A Handbook of Information Technology

Mark Charlton

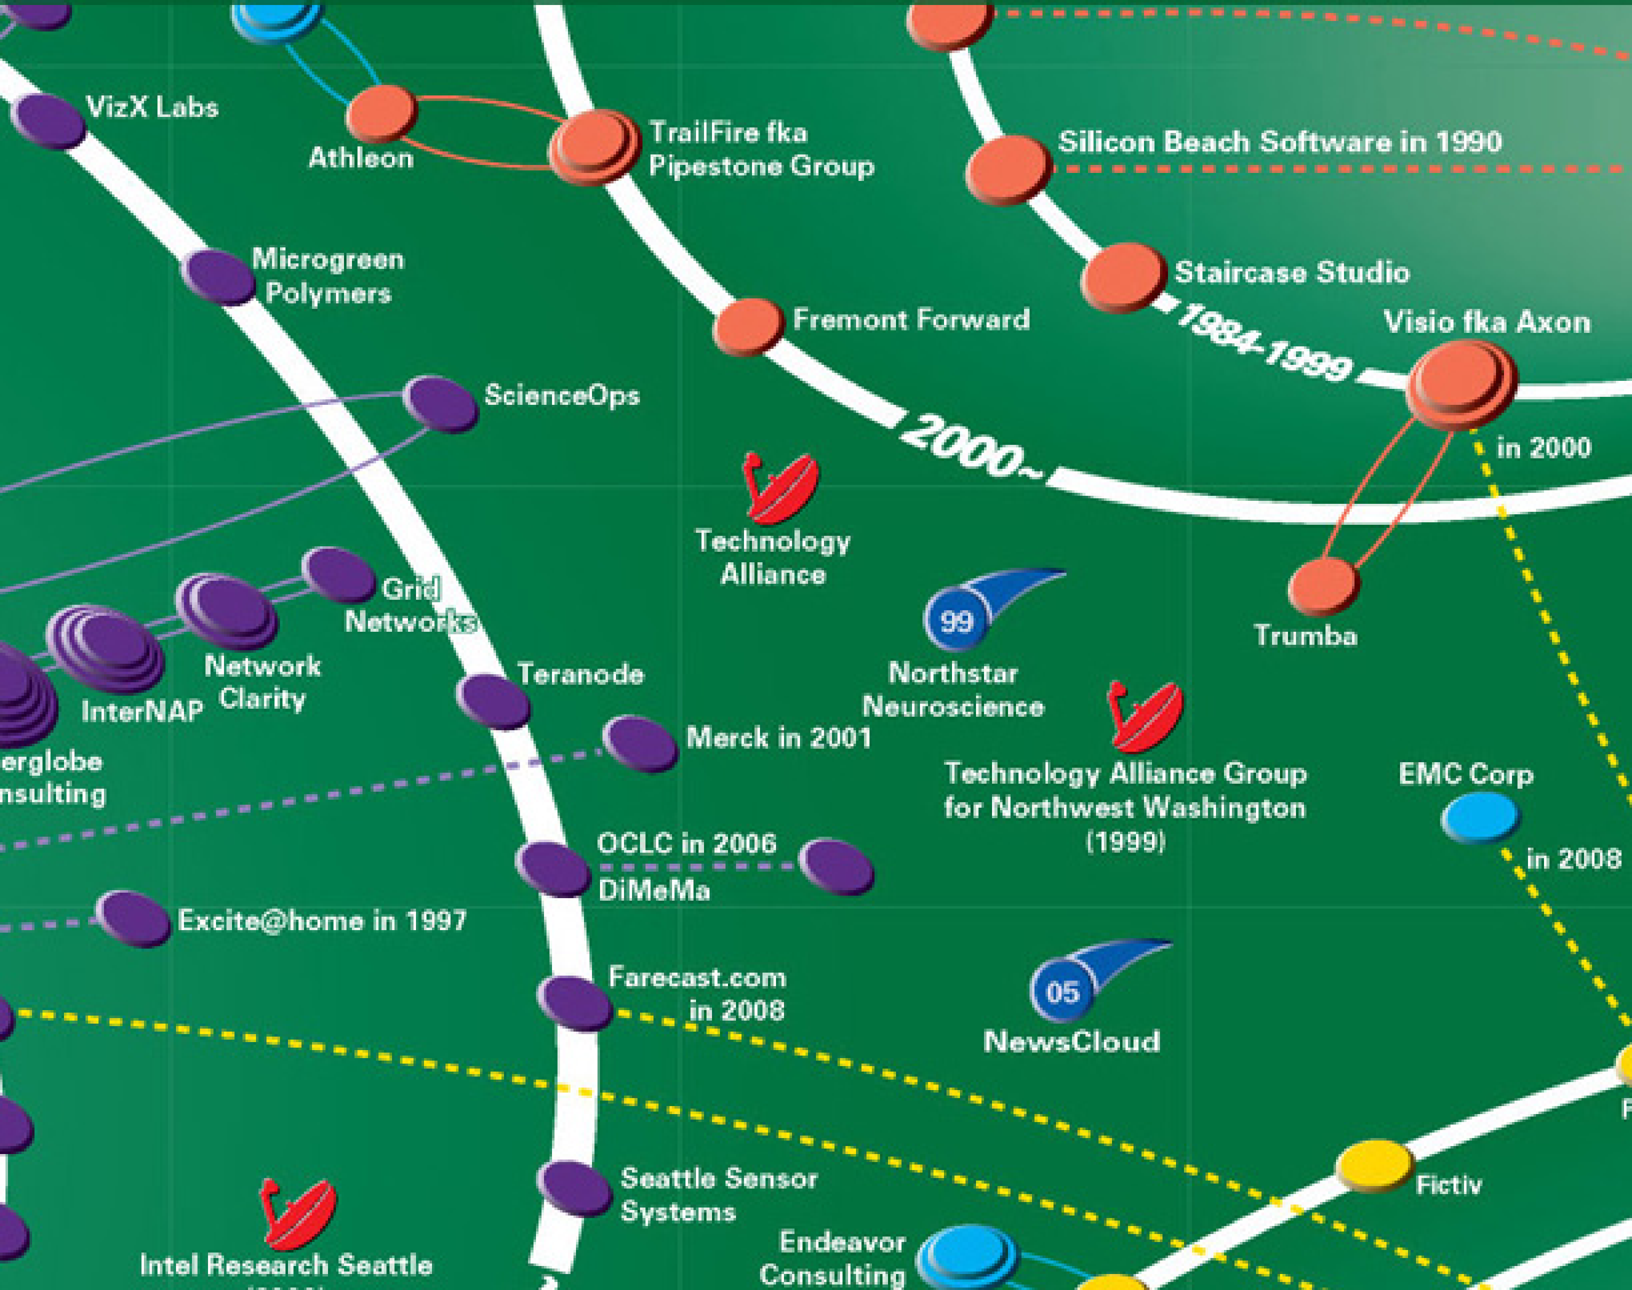

First Edition, 2009

ISBN 978 93 80168 55 5

Published by:

Global Media
1819, Bhagirath Palace,
Chandni Chowk, Delhi-110 006
Email: globalmedia@dkpd.com

Table of Contents

1. Introduction

2. What is Information?

3. Information, the Life Blood

4. What is Information Technology?

5. Organisation and Structure

6. Computers in Use

7. Computers for Individuals

8. Computer for Children

9. Academic Aspects

10. Community Service

11. Community Networking

12. Service Monitoring

13. Electronic Mail

14. Communication Network

15. Telecommunications

16. Various Networks

17. Model Open Systems

18. Future Prospects

19. Production of Computers

20. The Set of Instructions

21. Business through Internet

22. E-commerce in India

23. The New Generation

24. The Next Generation

ONE

Introduction

While mankind has developed myriad ways of applying and controlling power to dominate and shape our environment, through the use of — tools, weapons, machines, fuels, vehicles, instruments, clothing, buildings and roads, metals, plastics and drugs, agriculture, and electricity — the handling of information has lagged considerably, perhaps because the human brain is itself so remarkably powerful.

Until recently, there have been only three major developments in this area: the invention of written (or painted or carved) language, some five or six thousand years ago; that of simple arithmetic operations, using what would now be called a digital representation of numbers, about a thousand years later; and that of printing, about five hundred years ago. With written language, we get the capacity to make a permanent record of information and also to convey messages across space and time: storage, retrieval, and communication.

With digital arithmetic, we get the ability to perform accurate, repeatable manipulations of quantitative data. With printing, we can make many identical copies of the same record and so broadcast a single message to a wide and continuing audience.

Beyond these outstanding advances, until the last hundred years or so, the only progress has been in the engineering department with increasingly plentiful production of more powerful and reliable and efficient, faster and cheaper devices to implement these concepts. In the last hundred years, we have seen the rapidly accelerating advent of a technology so powerful, novel, widespread, and influential that we may indeed call it the Second Industrial Revolution. Its basis is electromagnetic, in many interconnected forms: photography, photocopying, cinematography, and holography; telegraphy, telephony, radio communication, radar, sonar, and telemetry; sound and video recording and reproduction; vacuum tubes, transistors, printed circuits, masers, lasers, fiber optics, and (in rapid succession) integrated circuits (IC), large-scale integration (LSI), and very large-scale integration (VLSI) of circuitry on a tiny semi-conducting 'chip'; and, finally, the bewildering variety of electronic digital computers.

The High-tech Revolution Vs Computers

The progress has been truly amazing. In only about 40 years, electronic communications and news media have become commonplace and indispensable. Computers have proliferated, becoming increasingly fast, powerful, small, and cheap, so that now there is scarcely a human activity in which they are not to be found, bearing an increasing share of the burden of repetitive information processing, just as the machines of the First Industrial Revolution have taken over the majority of heavy and unpleasant physical labour. Now, information can not only be stored, retrieved, communicated, and broadcast in enormous quantities and at phenomenal speeds; but it can also be rearranged, selected, marshalled, and transformed.

Until recently, these activities were the sole domain of the human brain. While creative, judicious, moral, and aesthetic choices are still best left to people, all the tedious and mechanical mental processes can now be relegated to the accurate, fast, and tireless machines. Any sequence of operations on information that can be precisely specified can be carried out without further human intervention or supervision.

At first, computers were the experimental toys of university researchers; then they became the tools of government establishments and giant corporations, huge, expensive, individually designed and manufactured, and beyond the reach of any but the wealthiest organizations.

People thought of the future in terms of machines of ever-greater speed and capacity; centralized behemoths would hold all the world's information in gigantic data banks, whence major decisions would issue to be imposed upon the populations at their mercy.

With the emergence of powerful, cheap, mass-produced computers-on-a-chip, the picture has changed radically. Now we see tiny computers everywhere: in wrist-watches, microwave ovens, electronic games, pocket calculators, cameras, typewriters, musical instruments, etc.

What used to be done, with few options, by intricate mechanical devices is now performed, with great flexibility and convenience and at much less expense, by the ubiquitous preprogrammed microcomputer.

The probable future has become one of millions of small yet powerful computers, controlling virtually every machine and appliance. These are distributed in every home, on every desk, in every workshop; many of them connected in a maze of small and large networks, much like the present telephone network. This enables individual computers to communicate, sharing information in a gigantic distributed data-base, and gaining, through distributed processing, computational power whose extent is yet difficult to gauge; all this following the individual requirements and choices of the owner or operator of each machine. Increasingly, we are confronted, not only with the results of the use of computers throughout industry, commerce, banking, advertising, science, the communications industry, newspapers, airlines, and hospitals; but with the realistic possibility of purchasing computer power for our own small enterprises, offices, and homes. This may be done in a variety of ways; but in all of them, the real cost of computation is constantly diminishing.

It is probably fair to say that the question of computerization is not "whether," but "when" and "how." We must choose whether to lease equipment or to buy it; whether to install terminals, connected to a computerized "service bureau," or a complete local computer system; whether to get a package of standard programs directed towards our kind of work, to commission new programs tailored to our special needs, or to learn programming and write our own; whether to go it alone or to share a system with a number of similar users (especially if they are in the same building); how far to take our first efforts at computerization; what to include and what to continue doing by hand. Computer programming is likely to become the literacy medium of the third millennium AD.

Elementary schools may well be teaching it before long, and we might be well advised to gain at least a smattering of knowledge of computers and of programming competence, especially since computer languages and programming environments are becoming increasingly helpful and friendly to the uninitiated user.

A computer is a machine for the automatic processing of information. Historically, this information was numerical, and computers were machines for doing arithmetic.

Unlike the simpler calculating machines, which can perform only one elementary arithmetic operation at a time, they need to be told what to do next (usually by suitable button-pushes); computers can be given a list of operations to perform (often with branching and repetitions, depending on tests of sign or value included among the operations), which they will then execute in proper sequence without further intervention. This sequence of instructions is called a *program*.

A digital computer stores its information in the form of words, finite ordered sets of digits, each of which can have only one of a finite set of values.

Considerations of simplicity, reliability and economy dictate that electrical engineers should design computers to consist of a great number of similar pieces of circuitry, each of which can only be in one of two states, usually denoted by 0 and 1.

Such *binary digits* (or bits) are the elements of which computer digital representation is built. A row of eight bits is called a *byte,* and the majority of computers have their storage of information organized in words of one, two, four, or eight bytes (8, 16, 32, or 64 bits).

The number of bits in a word is termed its *length;* if this is *k,* then the number of possible distinct pieces of information that can be stored in such a word is 2. In particular, four bits together can have 16 different contents, and these are standard binary representations of the numbers 0-15:

0000 = 0 0001 = 1 0010 = 2 0011 = 3
0100 = 4 0101 = 5 0110 = 6 0111 = 7
1000 = 8 1001 = 9 1010 = A 1011 = B
1100 = C 1101 = D 1110 E 1111 = F

with A = 10, B = 11, C = 12, D = 13, E = 14, and F = 15. These may now be viewed as the 16 possible digits of a representation (the hexadecimal, or hex), which is much more compact and humanly intelligible than a long string of zeroes and ones. For example, the byte 10110010 becomes "132," and the four-byte computer word 01001100011100101101000110001110 becomes the eight-digit hex word "4C72D18E."

The following are various parts of which any computer is composed. These are, essentially,

1. a *central processing unit* (CPU), which is the controlling and computing centre of the machine;
2. a *memory,* possibly of different levels, in which both *data* and *instructions* are stored;
3. a variety of *input* and *output (I/O) devices,* through which the machine communicates with the world outside it.

The CPU consists of an operation control unit (OCU), an arithmetic/logical unit (ALU), and a relatively small, very-fast accessible local memory (LM). The OCU keeps track of the memory location of the next instruction to be executed, and analyzes the current instruction, so as to activate the proper operation of a memory transfer, nonsequential jump (by appropriately changing the address of the next instruction), input or output of information, or computation (performed by the ALU), as is indicated by the instruction code.

The ALU actually carries out the elementary arithmetic operations (addition, subtraction or negation, multiplication, division or reciprocation) and logical operations (these being bit-by-bit operations, without carry' such as not, and, and xor; eg., not 1010= 0010, 1100 and 0101 =0100, 0110 xor 1010 = 1100) on the data given to it by the OCU. The LM receives the operands called for by the OCU and also the results of the ALU'S operations upon them. For example, the OCU may retrieve the factors of a multiplication from the main memory into a pair Of LM registers and instruct the ALU to multiply them and place their product in another LM register. Such registers are usually called accumulators, and they are normally double-length (since the product of two k-bit numbers is a 2k-bit number). Other LM registers are used for counting (eg. repetitions) and are called index registers; others hold intermediate values and are called buffers; and, finally, there are one-bit or two-bit registers which reflect the information on which tests are made by the OCU (for example, the occurrence of a carry, the vanishing or positivity of an answer, or the occurrence of an arithmetic overflow): these are termed flags.

Originally, the CPU was a sizable piece of electronics, hand-assembled and highly complex. With the advent of micro-miniaturization of circuitry, printing, and photographic techniques, and the mass production of components only the largest computers (mainframes) are built in the old way. Smaller systems generally have the entire CPU on a single chip. Among these, the name microcomputer is now applied to those with less than a million words of memory and a word length of one or two bytes; the name minicomputer applies to the larger machines, with two to four-byte words and one to a hundred million words of memory. (The smallest micro is probably more powerful than the big computers used by universities and industry in the 1950s.)

The *main memory* (mm) consists of magnetic or electronic components which store the information (both data and instructions) needed by the computer. The individual words are directly addressable from the CPU by number (rather like houses in a street), and their contents are retrievable in very short times, of the order of the operation time of the CPU (ranging from fractions of a *nanosecond,* 10^{-9} or one billionth of a second, for the fastest mainframes to several *microseconds,* 10^{-6} or millionths of a second, for the slower micros). This is often referred to as high-speed storage or random-access memory (RAM). While most of the mm is erasable and may be changed at will, some memory is used to store constants and often-used utility programs and is not erasable by the CPU: such memory is called *read-only memory* (ROM). Sometimes this is optional and can be plugged into the computer : this is called *firmware.*

Note that a computer with one-byte-long addresses can have at most 28= 256 words of mm; a two-byte address can reach 256^2 = 65536 words; a three-byte address can select any of 16777216 words of mm; and so on. In practice, most micros and minis have mm ranging in size from 2 14= 16384 to 218= 262144 words. It should be noted that 10^3= 1000 and 2^{10}= 1024. Because computers are so heavily slanted towards base-2 representation, it has become almost universal computer parlance to use the prefix *kilo* or K (which usually denotes a thousand units) to denote 1024 and the prefix *mega* or M (which usually denotes a million units) to denote $2^{20} = 1024^2$ = 1048576. Thus, we write 16K for 2^{U} and V4M for 2^{18}. Perhaps the commonest size of MM consists of 64K = 65536 words.

