocessor at its centre.

An overlay of intercommunication will then become necessary in order to coordinate the several sub-systems. One way in which this can be achieved is by using an additional microprocessor as a "taskmaster", and by permitting the various processors to communicate with one another only by means of queueing messages at this central facility.

There are three main implications of adopting such an approach:

1. This approach permits the system designer to consider the original problem in a very different way. This view may be closer to the "pure" nature of many problems, and less a matter of mapping every problem onto a single computer of fixed design and capacity.

Each solution would be allowed to grow "organically" to accommodate a particular problem. In a multi-microprocessor environment, any mix of computers and peripherals would be acceptable. The important factors to control would be the *interfaces* between processors and the *common database*. Clear and simple standards for communication and data definition would be required.

2. When the complete system was operating, it would be considerably more robust than a conventional centralised system running on a single computer. A hardware fault or a program error in one sub-system could not cause unpredictable repercussions in another if the interfaces were properly designed. Keeping the processing elements loosely coupled in this way would allow graceful degradation of the system and prevent complete shutdowns.

3. The ability to use very small task teams to work exclusively and simultaneously on individual sub-systems could cut system development time for the project as a whole. The teams could be physically separated because they would not be working on the same computer. Significantly, the teams would not be competing with one another for the same development resources. Less-qualified staff could be employed if the individual modules were kept simple. Externally-purchased packages of hardware and software could be utilised freely, provided that the interface and communications conventions were observed.

Unfortunately, there is no generally-accepted methodology available today for approaching system design in a modular way. It is conceptually and practically difficult to control and link a number of concurrent processes in separate computers, particularly if the processes need to exchange partial results. There is little experience to date with systems that employ true parallel processing. Techniques for accomplishing this kind of computing and its associated data communications need to be developed, proven, and disseminated before general progress can be made.

It may also be found that the high-level programming languages that are now established as standard for most commercial applications are unsuitable for developing modular systems. Today's analysts and programmers have learned a very restrictive centralised approach to computerisation, which is reflected to some extent by the programming languages that are used today. This tradition will be difficult to change.

STAND-ALONE SYSTEMS

The dramatic drop in processor costs has exerted strong downward pressure on the prices of computer peripherals (printers, VDUs, keyboards, tape cassettes, discs, etc.). These prices in turn have fallen considerably, although not to the same extent as the prices of processors. Part of this reduction in prices has been due to the incorporation of microprocessors into the peripherals themselves to replace custom electronics.

These reduced prices mean that a small but complete independent computer system with peripherals can be bought today for about $8,000, or can be configured from individual components for less. This price is continually falling. Such a system does not have much computing power or storage capacity, but it *is* a usable general-purpose computer. For a number of small applications, it is a perfectly adequate vehicle for productive data processing.

Known collectively as "small business systems", these computers are already well represented in the market, where many different brands are rapidly becoming available. Most are assembled by OEMs from standard components, and are packaged, fitted with some operating and/or application software, and sold at a good profit.

The principal target market is the vast number of businesses that are too small to have been able to afford their own computers in the past. However, these systems are equally well suited to end users of DP in large organisations.

The presence of such systems on the market means that users may now expect (reasonably or otherwise) to have their own dedicated computer systems that function independently of the corporate data processing activity. This notion may be thought in some quarters to be rather alarming, but it is technically straightforward and economically appealing, at least to the user. Such systems might offer the following advantages:

1. They would place the data processing tools in close proximity, both to the problem being addressed and to the people who work with and understand the problem. In some cases, the users themselves rather than DP specialists would control and operate the system.

2. The sharing of scarce computer resources amongst users would be unnecessary. The scheduling of jobs would be a less onerous task.

3. The systems would not be vulnerable to crashes of the main computer, nor could they cause it to crash.

4. A proprietary sense of "our system" amongst the users would encourage constructive use and general efficiency, and generate interest in further applications.

5. Users would quickly learn some home truths about the problems and capabilities of computer systems.

Under the label of "distributed processing", this direction of development has been gathering pace for several years. Once again, the minicomputer was the vanguard for these applications. The main benefits which this approach offers are not economic ones, but organisational and managerial.

There are, however, several obvious dangers and disadvantages:

1. Naive users may not cope well with the rigidities of computer systems. The scope for confusion and error is very large.

2. Systems would need to be made considerably more idiot-proof than most are today. If non-specialists were to use the systems, it would mean that a much wider range of possible errors would have to be defended against.

3. The organisation's central database, if accessed by these remote systems, would need to be stringently protected against misuse and corruption. It would also have to be made readily and simply available to remote computers.

Perhaps the greatest potential danger is not that the DP department might become superfluous, but that the organisation might find itself having to support a large menagerie of systems. Each computer might be supplied by a different vendor, might introduce incompatible hardware and unmaintainable software, and might not conform to the organisation's hard-won DP standards. Many users still do not understand this danger.

With cheap hardware flooding the market, the initiative for installing new systems may pass to the end user by default. This concept is a revolutionary one for centralised DP departments. Salesmen of small systems are already bypassing the central computer function altogether and attempting to sell packaged systems with hardware directly to the end user. Even if the user does not purchase such a system outright, he may wonder why his own DP department requires a lengthy development cycle and a large budget to develop the same application. It often puzzles a potential user why an application is apparently available off the shelf *with its own computer* elsewhere, for a fraction of the cost that is required to develop it internally.

SOFTWARE

In many ways, software is the key to the future use of microprocessors in data processing. In dedicated control applications, the software for a micro is likely to be specialised, trivial and static. In DP applications, in which the micro will be utilised as a general-purpose computer, it must have flexible and convenient software if it is to be used effectively. To some extent, these uses conflict with each other.

Software development in the early years of computing followed a familiar pattern. When first launched, most computers had very little operating software, usually a patchy operating system, and a simple language translator. As more of the new machines were installed, both the manufacturers and the users settled down to develop the large amount of operating software that the market required: assemblers, compilers, editors, trace/dump facilities, debugging aids, input/output handlers, and much more. The effort and expense involved in producing this software was immense, and was spread over years of the computer's service life.

Every time a new architecture was introduced, this cycle was repeated. The classic example of software obsolescence was IBM's introduction of the System/360 in 1964. The new computer put back the clock of operating system development for several years, even for the users of IBM's own previous computer lines.

Since this experience with the 360, most manufacturers have broadened and extended their ranges, rather than making obsolete the existing software base of an earlier computer line. The lessons of the 1960s have made the market much less keen to adopt completely new computer architectures than it might otherwise be, for the following reasons:

- Users' investment in applications software is so enormous that they cannot readily contemplate installing an incompatible processor.

- Few users wish to participate in the long, painful gestation of a brand-new operating system.

To evade these potential problems, many users have been content to settle on a "standard" mainframe. The price they pay is to continue to use an architecture (that of the 360 family) which is now more than fifteen years old, and which reflects the DP problems and approaches of a generation ago.

The introduction of minicomputers in the mid-1960s started another series of software development cycles. However, minicomputers differed from their mainframe counterparts in several important ways:

1. The mini manufacturers did not wish to enter the software business and had little interest in developing elaborate software tools for their computers. To them, software and support meant higher unit costs at a time when the trend was strongly towards cheaper hardware.