Almost all computer instructions comprise an operation code (usually one byte long, allowing 256 possible operations to be specified), followed by an operand reference (number, index, or address) of variable length (since some operations require more data than others; for instance, the STOP instruction needs no operand, so it is one byte long).

The *extended memory* (EM) is sometimes considered under I/O devices, both because it is often physically located outside the computer (while the CPU and the mm are usually in the same enclosure), and because its speed of access is much slower than the speed of operation of the CPU and is comparable with the range of speeds of I/O devices. Most read/write memory is magnetic (optical memory is read-only, and fast RAM is either magnetic or electronic), either in the form of tape, drum, or disk, coated with magnetic material, much like sound recording equipment; it is similarly erased, read, and recorded upon by "heads" which contain electromagnetic sensors/ polarizers. The cheapest (and most limited, in both speed and capacity) device is a common portable cassette recorder, with 1/4" tape cassettes. From this, one moves to specially engineered cassettes and recorders, and thence to high-speed 1/2" or I" reel-to-reel drives carrying thousands of feet of tape at very high speeds. Access times can be quite good for sequential access, along the tape, but random access time is poor at best, running to seconds or even minutes. Economy and a virtually unlimited total storage capacity (on numerous cassettes or reels; but only as many units as one has on-line tape-drives are actually accessible without human intervention) are the only advantages.

Disk Memory

When we wish for practically useful EM, combining large capacity with relative economy and speed of random access, we must turn to drum or disk memory; and, nowadays, the former have been practically replaced by the latter. Disk memory is of two types: floppy disk and hard disk, the first being the cheaper, slower, smaller-capacity option. Floppy disks are flexible, have diameters of 51/4" or 8", generally, and are removable from the disk-drive, so allowing one to build up an unlimited library of stored data. The information is stored on concentric circular tracks (not on a single spiral track, as on a sound record), on one or both sides of the disk. The number of tracks and the number of bytes per track vary (the density increasing with precision of engineering, and so with cost of the drive), but the total capacity of a floppy disk is in the range of 50KB to 1MB. The disks rotate at, typically, 300 rpm, and access time is governed by the time required to place the movable head on the right track, a fraction of a second, plus the fifth of a second taken by the head to traverse the circumference of the track, in search of a record; thereafter, consecutive bytes are accessed at some thousands per second.

Hard disks are rigid and have larger diameters. There are drives which range from anything from one to a dozen disks, rotating at about ten times the speed of floppy-disk drives (and so diminishing the access time of records in a track) with one or several heads.

Fixed-head drives naturally must have a head for each track (which costs more), but save head-movement time in random access. Winchester disks are movable-head drives with sealed-in disks, where the heads ride very close to the disk, cushioned by the layer of air between. In floppy-disk drives, the head actually rides on the disk, eventually wearing it out.

The capacity of hard-disk drive ranges from 10 MB to 100 MB in a single drive. Some movable-head hard-disk drives have removable disks or disk-packs, allowing for greater library storage.

The I/O devices are the computer's link with the outside world. In large mainframe computers, we see *paper card* (as in Hollerith or "IBM" cards) and paper tape readers and punches: the perforations in the paper carry the information. Increasingly in large computers, and almost universally in small ones, the main input is from the keyboard of a terminal. This is much like a typewriter keyboard, and depressing any key sends an 8-bit code to the computer.

When the computer is waiting for input from this terminal, it reads the code and interprets it as a datum; when it is busy, either the message is lost or it is held in a "buffer" for subsequent input (this depends on how the connection is made and what the computer is made to do). It is quite common for the computer to be connected to several terminals, all competing for its attention. This is called time-sharing. The computer cycles around the terminals, look for their several inputs while dividing its CPU time among them. The main output of the computer is to the display devices of the terminals; these are either video displays (cathode ray tubes, CRT, just like the screens of black-and-white or color TV sets; indeed, simple micros sometimes use ordinary television sets as display devices) or printers (in so-called hard-copy terminals). Of course, the computer may be connected to additional video displays and printers, of different qualities, as well as to plotters, a kind of printer for drawing graphs and diagrams. Many types and speeds of printers exist. Other input devices, such as audio amplifiers (receiving signals from radio tuners, record players, etc), video receivers and recorders, laser disk drives and a variety of scientific instruments, can all be classified as transducers (devices that transform physical quantities, such as position, conductivity, pressure, temperature, vibration frequency, or amplitude, into electromagnetic impulses) linked to digitizers (which convert such impulses into sequences of zero/one pulses). Output from the computer can similarly follow the reverse process, yielding visible or audible results, or the control of mechanical or electrical equipment. Thus, computers can draw pictures (often, moving pictures), make music and other sounds, and can control appliances, machinery, and whole manufacturing processes.

In order to connect remote terminals to a computer, use is often made of telephone lines; and, for this purpose one employs a device called a modem (for "Modulator/DeModulator"), which converts the computer's digital signals to and from telephone audio signals. This has a cradle shaped to hold the ear and mouth pieces of a telephone's hand-set. It is also possible to connect several computers in this way. This is called the formation of a computer network. Of course, computers may also be connected by cable, fibre-optics, or microwave link. It should be noted that terminals usually, and other I/O devices often, themselves contain computers of varying degrees of power and complexity; so that even a multiuser computer with a central CPU may still be seen as a kind of computer network in its own right. The computers in I/O devices such as terminals or printers are often referred to as peripheral processors (PP).

Many networks do not have a central computer at all; but are simply a collection of independent computers linked for the sharing of information and, sometimes, computing capabilities. Often, they permit the exchange of messages *(computer mail)* and the pooling of data *(distributed data-base).* They may also share a common bank of memory, accessible to all.

Finally, since the invention of computer networks, designers have been investigating the possibilities of computers made up of an array of CPUs *(multicomputer, parallel processors, distributed processing).* These new ideas are very powerful and farreaching: they will probably revolutionize our ideas of computers, and of their applications, in the next few years.

The various peripheral devices are connected to the CPU without examining how. In fact, this may be done in several ways. We can simply have a separate connection (or port) for each device, but this limits rather severely the number of devices that may be connected to the CPU. Another way is to have a single bus (or connection) to which any number of devices may be attached.

The information signal must then carry an appropriate address. The bus receives all signals, and individual devices (including the CPU) seek out and decode only those addressed to them. It is also possible to have a switching device, which receives addressed data and directs them to the appropriate recipient device, rather like a central post office.

The decision on what communication arrangements to adopt is made on the basis of considerations of cost, capacity, and speed. It is often the case that the several devices forming a computer or a computer network have their data coded in different ways.

It is then the job of the CPU(s) and PPs to share the work of interpreting signals into appropriate codes for each machine. This is broadly termed the problem of interfacing devices. Sometimes, the solution is to have a standard code or structure for the communications device (one meets the s-100 bus, the rs-232 serial port, the ASCII character-code, and so on).

Another interfacing problem arises from the difference in the rate at which different devices can send and receive information (this is measured by the baud rate, named after Baudot, the inventor of the first five-hole paper-tape code; one baud is one bit transferred per second; hence kilobaud, kb, and megabaud, Mb typical rates range from 200 baud to 200 kb). One solution is to send each piece of information (usually one character at a time, which takes 8 to 10 bits) only when the last has been acknowledged (this is referred to as a handshake); this is sure, but slow. Another way is to use a storage buffer in which a large batch of information is accumulated for fast transmission, thus not wasting the time of the faster device.

One last kind of choice must be mentioned: some channels of communication are serial (they transmit one bit at a time), while others are parallel (they can transmit a byte or a whole word at a time); the latter are obviously faster, more complex, and more expensive. When devices are connected by cable, the degree of parallel communication is exhibited in the width of a flat ribbon cable, carrying several wires, side by side, and in the number of pins in the plugs and sockets by which they are connected to the machines. Parallel transmission is a variation on multiplexing. What we have described is generally referred to as the hardware of a computer. By contrast, the programs that make the computer work, the "soul" of the machine, as opposed to the hardware "body," are collectively called its software. Inevitably, there came to be programs that were hard-wired (in the now outdated phrase) into the computer, in the form of ROM. These are termed firmware.

A computer without software is a helpless set-of circuits, and the expertise required to create the basic software that will bring the machine to useful life is comparable to that required to design the machine itself. Indeed, these days, computers are designed in cooperation between computer architects, who design what the computer will do, hardware engineers, who design how it will be constructed to be able to do it, and software engineers, who design and program the operating system that will run the machine. Beyond this, the computer will also need application software of many kinds, to enable it to do a variety of jobs, such as file-handling, accounting, statistics, payrolls, inventories, complex graphic displays, games, and so on. Typically, the application software is written (i.e. programmed) by the endusers (if they are sophisticated enough) or firms of consultants and programmers and system analysts, often called software houses. Sometimes, a software house will produce part or all of an alternative operating system, or an addition to an operating system, to make a computer more flexible or more efficient than the manufacturer's own software system allows it to be.

Computer Languages

The CPU of any computer is designed to accept and execute a specific set of *operation codes (op-codes),* ranging in number from a dozen or so to several hundred. Different makes and models of computers and microprocessors may have entirely dissimilar op-codes; but the operations that they represent are much more alike than different, both through functional necessity and historical development.

The interpretation of the op-codes is built into the hardware of the OCU (though sometimes the details of interpretation may be modified by the user through what is called *microcoding),* and it is part of this interpretation that the complete instruction being decoded contains a certain amount of further information (such as parameters, indices, and one or more memory addresses).

The "length" of a computer instruction (the number of consecutive memory words, or of bytes, occupied by it) may vary, but the OCU automatically adjusts to this. An executable program (sometimes called "object code") consists of a sequence of machine instructions, consecutively stored in the computer's memory and (with the exception of "jump" instructions) executed in the order in which they are stored.

The aggregate of possible machine instructions is called the machine language. In the computer, a program consists of a long string of binary digits (bits), usually written as 0's and 1's; and, of course, the same program would be interpreted quite differently (usually as nonsense) by any computer for which it is not designed. Indeed, the slightest error in a program almost always leads to an error in its output (usually a fatal error!). This state of affairs is sometimes expressed by the computer adage, "Garbage in; garbage out," or just "GIGO." Errors in programs are called bugs; and the tedious, odious process of finding and correcting such errors is termed debugging. It is estimated that, in the production of a working program, the debugging time may be two to four times as long as the time it takes to plan and write the program initially.

To give the reader a feeling for the nature of machine language, we present a simplified, fictitious, but typical, machine language specification. Our computer has two 16-bit accumulator registers (acc), X and Y, which may be coupled into a single 32-bit ace XY, with X holding the more and Y the less significant digits; these are attached to the ALU; and a program control register (pc) Z, also of 16 bits, attached to the OCU, which contains the address of the next instruction to be executed. Instructions will contain a 2-bit acc code a, with a = 0 referring to X, a = 1 to Y, a = 2 to XY, and a = 3 to Z. The memory consists of 216 = 65536 16-bit words, directly addressable with a 16-bit address, n.

There are four addressing modes, denoted by a 2-bit mode code d, with d = 0 referring to absolute address (address *n*\0 refers to the actual number *"n")*, *d* = I to direct address (address *n*\1 refers to word n in the memory), *d* = *2* to indirect address *(n*\2 refers to the memory address which is the content of the memory word with address *n)*, and *d* = 3 to relative address (n\3 refers to the memory address that is the sum of *n* and the content of the pc register Z, with any carry to the seventeenth bit ignored). We will write {{*a*}} for the acc denoted by a, *(n\d)* for the memory word with address *n* in mode *d*, and C[*x*] for the content of word or acc *x* (so, e.g., C[{n]2}]=C[C[n]], while C[{n\3}]=C(n+C[Z]). *A port code p* selects one of 8 *output ports (p=0,* 1, . . ., 7) and 8 *input ports (p = 8,* 9, . . . , 15), these ports transmitting 16 bits at once and having to be reactivated before each function (i.e., each output instruction sends out one 16-bit number, and the machine must wait for the output to be acknowledged before making another output; while each input instruction reads one 16-bit number, if an input is available, and the machine must wait for a new input each time). Finally, a bit code b refers to each of the 16 bits in a word (bit 0 being the least significant-rightmost and bit 15 the most significant-leftmost). In some cases, the codes a, d, p, and b are interpreted somewhat differently, depending on the particular op-code c.

As an example of a very simple computer program, we consider the solution to the following computer problem. Our computer is to be fed a sequence of one thousand 16-bit numbers at input port #9 (these may be keyed in by hand or fed in by a digitizer connected to some experiment). They are to be stored in memory words with addresses 5000, 5001, 5002, . . . , 5999. Their sum is to be computed and stored in address 6000 and output to a printer through output port #2. The program is to be stored beginning at address 0 in the memory.

Again, what is important about this example is not its detailed form, but the difficulty of its interpretation, and therefore also the difficulty of verification and debugging.

The programmer must deal with a mass of details that are of a purely mechanical nature and have no relevance to the problem being solved. "Higher-level languages" are attempts at making the computer understand the programmer's way of thinking, rather than forcing the programmer to think like a machine.

Higher-level languages

After some introductory description of an imaginary computer, and especially of its CPU, and the establishment of some essential notation. Presented with a set of arithmetic and other transformations required by potential users of a proposed new computer, the electronic and logical design engineers seek the simplest circuitry that will execute operations sufficient to generate all the required transformations. Circuits must be simple to be fast, efficient, reliable, and cheap. However, when an instruction has 16 or more bits, most of which should for the sake of efficiency, have some significant effect; the exact and complete specification or the action induced by it may well be somewhat forbiddingly intricate! And indeed, simplicity of circuitry does not usually lead to simplicity of use.

An ample of a program in this machine language, which takes 16 instructions and 28 16-bit words of memory storage, to read in one thousand numbers, store them in consecutive memory locations, sum them, and print out the sum. However, please note that the English description of what the program does (though much easier to say and comprehend) itself requires 103 characters, each equivalent to 8 bits; so that it is almost twice as long as the machine-language program. Indeed, the difficulty of the latter lies not in its length, but in its alien and opaque form of expression of what it does. The first improvement on machine language was assembly language. This is not very different from machine language, but removes the most glaring and trivial irritations. First, instructions may be labelled, so that one no longer needs to count lines to address jump instructions; one jumps to the label. This is also a great help when one wishes to insert additional instructions, either to correct or modify a program (every jump need not be changed). Second, memory may be labelled, so that numbers may be stored in symbolic addresses and retrieved therefrom, leaving it to a computer program to assign actual storage. Third, the entire program becomes relocatable anywhere in the computer memory. Finally, and perhaps more trivially, the op-codes may be replaced by abbreviated mnemonics, and the instruction layout may be relaxed, using punctuation marks. Of course, a program written in assembly language is no longer directly intelligible to the computer or executable by it. It becomes a piece of textual input, to a "translator program" called the assembler. Naturally, the program may be stored in the computer's memory, like any other piece of text; but, before it can be executed, it must be "assembled." The assembler finds appropriate space for the object code (i.e. the machine language equivalent of the input program) and for its needed storage space (often called variables). The assembly language instructions correspond, one-to-one, to machine-language instructions.

For example, we may denote memory locations by one, two, or three lower-case letters (both for labelling instructions and naming variables). A numerical address will still be allowed. Absolute addresses (i.e. actual numbers; d=O) will correspond to quoted numbers (e.g. "6000"); direct addresses (d=l) to variables named as stated; indirect addresses (d=2) to variable names placed in parentheses; and relative addresses will not be used. With relocatable programs and symbolic addresses, they are no longer useful). Accumulators

(X, Y, XY, or Z) will not be referred to symbolically, but only by these names. Instruction labels will precede the instruction and will be terminated by a colon (:). Spaces will be ignored. Op-codes will be denoted by strings of capital letters and terminated by a comma (,), which will also separate multiple arguments.

A sample specification of assembly language op-codes is shown in Figure 1. In the assembly language notation, A denotes X (a=0), Y (a= 1), XY (a=2), or Z (a=3). Italic letters (p, x, or b) denote numbers (possibly in quotes, to denote absolutes) or variables names (possibly in parentheses, to denote indirect addressing), and are interpreted according to the op-code (e.g. A cannot be XY in IN and OUT instructions; if A is XY in instructions with c = 1, 2, and 4 (with p > 5), the operand must be 32 bits long-either a double-length number, or an address referring to the first of two consecutive 16-bit memory words. Again, the reader must remember that it is not the particular details of our specifications that are important or memorable but the kind and degree of detail occurring.

Assembly language mnemonic	*Machine language instruction*
IN A, (p-8)	*c* = 0, *a* ** 2, *p* >*8*
OUT A, *p*	*c* = 0, *a* ** 2, p < 7
TO A, x	*c*= 1
FROM A, x	*c* = 2, *d* > 1
FROMA, B	*c* = 2, *d* = 0 -
CL, A	*c* = 3, *d* = 0
NEG A	*c* = 3, *d* = 1
NOT A	*c* = 3, d = 2
RNDx	c = 3, *d* = 3, a = 2
ADD A, x	c = 4, p = 0
SUB A, x	c = 4, p = 1
AND A, x	*c* = 4, p = 2
OR A, x	c = 4, p = 3
XOR A, x	c = 4, *p* = 4
LADD, x	c = 4, p = 5
LSUB, x	c = 4, p = 6
LAND, x	c = 4, p = 7
LOR, x	c = 4, p = 8
LXOR, x	c = 4, p = 9
CSHL A, P	c = 5, *d* = 0
CSHR A, P	c = 5, *d* = 1

(Contd...)