2. Initially, no one had any experience of writing programs for these new machines. Productivity of programming was low. Nearly all programs were coded in assembly languages.

3. The machines were cramped in both word size and memory space. This made program development awkward and caused much ingenuity to be expended in levering quarts into pint pots.

4. The computers were cheap in relation to the cost of developing the software they needed. This factor tended to discourage investment in high-quality software.

5. Minicomputers were used at first principally for dedicated industrial applications, which

were often quite remote from the mainstream DP activities. Frequently, they were programmed by engineers rather than by experienced computer programmers. These individuals tended to view the computer strictly as a means to a limited end, and did not develop or insist upon the kinds of software development tools which were commonplace in DP departments.

For these reasons, software for minicomputers remained rather primitive for a number of years. Eventually, however, it became necessary to develop more elaborate system software in order to satisfy competitive pressures in the market, larger minicomputers, and the use of minis as small mainframe computers.

It is central to an understanding of the microprocessor to realise that *all the factors described above apply to microprocessor software today.* There is also the further complication that the time required to develop a new microprocessor is much shorter than that needed to develop a new conventional computer. This has meant that the semiconductor houses, in their unceasing competitiveness, continue to introduce new computer architectures *every few years – a shorter time than is required to develop a family of high-quality software tools.* The implications are that:

1. At any given time, the newest microprocessors will have little development software available to support them. They will, therefore, be difficult to use effectively, even though they are more powerful than their predecessors.

2. The manufacturers are preoccupied with developing their next generation of micros, rather than with squeezing additional value out of the current lines.

3. No manufacturer wants to invest large sums in developing high-quality software for a device that will be obsolescent in a short time.

The recent experience of many users bears out the truth of these implications. Much of the available operating software, particularly that of the "personal computers", is badly written and poorly documented. No universal specification is applicable to BASIC, which is the *de facto* standard language for micros. As a result, each implementation of BASIC is slightly different from all the others, even on the same micro. Some operating systems are idiosyncratic, and their diagnostics may be far from informative. Disc-formatting conventions are not universally observed.

Most importantly, *software support scarcely exists.* The OEM suppliers of these small systems simply cannot afford to maintain expensive customer support organisations if they are to sell the systems in volume at a competitively low price. The manufacturers of the micros have no traditional experience of either writing or maintaining systems software, and would prefer to concentrate on what they do best: making cheap silicon products.

The environment for application software is, if anything, even worse. The prospective user of a business application package is confronted by a number of candidates. These may range from a completely bespoke implementation by a software house, including training and system support, for perhaps $50,000, to an application program on a cassette purchased in a neighbourhood electronics shop for $15, no questions asked.

The popularity of the "personal computer" based on the microprocessor has resulted in a growing supply of very cheap software packages, many of which are of dubious quality. Home enthusiasts, students, hobbyists, and moonlighting programmers are all contributing their home-grown programs to the many which already exist. Most of these programs are intended to entertain and interest other enthusiasts, but some are aimed specifically at the small-business market, where the microprocessor is having a major impact. Most users in this market are poorly qualified to make an informed selection of software.

The growth of this "cottage-industry software" poses some extremely awkward problems for the user of microcomputers, including:

1. There is no easy way to tell a good package from a bad one.

2. Hobbyists generally know nothing about good programming practices and disciplines. It is unlikely that anyone else will be able to maintain the programs these people produce.

3. Few people outside a business know enough about its day-to-day requirements to produce a correct and useful business application package. However, their ignorance does not keep them from trying to do so, or from selling the result of their work.

4. Program maintenance and support for a cheap package may be impossible to obtain from any source. Hence the low price.

5. The user may be compelled to alter his business procedures and methods to fit a computer package, rather than the other way round. The alternative may be a very expensive customised solution.

From these descriptions of the current state of system and application software, three points are clear:

1. The lack of useful, reliable software makes the microcomputer a poor choice today for many business data processing applications. The low cost of the hardware may be more than offset by the absence of a good environment for system development and by the doubtful quality and uncertain origin of the available application packages. Such good software as is available can be found only by trial and error and word of mouth.

2. The future widespread use of micros in business data processing will depend critically on the availability of good packaged application software. Some early "horror stories" are bound to circulate as badly-written packages are tried out in good faith by naive users.

3. Ironically enough, large user organisations are better placed to take advantage of the microcomputer than are small first-time users. Their greater depth of DP skills and experience will partially compensate for the deficiencies of the software. First-time users have no such compensation.

TOTAL SYSTEM COSTS

Conspicuously absent from this chapter is any suggestion that the microprocessor will make organisational computing much cheaper than it is today.

Of course, in some ways it will tend to do so. The rock-bottom cheapness of the hardware will be a positive factor. The modular approach to system building, correctly practised, will be another. But there is no strong indication that the overall cost of computing will decline significantly in the short term.

The essential concept is the *total system cost* involved in a computer application, from its inception to the end of its useful life. This cost is made up of the individual costs of several familiar components:

- Feasibility study and analysis.
- System design and specification.
- Programming/coding.

- Program and system testing.
- Acceptance testing and implementation.
- Documentation.
- User education and training.
- Live operation.
- System maintenance and modification.

These components are assembled differently by different organisations, but all of them occur in some form in most computer applications. The significant point is that the "processing content" of most of these activities is actually rather small, and for several of them it is nil. The total cost of computing for most organisations has tended constantly upwards, while hardware costs have for many years tended steadily downwards. This analysis confirms what most data processing managers already know: that the "people costs" of computing dominate the total system costs, and this dominance is becoming more pronounced.

What the microprocessor is likely to do is to emphasise dramatically the relative costs of people and hardware in computing. One unfortunate way in which this will happen is through the sale of cheap computers to inexperienced users, who will then find to their cost and dismay that these cheap devices and cheap systems are not the same things.

Nevertheless, the total system cost can be attacked at the following points:

1. More formalised system analysis and design methodologies may be developed to take advantage of the economics of the microprocessor. These may supplant the *ad hoc* techniques in use today, and shorten the crucial design stage.

2. Better standard programming languages may reduce the time required for programming and testing, and may make maintenance easier. There is considerable movement towards such languages today, much of it quite independent of the hardware manufacturers.

3. More stand-alone systems in user departments should mean less DP involvement with live running and user training. (Whether the user's own people costs should be counted in the total system cost is another matter.)

Even if these improvements fail to materialise, and system costs relentlessly increase, it is important to keep costs and benefits in perspective. Over the years, systems *have* become better and more useful. Some lessons *have* been learned. New techniques are available which qualitatively improve the systems that incorporate them. Progress may have been slower than many would like, but it has been fairly steady. The devolution of systems towards the user, and the proliferation of interactive, user-friendly systems, have made computers vastly more acceptable to laymen today than they were just five years ago.

The microprocessor will both encourage and strengthen these healthy trends. It will have a strong positive influence on the future of data processing. It will improve users' conceptual understanding of computing. It will not do these things by a direct assault on total system costs, for these are not within its influence. Instead, it will permit the computer to penetrate economically into a great many new activities which can benefit from its presence. Increasingly better computing value will be available for a given system cost.

CHAPTER 5

IMPACT

The preceding chapters have revealed that two important changes are occurring to traditional data processing because of the economics of the microprocessor and its associated technology:

1. Cheap hardware is changing the mix of computer equipment and is encouraging a different approach to system design.

2. The traditional relationships in some sectors of the computer marketplace are breaking down and re-forming along new lines.

The ramifications of these two changes are sure to be felt throughout the entire computer industry, and this raises the important question of just what these changes will be. This chapter looks cautiously into the near-term future (up to two years) to extrapolate current trends and to evaluate their impact. A shorter forecast horizon would be useless, a longer one foolhardy in view of the current pace of developments.

THE MARKET FOR MICROPROCESSORS

As we discussed in Chapter 3, the larger a market there is for a particular microprocessor, the cheaper it becomes. Thus, market forces always tend to favour a mass application over a special-purpose one. If commercial data processing requires completely different microprocessors from those that find their way into widespread domestic use, it can expect to pay much more for them.

This point is clearly illustrated by considering every potential DP application for micros: small computers, network elements, and so on. It would be difficult, on this basis, to estimate a worldwide requirement for more than one million microprocessor units. By comparison, using just one microprocessor in every telephone would require more than 100 times as many micros in the USA alone, and they would probably be much simpler micros. This comparison does not mean that there is not room in the market for both kinds of applications. But it does mean that it is very unlikely that a dozen different manufacturers of 16-bit microprocessors can all be profitable if they depend on the DP market alone to provide mass applications for their most complex products.

Using the microprocessor in a general-purpose programmable computer invites a more comprehensive instruction set and a larger word size. Using it as a dedicated processor of simple data and transactions does not. At the moment, the former type of application is very much in vogue, owing to a pent-up demand by small organisations for cheap computers. But once this demand has been satisfied, there are very few large-scale applications of micros in evidence today which would benefit from more-complex microprocessors. Figure 8, on the next page, provides a ranking of some current 16-bit micros. This does not mean that such applications will not be found, but it does mean that the progression from simple micros to more-complex ones is neither automatic nor inevitable.

Another factor that strongly opposes new micro designs is the desperate need for a *standard device* in the market. Currently, there are some four or five different 8-bit micros that could

Figure 8 An approximate ranking of some 16-bit microprocessors according to raw power

Power

- Motorola MC68000
- Zilog Z8000[2]
- Intel 8086
- Western Digital MCP-1600
- Zilog Z8000[1]
- Texas Instruments TMS 9900
- Fairchild 9440
- Data General mN601
- National Semiconductor INS 8900
- Texas Instruments SBP 9900
- National Semiconductor PACE
- General Instruments CP 1600

1 40-pin version, 64Kb max. addressing
2 48-pin version, 8Mb max. addressing

(Source: *Mini-micro systems*, January 1979)

reasonably claim to be the most popular, but it is rare to meet someone who has had experience of more than two of them. This chaotic state of affairs should represent a strong barrier against further new designs. Yet a number of 16-bit micros have already entered the market, and although some of them offer a crude compatibility with their 8-bit predecessors, effectively they are new designs.

Just as the IBM 360 and the Digital PDP-11 architectures have become *de facto* standards in their respective markets, so a single microprocessor architecture, or at worst a single design for each of the popular word sizes (4-, 8- and 16-bit), should emerge. Although the semiconductor manufacturers will initially resist the appearance of a standard micro, and although the architecture finally chosen may, like the 360 and PDP-11, be objectively inferior to others on offer, *de facto* standardisation must occur.

One interesting way in which this might come about would be for the market to settle on a single-chip version of either the Digital PDP-11 or the Data General Nova. Leaving aside the question of whether the architectures of these minicomputers are appropriate to a micro, the idea is appealing because of the very large amount of software which already exists for both, and the general state of knowledge of these devices and their capabilities. Both Digital and Data General already have microcomputer versions of these minicomputers, the LSI-11 and the microNova respectively. These are single-board rather than single-chip devices at present, and, consequently, they are rather expensive. Neither company is a significant manufacturer of semiconductors.

In this case, however, the interests of the market and of the suppliers appear to diverge. Although the market might benefit considerably from adopting as a standard an architecture that is already widely known, a flood of cheap LSI-11 chips could swiftly undercut the market for some of Digital's larger and pricier PDP-11 minicomputers. It might also reduce the lucrative sales of Digital peripherals by opening more of the PDP-11 market to plug-compatible competitors.

Both Digital and Data General, therefore, appear to be treading carefully in marketing these devices. So far, most of them seem to be sold into up-market OEM equipment. This pattern would change if another semiconductor house produced and sold, say, a $10 PDP-11 chip. This is quite feasible technically, and there is litigation in progress at the time this report is being written to forestall just such a move. The outcome could be very significant for the early adoption of a standard 16-bit micro.

The market has been slow to accept the new 16-bit designs. This reaction suggests that the increasing complexity of new microprocessors is slowing down or perhaps even reaching a natural limit. While the manufacturers appear to be surprised at this development, there is no reason why they should be. There appear to be three natural markets for microprocessors of different sizes:

1. 4-bit Simple device control (watches, calculators, cookers) requiring small programs with numeric displays.

2. 8-bit More complex device control (factory systems); information handling and data communications (word processors, modems).

3. 16-bit Simple computing applications (minicomputer replacements).

As micros become more complex, they lose some of their direct applicability to simple tasks. Yet it is these commonplace tasks which are likely to provide the biggest markets for micros and the greatest profits for their manufacturers. (Figure 9, on the opposite page, shows historical and predicted sales of microprocessors by word size.) We, therefore, expect to see fewer and fewer new architectures on the market, and very little pressure to expand beyond 16 bits. Instead, there will be a trend towards the single-chip microcomputer with on-board memory and interface circuitry. This is the simplest development that will reduce assembly costs and improve overall reliability.

Until recently, the semiconductor industry has been able to produce whatever was technically feasible, and the market would buy it. With microprocessors and their associated software costs, this environment is changing. The silicon industry now finds itself deeply involved in systems, instead of just in simple components. As a result, there are likely to be fewer semiconductor houses producing microprocessors in two years' time than there are today.

SMARTER PERIPHERALS

Almost every manufacturer of VDUs now incorporates microprocessors into his products. At least one VDU on the market has three separate micros in it. Until recently, the distinction between "smart" and "dumb" VDUs implied a significant difference in price because of the extensive special electronics the former required. However, by incorporating microprocessors and other IC components, the quantity and the cost of electronic hardware can be made to be nearly the same for both types, the functional differences being accounted for by software.