(Contd...)

ASHL A, P	$c = 5, d = 2$
ASHR A, P	$c = 5, d = 3$
JMP, x	$c = 6, p = 0$
JPCZ A, x	c = 6, p = 1
JPCNN A, x	c = 6, p = 2
JPCNZ A, x	c = 6, p = 3
JPCP A, x	c = 6, p = 4
JPCB A, 0, *b*, x	c = 6, p = 5
BIT A, 0, *b*	$c = 6, p = 6$
BIT A, 1, *b*	$c = 7, d = 1$
BITNOT A, *b*	c = 7, d = 2
STOP	c = 8
INTER	c = 9

Figure 1: Typical assembly language mnemonics

Consider Figure 2. Here, numbers are given in decimal notation. Note the tremendous improvement in direct legibility of the program in comparison with the binary strings. The assembler is left to figure out where to store the sum (the variable "sum" which the machine language program put into address 6000) and the current address of the listed numbers (the variable "w," which the machine language program puts at location 28 if the program begins at location 0, necessitating counting the words occupied by the program; if the program were to be modified or relocated, this would have to be changed, and would be an unnecessary source of programming errors); and similarly, there is no need to count words to determine the addresses of "ret" and "out" (placed by the machine language program at 5 and 24 ' as it happens). As an exercise, the readers may wish to try to modify the program so as to allow for the possibility that the sum of the 1000 numbers might occupy more than 16 bits. In machine language this requires no less than 12 changes and seven shifts of memory location.

```
      CL X; FROM X, SUM; TOX, '5000'
ret:  FROM X, W; SUB X '6000'; JPCZX, out
      INX, 1; FROM X, (W); ADD X, sum
      FROM X, sum
      TO X, W; ADD X, T; JP, ret
out:  TO X, sum; OUT X, 2; STOP
```

Figure 2: The summation program in assembly language

In assembly language only the substantive changes (from the programmer's point of view) need be made. The amended program is given in Figure 3. Here we note that only seven changes are needed; but observe that "sum" becomes a double-length address-just another piece of bookkeeping handled automatically by the assembler.

```
        CL XY; FROM XY, SUM; TOX, '5000'
ret:    FROM X, W; SUB X '6000'; JPCZX, out IN X, 1; FROM X, (W);
        TO XY. sum ADD XY. (w);
        FROM XY, sum TO X,W; ADD X, T; JMP, ret
out:    TO XY. sum; OUT X, 2; OUT Y. 2: STOP
```

(Underlined material represents changes from previous program).
Alternative version would have third line replaced by:

```
        CLX: IN Y 1, FROM Y (W); LADD, sum FROM XY sum
```

Figure 3: The summation program modified for long sums

Of course, a program in assembly language still represents, instruction-by-instruction, a program in machine language, and therefore still has the defects of "walking the road in extremely short steps," as we put it earlier.

All that shuffling of numbers in and out of the accumulators, counting and restricted testing, is not natural to us, and its mechanical nature suggests that an improvement is still possible. Thus it was not long before programmers devised much more humanly natural languages, which are generically called *higher-level* (or *algebraic) languages.* Some of the more common ones are ADA, ALGOL, APL, BASIC, C, COBOL, FORTRAN, LISP, PASCAL, PL/I, RPG-II, SNOBOL, SPITBOL, WATFOR, and WATFIV. There are more; and all come with a variety of versions and dialects.

The characteristic that puts a language into this class is that one higher-level instruction translates into several assembly or machine language instructions. Beyond this, the languages differ according to the kind of program they are intended for.

While most languages will perform most tasks, FORTRAN was clearly intended to do scientific and engineering computations, COBOL to perform business data-processing, LISP to manipulate lists, and SNOBOL, strings of text. As a final illustration of the power and intelligibility of such languages, we present higher-level versions of our summation program in BASIC, WATFIV, and PASCAL. Once more, note that the details of the individual languages do not matter to us at this point (though, if the reader ever intends to program, then some language will have to be thoroughly understood; but assume that many of my readers have no such intention); it is their general flavor and appearance that is noteworthy; and the casual readability of the samples presented is their salient characteristic.

Of all the hundreds of programming languages devised, it is safe to say that more lines of program have been written in FORTRAN or COBOL than in all others combined, though this reflects their age rather than their desirability. The fast-gaining runner-up must surely be BASIC, which comes with every micro and most minicomputers; though some would say that BASIC, which perpetuates the style of FORTRAN, spoils any chances of producing really good programs and should be replaced by PASCAL. Fashions come and go in programming languages, not always fully based on rational arguments, and personal preferences vigorously touted.

Given a higher-level language, it is necessary to have a translator program to turn it into machine (or more frequently assembly) language. Such translators are of two kinds.

An interpreter operates at execution time: beginning with the first instruction of the stored higher-level program, it translates it into one or more machine language instructions, which it proceeds to execute; then it reads and translates the next higher-level instruction and executes that; and so on. If a jump is encountered in the source program (as the higher-level language program is often called), the interpreter goes to the next instruction in accordance with the jump and proceeds to read, translate, and execute it, without taking into account whether it has already encountered it. Thus, the same source instructions will have to be passed and translated a thousand times (a very time-consuming and wasteful procedure). But the (usually much longer) object program need not be stored, and interpreters themselves tend to be shorter and simpler programs. This is why a BASIC interpreter is a natural adjunct to a microcomputer, which tends to be a little cramped for space; especially since many microcomputer users are riot interested in extremely lengthy computations.

By contrast, a *compiler* is a program that translates a source program into an object program, the former being in higher-level language and the latter in machine or assembly language.

The program is translated as it sits in memory, not during execution; so that the problem of repeated translation is avoided. Indeed, there are many so-called "optimizing compilers" which pass through the compiled object code several times, eliminating redundancies and inefficiencies to produce faster object programs (something that would be impossible for an interpreter to do). Speed is greatly increased, in general, at the expense of space in the computer memory. Of course, once a program has been compiled and proven to be free of bugs (debugging is far easier in source code than in object code), the source program may be stored for future use and only the object program kept in main memory. Compilers are available for almost all higher-level languages on almost all machines, and are essential for extensive applications.

Computer Software

The programs, or software, available for computers fall into two classes, *system software* and *application software*. The system software comprises those programs that are considered indispensable to the general operation of a given computer system, forming what is often termed the operating system of the computer, and generally are supplied by the manufacturer, though alternative operating systems are sometimes available from software houses. The applications software includes all those programs that are needed by one computer user but not by another, even if such programs are widely required, and while some manufacturers will sell applications software (as optional additions to their systems), more often this is obtained from software specialists.

We begin with the *operating system.* This may be divided into a kernel of absolutely indispensable programs and a shell of almost indispensable, so-called utility programs. The kernel (or nucleus) is also sometimes termed the monitor. To understand why this kernel program is needed, we consider what happens when we type a character (say an "A") on our keyboard, as part of a message to the computer. Because our timescale recognizes perhaps 1/20 second, while the computer may deal in units of 1/1,000,000 second, if we want to type an "A" we may find ourselves actually transmitting some 50,000 A's to the computer, before our finger leaves the key! If the keyboard mechanism were to transmit the A for only one microsecond, on the other hand, the computer might will be busy elsewhere and miss it altogether. Therefore, it is necessary to establish a protocol (in the usual diplomatic sense, a formal structure for the orderly passage of information) or "handshaking" procedure. This is often done by sending the A on one line (or lines) and simultaneously sending a pulse along an auxiliary line, which sets a "flag" bit in the CPU. When the computer is ready to receive a new character, it checks the appropriate flag bit over and over, until it detects that it has been set. Then it reads the character, resets the flag to its "null" state, and sends a signal pulse back to the keyboard (or its controlling circuitry), telling it that the character has been received. The flag is always open, so the transmit pulse need only be sent once by the keyboard; but the character continues to be sent until it is acknowledged by the computer. Thus a keyboard character will neither be missed nor read repeatedly. In addition, either the CPU or the keyboard controller should send the "A" to whatever device displays the typed characters.

If the particular keyboard were to be the only input device ever to be connected to the CPU, it would be best to incorporate all this in the circuitry of the CPU; but since there may be several different inputs to the CPU, and several types of terminals may be used, coming from different manufacturers, it is found preferable to have the I/O procedures made part of the operating system. In addition to the keyboard input already described, there will be short machine language programs for input from any other devices attached to the computer and corresponding programs for output to display screens, printers, and so on.

The orderly operation of a computer is subject to interruptions (or *interrupts,* as they are known in Computer jargon), either because a program contains an instruction that at some point, is not executable (such as a division by zero or a reference to a nonexistent memory location), or because someone needs to terminate execution for some overriding reason (such as the decision that a program is in error and is generating garbage or is in an infinite loop). Handling errors and interrupts is another function of the kernel of the operating system, as is the start-up (or "bootstrap") procedure which initializes operation. If, as is often the case, several users are connected to the computer, then the operating system must handle the tasks of Job-scheduling (and Job accounting and billing, if this is appropriate), and the allocation of storage, in main and extended memory, and of other resources (such as printers or communication lines), and the management of *time-sharing* (as between several terminals). Another task that is handled by the operating system is the management of *user files,* and their transfer between main memory and extended memory (such as disk). Indeed, one sees frequent reference to *disk operating systems,* such is the importance of this function of the operating system.

Another function of the kernel is to provide protection and security to users and to itself, both from authorized users encroaching on forbidden territory and from unauthorized users attempting to use the computer.

Ideally, each user should, on giving the correct access code (such as a "password")—this feature did not exist on the first computers, and is still absent in personal microcomputers with only one user/owner-who has access strictly limited to the parts of the computer allocated to him or her by the operating system, but in such a way that the impression is maintained that he or she is the only user present.

Another very congenial aspect of this concept is that some operating systems conduct their memory management function in such a way that the user need make no distinction between main and extended memory; this is referred to as a virtual memory system. The kernel program has to be able to display, move, modify, and search, at least the main memory, and to initiate execution at any given address, or activate any of the peripheral devices. Using these functions, by means of appropriate commands in the "operating system language," together with its file-handling capabilities, the system can load, compile, or execute any program stored in its memory (main or extended) written in machine, assembly, or higher-level language, provided that a suitable compiler or other translator is available to it.

We now turn to the utilities provided by most operating systems. First, we have an assembler (and perhaps also a *disassembler,* which translates a program written in machine language into the more intelligible assembly language; as well as a *macroassembler,* which allows the user to define his own macroinstructions in machine or assembly language), together with a selection of translation programs, either interpreters or compilers, for the higher-level languages that the user wishes to employ.

The smaller microcomputers will provide an interpreter for some dialect of BASIC, since this a relatively simple language to learn and to interpret into machine or assembly language; beyond this, one must pay for additional languages. There may also be a variety of TRACING and DEBUGGING utilities, according to the cost, size, and sophistication of the computer and its operating system.

There will be facilities for linking or chaining programs together. While the kernel will contain the rudiments of a file-handling system (to create, destroy, list, locate, and transfer files), there will also be utilities for further management of these. Such programs will sort, find (among other files), search (in a given file, for information specified), transform, edit, and combine files. Indeed, a good program editor can enormously facilitate the rapid and painless creation and modification of files, which include both text and programs.

The line separating the kernel from the shell of utilities is purely conceptual and far from sharp. Similarly, the boundary of the entire operating system is far from definite. What is available beyond the essentials mentioned above will be called part of any decent system by some, part of a compiled language by others, and just applications software by still others, depending on their point of view. The ability to handle a variety of data structures (such as arrays, lists, strings of characters, trees, queues, and stacks) may come from the use of a suitable higher-level language or from an extended "system development" utility package.

Similarly, languages intended for scientific and engineering applications usually handle floating-point and multiple-precision arithmetic, though this may be provided even in computer hardware. The same applies to routines for computing, for example, sines, cosines, logarithms, etc., and pseudorandom numbers (the last beloved of, computer games inventors!).

A *data-base management system* (DBMS) may be just a glorified file-handling utility; or may be an elaborate program, cross-indexed and relational, with its own language of special commands, for answering any conceivable question about a large amount of intricately structured data.

As is often the case, you get what you pay for, both in money and in memory space; and vendors' claims have to be carefully scrutinized and verified. It is advisable not to overbuy one's capabilities beyond one's needs.

Similarly, a graphics package may simply allow one to produce passable graphs, histograms (i.e. bar-graphs), pie-charts, and perhaps games, with a resolution of some 200 by 300 delightfully named pixels; or it may allow you to resolve perhaps 2000 by 3000 pixels, in a variety of colors, with the ability to draw complex three-dimensional shapes, properly shaded, illuminated, and textured, and move them by commands in a special language. The latter systems are a lot of fun to operate, and can be most helpful to draughtsmen, animators, film designers, and simulator-trainer designers; but they cost a bundle, and they require quite powerful computers to hold and run them.

Another offshoot of the file handier is a cluster of programs for generating reports, journal ledgers, and accounts of all kinds, as well as forms and mailing' lists (the latter leading to a kind of DBMS in which mailing lists are matched to interests and characteristics of the individuals itemized).

Again, there are statistical packages of various degrees of sophistication, from a mean-variance-covariance calculator, to systems able to analyse very complex sets of data by elaborate techniques, using a whole statistical computer language. What has been described as "the most popular program ever written" is usually given the generic name of an *electronic worksheet,* though the trade name of the first such program, Visi-Calc, like Kleenex and Vaseline, has become almost generic.

A table is presented on the video screen and each entry is allocated either a numerical value or a formula relating it to other entries. When the data are sufficient, the resulting numbers are displayed. When an entry is changed, it and all entries depending on it are altered accordingly at once. This provides a representation of a given situation surpassed only by a graph in its impact, and a first-class planning aid. Finally, among the borderline system/ applications software utilities, will put what is usually called a word-processing package (though would prefer either character or text processor). This is an extension of an editing utility, in which the text may by "scrolled" up and down on the screen, edited in the usual ways, but also more specifically for producing letters, articles, reports, and other written copy.

Margins may be set and the text right, left, or double justified, or set up in multi-column pages, with page numbering, indexing, and even the use of different sizes and styles of typeface, in the most sophisticated systems. Here again, cost rises steeply, and one should buy only what one will need.

Beyond these programs, there lies an endless variety of unquestionable applications programs. There are programs for ballistics, boat design, analysis of molecular structure from x-ray diffraction data, tabulation of Bessel functions, simulation of naval battles and economic cycles, etc.

Most of these are not in the market and only work on one machine; but there are very many programs available in the open market, produced or distributed by software houses, with various levels of efficiency, and sophistication. *Programming* is a fascinating, intricate, rewarding, but unforgiving and at times infuriating occupation. We will encourage all of you who have the time to try it. If you have a computer at your disposal, it would be a shame not to learn a simple language, such as BASIC or PASCAL and try your hand at writing a simple program or two. The sense of achievement when you have a working program is great; perhaps because the process is addictive and consumes much more time than you would believe possible.

TWO

What is Information?

Information is one of those misunderstood concepts. Yet it is at the same time one of the most used. We go to the bus station to seek information from timetables. We obtain information from government offices such as the DSS and the DVLC. Banks, supermarkets, leisure centres, libraries and even the police ply us with more. Newspapers, television and the radio present us with their own ideas of what information should be. In many instances, however, we may not agree with their conclusions. McLeod termed its subjectivity as being one person's junk and another's treasure.

We all understand and deal with this accordingly when extracting news from the media, for example. Yet many of our organizations consider information to be something more. They see it as a vital resource, to be managed like any other valuable resource. How it is used and disseminated through the available technology can determine how efficient, and indeed effective, an organization is. Peter Drucker sees information acquired in a systematic and purposeful way as enhancing an organization's productivity.

Information is important, we cannot operate without it. But more than this, we are discovering that our ability to process it by increasingly sophisticated technological means is fundamentally changing the way that employees perceive their organizational environment. The consequences of this could be either to break down established functional controls or indeed to enhance them by becoming super-efficient. Much will depend upon the characteristics of the organization before implementation.

As firms become increasingly 'information dependent' we need to identify the perceptual relationship between ourselves and the information we use. How, for instance, can we possibly design an adequate information system if we do not understand the nature of information? The answer to this lies, in part, with the conventional wisdom of the day. Managers, like everybody else, develop their views through exposure to established ideas. Thus, the way in which they understand information will influence the way in which they treat it. Such conventional wisdom is coloured by the technicians, on the one hand, and the user's own experience on the other.

Anecdotal prescriptions abound as guidelines to the way presented for all who care to listen. These arise from a blanket of professional and academic thinking which surrounds the business environment and provides remedies for action. We could categorize all these ideas into two bodies of thought, two paradigms. The first could be termed the resource-driven paradigm. This is because its central theme in understanding information is the continuity and consistency of the information itself. It is very much in vogue at present.

The second body of thought is the perception-driven paradigm. Information is seen as an abstract concept, the product of individual perception. It is a temporary phenomenon and as such belongs only to the receiver. The difference is not merely one of academic debate. Managements adopting one or the other can affect the design of their organizations. If information is considered to be a resource then resulting systems are usually more centrally controlled, the assumption being that all information is corporate property. Whereas information considered to be personally owned is seen as being outside the formal structure.

Within the framework of this paradigm the view of information is coloured by its use as a resource. Like any other resource it can be tapped at any time with the certainty of achieving a predictable value from it. Information is regarded as unchanging, and therefore can be easily accommodated into a firm's formal procedure. There are a range of propositions available which seeks to explain information. Each proposition has a consequence for organizational design. Listed below are those major themes and what their implications might be on business.

Herbert A. Simon took the view that information, along with energy, constitute the two basic currencies of our society. We need energy to breathe, to move, to think, to live. But this alone is not enough. We need to know when to breathe, when to move and how to think. Information provides us with that knowledge. To adopt this perspective is to interpret information as an independent entity. It is all about us waiting to be picked up and used. There may be minor variations in its interpretation by different individuals, but consistency can be easily achieved through better training.

Specific sets of information, such as finance information, are given a unique value by organizations which are usually based upon departmental rather than individual needs. Particular information is thus attached to particular departments. This allows information to be generally accessible because perceptions of it are bound by the formal departmental framework. Information is then tied into that framework and cannot be used legitimately within any other context, thus guarding its consistency. What is excluded, therefore, is an important role for the individual receiver. This is not to suggest that information cannot be perceived by individuals. Psychologists such as Anderson and Bower have long been aware of the importance of information reception in determining an individual's behaviour.

However, information can be received automatically without an individual having to understand it. For example, through training an employee could process production statistics without knowing what they mean. In this case it is the functional process (through its formal procedure) that is acting as the receiver rather than the individual. An accounting process, for instance, responds to specific inputs without anyone having to understand them. A motor car responds to certain information transmitted through mechanisms such as the throttle, steering wheel or brake. Both of these examples are designed to act as receivers to specific information.

By implication, information can be transmitted to any receiver that logs into the particular transmission. In the same way as a radio picks up a station when it is switched on, the right form will pick up the right information. And like radio waves, information is omnipresent, only needing the correct tuning. Such a view allows information the continuity and consistency required to be regarded as a resource. It also allows a further proposition.

The design of any formal information system must assume that information does not change during transmission. It can, however, be transformed from data or other information sets before transmission. Within such a framework it is the transmitting functions (departments, processes, etc.) which are considered to be the main determinants. The information receiving functions (also departments, processes, etc.) are designed to link into the operational needs of the transmitter functions.

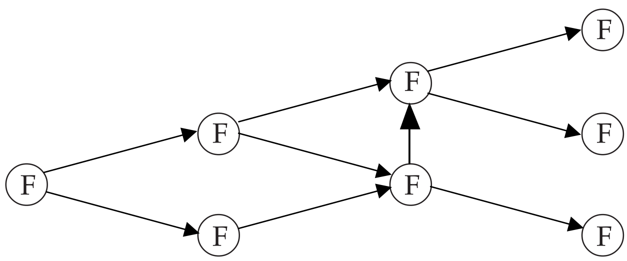

F = functions

Fig. 1. An information network between functions

Each function, however, can be both transmitter and receiver of information, and together with other functions will form part of an integrated network (Figure 1). The functions are tied by their simultaneous roles as information receivers and transmitters, thereby being prevented as much as possible from deviancy in their information usage.

Consequently, business organizations are driven by those functions which transmit information. The higher the value that a particular set of information is given, the more important will be its relating transmitter function. For example, the finance department could both transmit and receive information. Indeed, each job function within that department could also do the same. A cost clerk, for instance, receives information from the factory floor which he then processes into other information and transmits to some other source, perhaps management. The value of the information he processes will depend upon the utility his organization has placed upon it. If the utility is high, then the prestige attached to the transmitting function will also be high irrespective of job complexity. In the same way, if a department has a great number of important information transmitters within its control, then it too will be prestigious and possess influence. Why certain sets of information are prestigious and others are not will depend, amongst other things, upon the organization's political processes.

Perception-driven Paradigm

The perception-driven paradigm does not consider information to be a resource. Individuals or groups are seen to own their information. What belongs to the organization are the data sets, the facts and figures specific to certain functions. For instance, data on how an organization is performing in its market quite obviously belong to that firm. How well that data is used is dependent upon the individual's own competence in interpreting it. Management can control this indirectly through training, but they cannot directly control the thought processes which turns data into information. Once again, there is a collection of propositions or ideas about information which formulate a general approach.

Adherence to this paradigm, as in the previous paradigm, may not necessarily be by conviction but rather by default because people not treating information as a resource automatically form part of this group. However, there are others who are members by conviction. Information is said to be receiver-dependent, in that, any set of information can be considered to be information only if it is recognised as having value by a particular individual. Thus, a river flooding may be information to someone who lives on its banks but of no consequence to someone living several miles away.

It would be difficult, therefore, to understand information as a group or organizations property since its determination is subjective. Management teams have implemented computer technology and in so doing also developed their information sets into a resource. They experience problems, usually some time after a successful launch, and then abandon using their technology in an integrative manner, returning to the 'islands' approach (that is, a computer system, perhaps physically integrated through networks, etc., but perceptually segregated into different functional areas). Information is no longer seen as a resource; indeed its treatment as such is often seen as the culprit and not the technology itself.

Both paradigms recognize a difference between data and information. Computer personnel, for example, have long known of the distinction. Information is seen by them as processed data. The two are interrelated, data being the input to a process, and information the output. In the same way that an output of one system can form the input of another, information from one system can form the data of another. This is consistent with treating information as a resource. Perceptually-conceived information on the other hand is not the consequence of a formalized process. Data are a factor, but then so are other aspects such as individual traits, culture, structure and political processes.

The relationship between data and information is, therefore, not so strong. This contradicts the objectives implied by the creation of management information systems and data processing. The former allows the individual access to established information, whilst the latter converts data into information. Since the production of information cannot be formalized there is no point in trying to achieve it. Systems should, therefore, be designed to produce the most accurate and readable data in order to enhance easy conversion into information.

When Ligomenides describes information as a fundamental force, he should be referring to data. Levitan, on the other hand, is correct in claiming information to be dynamic and continuously evolving. The logic of these two views can be seen when data and information are differentiated. It is only data which is time-independent, and therefore unchanging and static— a fundamental force. By this we mean that data retains a constant empirical value: a fact is a fact and cannot be altered. Information, on the other hand, can be altered and changed. It is time-dependent and thus its value can alter from one moment to the next. It is dynamic and evolving because it is determined by an individual's perception.

It is, indeed, time and utility which transform data into information. If at a particular time a set of data is useful to an individual then it can be described as information. Equally, at another time the same set of data may not have utility and therefore not be information. The data themselves have not changed in any way, remaining always constant. What has changed is the individual's perception of those data. Information is in reality the consequence of a complex psychological process which transforms perceived data, into usable thought inputs.

It is, therefore, data which the individual receives from the environment, the brain which transforms these data into information. No formal system can directly pre-empt this stage; thus no formal system can transform data into information.

Data can be transformed into different data, transmitted and received. For example, production data submitted to the cost department and processed into second generation costed production data, which are in turn submitted to the financial director. Information is transformed, in part, from data; it can only be received and never transmitted. The financial director may attach value to costed production data and thus receive it as information. If he/she passes it on, it will once again be as data, perhaps refined into third generation data. There cannot be information systems because information is so transient.

The individual's subjectivity, which can be developed by training and formal procedure, determines what is information. At another time that perception may be different. Information systems cannot work in this way. Their assumption must be that information is not transient but consistent.

Quite obviously information systems do exist in one form or another. They may not be what they claim to be but they are nevertheless working. Is it worth arguing the difference? The answer to that lies within ourselves. Our actions and behaviour, and consequently our organizations, are governed by our perceptions. If, therefore, we understand information to be permanent rather than transient then how we organize ourselves will be coloured by that. We will design systems which use information as a resource and expect it to be totally consistent.

Structural Implications

The difference between these two paradigms may seem subtle, and in terms of information usage it may well be. However, NBS research indicates that the effects upon organizational design can be major. Structural characteristics consequent upon individually-owned information, for instance, can adopt two extremes. On the one hand they can be organic and informal (adhocracies) whereby all members are seen as contributing equally, each possessing a particular range of techniques to interpret data, whilst at the other extreme there are formal and mechanistic organizations (bureaucracies). The rigidities of structure stifle any use of information as a resource, so by default such processes are left to the individual. These contrast with the resource-driven paradigm in which the structural characteristics of organizations are seen to be centrally orientated, with spheres of influence at the periphery, their importance dependent upon the perceived value of particular sets of information.

Wider Implications of Information

The development of an organization oriented toward the resource-driven concept of information is well known, and many large firms are aligned in that direction. A sophisticated network is established, more than likely computerized, accessing and distributing information. Particular attention is paid to ensuring standardization of information sets. For example, a piece of costing information is so devised that it can be used by accounts, marketing and production. These organizations treat all information as their property, attaching to it a value like any other asset. This denies the employee any sense of ownership and tends to detach them from the work process with consequential inefficiencies evolving. It is asserted, that a motivated staff is a contented and efficient staff.

One important way to motivate is to give the individual responsibility within their job. Such responsibility is manifest in how much ability a person will be given to determine the information need of their particular function. To standardize information in the way that many corporate systems do, takes away that fundamental element of responsibility. Levitan, on the other hand, would claim this to be an inevitable consequence of increasingly sophisticated corporate structures. In addition, Karen Levitan and others imply that information which is not used as a resource is inferior. On the contrary, one could argue that giving the individual ownership could result in a greater entrepreneurial drive.

Standardization is the killer of innovation. It is difficult to conceive of any substantial organization, however, where such a degree of individuality could be tolerated. Their size demands a need for cohesion through standardized practice. Therefore it is only in the smaller establishments where such freedom would work. Large organizations which do not treat their information as a resource still exist. Their mechanisms of control, however, are not informal or organic. Indeed, the realities are the converse. Individuality is stifled by greater control, rather than less. No-one is trusted to use information for the organization's good without formal control procedures. Classic examples of such organizations are to be found within the civil service or large quasi-government institutions such as the post office.

It has been suggested that there is a relationship between the way in which a firm uses its information and its structure. A matrix can be established differentiating organizations in terms of how they treat their information and the consequential structural characteristics. If, for example, information is plotted against size (another important contingent of organization) four major structural types could be identified. Each is the outcome of the way a particular firm uses its information.

On the vertical axis of the matrix are the two extremes of information. On the horizontal axis are the two extremes of organizational size (Figure 2). Both information and size are considered to be contingents, that is variables which affect the characteristics of an organization. The four possible combinations of these two variables determine four organizational structures. Thus, each structure is an ideal for each of the four combinations. For example, a managerial structure is the appropriate structure for a large firm using information as a resource.

Resource	Managerial	Technical/ Professional
Treatment of information	Bureaucratic	Entrepreneurial
Perception	Large	Small

Structure parameters

Managerial	: larger in size with central/top control but with more departmental integration
Bureaucratic	: larger in size with formal structure and a diversified control
Technical/professional	: smaller in size with a professional body and formal structure
Entrepreneurial	: centrally controlled but with less formal structure and smaller in size

Figure 2: Information/structure matrix

Research into the validity of such a matrix is by no means complete. There is evidence from the NBS regional survey, for example, that such a correlation does indeed exist. In general, the more successful managerial-style firms tend to treat their information as a resource, whilst the more successful entrepreneurial organizations tend to treat their information as a perceptual phenomenon, and so on. It is not yet certain whether the matrix is predictive, that is implying combinations which are not successful. There is sufficient evidence, however, to support the claim that information usage affects organizational structure.

The issue of information is much more than whether or not it is a resource. Information is indeed the essence of an organization, shaping and determining its structure.

For example, bureaucratic organizations could hot effectively treat information as a resource because their mechanisms could not control it that way. A technical/ professional organization, on the other hand, is better able to handle its information as a resource because of its structural characteristics.

The implication is that there should be a natural balance between the design of an organization and its information structure. Where this has not been achieved it is because managers are not matching the structure of their organizations with an appropriate usage of information. NBS research has also shown that the implementation of computer technology can worsen any potential for imbalance. This is perhaps because of two factors.

Firstly, the characteristics of the technology itself do tend to enhance resource-driven rather than perception-driven information. Secondly, the greater volume of information generated makes the problem more apparent.

If the matrix in Figure 2 is valid then there is neither a right nor a wrong way to control information: much depends upon the individual circumstances. Resource-driven information and perception-driven information are opposite extremes of one continuum.

Management teams were grasping eagerly at the new computer technology as a solution to their efficiency problems. There were at that time forces in the environment—market, political, social or economic—with which the traditional methods of organization could not cope; it was believed that the computer would provide all the answers. In the space of that time many colleagues began to discover that the impact of computer technology seemed to be contradictory. On the one hand, it did tend to improve efficiency, whilst on the other it was not always so certain that an improvement in effectiveness followed. It seemed that underlying, and as yet not understood, forces were at work.

Departments employing the same technology were more or less successful than others. Business organizations with all the ingredients of success and implementing computers to further fuel their growth quite often experienced failure a few years later. Since those early days there has been considerable discussion and research on the impact of information technology. It seemed that the perceptual aspects of organization were not merely misunderstood but often ignored, and thus we were not able to understand fully the implications of introducing information technology. We now realize that not only does it have an impact on the way we run organizations, it also has an impact on our behaviour. There is a readily understood physical aspect of any firm and a more abstract, perceptual aspect which arises from the way that people behave as a collective. Information technology is not the only agent in this process, but because of its particular characteristics it can have an extraordinarily powerful effect on both these aspects. One of the two main objectives of this book, therefore, is to underline the importance of the relationship between, information technology (and to a lesser extent general technology) and its organizational structure.

The two are inseparable. Models that seek to understand either technology or organization must do so within the framework created by this relationship. Otherwise an essential factor leading to an appropriate analysis will be missed. The other main objective is to produce a model which can do this. Unlike many of its predecessors, the model will not concentrate on the functional or hierarchical nature of organizations but rather on the dichotomy between the physical and perceptual aspects. This exists whether or not there is information technology present, but in many cases the implementation of such technology can exacerbate potential dysfunctional forces existing between the two. In effect there are two organizations—a physical organization and a perceptual or virtual organization.

The impact that information technology can have on the sometimes tenuous relationship existing between the two can be quite profound. The model explores thoroughly this relationship. It proposes virtual organization as the antithesis of physical organization. That is abstract, unseeing and existing within the minds of those who form a particular organization. The point being that because the framework of virtual organization is often subjective and open to many different perceptual interpretations, it is difficult for anyone to establish an appropriate response to it. This makes it no less real; nor does it reduce the need to respond to it. Left unattended the effects of virtual organization can be quite catastrophic, especially if stimulated by the impact of information technology. This has considerable significance for the practitioner, especially those who have introduced, or are thinking of introducing, information technology. The model suggests that for a firm to survive there would ideally exist a dynamic equilibrium between the two organizations which has been termed the organizational balance. Anything which upsets this balance, such as the introduction of information technology, has to be compensated for by proactive action of management. If managers do not understand, nor indeed perceive the existence of virtual organization, then they will not be able to take the appropriate measures. Long-term structural weaknesses or even total failure of the business can result.

Demonstrating the existence of virtual organization can be problematic. By definition it cannot be observed directly since it exists only in the mind. At best one could ask a selection of staff for their impressions and make certain deductions, but these are likely to be dubious and tainted with the researcher's own perceptions. As such there has deliberately been no direct empirical research to establish the foundations of this model. Support for its validity has been sought in three other areas. Firstly, there is indirect empirical data gained from several research projects conducted at the Newcastle Business School. This comprises a study over two years by Stuart Maguire and myself into eighty-five small to large firms in all sectors of the North East. We examined particularly the strategy that firms employed to implement and control their information technology.

The data were gathered through a series of observations, interviews and involvement in the process. Other data were obtained from one or two firm studies on the impact of information technology by a group within the Newcastle Business School called AMBIS (Advanced Manufacturing and Business Information Systems). The argument for using these data rather than dedicated data is that they should be less contrived. Secondly, there is support to be gained from relevant personal experience. Although many would find this approach invalid because it does not comply with traditional or conventional ideas about research data, there is an increasing acceptance that such data can prove useful. In the social sciences, for example, Duke argued that such an approach (amongst others) would make a valuable contribution to the understanding of social experience. It is accepted that it cannot provide the only data, but the data that it does provide should be considered no less valid than any other.

We return to a more traditional style in support of the main arguments. That is, rather than gathering empirical evidence first and then developing a model, we should develop the model first from experiential instincts and then support it with data. This is put forward in the belief that to pursue the traditional methods of data-first research would not produce the generic model necessary to achieve the objectives set above. We have become unimaginative in our ideas about organization because we have been blinkered by too many sets of data.

How the book is structured in terms of presenting a logical argument takes on a greater significance than if it were based merely upon direct empirical data. It requires presentation of basic argument, counter-argument and a compromise position which leads the reader forward to the logical conclusion. This is the dialectic which involves thesis, antithesis and synthesis.

This overview, therefore, should not only cover the background to the book's major arguments and supportive material, but also provide the reader with a basic chapter plan.

To explain the impact of information technology we must look beyond the functional aspects of business organizations. The importance of information is reflected in its ability, through the appropriate technology, to corrupt departmental boundaries. This chapter, therefore, proposes that it is better to understand an organization as a network of information flows, which is itself set within other network flows of region, market, industry, and so on. Such a perspective is underlined by many firms accessing information beyond their established functional boundaries. Within the organization, this could involve different departments accessing a common pool of information from the company database.