A useful extension to the VDU is the stand-alone data entry station, capable of fairly complex

Figure 9 Worldwide microprocessor and single-chip microcomputer sales

	1973	1974	1975	1976	1977	1978	1979*	1980*	1981*	1982*
4-bit:										
Annual	0.1	0.3	1.0	2.0	4.1	9.0	15.0	34.4	95.4	207.0
Cumulative	0.1	0.4	1.4	3.4	7.5	16.5	31.5	65.9	161.3	368.3
8-bit:										
Annual	0	0.01	0.6	1.5	3.0	6.0	8.0	13.0	20.0	35.0
Cumulative	0	0.01	0.6	2.1	5.1	11.1	19.1	32.1	52.1	87.1
16-bit:										
Annual	0	0	0	0	0.2	0.3	0.8	1.6	3.6	6.0
Cumulative	0	0	0	0	0.2	0.5	1.3	2.9	6.5	12.5
Total:										
Annual	0.1	.31	1.6	3.5	7.3	15.3	23.8	49.0	119.0	248.0
Cumulative	0.1	.41	2.0	5.5	12.8	28.1	51.9	101.0	220.0	468.0

* (Estimated) Sales are in millions of units

(Source: *Computing,* 5 April 1979)

editing of keyed input. No on-line links to a central computer are necessary if sufficient local magnetic storage is provided. Such "intelligent terminals" are already popular in companies that must collect and edit data from widely scattered sources, but cannot justify an expensive on-line data network for doing so.

Other peripherals will also benefit from the microprocessor, but in a different way. Printers and discs currently require substantial CPU overheads in order to supervise and control their physical and logical functions. Monitoring error conditions, performing automatic re-tries after each failure, and collecting performance statistics are all tasks which historically have been delegated to the operating system for want of a better place.

A superior approach is to convert printers and discs into self-contained sub-systems, each under the control of a microprocessor. All device-dependency can then reside with the micro. So can all handling of the physical operations of the device, and also extensive self-diagnostic procedures. Communication with the main computer could then be restricted to high-level program requests for specific peripheral functions, followed by bulk transfers of data between the sub-system and the main computer.

This approach would have several advantages:

1. Diagnosing faults would be much easier. There would never be a question of where in the system an error had actually occurred.

2. The main computer would be freed from a great deal of tedious housekeeping and consequent CPU overheads.

3. The main operating system could be made much simpler in structure and much smaller in its memory requirements.

4. A malfunction in a peripheral could not crash the main system, nor would the peripheral be made inoperative by a main system crash.

5. By communicating at arm's length with the main computer through a high-level protocol, the peripheral would be truly plug-compatible. If this protocol were a universal standard, a different peripheral could be substituted for an existing one, without the need for any changes whatsoever to the main computer's programs.

6. From the point of view of the main computer and its application programs, the physical and logical attributes of the peripheral device would be completely separate.

Of course, there is no restriction that the peripheral be a single piece of hardware. It could be *any* self-contained device that is required as a passive utility by an active program in another machine. An interesting candidate for this role is the *database processor.* This could take the form of a computer (perhaps even a large mainframe) holding and maintaining a conventional database, with or without the help of a DBMS. This computer would be available to a number of other computers that required access to the database. All accesses could be strictly controlled by this "back-end processor", and requests for data could be made at a very high level without regard for the physical structure of the files. Such a utility would be extremely valuable as a node of a large network.

NETWORK LINKS

The task of linking computers into data communications networks is currently an extremely complex one. There are a number of so-called standard architectures for building networks, but none of them is universally accepted. Several are obviously intended to lock the user into relying permanently on a particular hardware manufacturer.

A computer that requires access to a network could view that network as a sort of peripheral. If all the machine-dependent software and low-level protocols were programmed into a microprocessor, then that micro could be used as a standard network interface for every computer attached to the network. This development would have the following effects:

1. The computers could talk to the network using only the highest-level protocol and leave all lower-level communication to the micro. This approach would greatly simplify network software design.

2. Each micro would talk only to other identical micros and to the one device for which it was the interface. This device could be a computer, but it might also be a peripheral, such as a printer. This lavish use of standard interface computers would be economical, because of the cheapness of the micro.

3. A great deal of software would be moved downstream from the nodal computers to their interface micros. Therefore, this software would need to be written only once. It would not have to be implemented on every kind of computer that was attached to the network.

While the concept of using an intelligent network interface computer is by no means new, realising it on very cheap hardware could mean that low-budget internal company networks could suddenly become economically attractive. This application must represent a prime target for a well-constructed package to be marketed to medium-sized computer installations. It raises the possibility of a private data "ring main" in an organisation, into which every user who required access to the corporate database could simply plug his computer or terminal.

LANGUAGES AND SOFTWARE

Microprocessor applications are rapidly sorting themselves into two distinct categories:

1. Traditional DP applications in which the micro is the centre of a small general-purpose computer system.

2. Specialised applications in which the micro is a dedicated component within a larger system.

The requirements for computer languages roughly follow this division.

In the first category, BASIC has dominated the applications market so far. BASIC is very unsatisfactory from several points of view, but it has the enormous advantage of being easy to learn and use, if not to read and understand. Because BASIC was originally devised to enable schoolchildren to use a computer, it is not very surprising that many programmers have mastered it.

COBOL has been developed for micros by some manufacturers and OEMs, and is now available in various implementations on many of the popular 8-bit micros. Its use on such small machines is uncomfortably cramped, and it often lacks some facilities that most programmers would take for granted on a large mainframe. However, it preserves some compatibility with larger computers, and the language is adequate for simple commercial applications, particularly those that require screen formatting.

It is very unlikely that either BASIC or COBOL will be displaced from commercial use of the micro, because it is this simplicity and superficial compatibility that will allow the micro to be deployed in business applications. The commercial DP market is rather reactionary and does not adapt easily to a new language.

For other applications, however, the outlook is very different. In process control, data handling or device supervision, BASIC, FORTRAN and COBOL are very awkward, and greatly restrict the flexibility of the micro by under-utilising its innate capabilities. Assembler languages are generally used in these applications, plus special real-time languages such as PL/M and CORAL.

To avoid the need to learn a new instruction set and new architecture with each new micro design, a standard high-level programming language is needed for non-commercial applications. There is currently a strong impetus towards adopting PASCAL as this language. PASCAL is a descendant of ALGOL, which it resembles. Originally developed in an academic environment, PASCAL is currently enjoying a wave of great popularity amongst micro enthusiasts. Universities in particular are promoting it, and an international standard may not be far away. All this activity will not necessarily save PASCAL from the fate of PL/1, but it stands a good chance of success. Its prospects are good because it is not aiming to replace an established standard, but rather to impose some order on the Tower of Babel which is the state of microprocessor languages today.

If the micro is to penetrate the commercial systems market on a large scale, it will require excellent application packages as a vehicle. Otherwise, expensive custom software will have to be added to cheap hardware, resulting in an unappealing combination for the first-time user.

Some software and systems houses have specialised in developing such packages for minicomputers. They may move down-market to the micro with the same packages, albeit at greatly-reduced revenue per sale. As we noted earlier, a significant barrier to developing comprehensive software is the constant appearance of new micros in the market. As this process slows down, the potentially extended market life of a good package will make it a more attractive candidate for development. Many copies of each package will have to be sold in order to get the price down to a level acceptable to the lower end of the market. An example of success in such an approach is the CP/M microcomputer operating system, of which over 100,000 copies have been sold for about $100 per copy.

There is a fundamental question of the degree to which an organisation will alter its internal procedures in order to accommodate a software package. If the alternative is an expensive bespoke system, then the choice should be a simple one of economics, but the costs of disruption and change are difficult to quantify. The most probable course is that small companies, lacking an alternative, will use packages and adapt to them, whereas larger organisations will continue to expend resources on customised systems, whether they really need them or not.

THE COMPUTER INDUSTRY

There is tremendous turmoil in the computer industry today as a direct consequence of microelectronic technology. It is easy to see why this is so.

Every manufacturer across the entire range of computers — micros, minis and mainframes — is affected by the plummeting costs of electronic hardware. This means that for constant performance, the price of a computer (and with it, the manufacturer's profit margin) must fall. Manufacturers have reacted to this phenomenon by migrating up-market, where there is greater added value and greater unit profit. In general, the three classes of computer manufacturers have behaved as follows:

- The micro manufacturers, instead of selling only chips and boards, are marketing small business systems that are indistinguishable from small minicomputers.
- The minicomputer manufacturers are enhancing their top-of-the-line products with virtual memories, multiprogramming operating systems, and elaborate ranges of peripherals, and are selling these, in effect, as mainframes.
- The mainframe manufacturers have more of a problem because there is nobody above them whose market they can invade, so they are badgering their existing customers to upgrade and enlarge their existing systems. Some have even turned the tables and are pushing vigorously downwards into the minicomputer market.

Some users who have suffered in the past at the hands of arrogant computer manufacturers may relish their discomfiture today. But there are important consequences for the user as well as for the manufacturers:

1. The manufacturers will have to adapt to new markets and different applications from those to which they are accustomed. This process is sure to result in their making some mistakes. A manufacturer who is forced to go into a new line is unlikely to service it as well initially as he did his traditional lines. His product or support may be off-target, and the user will be the first to know it.

2. There are signs that the micro manufacturers and some of the mini manufacturers do not yet fully understand the commitment they will have to make to software and customer support, which were formerly the exclusive province of the mainframe manufacturers. This deficiency could rebound against them if the market is unwilling to accept a lower standard of support than it expects.

3. The marketing and support structures of some companies are inappropriate to the new markets they are trying to enter.

The likely outcome is predictable but difficult to pinpoint precisely. Smaller revenues from cheaper products mean smaller profits and ultimately less room in the market for competitors. What is uncertain is who will disappear first.

An early sign of this revenue shrinkage is the intensified price competition in mainframe computers. It is very significant that IBM has signalled with its 4300 range that it will now compete actively on price, and not merely provide a price umbrella under which the other manufacturers may shelter. This move made very bad news for some of IBM's competitors.

Although it is not directly threatened yet, the mainframe computer industry is probably the most vulnerable sector in the longer term. There are three reasons for this:

1. The new economics of hardware increasingly favours smaller computers.

2. The rationale around which mainframes were developed — the need for expensive high-performance processors to be centralised and shared out — no longer applies.

3. New approaches to system design are likely to emphasise cheap, dedicated, physically-distributed processors.

THE COMPUTER SERVICES INDUSTRIES

The computer services industries in the UK represent a considerable investment in hardware and skills. These industries consist of computer bureaux (both batch and time-sharing), software and systems houses, and consultancies. All of these will feel the winds of change, some much more than others.

The computer service bureaux will be the most dramatically affected. This industry, as it was originally conceived and as it has operated for many years, appears to be doomed. Economics and users' preferences will no longer favour running large batch programs on an expensive bureau computer. Similarly, a small in-house interactive computer should be able to compete with most time-sharing bureaux in both cost and ease of use.

If service bureaux expect to survive, there are two paths they can take:

1. Selling hardware and software
   Instead of tying the customer to its computer, the bureau can sell him his own computer and the special programs it has developed. The advantage is that the successful bureau already has an effective marketing network and a good knowledge of its customers' needs, both of which it can put to immediate use. The disadvantage is that a continuing source of revenue is in effect converted into a one-time payment. Nevertheless, some bureaux are already adopting this strategy and are becoming OEMs of a sort.

2. Selling very specialised computer services
   If a bureau has application packages which are very complex and much sought after, or if the customer's applications must run on a super-computer of some kind, then the bureau may enjoy a natural monopoly which is immune to invasion by an army of tiny computers. However, this will be the case for only a fraction of bureau applications.

Software and systems houses are comparatively flexible, and will not be as susceptible as bureaux to sudden obsolescence. However, they must beware of offering expensive software to run on very cheap hardware. They will need to concentrate on producing better packages than the OEMs, and most of them will find it irresistible to move into hardware sales as well. Some have already done so.

Commercial users in future will shop for complete hardware/software systems that will be much cheaper than they are today. If software houses can produce truly general and adaptable application packages, they will be able to compete very effectively with other OEMs. However, the market will encourage a large number of sales at a low unit price, a strategy which is not familiar to most software houses. Significant internal changes to software houses will be needed.

These houses have a considerable advantage in that they are accustomed to providing education and "hand-holding" for their customers. Software support and customer support will be in great demand in future, because the hardware suppliers will not be able to provide enough of them. This support will probably be charged for separately from the software itself — IBM has already started this trend.

Consultancies have an inherent ability to lean with the prevailing wind of the moment. They are usually able to change their preferences at short notice to accommodate both genuine technical advancements and short-lived DP fads.

The danger to consultancies in future is that their fees for advice may become a significant fraction of the cost of the computer system under study. Some clients (particularly first-timers) will balk at this apparent imbalance. They may decide that the risk of doing the wrong thing carries a smaller cost penalty than buying consultants' high-priced advice. Some consultancies have already suffered from this phenomenon.

Consultants will, therefore, need to specialise in two ways:

1. They must perform assignments for larger organisations that are already heavy users of DP. Their advice to these users will be focused upon conceptual alternatives and the design of major new systems. Many organisations with large DP departments carry a staff of skilled analysts and programmers but have very few first-class conceptual thinkers. Such individuals are difficult to recruit and retain. Consultancies could provide a very cost-effective way of making these skills available to a number of organisations.

2. They should accept the challenge of educating clients' top management in the concepts of data processing. Very often, senior British managers refuse to devote time to internal computer appreciation and training courses offered by their own DP departments or by computer suppliers. A consultant who is skilled in making such presentations can frequently break down this barrier, because the manager is less afraid of revealing his ignorance to an outsider. He is also likely to listen attentively when the advice is costing his company a lot of money.

SKILLS AND EDUCATION

It is probably true that the degree of success achieved in applying computers to practical problems today is more limited by education and imagination than it is by the tools available. With today's knowledge, attitudes and skills, it is doubtful whether the computer industry could deliver much better computer solutions to difficult problems, regardless of how many computers and programmers were available. Human resources are the most limited resources of all.

We suggested in Chapter 4 that the microprocessor might provide the economic basis for a new approach to complex problems. In the past, many "new approaches" have fizzled out, leaving DP in its traditionally limited role. If such a failure is to be avoided this time, then the question arises as to what skills and attitudes must be cultivated in order to utilise fully this new tool.

The experience of some users who are already experimenting with microprocessors suggests that no completely new skills are required, but rather a rearrangement of the old skills is necessary.

The most valuable skill is the ability to perceive a complex problem in terms of its functional parts, each one separable from the others and amenable to a relatively simple solution. This ability appears to be an innate gift in some individuals, but others have successfully acquired it through experience.