Externally, employees can gain entrance to public databases such as Prestel, which are run by autonomous bodies. The drive is towards centralization and hence better management control. Paradoxically, the opposite is being encouraged, whereby every user has the potential to access information which enables them to operate more efficiently, and also to do so outside the control of their manager. Under these conditions it is difficult for managers to control not only who reads and writes what but also, and perhaps more importantly, who knows what.

Large amounts of information can be processed and communicated via a network without managers ever being able to determine whether it runs contrary to their own policy. Security has, as a consequence, become an extremely important issue. The more the dissemination of information is allowed, the more likely control mechanisms will become corrupted.

In simple terms, data are facts, whilst information is the consequence of interpretation. What is information to some is not to others, and it is the subjective interpretation of what is and is not valuable which has caused managements considerable problems. The control parameters on the organization's database are based upon managers' own perceptions of its value as information which may not coincide with that of a potential intruder or misuser. But more than this, individual employees' perception of their information stock widens because of their greater accessibility to the information network. As a consequence, their perception of their own job parameters also widens beyond what management might think desirable.

THREE

Information, the Life Blood

Information is the life-blood of a community information service and the information file is at its heart. Unless the heart is sound and continually pumping a supply of regularly renewed and fresh information into the system, it will not function at its best. Therefore, it is important to give extra care and attention to planning the resources needed to set up a sound information base and a workable system for keeping it up to date.

Some basic things you will need to consider are as follows:

— How is the information to be collected?
— What method is to be used for processing and storing information?
— How is the information to be retrieved?
— What size is the file likely to be?
— What areas, geographical and subject, is the file to cover?

In addition, you will need to consider how the information is going to be disseminated, if it is not just for internal use; but that is the subject of the next chapter. The kind of system you choose and the degree of its sophistication will depend on the nature of your information service, the size and complexity of the information you need to store, the staff, and the financial resources available. When designing a system, make sure it is readily understandable to all those using it-volunteers and public included-and not just to professionals.

Coverage

In identifying the need for your community information service, you should have reached a decision already on the community to be served. This may be either a geographical community or a community of interest or both. Either way, you will still need to make a decision about the extent of your information base. For example, a neighbourhood information service must decide what information to collect on the wider area-town, county, region, nation-outside its local community. An information service aimed at a fairly wide area may need to limit its interests to a particular sector of the community or to a particular range of subjects, e.g. the low paid, those receiving income support, or those with debt problems. Even an information service in support of a local campaign may benefit from collecting information on similar campaigns or problems country-wide.

Types of Information

The information base of your service is most likely to include the following types of information:

Soft information— details of clubs, societies, organizations and services; individuals; events, etc. Usually this information will not be available in a published form, or at least not in sufficient detail, and it will be necessary to make a conscious effort to collect it. This information may comprise the major part of the resources file.

Hard information— factual information on a specific subject e.g. benefit rates for single-parent families, how to get legal aid, how to change your name. This information will be available in a variety of Banns, some ephemeral such as leaflets, pamphlets, booklets, broadsheets, posters, periodical articles, etc. Such items are usually kept in a vertical file, in storage boxes or similar receptacles and accessed via some form of index or classification scheme. Other hard information may be in the form of books, multi-volumed reference works, law books, or even audiovisual forms. If there are many items of this type, they will need to be classified and shelved.

Supplementary information— information already produced in a particular format-directories, handbooks, diaries, annual reports, constitutions, newsletters-of organizations appearing in your information file.

'Soft' Information

It is highly unlikely and undesirable that an information service should need to 'go it alone' in the collection of information. Such is the volume, complexity and variety of sources of information in present-day society that, unless your service has very narrow terms of reference, it will be virtually impossible for it to collect and keep up to date all the information that is required.

Therefore, it is important that you first of all identify the information providers, support services and 'gatekeepers' in your community and establish effective contacts with them. These links may already have been forged in the process of conducting a community profile. They now need to be fostered and strengthened, so that there can be a mutual exchange of information. These contacts will also be able to provide you with useful feedback from the community as to the success or otherwise of your service. The network of contacts can be maintained on a fairly casual basis as the need arises, or you might want to formalize the arrangement. Some ways of doing this include irregular 'get togethers', informal luncheon clubs, regular meetings with agendas and minutes, or the circulation of a newsletter or bulletin. Other activities that could develop out of such meetings are joint collection of information, shared publicity, compilation and publication of directories, information handbooks and leaflets, training, and general discussion of common problems. Collecting information is a time-consuming process and there is no one method of going about it. You will probably have to use a combination of several techniques to build up a satisfactory information base. One thing definitely to be avoided is duplicating work that has already been done. So first of all identify the following:

Existing information files. Contact all the other community information services, council departments and organizations who are likely to maintain information files and ask if they would be willing to share this information. Try to offer something in return, either an exchange of information or some other help you can provide.

Local directories. There are several types of directories that may be available in your area and which are useful sources of soft information: (a) local government authorities often produce directories or town guides of their area which contain a certain amount of local information or they may produce a guide to council services; (b) telephone directories, especially Yellow pages which have an alphabetical subject arrangement; (c) area directories, such as those produced by Thompson's Newspapers, which are similar to Yellow pages but also contain a section of community information-your local newspaper may produce a directory of local services, sometimes in conjunction with local information groups; (d) citizen's guides may also be produced by the local newspaper as a supplement to an existing newspaper; (e) CVSs or RCCs sometimes produce directories of voluntary organizations in their area or you may find these listed in your local CVS/RCC annual report; (f) directories covering a special subject, such as accommodation, halls for hire, etc.; (g) directories aimed at a particular interest group or groups e.g. people with disabilities or senior citizens.

National directories are useful not only for the wider network of services but also for identifying local offices of a national organization.

Other methods to use in building up your information base and maintaining its currency include:

Looking and listening. A lot of information can be gathered by simply walking around your community, looking at notice-boards, picking up leaflets, or attending open meetings, fairs, fetes and other community events. This method of collecting information is unpredictable and requires more than a 9-to-5, five-day week.

Contact with individuals. People are a major source of information in any community and there are always those to whom others naturally go for information. Try to identify such people and enlist their help, possibly by offering something in return-bartering is a time-honoured system! Direct contact with individuals in other information services and organizations is often a better way of eliciting information than impersonal methods and also paves the way for further cooperation. Pay regular visits to such centres.

Scanning newspapers, magazines, newsletters, etc., can turn up details of new organizations; changes of personnel, premises or hours of opening; new services introduced; and may highlight problem areas where information may be needed.

Other documents, including council minutes and agendas, planning regulations, annual reports (ask to be put on mailing lists for these), leaflets and manifestos.

Media publicity Local press, radio and cable television stations may be prepared without charge to put out a call for organizations to contact your service. If not, then consider inserting an advert or, where there is, a local directory produced, have a tear-out slip included for organizations to fill in and return to you. As a result of these various methods, an amount of raw material will have been obtained which may be of use to your information service. How accurate, up to date and complete the data is will depend very much on the reliability of the source. In most cases, unless you have every confidence in the source, the information will have to be checked before it is entered in a more permanent form in your resource file. So, at this stage, simply record it onto say a scrap 5in x 3in card and put it into a temporary file.

The information you have obtained will be of several different kinds but the most common will probably be that relating to organizations, clubs and societies. This is the most difficult type of information to collect since it is ever changing and therefore unrelenting effort will be required to achieve even reasonable completeness or accuracy-100% will be impossible. However, with a good system for collecting and updating the information, you should be able to build up and maintain an acceptable file.

A lot of staff time will be needed to carry out this work, so first of all investigate whether there are other local information services that need this type of information and would be willing to help in its collection. It may be possible to collect the information through an umbrella group, where the work can be framed out so that it is not too great a burden on any one organization. Alternatively, a joint approach can be made on behalf of a number of organizations.

Ideally, the most effective way to collect information of this kind is to pay a personal visit to each organization or service, but rarely will that be possible without extensive staff resources or a very small community. The cheapest and quickest way to collect this information is by telephoning. Prepare a standard list of questions to ask so that entries in your information file follow the same format.

Telephoning is only practicable if you have a reasonably small number of organizations to contact, say up to one hundred. If visits and telephoning are ruled out, then the information can be collected initially by postal questionnaire, with a request for up-to-date information sent at least once a year and also when it is known from other sources that changes have taken place. A covering letter should be sent with the questionnaire or update request, explaining "Why the information is needed and asking for cooperation. If the information is to be passed on to other organizations or commercial bodies, either free or for payment, this should be drawn to the attention of the recipient of the questionnaire, so that they have an opportunity to decline to give the information or insist that it is not so disseminated. If the information is to be stored in a computer, then you will need to meet the requirements of the Data Protection Act.

If it can be afforded, include with your questionnaire some means of returning it free of charge, as this will significantly improve the response rate. Separate forms may be required to collect different kinds of information. The following is a suggested list of headings which you might need to include on forms for collecting information from two of the most common types of organizations.

Clubs and Societies

- Name: the name by which the club is best known, plus the full name if different. Indicate relationship to larger body, i.e. branch, regiment, lodge, etc.
- Secretary's name, address and telephone number or those of a similar person if the organization has no secretary.
- Place, day and time of meetings.
- Purpose of organization if not self-evident from the name.
- Eligibility: any restrictions by age, sex, ethnic group, occupation, status (single, divorced, widowed) etc.
- Subscription: any charge to become a member or to attend meetings.
- Annual general meeting: the date of an organization's AGM can be a useful indication of when to send out update forms as this is when officers or policies are likely to change.
- Other officers: names, addresses, telephone numbers of treasurer, president, chairman, publicity officer, membership secretary, etc.
- History of the organization including special events or persons associated with it.
- Publications: newsletter, diary, commemorative brochures, annual report, etc.
- Date the information was obtained and/or last updated or checked.

Agencies and organizations providing a service

— Name: popular name and full name, e.g. CHAT (Come Here and Talk), SHAC (Shelter Housing Action Centre), plus relationship to parent body where necessary.
— Address: street address and postal address, Post Office, box number if used.
— Telephone and fax number, including-service number, administration number, after-hours or hotline numbers.
— Contact person: personal name, title and address, if different from above.
— Hours of opening: days and times; note if seasonal, e.g. holiday period, term time.
— Services provided: advice, counselling, practical help, etc.; types of enquiry the service would like to have referred to them.
— Eligibility: age, income, sex, residency, status.
— Application procedures: walk-in, appointment or waiting list; what papers or documents need to be brought.
— Cost: free, fees, means-tested, donations; any facilities for payments to be spread over a period; help available from local or central government, charities, etc.
— Geographical area served: neighbourhood, city, county, region, adhoc area; no geographical restrictions.
— Branch offices: extension bureaux, mobiles, surgeries, etc.; include hours of opening, routes and times of stops.
— Director, administrator or executive director of the service; name and telephone or extension.
— Volunteers: does the service use volunteers and for what purposes? Method of recruitment.
— Publications: directories, handbooks, leaflets, annual report, etc.
— Funding: local government grant-aided, donations, etc.
— Special facilities: foreign languages spoken, access and facilities for people with disabilities, photocopying or fax service, escort, advocacy, etc.
— Transportation: how to get to the service, e.g. bus route numbers, underground line and station, etc.
— Date information was obtained and/or last updated or checked.

These two lists represent most of the facts that a community information service might need to know about an organization or service. In practice, depending on the scope of your service, it may not be necessary to include all the fields in the questionnaire. It is a good idea to get the person filling in the questionnaire-to-sign it, thereby giving the information service some-protection against claims of giving out wrong information—a not uncommon occurrence, especially when the information is committed to print.

There are other types of 'soft' information that your service might need to collect, for example halls for hire, places of local worship, accommodation, 'What's-on' events, local industry. The procedure is much the same:

— Decide what information you need to know.
— Devise a standard format or list of questions to ask.
— Identify possible sources of information.
— Decide which method or methods to use to collect the information: telephone, personal visit, postal questionnaire, press advert, etc.

How to Store 'soft' Information?

There are various methods that can be used for storing information:

A list is easy and quick to consult, can be photocopied for clients to take away and can be faxed, but has very little flexibility for inserting and updating information.

A strip index provides flexibility, can easily be photocopied, and is available in various forms: large binders holding many pages, smaller folders containing just a few pages, address books, wall-hung panels, rotunda units, etc. Strips are available in several colours, which allows for a simple categorization, and various widths. Even so, they are limited as to the amount of information that can be contained on them. It is also not quite as easy to insert new strips as the manufacturers claim.

Cards are still one of the most commonly used systems for storing information. They are infinitely flexible, easy to insert and update, can be sorted into any preferred order, and require very little technical know-how. The most popular sizes are 5in x 3in, used mainly for indexes or temporary information waiting to be checked, 6in x 4in, and 8in x 5in, which are large enough to contain sufficient information about most organizations or services and to be preprinted with a standard grid for recording the information. There are more superior types of card files available, such as the Rondofile or Rotadex systems which use custom-designed stationery and therefore tend to be more expensive. With systems that use standard cards, you can-if you are hard-up or want to save trees-use the reverse of scrap cards. Card files have a number of disadvantages: they are not so easy to carry around, they arc not easily reproducible if someone wants a copy of a particular section, multiple access (e.g. by name, subject, area, service, etc.) requires duplication of cards, and the updating process can be tedious.

You may still find around edge-punched cards which offer a primitive type of multi-access to the information on the card-through the strategic use of knitting needles or rods. Personally speaking, I would save up and get a computer.

Loose-leaf binders with one sheet for each entry, sometimes called 'sheaf' binders, have a similar flexibility to cards. They are slightly slower to add to and amend but are more portable. They are suited best for alphabetical arrangements as they are difficult to guide.

Microcomputers now have the capacity to handle immense information files that can be accessed by any number of predetermined variables or even free-text searching by keyword, depending on the software used. With a printer attached, all the information or sections of it can be printed out on demand or to provide a master for printing by conventional methods. Personal computers now begin at prices that bring them within the range of most community information services. However, there is danger in automatically seeing the computer as the best solution for maintaining your information file, as this may not always be the case. A lot depends on the size of your file and its complexity, what you want to do with the file, and to what other uses you want to put the computer. The importance of the computer in, information and advice work is such that it merits a chapter to itself, and you will find these and other aspects dealt with in more detail in Chapter Eight.

Filing System

Whichever method is chosen for physically storing the information, except for that by microcomputer, it will be necessary to decide the best way of arranging entries so that the information can be retrieved swiftly and accurately. In order to ensure that the total resources of your information service are used to the full when answering an enquiry, it will help if the system used to organize the information file can be integrated with that for the vertical file and the book-stock. This point needs to be borne in mind now, although it will be dealt with more fully later.

The bulk of the information file is likely to comprise 'soft' information -which you have collected about organizations, clubs, societies and services and, to a lesser extent, about individuals who are a valuable source of information or advice, plus items of 'hard' information not available or conveniently available in printed form. There are several ways in which you might want to retrieve this information-by the names of the organizations, by subject (i.e. area of interest), by clientele, by the type of service (e.g. counselling), or by place. In practice, the 'clientele' and 'type of service' approaches are usually catered for through the subject file rather than, as separate files.

Organizations file: The master file for the whole system will usually be an alphabetically arranged sequence of entries by name of organization. It is generally recommended that the full, official name of the organization is used, with references from other forms of the name. In practice, it does not really matter as long as there are references from the alternatives to whichever form is chosen. There is a good argument for putting the main entry for an organization under the form of name by which most people refer to it, even if this differs from or is an abbreviation of the official name. This is especially so where the file is to be used by the public.

In addition to details of organizations, each entry in the master file may also contain certain 'housekeeping' details relating to the maintenance of all the information files, plus references to appropriate headings in the subject file, where relevant information can be found, and a classification number.

Subject file: Most enquiries received by a community information service are likely to be about the need for a service or activity or for help with solving a problem, rather than for information on a specific named organization. So the information you have collected about organizations, services and individuals will also need to be accessible by the services and activities they provide. There are two ways this can be achieved, by either using a subject index or having a subject file of organizations. In both, you will need to choose or adopt a set of subject headings which will adequately describe the information on file or elsewhere in the system and the interests of your clientele.

If you are compiling your own list of subject headings, choose terms that are in common use by your clientele, especially if they will be consulting the files directly, rather than the 'official' term, e.g. OUT OF WORK or ON THE DOLE instead of REDUNDANT. Always refer from alternatives not used to the chosen term.

A subject index is rather like the index to a book. It is an alphabetically arranged file of subject heading cards on which reference is made to where information on that subject can be found. The cards themselves do not contain information. The subject index can contain references to organizations, individuals and other supplementary material such as pamphlets, periodical articles, audiovisual material, books, etc.

A subject file consists of an alphabetically arranged set of subject heading cards behind which are filed copies of the organization cards appropriate to that heading and possibly cards containing items of 'hard' information. An alternative is to arrange cards according to a classification system either devised by yourself or adopted from another service, such as the NACAB classification scheme.

Place file: Where an information service covers a number of distinct towns, villages or neighbourhoods, each having a similar range of organizations, it may be helpful to have a file that can be accessed by place. There are a number of options. You could sort the organizations file initially by place, especially if it contains few organizations whose responsibilities extend to the whole area covered by your service. Alternatively, if most of the organizations are prefixed with the name of the place (e.g. EXVILLE ATHLETIC

CLUB, EXVILLE COUNCIL FOR THE DISABLED), the alphabetically arranged master file will automatically bring together those from the same location, provided inverted headings have not been used.

Some arrangements will have to be made for organizations which do not conform to the pattern, either because the place-name does not feature at the beginning or because it is not contained in the name of the organization at all. In these cases, an additional card can be made out with either an inverted heading (e.g. EXVILLE, ROTARY CLUB OF) or the place-name can be prefixed.