The second important skill is an empathetic one, which perceives a specific user need in the form of a set of mixed manual/computer procedures which are both useful and easy to computerise. This skill will be much more important than it is today, because new systems will be far more interactive, user-intensive and immediately functional than are most current ones.

Traditional batch systems require very little knowledge of moment-to-moment user needs compared with interactive systems. There will be a distinct shift in emphasis towards considering the computer as a handy extension of the user in his normal work. This means that the analyst will have to cultivate an intuitive grasp of the potential of the human and the computer working cooperatively on each facet of a problem.

Significantly, neither of these abilities is intrinsically a "computer skill" in the traditional sense. It is likely that gifted individuals will emerge from user departments, as well as from DP departments, to be the system designers of the future. We can expect that the conventional programming and coding skills of today will be greatly de-emphasised in future. One of the goals of modularity and packaged systems is to ensure that there is less software being written in future than there is today.

It is characteristic of the British educational system that very few individuals emerge from either a secondary school or a university with a working knowledge of computers. This deficiency poses three serious problems:

1. If this knowledge is to be obtained at all, it must be imparted by employers at their own considerable expense.

2. Some of these individuals will rise to influential positions in society, but will bring with them no insight into how computers could contribute to their fields.

3. The public's fear of and antipathy towards computers will be perpetuated.

It is essential to future economic success that all or most pupils receive at least a rudimentary introduction to computing principles. Given the heavy bias towards the arts at the expense of technology in much of British university education, it would be best if this instruction took place during the secondary school years and was made a compulsory part of the programme.

Countries which have adopted such an approach are already reaping the benefits. For example, in the USA, the first students to receive some early computer training in secondary school are now completing their university educations. University staff members report that engineering graduates who were exposed to computers *before* they entered university are bringing a new and more productive attitude to engineering. It appears to be important that the computer be an accepted tool, freely available to the student, rather than a novelty introduced late in his education.

There is no reason why this advantage should be limited to engineering. Science, medicine, law, administration, and even law enforcement could all benefit greatly from such an attitude. The important point is that the computer training must take place early in a student's education. With hardware costs declining, there should be little difficulty in equipping schools with cheap interactive computers. The biggest problem will be training the teachers, both at school and university levels.

DISTRIBUTION AND DECENTRALISATION

Micros will compel many an organisation to re-examine its reasons for deploying its DP resources in the way that it has. The results of this introspection may be surprising.

There is no single best approach to DP organisation within companies, but there are certain basic concepts that are common to most companies. The distribution of data processing in a company can be envisaged as having three dimensions:

1. Hardware and operations

   This dimension is concerned with the physical placement and running of machines.

2. Development resources

   This dimension includes programs, systems, and the people who invent and develop them.

3. Management and control

   This dimension includes the integration of results, the protection of common data, and the maintenance of standards and policies.

There is room in these three dimensions for virtually every style of corporate DP management. For example, it is possible to imagine a company which controls very tightly dimensions 2 and 3 above, but which installs dedicated computers (not locally programmable) throughout the company. Its computing could be said to be very distributed but highly centralised. The opposite extreme would be a company consisting of highly autonomous subsidiaries that prepare their own systems and programs, but which are compelled by group policy to run them all on a central computer. This would be decentralisation without distribution. (These issues and others are addressed in Foundation Report No. 18, "Strategic and Policy Issues in Distributed Processing".)

The microprocessor will not alter these dimensions. What it will do is to make it economically possible for each company to adopt a pattern of DP management which best suits its users, its policies and its corporate style. As a result, many companies are likely to embark on ambitious new programmes of development in directions they could not seriously have contemplated before now.

CHAPTER 6

A COURSE OF ACTION

This chapter presents a summary of the main findings arising from our study of the microprocessor in data processing. It then suggests some specific steps that management services and data processing departments may take in order to broaden their experience of these devices and prepare themselves for the changes which may soon come about.

MAIN FINDINGS

1. The microprocessor, made possible by very large-scale integrated circuit technology, has begun to infiltrate traditional data processing activities. In its present form, it is mostly confined to stand-alone microcomputers which have limited capabilities and restrictive development software. Alternatively, it may be purchased as a single-board computer, which must then be interfaced with other devices to form a workable system of some kind.

2. Because they are very cheap, microprocessors permit a limited but useful amount of computing power to be brought to bear economically on problems that do not warrant a major investment in computing. They encourage users to allocate computer hardware physically to those locations in the organisation where it is immediately useful. In many cases, micros will find their way into user departments. Thus, users may become more directly involved in designing and operating the computer systems that affect them.

3. Micros make possible an approach to system design based on separate but cooperating small computers, each one handling a well-defined process within a larger framework. Computer systems in future can thus reflect the requirements of the organisation rather than those of the central computer.

4. Falling hardware costs mask the fact that the total system costs of data processing have continued to rise. Software costs, in particular, threaten to take the naive user by surprise. Some very poor application software is finding its way onto the market. For large users of conventional mainframes, the investment in existing software is so great that they may resist redeveloping their systems whatever the attractions of doing so.

5. Although their capabilities are very limited at present, more powerful micros are on the way. These include a number of 16-bit designs which are in effect conventional minicomputer CPUs. When the market chooses a standard architecture, a large amount of supporting software will quickly follow, and the usefulness of the micro will be suddenly multiplied. A limit to increasing complexity may already be in sight, because the principal high-volume applications seem to be concentrated at the low end of the market and require only simple microprocessors. Future developments will probably concentrate on integrating more separate functions onto a single chip, rather than on designing more powerful processors.

DECIDING TO TEST THE WATER

With these findings as a background, a large organisation should seriously ponder the question, "Why should we become involved with microprocessors now?" There are at least six good reasons why it should:

1. "Because it's there." The microprocessor is *the* topical subject in DP today, and, as such, its advantages and pitfalls should be familiar to any company that is heavily committed to DP.

2. The micro multiplies the potential for introducing useful DP applications into more user departments. If DP does not take the lead in this process, the users may do it for themselves.

3. Most companies have a waiting list of applications which are pending development by DP. For certain kinds of applications, micros can help to cut this list down quickly and economically.

4. The hardware investment required to make a useful start with microprocessors is not large in relation to most DP budgets.

5. Many professional DP staff are happiest when they are involved with the latest products of technology. Rightly or wrongly, they prefer employers who help them to keep their skills up to date. In this respect, the micro is very much at the leading edge of technology today.

6. As with any new device, there exists a learning curve for using microprocessors that favours an early involvement in order to ensure good results later. Deferring this involvement means postponing the time when useful applications can begin to emerge.

Point 2 above is especially significant. Unlike earlier computers, the microprocessor is capable of entering a large company in a number of ways, not all of which fall automatically within the scope of DP's formal authority. As microprocessors are incorporated into more and more kinds of products, and as the boundaries between computers and other devices (for example, word processors or electronic PABX systems) become less and less distinct, DP will find itself increasingly on the defensive, unless it expands its competence to include the microprocessor.

Furthermore, pure DP applications of the micro in industry are likely to comprise only a minority of the total. Crucial applications will be found in the production line, the warehouse, points of sale, and other "sharp-end" locations in the organisation. However, DP staff may be needed to identify these applications and to develop and integrate them. Other people in the organisation may lack the analytical skills and outlook necessary to exploit these opportunities. DP departments must, therefore, equip themselves with practical knowledge of the microprocessor in preparation for these new developments.