The third option is to have a separate file of organization cards arranged initially by place and then alphabetically by name of organization or by subject.

Filing alphabetically

Filing alphabetically is not quite as simple as ABC, as any librarian will tell you. There are two recognized methods, known as 'word by word' and 'letter by letter'. In word by word, entries are initially sorted by the first word in the heading, then headings that begin with the same word are sorted by the second word, and so on. Another way to describe this method, which may be more helpful, is to treat spaces between words in a heading as an imaginary letter coming before 'a' in the alphabet, then sort letter by letter.

In letter by letter, you simply ignore any spaces in the heading and sort one letter at a time from left to right. Here is how a small group of headings would look sorted by the two methods:

Word by word	*Letter by letter*
DO SOMETHING! BOYS CLUB	'DOG AND BONE' DARTS TEAM
'DOG AND BONE' DARTS TEAM	DOGSBURY RESIDENTS'
DOG SNIFFERS ANONYMOUS	ASSOCIATION
DOGS HOMES	DOGS HOMES
DOGSBURY RESIDENTS'	DOG SNIFFERS ANONYMOUS
ASSOCIATION	DO SOMETHING! BOYS CLUB

There are several other niceties to do with filing alphabetically but we only need mention here that hyphens are treated as spaces in word-by-word sorting and numerals are spelt out as said, so that 1900 (the year) is considered for sorting as Nineteen Hundred and 1,900 as One Thousand, Nine Hundred.

File 'housekeeping'

In addition to the information needed to answer enquiries, there are other items of information to do with the maintenance of the file which may usefully be included on each entry, such as the following:

— the date the information was obtained, last updated or last checked, which will indicate not only the degree of reliability to be placed on it but also when to update;

— the date a questionnaire or update letter was sent - a check for chasing up non-returned forms;
— additional contacts or other information which does not fit into the main part of the entry;
— subject headings used in the subject file;
— place headings used in the place file, if not obvious from the name of the organization;
— feedback-comments from users of the service;
— 'tracings', i.e. a list of other cards that refer to they can be traced and amended or withdrawn when necessary.

The amount of this 'housekeeping' information you need to include will depend on how sophisticated you need your file to be. The bigger the information service and the larger the file, then the more likely it is that you will need to introduce a systematic procedure for processing information. In such a case it is usual for the housekeeping information to go on the reverse of each entry, using a standard grid like the one illustrated here. It is arguable that if your information file requires this degree of sophistication, you ought seriously to consider keeping it on computer.

For a small community information service, it may be quite adequate just to note the date when the information was collected, checked or updated.

Updating

By their very nature the entries in your database will be changing continually. Hours of opening, meeting places, officers, membership fees, subscriptions, charges, etc., are all susceptible to frequent change. Therefore, it is important to have a system for regularly updating each entry in the database. There are several ways of achieving this but, whichever method you adopt, it must be regular and ongoing.

Interim updating

All the entries for organizations should be checked at least once a year to ensure they are still correct. In between times, however, new information will be brought to your attention by various means, including word of mouth, newspaper reports, organization newsletters, company reports and by direct contact with the organization itself. If the source of information is reliable, it can be substituted immediately for the out-of-date information in the database but, if not, it should be noted and a further check made to verify its accuracy. A simple way to add new information to an entry in a card file is to write it onto self-adhesive labels or slips (obtainable at most stationers) which can be stuck over the original information. Even though an entry has been updated in between times, it should still be checked formally once a year.

Annual updating

Once a year is probably a reasonable time-span for updating your database but a shorter interval may be necessary if the information content changes frequently. You will need to decide how the annual update is to be carried out. Some organizations like to update the entire file at the same time, so that they can say the file was as accurate as possible at a particular date. However, this does create a tremendous amount of extra work for a period of several weeks and you may feel it necessary to get extra paid, work-experience or volunteer help. Another disadvantage of this method is that it captures a picture of organizations at a set moment in time. Some organizations may be just about to make changes, and your records for these would be out of date for almost a year.

Alternatively, you may decide to update on a continuous basis by using either the date the record was added to the database or the organization's annual general meeting (AGM) date as a trigger. This is easy to do if your database is kept on a computer, since a date field can be used to produce a subset of records and, by linking the database with a word processing package, a personalized standard letter can be printed out for sending to each organization. It is less easy with a card or any other kind of file. You can get little coloured metal tags to clip on to the top edge of cards, a different colour for each month, or you could use an appropriate symbol on the heading of the card, such as a coloured dot, a letter or number. Tags can make the file look a bit messy and are easily knocked off if it is in constant use; symbols require the whole file to be checked through.

It is better to update records shortly before or after each organization's AGM, since most organizations, if they are going to change their officers or constitution, generally do it once a year at the AGM. This is so because many organizations do not know their AGM date a year in advance, it is enough just to indicate the month in which the AGM is usually held. For those organizations and services which do not have an AGM, you will need to use the date the information was added as the trigger.

When writing to organizations to update their information, it is better to send them a copy of their record with a request to indicate any changes, rather than sending a blank questionnaire. This saves the secretary unnecessary time in repeating information that has not altered.

An alternative to writing to each organization is to phone, but this is time-consuming and costly, and only feasible if the number of organizations to contact is fairly small. Avoid sending out updates at times when people are likely to be away, such as during the summer. There are certain times when it is more appropriate to update other kinds of information, e.g. social security benefit rates (which usually change in April and September), adult education classes or school, college and university courses which change term-wise.

When an updated entry is returned it should be checked to see if the subject or place headings still apply. Minor changes can be recorded on the master card and any other cards in subject or place files; major changes may require a new set of cards to be produced. For those with computerized databases the process is simpler as just the one record has to be updated. Don't forget to change the date the information was last checked or updated.

'Hard' Information

As well as the database, most information services will need a certain amount of 'hard' information to answer enquiries. Short, unrecorded or inadequately recorded items of 'hard' information, as we have already seen, can be incorporated into the database. However, the majority of hard information will usually be found in one or more of the vast range of print forms, starting from locally produced free broadsheets to extremely expensive multi-volumed loose-leaf reference works.

Within the scope of this book, it is not possible to go into detail about the sources of this material, since they will vary considerably according to the type and subject range of the service you are operating. Community information changes rapidly so any printed sources should be treated with care. However, you might find useful a chapter on the sources of community information that I contributed to Printed reference material and related sources of information, 3rd edn, Library Association Publishing, 1990, which concentrated on organizations and on those items which are regularly updated.

The following framework for collecting community information was originally devised by Grainne Morby, who at the time was working for the Community information Project, with some additions of my own:

Distinguish between subject areas

In other words, identify the main topics into which the subject scope of your information service naturally divides. Throughout this framework I will take as an example an information service for which the main area of interest is housing. It is sometimes useful to apply a consistent criterion when dividing up a subject, although obviously this would not be so easy for an information service whose scope covers as broad an area as, say, 'community information'.

However, taking housing as our example, you could decide to divide it by type of accommodation, e.g. owner-occupied, private-rented, council-rented, New Town, Housing Association, institutional accommodation, tied accommodation, mobile homes, etc. In most cases you would probably need an additional category for subjects that cut across all or more than one category. In the case of housing, this might be 'squatting' or 'homelessness'.

Since the material you want to collect will usually be written to meet a particular need or needs, it is advisable to collect information around those needs rather than for fitting into a theoretical framework. The kind of broad client groups that might be identified in the field of housing are landlords, tenants, owner-occupiers, squatters, transients, elderly people, people with disabilities, the homeless. Then, one could identify client groups who share a common problem or need associated with their particular type of accommodation, such as eviction, dampness, house repair, renovation, harassment, mortgages, rent arrears, planning permission, redevelopment, etc. Once these subject areas and client needs have been identified, it should give you clearer guidelines when you come to select material from the various sources identified in the next part of the framework.

Print sources
Official
(a) National
- *pre-legislation:* political party manifestos and policy documents, speeches by Ministers, the Queen's Speech at opening of Parliament, White Papers, Green Papers.
- *Legislation:* Bills, Hansard reports of debates in both Houses of Parliament, Acts, Statutory Instruments.
- *Guidance to local statutory bodies:* circulars and letters from Ministers, guidance notes, reports of Inspectors.
- *Reference books:* collections of legislation, e.g. Statutes in force, Legal aid handbook.
- *Periodicals:* from government departments and agencies, e.g. Department for Employment gazette.
- *Guidance to the public:* leaflets, posters, videos and other A/V material.
- *Ombudsmen*.

(b) Regional
- Health authorities, British Rail.
- Privatized utilities: gas, water, electricity, telephone.
- Consumer watchdog bodies: Post Office Users' National Council (POUNC), Office of Gas Supply (OFGAS), Office of Electricity Regulation (OFFER), Office of Telecommunications (OfTel), Office Of Water Services (OFWAT).

(c) Local
- *Local authority policies and decisions:* agendas, minutes, standing orders, bylaws and regulations.
- *Guidance to the public:* handbooks, leaflets, etc. from local authorities, area and district health authorities, Community Health Councils, local offices of government agencies.

2. Non-official

(a) National

— Reference books and handbooks from commercial publishers.
— Voluntary organizations.
— Pressure groups, self-help groups.
— 'Umbrella' organizations, e.g. National Council for Voluntary Organizations, Federation of Independent Advice Centres.
— Professional bodies, e.g. The Law Society.
— Trade associations, e.g. National Association of Retail Furnishers.
— Educational bodies, e.g. National Extension College, Open University.
— Trade unions.
— The media.
— Practitioners in the field who publish.

(b) Local

— Claimants' unions.
— Law centres, other specialist advice centres.
— Neighbourhood and generalist advice centres, resource centres.
— Local 'umbrella' groups.
— Pressure groups, campaigning groups, action centres. Voluntary organizations.
— Local media: newspapers, community newspapers, radio, cable television.

(c) New technology

— Video-text systems: Prestel, teletext, private viewdata.
— Online computer databases, e.g. Volnet
— CD-ROM, videodiscs, etc.

This list is not intended to be exhaustive but will give you some idea of the range of sources in which to look for information. Howard Matthews, in his book *Community Information* has drawn up a set of criteria for selecting community information materials. It is aimed particularly at library community information services but may be generally helpful for anyone setting up such a service.

1. Selection should be done by those who know the community best.
2. Material should relate to specific, local, identified needs.
3. All material should be current.
4. All material should be written at a level appropriate to its intended use.
5. Material should be concerned with the practicalities of problem solving.
6. All new material should either fill a gap, offer a new viewpoint, or promise something better.
7. Material should take into account the level of use the client will make of it.

Obtaining publications and other material

Once the material you require has been identified, you will need to obtain it. Some material will be priced, in which case- you will need to order it, and some will be free.

Priced publications

The procedure for obtaining priced publications, particularly in the field of community information, is not as simple as you might suppose. Since a lot of the material may well be fairly cheap and the small or community publishers, booksellers and library suppliers are reluctant to obtain it as the profit margin is likely to be too low for the amount of effort involved. The alternative is to order direct but this also has its difficulties as many small publishers now insist on money with order, and this can cause problems for some organizations through the need to raise a cheque or limitations on the amount that can be spent from petty cash.

A few national voluntary organizations, e.g. Age Concern England, operate a standing-order system for their publications-this is one solution, although you may need to accept blanket coverage of the organization's output. In the larger cities, you may find a community or alternative bookseller who is prepared to supply this material to order. A directory of alternative bookshops is published annually by the magazine *Radical Bookseller*. Whatever the method of ordering, it is important to keep a record of what has been placed on order and from whom to avoid duplication of titles and as a check on supply. For each title make out an order or slip. This can be either a scrap 5in x 3in card or one specially printed with a standard grid. The following indicates the kinds of information that need to be recorded:

- — title of publication-more important than author for this kind of material, since authorship is often unclear;
- — author-might be an individual or an organization;
- — price, date of publication and, if a periodical, frequency;
- — supplier-this may be a bookseller, library supplier or the address of the producing body;
- — date of order;
- — source of information about the item;
- — location, if you are ordering for more than one centre.

You may find it helpful to use the order cards later as the basis of a catalogue of the publications in your information centre; in which case., leave space at the top of the card for either a subject heading, classification number or filing code. After an order has been placed with a supplier, the order cards are filed, preferably by title, in one alphabetical sequence known as the 'order file'. When a publication is received, the order card is removed from the file and accompanies the publication to the next stage if it is to be used later as the catalogue card. If not, you may find it useful to file the order card in a 'publications received' file, until such time as the publication is permanently recorded in the system, after which the card can be thrown away.

Periodicals

Publications which arrive at regular intervals—weekly, fortnightly, monthly, quarterly—will not require an order card but you will need to keep a record of each title taken. This should include the following items of information:

— title of periodical;
— frequency;
— supplier;
— when subscription is due for payment;
— a grid to record receipt of each issue;
— instructions for disposal of back copies.

For those using computers, there are software packages available for periodicals management but these are aimed at organizations who take an extensive range of periodicals and such packages are not cheap. If you have a database management package on your computer, you would do better to devise your own periodicals database.

Free material

This can vary from quite substantial loose-leaf binders plus updates to leaflets, posters and bookmarks and may be required in bulk for clients to take. Since this is likely to be a frequent type of request, it may be quicker and simpler to duplicate a standard letter to send to organizations, with space left to fill in address, title of item, format (leaflet, booklet, poster, etc.) and number of copies required. It is advisable to keep a record of what is requested, if only as a check on whether it has been received or not.

A lot of the sweat has been taken out of identifying free materials, and addresses from where they can be obtained, through the setting-up of Camden Libraries' Free Leaflets Information Service (FLIS). For an annual fee (E120 in 1992) subscribers receive a copy of the Frills directory of free information leaflet suppliers, a monthly batch of new and updated leaflets and amendments to the Frills directory, and details of a model system for organizing leaflets. Details of this service can be obtained from Community Information Services, Camden Leisure Services Department, St. Pancras Library, 100 Euston Road, London NWI 2AJ. A similar, but not so extensive, service is provided by the NACAB as part of its Information Service, but you would need to have approval to subscribe to the full pack.

Organizing hard information

We have already seen that hard information can be found in a great variety of forms. For the benefit of organizing it, the following categories can be identified:

Material in book form for use by clientele or information workers will be best arranged on shelves in the broad subject groups identified when acquiring the material or according to a classification scheme.

Ephemeral material for use by information workers is best stored in a vertical file using the same subject headings as those in the subject sequence of the information file, though probably with more subdivisions, or using a classification scheme. Alternatively, the material could be kept in file boxes on shelves using broad subject headings.

Ephemeral material for use by clientele on the premises is best kept in file boxes, ring or display binders and ideally interfiled with books. Vertical files are not recommended for public use as they tend to be a deterrent. Another way of dealing with ephemeral materials is to gather them together in packs on particular topics and, for maximum impact, to display them face outwards on sloping shelves.

Ephemeral material for the public to take away can be displayed in special leaflet dispensers, on a sloping surface or on tables. Ideally, they should be displayed in broad subject groups or by originating organization, e.g. Benefits Agency, Department for Employment, but this is rarely possible because even within organizations leaflets are produced in varying sizes.

Supplementary Information

Many organizations produce printed material in the course of their work or activities. The most common types of material are directories or address lists of members (individuals or constituent groups), newsletters, events sheets or cards, annual reports, constitutions, posters, leaflets, funding appeals, campaigning literature, advertising brochures, balance sheets and commemorative booklets (anniversaries, etc.). All this material is potentially useful as back-up information to the database but there is no sure way of obtaining it. When sending out the original questionnaire to organizations, you can ask to be put on their mailing list for this type of material, if one is maintained.

Some organizations that operate on a shoe-string budget may make a small charge for their mailing list facility. Even when you are put on a mailing list, regular receipt of material may depend very much on the enthusiasm or efficiency of the secretary of an organization. Often the best solution to obtaining material is to maintain regular contact with organizations and to pester them continually. When material arrives, it should be dated, so that you know roughly how old it is, and carefully filed for future use. The simplest method is to have an envelope or folder for each organization. Write on the outside of the envelope the name of the organization as used in your database and arrange envelopes alphabetically in boxes or a vertical file. You should also add to this file any cuttings from newspapers or magazines concerning those organizations. The file should be checked regularly, at least once a year, to remove out-of-date material. When writing to organizations for updated information for your database, ask them for any new literature.

Before throwing away the withdrawn material, check to see if it would be of use to some other organization or section of your service which collects material on local history.

Classification

It is not my intention to go into any great detail about the classification of community information materials. For most small information services, it will be quite adequate to arrange material by broad subject groups subdivided according to the form in which the clientele's needs and problems are presented. For ease of labelling vertical files, book spines and boxes, a simple notation can be used based on the initial letter or letters of subjects. For example, Cambridgeshire Libraries' community information service covers 21 categories:

Ben	Benefits	Gov	Government
Bus	Business	Hea	Health
Car	Careers	Hou	Housing
Corn	Community	Law	Legal rights
Con	Consumer	Mon	Money
CR	Community relations	SC	Senior citizens
Dis	Disability	Tra	Transport
Edu	Education	Ump	Unemployment
Emp	Employment	Worn	Women
Env	Environment	You	Youth
Fam	Family		

Some libraries also include as separate categories: Death, Equal opportunities, Fuel (or Energy), Gay rights, Leisure, and Trade union rights. These categories can then be broken down into smaller subcategories as necessary. Taking the earlier example of Housing, a possible range of subcategories with a mnemonic notation might be:

A	Agricultural accommodation	Re	Rented accommodation
H	Homelessness	ReC	Council housing
I	Institutional accommodation	ReH	Housing Association housing
M	Mobile homes	ReP	Private rented
O	Owner-occupied	S	Squatting

Numerals could be introduced for more detailed subdivisions, e.g. O1-Buying and selling a house; O2-Home insurance; O3-Planning applications; O4-Development; etc.

An effective method that can be used where you have only a small number of broad categories (no more than ten and preferably less) is colour coding. Allocate a colour to each subject and attach an appropriate coloured sticker or tape to the spine of each book or file box. Colour coding can be used in conjunction with notation to provide more detailed subdivision of a subject. When a community information service grows to a size where broad categories are not adequate to 'contain all the material and allow efficient retrieval, then it may be necessary to consider using a classification scheme.

Two choices are open to you, either using or adapting a ready-made scheme or constructing your own. There are a number of general classification schemes used in libraries, of which the most well known is the Dewey Decimal Classification. These schemes have been constructed to organize the whole of knowledge and subjects are usually arranged on philosophical or logical principles. Consequently, they may not be sufficiently detailed or treat subjects in the way your clientele express their information needs. Most schemes have a degree of flexibility and it might be possible to adapt them. The NACAB has its own scheme for arranging materials in the broad area of community information, but it hardly merits the description of a classification scheme. However, it is sufficiently detailed and client oriented to be worthy of consideration.

FOUR

What is Information Technology?

In so far as control over information forms the basis of managerial authority, the technologies which process it are inherently political. Indeed, business organizations are under pressure to change their decision-making structures as a direct consequence of information technology implementation. These 'subversive' effects on organizational structures are not immediately obvious to management. They arise as the new technology changes the established work practices and consequently creates different relationships on the organization's information network. Management is then faced with a dual threat:

1. A new pool of knowledge outside its formal control structures. Employees are able to access information without hindrance.
2. Power is shifting from line managers to those who control the information technology, such as the technicians in the DP departments or knowledgeable users.

Management's response has been weak and incremental. The control of information generated by computer technologies increasingly reflected the new power structures that these technologies established. On the one hand, functional control (allocation, budget, etc.) remains with general management. On the other, operational control (how and what is processed) is left to the technicians in the DP departments. The pervasive nature of the technology spreading through all levels of the organization established those who controlled it as important powerbrokers or gatekeepers. Newcastle Business School research has shown that senior managers are often unwilling to become involved. They served their apprenticeship in computerless or, at best, early mainframe-driven organizations and therefore do not always fully understand or appreciate the new technologies. Despite the obvious need for new control structures the traditional methods are adhered to.

The dichotomy between the established management structures and the changed (and changing) technological structures is therefore maintained and often worsened. Such opposing forces create a complex environment and place considerable constraints on management action. The dichotomy quickly becomes institutionalized. Lack of management action causes other user employees to perceive themselves as powerless in any dealings with the new technologies. This serves to enhance the influence of the technicians and is built into the organization's structures. Business action then tends to become increasingly inflexible; user managers are reluctant to take risks in an environment they do not understand, and many operational decisions are left to the technicians. This area of decision-taking is also built into the firm's task structure. Meanwhile the new networked systems demand increasing flexibility in order to work but are instead being faced with increasingly inflexible and decision-shy managements.

To appreciate the nature of this threat and how profound its impact can be on the organization, it is essential to grasp the characteristics of technology, in particular its integrative nature with any social entity whether it be an organization or something else. Explore this area, but it also then distinguishes between different types of technology, highlighting the perceptual characteristics of information technology. It concludes that to control information technology we must realize that we cannot treat it as just another piece of machinery.

The relationship between technology and an organization has, indeed, been thoroughly explored. It is obvious that a firm's structure is characterized by the many forces emanating from its operating environment. What is perhaps not so obvious is that technology is one of those forces. This is because it is quite often seen as a locally owned, physical resource. The machine on the shopfloor is just that and not part of a more complex social process. Within the framework of this perspective technology is a tool, totally predictable and controllable by management.

The physical technology alone (e.g. machinery or computers) will not explain how organizations are affected by it. As Weizenbaum suggested, a computer or telecommunication system (or any form of technology) is merely junk without the people to relate to them. It is how those people relate to their technology, how they use it, their perceptions of it, and so on, that are important. These aspects govern not only the technological process but also the characteristics of the organization structure. In the same way, structure should not be seen merely in a physical sense. The dividing of a firm into departments and functions, although of interest, does not capture the essence of an organization, which is to be found within the minds of its employees. The relationship between technology and structure is thus not physical but perceptual because it is within people's minds that the links are made. Our understanding of such techno behavioural forces is not great. This is partly because of the hegemonic view of business organizations as hierarchies in organization theory. It would seem, therefore, that we have not as yet learnt to examine organizations in an appropriate dimension which can highlight the effects of information and its technology.

The traditionalist's view of an organization is typically presented by Schein as a rational co-ordination of activities. This emphasizes the hierarchical and functionally-bound aspects. Butler, on the other hand, discussed the lack of understanding of such models determining that all the interrelationships are not adequately identified. They tend to portray firms as static entities without the turbulence of opposing forces. In particular, technology is seen as a side issue, if at all, and not central to the fundamental mechanisms of an organization.

A more appropriate model in Butler's mind is Udy's, which puts technology at the centre of an interactive environment of production objectives, social setting, physical exigencies and work organization. Technology, not the structure, sets the constraints ('trade-offs' as Butler calls them) upon an organization. Thus, a certain structure can be considered suitable for a given technology. For example, a company using mass production techniques, such as Ford, will best be structured to suit the needs of mass production technology. The relationship does not end there, however, as in Pugh's model, because he suggests that it is technology which drives the organization.

Where this model fails is in not identifying with sufficient clarity the perceptual aspects of an organization's technology. It is only within this domain that useful links can be established between technology and structure. The perceptual aspects of structure must also be emphasized for this to be successful since within any physical domain technology can only be seen as a resource and structure as a functional framework. It will then demonstrate that employees' perception of their structure determines the consequent structural characteristics. This chapter also suggests, however, that the relationship is two-way, in that an established structure can itself influence and change the perception of the employees. The impact of information technology is then again introduced. Already it has been established that such technology can change the way that individuals perceive their environments. A link with the present chapter is that since an employee's perception of the environment can be changed through information technology, his/her perception of the firm's structure must also change. The post-implementation structure that emerges may be profoundly different from any previous structure.

This control-theoretic view of information technology's impact is particularly relevant to the devolution of management authority. The impact of the new technology may radically affect the power structure or the political relationships between management and staff (but not the functional relationships) thus creating an imbalance between the perceptual and physical aspects of the organization. Through the perceptual domain, information technology has also infiltrated the cultural base of organizations and has thus had the potential to alter cultural objectives, which are the consequence of existing value systems. For example, it may be the formal objective of one firm to produce good quality products at the expense of higher profits. This objective has arisen from the organization's culture, but any form of technology, especially information technology, is not sensitive to cultural requirements.

Managers may try to fit the technology into their companies but it is generally within a framework provided by the physical domain. The objectives suitable for the technology do not always match existing cultural needs. To seek an explanation of the technological impact, therefore, we should no longer look at organizations in terms of departments and divisions. Any boundary established by such a definition may be rendered irrelevant. It is more sensible to consider an organization as a network of perceptual maps, dependent on an information network which is itself part of a wider network within the region, market, industry, and so on.

The need for such an outlook is underlined by the number of firms accessing information beyond their established functional boundaries. Internally, this involves different departments accessing a common pool from the company database; externally, employees gain entrance to public databases such as Prestel. It is increasingly difficult to understand these organizations as physically controllable entities. The problem, however, extends beyond management. If it is accepted that our models of organization are also too limiting because of their inability to highlight the dichotomy between the perceptual and the physical, then we must develop new models, which is one of the major proposals of the book. In order to do this we must first establish an appropriate methodology. The analysis which must be pursued is wide-ranging. It extends beyond the convenient boundaries that traditional theory has established. The tidy logic of an organization reaches only into the physical domain, beyond that is a vast perceptual arena which can be understood only as a network of relationships.

The Dynamic Element

Before its presentation, however, there are other factors which must be addressed. The point has already been made that an organization is a fluid and ever-changing entity. The components so far discussed form the framework of the organization but they are no more than boxes in an analyst's model. Other forces exist which breathe life into these concepts. They can be termed as the dynamics of an organization. People are the organization, and as such their perceptions, aspirations, alliances and conflicts form an important dynamic. It suggests that the outcome for any organization is not through a rational process, although that may seem to be apparent within the physical domain, but rather as a consequence of perceptual interaction between individuals or groups. These can be dysfunctional as well as functional.

For managers the issue is how such diverse forces can be controlled. They can only directly access the physical domain and therefore have to attempt to manipulate this in some way to have any impact upon the perceptual domain. One area, however, which allows a controllable linkage between the two domains is the information network. The quality, quantity and content of information, amongst other factors such as historicity, will determine an individual or group perception. On the other hand, the network itself can be physically manipulated thus affecting the quality, quantity and content of information. Its control is therefore of considerable importance to management.

Its particular emphasis is on the use of computer technology to generate information. Other forms of generated information through manual methods can also be controlled as information management, but because of the computer's ability to generate so much information other important forces come into play which are not present in a manual system.

The ways in which employees regard the characteristics of their information will determine in part the nature of their information network and ultimately the organization structure itself. Management must be aware of this and develop strategies through training and computer technology implementation which are sensitive to this. Information, however, is only part of the process. On its own it is not particularly a dynamic of the organization. Individuals may receive information and do nothing with it. A firm, therefore, does not change because of information. Information must be used. Action or behaviour is the consequence of information being used in the physical domain. In the perceptual domain, information being used is manifest through the decision process. There are no actions without decisions, no matter how minor or simple they may be, so that an organization, which can be analysed in terms of its actions, is driven by its decision processes.

In any context there is no such thing as an individual decision, even if the person involved is on a deserted island. Especially within a complex social process such as an organization, a decision is based upon a range of perceptual and physical influences. A person's experience coming to bear in a decision, for example, is the consequence of previous interactions with other people and institutions. Therefore one cannot fully control the decision process by controlling those individuals who are seen to own it. Their actions might be constrained but the source of their decisions is to be found in the perceptual domain. This is particularly important when such processes are affected by computer technology. Managers tend to develop mechanisms around the physical aspects of the technology and ignore the wider perceptual impact. It then places these ideas in the context of the other major components of a firm, in particular the information network, by highlighting the close relationship between a decision and information.

We have so far developed the framework of an organizational structure dominated by two domains: the physical and the perceptual. In a similar fashion to matter and anti-matter, their forces to Conflict and contradict one another. Management must maintain the narrow ground between the two, because it is only here that long-term stability is established. At this point the forces of both domains are in equilibrium, which for the purposes of the model has been termed balance. The achievement of such a state is one of the key outcomes to be identified.

The domination of one domain over another could lead to long-term instability, the degree of which depends upon a range of factors unique to specific organizations. For example, if there is no balance and an organization is dominated by its physical domain, then its short-term survival, at least, will depend upon how effective that domain is in dealing with the firm's particular environment. In the longer term, however, all firms must achieve an equilibrium whether the dominant domain is in tune with the environment or not. This is because the cost of not achieving equilibrium will be manifest in the dysfunctional relationship between the two domains. Forces will come into play that will undermine stability. It shows a complex social entity driven by an abstract relationship between two sets of forces. Much of this arena exists within the minds of the employees, of course, but like other societal phenomena its logic extends beyond them.

The forces can exist independently of particular individuals in that they are perpetuated by culture, technology and training. Individuals may come and go, some people may contribute more than others, but the logic of the organization as a separate entity remains consistent. It is proposed that an, organization can pass a significant point, the virtuality point, whereby its survival does not depend upon a particular set of individuals. This is a different world from the one portrayed by established theory. It comprises many of the same elements, such as structure and technology, recognized by conventional wisdom, but it is the perspective, the way in which those elements are arranged, that is different.

The emphasis is upon a force field relationship rather than the aspects of one domain or another. By having this different framework, we can also analyse the effects of change in other ways. In particular, the impact of information technology is shown to have a far wider effect upon organizations than the traditional models may lead us to suppose. Much of it is common sense, gleaned from the working environment and from those in it. Nevertheless, the model is placed before all management as a challenge: to use the ideas this book presents as a basis for analysing the worth of their information technology.

FIVE

Organisation and Structure

The explanation of organizational structure was left to two mainstream approaches. The first was the economic model whereby the business organization was subsumed into the market. Structure was thus defined as an extension to the market. The second was through the Taylorist and management science movements, where organizations were seen as no more than vehicles for management control. In neither case was sufficient analytical sensitivity provided to appreciate the close-relationship between structure and its technology.

Today, different ideas about organization are prolific. The traditional views are still around but there are many others which have greater sophistication and which span a range of academic disciplines. They can be condensed into three major perspectives or metaphors. Scott, for example, identified these as follows. Firstly, organizations are rational systems, based upon some collective rationale such as profit maximization or growth. This perspective also embraces the two approaches already mentioned above. Secondly, organizations are natural systems; arising from the behaviourist movement of Elton Mayo whose perspective views organizations as a living organism. Thirdly, organizations are open systems which are seen as processes being open to the environment.

Watson, on the other hand, identified 'waves' of organization theory. The first sees organizations as machines, portraying the view that, like machines, they are predictable. The second wave sees organizations as organisms to demonstrate that they are as complex and unpredictable as a living organism. The final wave sees organizations as processes, thus highlighting their dependence upon interactive behaviour. Others, such as Gareth Morgan, have gone further in identifying metaphors of organization (brain, machine, psychological prison, etc.) that we use to aid our understanding of them. Whichever perspective is supported, they all agree that organization structure provides some form of framework.

Mintzberg's book places considerable emphasis on structure and his definition proposes it as the summation of the ways in which a firm's labour is directed and coordinated into tasks. Although this definition's simplicity is appealing, it leads us to view an organization in terms of a division of labour and the co-ordination of management control to maintain that. It is an outlook typically portrayed by the firm's organization/ management chart, which identifies quite precisely who is responsible for what and to whom. This merely lists a functional relationship, which although important, reveals only part of the organization structure. Handy hints that there may be more when he talks 6f linking mechanisms between roles and co-ordinating structures. Although the elements of line control and co-ordination are still present, Handy extends his definition to cover the perceptual domains of an organization. This can lead to a more open definition. The concept of linkages increases our awareness of both interdependence and independence between factions; of how resource groupings can either contradict or complement other resource groupings. Linkaging is, therefore, necessary to cement the otherwise contradictory elements to prevent dispersal, or to reaffirm the complementary to prevent them becoming contradictory.

Idea of a Boundary

Linkaging, therefore, provides the organization with its completeness, establishes a boundary (perceptual or physical) which can be identified as being the domain of a particular organization. The idea of a boundary is an important aspect of structure since it enables us to determine how integrated the structure is with both its internal and external environments. For managers also, the boundary is useful since it enables them to identify their area of control. However, with the increasing implementation of information technology the issue has taken on greater significance. Boundaries can no longer be so easily identified. Modern technological systems directly link organizations into other organizations, thus merging in part their established boundaries. For example, a major retail outlet requires all its suppliers to link into their own stock control network. Those firms have thus become tied into their customer. In this situation it is difficult to determine where the supplier's boundary ends and the customer's begins.

Robbins proposes a traditional view of the relationship between a structure and its boundary. He determines the structure being defined by its boundary in the same way that a human skeleton defines a person. However, by equating boundary with human skeletons he has missed an important point. The body of a person is a finite concept in that it is spatially bounded within the physical domain.

On the other hand, an organizational structure, although providing a framework, is as much perceptual as physical. That is, it exists because its members perceive it to do so. To someone who does not understand the concept of an organization, the distribution of its physical resources has no significance and thus they will not see an organization. This problem is not as remote as it may appear. It is common, for example, for small business organizations to share resources such as office space, equipment and even staff, therefore possessing the same physical structure. There is no suggestion, however, that they are the same organization. Their differentiation is thus established by the perception of the actors (employees, suppliers, customers, etc.) and not the physical aspects. Organizations are not like people even though they comprise them. They are not necessarily spatially bounded; their buildings, for example, do not represent their outer edges. Indeed, some organizations have no formality about them or physical structure, such as some political or religious institutions, but they are just as real as a small business enterprise. Organizations are not bounded by a temporal framework in the same way that people are. Individuals have their historicities which influence their behaviour, but they are creatures of the here and now. Their socialization has made them believe that all there exists is the physical domain which is governed by the present. The operational levels of organizations are also bounded by time and the physical domain. However, the structure is not. It exists in the minds of its members and therefore is in the present, but it is also equally valid in the past because its existence is just as dependent upon past interaction.

For example, a firm's production process, which forms part of its structure, is physically bounded by time. It can produce only so many items within an hour or a day. This is its operational level—the dimension of doing. But the same process is time independent because how it can produce is and will be determined by decisions taken in the past, the present and the future. Structure, therefore, has to be more than a static framework, of functional and co-ordinative hierarchy, because organizations themselves are fluid, perceptual entities and as such could not be supported by a rigid framework. Robbins' model of complexity, formalization and centralization does suggest a certain fluidity in structure, but only to the extent of differentiation between particular organizations and not within the unit itself. In other words, Robbins recognizes an infinite variety of structures befitting the mode of operation, but once set they provide a rigid framework in the short term which can be changed only in the longer term.