A SURVIVAL GUIDE FOR MANAGEMENT SERVICES

There is a methodical process by which management services departments can obtain this expertise in a reasonable time and at modest cost. This process consists of three distinct stages:

1. Getting the right people and making them aware

   Getting the right people does not mean recruiting microprocessor experts from outside. There are not many of them about, they are difficult to find, and the idea is to *build* skills, not to buy them.

   The right people should all come from within the company. Initially there should only be a few of them – no more than four – and they should not consist entirely of "old hands". A mix of people of different skills and ages is probably best, including someone who is at home with hardware, and another who is a competent assembly-language programmer. All should have the ability to "think small", and they should enjoy working within tight constraints. Pains should be taken to avoid giving other staff the impression that an elite team will be doing all the interesting work.

Their initial goal should be to become familiar with current developments in microcomputers. This should initially be a part-time assignment. A good deal of reading will be required including reference books, computing journals, hobby magazines, and manufacturers' product literature. External courses may be useful as well. By the end of this stage, the team members should be very keen to acquire some microcomputer hardware.

2. Obtaining the tools for development

   As we have already noted in this report, there is a bewildering array of hardware on the market already, and more is appearing all the time. Making a good choice is, therefore, difficult.

   There are two basic approaches to microcomputer hardware: through small business computers or through microcomputer development systems (MDSs). The former are highly packaged with languages and some application software, and they are intended mainly for small-scale commercial DP use. The latter are in effect sets of specialist tools with which to develop *any* kind of microcomputer application, particularly real-time and industrial applications. Most MDSs are sold by the micro manufacturers themselves.

   The small business system is adequate for developing most DP application programs. However, if system software design or hardware experimentation is contemplated, then an MDS is essential.

   The choice among various small business systems should be made with a view to the *development software* which is available on each. Development software for micros still has many limitations, but these can be minimised by careful evaluation and shopping. The comments of other users are very helpful during this stage. An important consideration is that the hardware needed for testing and debugging some application programs may be considerably larger than that required by the application itself.

3. Choosing and performing a pilot application

   A successful pilot application is a most important stage in a company's first experience of microprocessors. It not only serves as a proving ground for the new skills that are being developed, but can also be used to combine DP skills with users' special knowledge in a highly-visible team effort. The benefits arising from a successful project may far exceed the limited objectives of the pilot project itself.

   Selecting the right application is, therefore, the most important task of this stage. The application should have these attributes:

   - It should meet a genuine user requirement.
   - It should be strictly limited in scope but not trivial, and should have clearly-established and well-understood objectives.
   - It should require only one terminal, and should have small data volumes.
   - It should involve an interdisciplinary team which includes the user.
   - It should be achievable without external help in a reasonable time, a few months at most.

During the project, the team should report to a senior manager, preferably one outside the data processing function. Control of the project should be tight, but not bureaucratic. Progress should be carefully monitored and recorded, and problems should be brought into the open as soon as they are identified.

At the end of the project, a successful application will meet the following criteria:

- It will be clearly seen to be a success in quality, time and cost.
- The user will have a productive new application which meets a real need.
- The methodology used will be repeatable on other projects.
- The involvement of DP and user staff will have been productive and reassuring.
- New ideas for other applications will have emerged during the project.

Following this experiment, the management services department should have not only a good idea of the capabilities and limitations of the microprocessor, but also an inkling of where the micro can best fit into the company as a whole. Different organisations are finding completely different uses for the micro in DP, thus reflecting both the versatility of the device and the ingenuity of users and systems people.

CHAPTER 7

CONCLUSION

It may well be that in a few years' time no distinction will be made between a microcomputer and any other kind of computer. Programming one machine will be exactly like programming any other, with the same development aids available on all. The only clues to the internal structure of the computer may be its price and the uses to which it is put, thus reflecting the key parameters of *economics* and *utility*.

If this comes about, then 1979 will have been one of the last years in which the computer was an important consideration in its own right. In 1979, humans still queued up and subjected themselves to considerable inconvenience in order to obtain a small slice of this expensive and scarce resource.

The writing is already on the wall for this quaint approach to computing. The microprocessor and its associated technology have begun the irreversible process of cutting the computer down to size — and with it, some of its high priests and acolytes who have profited so greatly from the computer's remoteness and expense over the past twenty years.

If the much-predicted "computer revolution" of the 1960s never came about, it may have been because the computer was inaccessible to those who might have led the revolution. Until recently, access to the computer has been the exclusive right of a privileged caste of data processing experts, many of whom owed their primary allegiance to their profession, rather than to their employers. Sometimes their "expertise" consisted largely of their exclusive knowledge of the awkward and inconvenient languages and operating systems which had been devised for utilising the computer. Little understanding of the organisation and its problems was assumed or expected of these experts.

In 1979, however, change is in the wind. In more and more companies, the initiative for new applications is passing to the front-line user who has genuine needs and no interest in computers as such. This transition will not be an easy one, and it must be expected to be fraught with confusion, recrimination and occasional disasters. But it will eventually produce a new kind of user: one who is knowledgeable, technically sophisticated and very demanding.

Once this happens, the *real* computer revolution just might come about. If it does, it will owe less to the giant, million-dollar, water-cooled son-of-ENIAC than it will to the humble pocket calculator.

BIBLIOGRAPHY

(References marked * are particularly recommended.)

1. ANON. "Applications are chasing the chips." *The Economist,* 24 March 1979, 119-120.

2. ANON. "A programming scheme that failed." *Business Week,* 28 May 1979.

3. ANON. "Chips lead to close control." *New Scientist,* 19 April 1979, 189.

4. ANON. "Communicating computers ring the changes." *New Scientist,* 3 May 1979.

5. ANON. "Computers learn to speak, but are still hard of hearing." *The Economist,* 9 June 1979, 105-6.

6. ANON. "Computers rush to talk PASCAL." *Business Week,* 18 June 1979.

7. ANON. "High-level languages: the next technology." *Microprocessors and micro systems,* vol. 2, no. 5, October 1978, 286-9.

8. ANON. "MicroAde concepts and facilities." CAP Microsoft Limited, 1978.

9. ANON. "Microelectronics, the new technology." Department of Industry and Central Office of Information, Colibri Press for HMSO, 1978.

*10. ANON. "The microprocessor: A revolution for growth." *Business Week,* special report, 19 March 1979, 42BI-42WI.

11. ANON. "The year of LSI and the maxi machines." *Computer Weekly,* 12 January 1978.

12. BAKER, K.D. "Functional decomposition on multi-microprocessor systems." *Radio Electronics Engineer,* vol. 47, no.11, November 1977.

13. BANKS, M. "Going where no computer has been before . . ." *Computer Weekly,* 4 August 1977.

*14. BARRON, I.M. "The future of the microprocessor." *Microelectronics,* vol. 8, no. 4, April 1978.

*15. BARRON, I. and CURNOW, R. *The future with microelectronics.* London: Frances Pinter Ltd., 1979.

16. BASS, J.C. "When the chips are down . . ." *Sunday Telegraph,* 24 June 1979.

17. BECKETT, B. "Decade of the micro." *Computing,* 17 November 1977.

18. BLOCH, A. *Murphy's Law and other reasons why things go wrong.* Los Angeles: Price/Stern/Sloan Publishers, 1977.

19. BRESLIN, J. and TASHENBERG, C.B. "Distributed processing systems — End of the mainframe era?" *AMACOM,* February 1978.

*20. CANNING, R.G. and McNURLIN, B. "Micros invade the business world." *Datamation,* August 1978.

21. COSSERAT, D. "The true costs of microprocessing." *Microprocessors,* vol. 2, no. 3, June 1978, 115-8.

22. CURNOW, R. "The future with microelectronics." *New Scientist,* 26 April 1979.

23. DOERR, J. "Low-cost microcomputing: The personal computer and single-board computer revolutions." *Proceedings of the IEEE,* vol. 66, no. 2, February 1978, 117-130.

24. DOZIER, H.W. and GREEN, R.S. "Single-chip computer expands its memory." *Electronics,* vol. 51, no.10, May 1978.

25. FALOR, K. " 'Distributed' doesn't mean 'random'. " *Datamation,* August 1978.

26. FARRELL, E.P. and KANELLOPOULOS, N.G. "Debugging aids for microprocessor systems." *Microprocessors,* vol. 2, no. 2, April 1978.

*27. FLANAGAN, D. (editor). *Microelectronics* (a Scientific American book). San Francisco: W.H. Freeman and Company, 1977.

28. FLINT, D. "Trends in database management systems." Butler Cox Foundation Report Series No.12, June 1979.

29. GOLDSMITH, P. "Software: the key to power for the user." *Computing,* 12 January 1978, 12.

30. HAMMER, M. and ZISMAN, M. *Design and implementation of office automation systems.* MIT Laboratory for Computer Science, 1979.

31. HARRISON, P.G. "The technology race in microprocessor application." *Euromicro Journal,* vol. 4, no. 1, January 1978.

32. HEALEY, M. "Scene is set for 'cowboys' to move in." *Computer Weekly.*

33. HUGHES, P. and DOONE, T. "Multiprocessor systems." *Systems International,* vol. 6, no. 1, February 1978.

34. JENSEN, K. and WIRTH, N. *PASCAL user manual and report.* Heidelberg: Springer-Verlag, 1975.

35. KEWNEY, G. "Growing pains of the soft machines." *New Scientist,* 22 March 1979, 938-941.

36. McLEAN, J.M. and RUSH, H.J. "The impact of microelectronics on the UK: A suggested classification and illustrative case studies." University of Sussex, Science Policy Research Unit, Occasional Paper Series No. 7, June 1978.

37. MANCHESTER, P. "Answers for the unasked questions." *Computing Europe,* 15 June 1978, 18.

38. MARAZAS, G. "Microprocessors heed the call to supervise I/O." *Electronics,* vol.51, no. 3, February 1978.

39. MITCHELL, R.E. "Third generation myopia." *Datamation,* September 1978.

*40. OGDIN, C.A. "Sixteen-bit micros." *Mini-micro systems,* January 1979.

*41. PURCHASE, A. "Small business computer systems in Western Europe." SRI International, 1978.

42. RAO, G.V. *Microprocessors and microcomputer systems.* New York: Van Nostrand Reinhold, 1978.

43. RIGO, J. "Harry's micro: The birth of a department." *Computerworld,* 9 April 1979, 28.

44. ROBARDS, M. and GARNER, T. "ACRE: A real-time distributed system." *Microprocessors and microsystems,* vol. 2, no. 4, August 1978, 195-203.

45. SEIDMAN, S.H. "Hobby and game markets fade . . . is there life after kits for retailers?" *Datamation,* April 1979.

46. SHELTON, C.D. "Microcomputer design." *Wireless World,* vol. 84, no.1505, January 1978.

47. SOUCEK, BRANKO. *Microprocessors and microcomputers.* USA, 1976.

48. STEVENS, R. "Driving the S-100 bus to a micro future." *Computer Weekly,* 23 November 1978.

49. SUMNER, C. and LEVY, W. *The affordable computer.* AMACOM, 1979.

50. SUTTON, C. "Taking the mystery out of the micro." *New Scientist,* 17 May 1979.

51. THEIS, D.J. "An overview of memory technologies." *Datamation,* vol.24, no. 1, January 1978.

*52. UTTAL, B. "The gentlemen and the upstarts meet in a great mini battle." *Fortune,* 23 April 1979.

53. WHITBREAD, M. "Design techniques with microprocessors." *Microprocessors and microsystems,* vol. 2, no. 5, October 1978, 263-7.

Abstract

Report Series No 15

Management Services and the Microprocessor

by Edward Goldblum

September 1979

The microprocessor or "silicon chip" has rapidly become the most topical subject in computing. Everyone knows – or thinks he knows – that the micro will have a great impact on the way we live. But few people have considered the impact it will have on data processing and information systems, the activities to which it is most closely related.

This report first describes the historical background of the modern integrated circuit, showing that it has been economics rather than technology that has determined the pattern and the pace of developments. Both the production and the use of microprocessors are analysed in terms of their key economic factors.

The report then describes the new opportunities in data processing that are presented by the introduction of very cheap microprocessors. These opportunities do not lie along the traditional lines of major developments in computing. Instead, they follow and enlarge upon the trends which were begun by the minicomputer.

The advantages and disadvantages of these new opportunities are discussed in detail, including the impact on various sectors of the computing industries, the threat to the centralised data processing function, and the concept of total system costs. The key role of microprocessor software is highlighted.

The report concludes with a practical guide for the management services department which wishes to make a constructive start in microprocessors without committing itself to a major expenditure of resources.

The Butler Cox Foundation is a research group which examines major developments in its field — computers, telecommunications, and office automation — on behalf of subscribing members. It provides a set of 'eyes and ears' on the world for the systems departments of some of Europe's largest concerns.

The Foundation collects its information in Europe and the US, where it has offices through its associated company. It transmits its findings to members in three main ways:

- As regular *written reports,* giving detailed findings and substantiating evidence.

- Through *management conferences,* stressing the policy implications of the subjects studied for management services directors and their senior colleagues.

- Through *professional and technical seminars,* where the members' own specialist managers and technicians can meet with the Foundation research teams to review their findings in depth.

The Foundation is controlled by a Management Board upon which the members are represented. Its responsibilities include the selection of topics for research, and approval of the Foundation's annual report and accounts, showing how the subscribed research funds have been employed.

The Butler Cox Foundation

Butler Cox & Partners Limited
Morley House, 26-30 Holborn Viaduct, London EC1A 2BP
Tel 01-583 9381, Telex 8813717-GARFLD

SISDOCONSULT
20123 Milano – Via Caradosso 7 – Italy
Tel 86.53.55/87.62.27

Akzo Systems B.V.
Velperweg 76, Arnhem, The Netherlands
Tel 85-662629

Butler Cox & Partners Limited
216 Cooper Center, Pennsauken, New Jersey 08109, USA
Tel (609) 665 3210

Printed in England