Organisation's Structure

Watson's ideas about an organization's structure provide a better match for a considerably more complex form of structure than so far developed. He does not see structure as a given framework but rather as a fluid, abstract entity which is the outcome of ever-changing employee interaction. Such a view can present no definitive solution to the quest for perfection, because the outcome is dependent upon so many variables.

With such a cost comes also a benefit. There is a freeing of structure from the constraints of 'conventional wisdom', by allowing greater flexibility in organization design and therefore enabling a more appropriate framework to be established for the analysis of the impact of other forces such as technology. Watson's words reflect a school of thought which considers the only reality of an organization is to be found within our minds. The structure we put on that organization is also within our minds. It will change as people's actions and perceptions change, and can never be completely controlled. Innovating forces like information technology, for example, can have the greatest impact upon an organization's structure through the minds of its participants. Structural change could occur through differing actions/perceptions of the employees, which may be beyond the control of management.

Structure must be seen to be more than a rigid framework. It may well be that some organizational structures have become increasingly flexible in the past few years as a consequence of their information technology systems. However, it would also seem likely that structures have never been as rigid as many would suggest. Some are more stable than others within a given period. Others are more responsive to certain situations such as crisis or particular market conditions. They may all have areas of rigidity within them which can subsequently cause problems. Whatever their apparent characteristics they are manifest not because of the way we organize our physical resources but rather as a consequence of the way we relate to those resources. Structure is indeed in our minds. As such any boundaries which are obvious to us are there only because we perceive them to be. If we then introduce a technology with a strong perceptual impact, such as information technology, it is not surprising that a change occurs in the way we view our organization. Such changed perception may lead to a changed structure.

The idea that organizations are perceptual maps has implications for their analysis and design. How we view them, and hence how they manifest themselves, will depend upon the particular organizational model we use. Thus, in a self-determining way we align our particular organization's characteristics, as far as possible, to those found within the model which matches our own opinions. For example, if we believe along with our colleagues that the most realistic metaphor for an organization is a machine in terms of formality and controllability, then our behaviour would be governed by that belief. The organizational structure would then be affected and to a limited extent changed by such perceptions if there is not a match. Problems can occur, of course, if these perceptions are not fully disseminated throughout the organization. The models of organization we have, therefore, are themselves affecting the environment which they are attempting to analyse. From this aspect alone it is worthwhile determining what these models are. However, we need to also consider them in terms of their suitability for the impact analysis of information technology.

The mechanistic view of organization is the first wave of theory, which comprises the metaphor of organization as a machine. Like a machine, its characteristics are determinable and controllable. The degree of its formality is not considered because all actions and behaviour within it are considered to be purposeful and co-ordinated. In other words, organizations are rational processes. Therefore the underlying assumption is one of an organization as a formalization of a need to meet certain objectives. Such objectives are seen as rational in that they are based upon an obvious range of criteria suitable for the type of organization. Thus, a business organization's objectives could be to maximize profits, to grow or to dominate a particular market or industry.

The emphasis is on goal attainment and its management. The profound question presented by this metaphor is, why are some managements better at achieving their objectives than others? This metaphor evolved from the economic model in which organizations are perceived to be unitary entities pursuing the same collective objective, controlled and driven by market forces. As such, organizations do not have any analytical value other than as a unit of the market. From this, however, four softer perspectives developed, each a response to the demand of business managements at the time for better understanding.

The first is Taylor's scientific management. The second is administrative principles proposed by people such as Henri Fayol. The third is Weber's theory of bureaucracy, and lastly Simon's discussion on administrative behaviour. There are many books which cover these areas quite adequately. Apart from an introduction below, these will not be dealt with in detail. The emphasis will be placed on examining their approaches to technology.

The founding father of the scientific movement is claimed to be Frederick Taylor and subsequently developed by others such as Lillian Gilbreth. Their perspective is an extension of the economic model by recognizing that organizations are not determined solely by market forces but can be moulded by management activity. They prescribed a methodology for action through a scientific analysis of tasks in a firm so that formal procedures could be established for optimum efficiency. The approach, in short, is prescriptive in that it presents patterns for management behaviour. Although lacking in fundamental academic support the traditions of scientific management are still reflected in many of our management training programmes.

The 'modern' view is more developed in its recognition of greater uncertainty but nevertheless still portrays a controllable environment. The elements of that environment are also controllable, Thus, the notion of technology is coloured by the certainty of its predictability. Technology in all its forms is a management tool. The original exponents of scientific management thought this, as do their followers. There is no relationship between a structure and its technology because neither is seen as a social process but rather as building blocks for management use. Thus, the impact of innovations such as information technology is not seen to affect the organizational structure as such, but rather the relationship between management and workers: any impact is regarded as an input to be controlled.

Innovation, therefore, changes only the way in which management operates (although this might have an indirect effect upon the organization itself). Impact assessment of information technology is in terms of functional analysis or efficiency criteria (how good or bad its contribution is to management objectives), the assumption being that the technology will always fit into existing organizational structures with the appropriate implementation strategy. The emphasis, therefore, is on performance measures and the ability of information technology to meet them. There is a range of literature which explores the scientific management viewpoint of information technology. Its presentation is prescriptive and consultative.

The remaining three perspectives within this metaphor are dealt with under the one heading. They are the models initiated by Fayol, Weber and Simon. Their collation is not because the themes they support are necessarily in agreement, although they do cover broadly the same principles, but rather because their implications for the analysis of technology are similar. The major difference between these and the previous approaches is that scientific management sought to establish the organization as a management phenomenon whereas Fayol, Weber and Simon sought to establish the organization as an administrative process.

The immediate implication of this is that, as an administrative process, the organization is less controllable and less predictable. It is no longer a tool but a set of administrative interactions. The objective of this group, therefore, was to identify such interactions and then to determine their appropriate control. Fayol was perhaps the first to recognize the importance of the administrative process in understanding organization. However, his work was not translated until the late 1940s and thus not widely known.

A sociologist, Max Weber, was also developing an ideological set of administrative principles. He identified bureaucracy as being the ideal form of organization. The formal and impersonal structural characteristics allowed for the most appropriate administrative processes and for an efficient achievement of objectives coupled with a certain humanity. Although advances were being made in other areas, it was not until the late 1950s, with the appearance of Herbert Simon's ideas, that advances in this area were made. He sought to portray a more open administrative process in which the actors are not machine-like in their rationality but more human in pursuing their own interests and in many ways ignorant of their total environment. He saw their rationality bounded by the complexity of the environment and thus their decision-making was not perfect or indeed predictable. Simon, along with others such as Cyert and March, perhaps provided the link between this metaphor and the second wave. They viewed people as being less predictable than did their forebears, but still nevertheless determined the organization itself to be of machine-like stature.

Perhaps by default, technology is, therefore, more integrated into the administrative mechanisms of the organization. There is concern about technological fit, thus the approach is directed away from one of expert handling to one of greater integration. This is particularly evident in this view's handling of information technology. The literature which reflects it concentrates upon the user end: it is here that a controllable interface exists between the administrative process and the machine. A classic example of this view is to be found in a book by Whisler, although it is becoming outdated. In the 1980s, with the increasing power of computer technology, many other books found their way onto the market. They all sought to merge the computer with the work processes of the organization rather than to allow it to dominate them, but stopped short of seeing it as a social process.

Peter Blau determined that it is irrational to view an organization as rational in the way that the machine metaphor does because it does not take account of the non-rational behaviour of people. Such a perspective summarizes the natural or organismic metaphor, in that organizations are seen as complex, living entities which do not have machine-like qualities. They are no longer reified, cohesive units which act in harmony and with total rationality towards a common set of goals. In the same way that a living organism is a collection of competing and interacting forces, some living and some not, so too the organization comprises individuals and social forces which compete with and contradict each other.

The consequence is that the organization cannot be so easily prescribed for management action, being at best a rather fuzzy area for common agreement. Technology, therefore, as an element of organization, cannot be so predictable in terms of a resource. Neither is there agreement on how to model the metaphor. With some justification it is claimed that Elton Mayo provided a foundation with his Hawthorne experiments. As social psychologists, their contribution was in recognizing the importance of behaviour in organizational design. Their work initiated the behaviourist school for the explanation of firms.

The offshoots are many and are still very much in vogue today. Since Weber's work the sociologists have also been major contributors to this metaphor, particularly through the work of Selznick and Parsons. Rather than placing the emphasis upon the people element of an organization as the behaviourists did, these and other theorists who followed were interested in the complete entity and how it worked. In true functionalist tradition they identified various component parts and studied interaction between them. Selznick highlighted the relationship between these and the organization's employees, whilst Parsons developed a dynamic model in which people were an integrated part. The technological tradition which grew from this metaphorical approach is less prescriptive and more analytical than that of the machine metaphor.

The earlier ideas are typically represented by the works of Gouldner or Wedderburn and Crompton. Theirs is mainly an examination of production technology and its relationship with the workforce. The structure of an organization was a consequence of the way its workforce produced, in particular of the relationship existing between workforce and technology. However, a more formalized development of this took place in the late 1950s and early 1960s which was influenced in part by Elton Mayo and also by considerable empirical research. Although not labelled until much later, the approach became known as contingency theory. In contrast to the classical approach (whether the machine or natural metaphor) which sought to achieve a universally best structure for a business organization, the contingency theory proposes that there is no one best solution.

Different situations demand different organizational structures, which are governed by forces of change known as contingencies. Thus, the prevailing contingent would shape the corresponding organizational structure. For Woodward, after extensive research of over a hundred firms in Essex, the most important contingent is technology. If, for instance, a firm has a mass production technology, then its structure would develop in terms of the other contingents, such as size and culture, which would best suit that. So mass production technology, for example, by its very nature determines that firms generally be large with formal cultures. Quite apparently, therefore, technology is seen by Woodward as a determinant of organization. Although others, such as Pugh and the Aston Group, do not agree with her upon the extent of the power of technology to determine structure, there is a general consensus on the nature of technology itself. It is no longer a resource as such, but as much an element of organizations as structure, people and their culture.

Information technology is perhaps not addressed as positively as it is in the previous metaphor. This is because there is no formal distinction made between general technology and information technology. Both are social processes and therefore to be treated as such. That is not to say there are not writers on the subject. Miles *et al.,* Anderson and Mumford and Zuboff are typical of the works produced. Information technology is analysed in terms of its ability to integrate into society or business. In general, the conclusions seem to be doom-laden, predicting dire consequences if we do not get it right.

Organizations as Processes

This metaphor describes the organization as a complex conglomerate of interaction which does not necessarily produce a clearly identifiable form. An organization exists because people are told it does by management teams, corporate identity, and so on; although, in reality, apart from resource allocation, members/employees act very much for themselves as individuals or within groups, occasionally forming alliances to get things done. It is the process of doing (production, administration, etc.) which is the major input to organization structure, so that the organization is a series of different or interlocking processes. Its structure is shaped by those processes because it is the arena in which they are actioned. As with the other two waves there are disagreements, in particular over the identity of the process. Some put forward the decision process as a dynamic which determines the organizational structure. Others view the organization as a political arena in which structural adjustments are made through political mechanisms such as ownership of power. Yet others, such as Johnson *et al.,* regard these processes in terms of their relationships with one another and how open they are to external influence. Contingency theorists Lawrence and Lorsch, on the other hand, determined a particular relationship between the organization and its environment as being the important process, whilst Weick highlighted the processes between the individual and their social formation.

In general, organizations are no longer seen as hierarchical and solid (albeit at times fuzzy) frameworks. Structure is considered not to be as important as the processes themselves because it is a transient thing, created and remoulded by the whim of changing interactions. Watson described this organization as being formed in the minds of the employees through their perception of its interactions. The process or open systems metaphor, therefore, shifts the emphasis of analysis from the design of the organization to the functioning of its elements. It allows these elements (e.g. technology, politics, decisions) a far greater impact than the previous two metaphors. For technology in particular the approach is significant. In the machine metaphor, technology is not part of the organization, whilst in the organismic metaphor technology is part of an organization's social process and a determinant of structure. In the process metaphor, technology can be a process in its own right, determining the perceptions of individuals and affecting their behaviour. Its ability to affect the characteristics of an organization are thus considerably increased because the mechanisms involved (i.e. employee perception) form the essence of an organization. These ideas have been most prevalent in the analysis of information technology. The relationship between the behaviour of people and the technology they use is portrayed by the socio-technical systems theory. This applies a systems methodology to the analysis of information technology and how it affects the way we organize. Its roots can be traced back to Mayo and his Hawthorne experiments in terms of the paradigm's behavioural aspirations. It was Mayo who first formally stated the link between our behaviour and organizational design. But it was Trist and Bamforth who galvanized those ideas into understanding the impact of technological innovation.

With the growth of computer technology as the dominant innovative technology it was, therefore, not surprising that Trist and Bamforth's methodology was quickly taken up by others to analyse its impact. Burns and Stalker produced an impressive model based upon empirical research describing a close behavioural relationship between our technology and the way we organized. If the structure of the firm did not suit the technology employed then pathological responses evolved on the part of management.

Seeking a Solution

Just as firms are not merely a collection of resources, human or otherwise, nor are they merely a collection of processes. As with all social entities, they are a combination of physical and perceptual: but it is how the relationship between the two is understood which is important. Quite obviously many organizations do have buildings which figure strongly in the image they project but also quite obviously the influence of the organization does not end at a brick wall. They do comprise physical beings constrained by their physical environment but beyond such environments they have resources which can be both physical (machinery, etc.) and perceptual (knowledge). Any complete model of organization must accommodate both domains. Organizations, however, must be more than the summation of individuals, resources and perceptions within them. Unfortunately the metaphors discussed describe no more than this. Yet organizations tend to develop some sense of being in their own right. This is not in terms of a conscious, spatial positioning, as any individual would have, but is rather a consequence of history.

Continuity of an organization, for example, cannot be solely explained as the perception of individuals or the particular allocation of resourcing. There is a logic within it which outlives any temporary input an individual may care to make. It is indeed a process, but one which is at the same time both independent and dependent. For example, previous decisions made by management can constrain and develop the organization's structure in future years. Both the physical and perceptual aspects of the decision as real as each other. One could be seen in the production line whilst the other, although not seen, is ingrained within the management culture. Together they form ail inheritance passed onto subsequent management which constrains their action. It is in this sense that an organization can be said to be existing independently. Decisions and actions are taken all the time which can then exist beyond their initiators and add to a bank of knowledge within the firm. To a certain extent we are referring to the organizational culture, but culture's form relates only to the rules for action and does not express the logic of the actions themselves, as does structure. Structure can thus be determined by neither domain alone. More probably, it arises as a consequence of a relationship between the two domains continually interacting with one another through a historical process which in turn feeds the logic of the organization. In the same vein, information technology (or technology in general) cannot be treated as either a physical or perceptual phenomenon.

The way in which each of the three metaphors provides analytical tools for the assessment of its impact restricts our being able to see the whole framework. Information technology does have a strong physical impact: it can, for example, turn an inefficient firm into an efficient one. The criteria used to assess such a change are based in the physical domain in that there can be more information output in terms of management reports or invoices billed, and so on. It can also change user behaviour, as we see in socio-technical systems, so that, employee perceptions of their tasks or indeed their firm change. In many models the two domains sit either separately or uncomfortably side by side.

However, it is not so much the consideration of the individual domains that is important (although obviously for line managers it cannot be ignored), it is the relationship between them which must be analysed and the subsequent effects such interaction has upon organization design.

SIX

Computers in Use

Recent years have seen both a decline in the price and a rise in the power of personal computers. You can get a fairly basic word-processing kit, including a monochrome monitor, dot matrix printer and word-processing software, for around £400; Whereas £1,500 would buy you a powerful and sophisticated IBM-compatible personal computer (PC) with a built-in integrated suite of software (word processing, database, spreadsheet, graphics, communications, etc.), DOS and Windows operating systems, colour monitor and a bubblejet printer.

There may be a temptation for a community information service to assume that a computer is now a necessity for its work, before asking the obvious question: 'What do we need one for?' Consideration also needs to be given to the question of who is expected to use the computer. If your service relies on volunteers, they may have an aversion to new technology or simply not work long enough to become familiar with and confident in using a computer. Before plunging into the murky waters of information technology (IT), you should first of all ask the kind of basic questions that needed to be addressed when setting up your service, such as 'Why do we need IT', 'What functions do we want it to perform?', 'What system is going to meet our needs?' 'Who is going to use the system?', 'Where can I get advice?', 'What training will be required?', and so on. There are many reasons why an information service might want a computer but they may not necessarily all be the right reasons. Some wrong reasons are:

— everybody else has got one;
— to improve the image of the service;
— to sort out our database /statistics/ financial accounts, which are in a mess.

This latter brings to mind the computer cliche 'garbage in, garbage out', namely that if you put rubbish into a computer, you can expect to get rubbish out. Before inputting any information into a computer, you must first have a logical and organized manual system. The computer won't do it for you.

Some right reasons, might be:

— to provide faster and more timely access to information;
— to facilitate access to a wider range of information;

— to provide more sophisticated access to information;
— to enable more people to have direct access to information.

Before making a decision on what equipment to purchase, consider what functions you want the computer to perform. These might be one or several of the following:

— database management;
— recording statistics and, possibly, producing these in graphic form, i.e. as bar or pie charts, graphs, etc.;
— running welfare benefits and other programs;
— running a local view data system;
— online access to other external databases;
— word processing;
— financial accounting;
— desktop publishing.

If you require most of these functions, a fairly powerful machine would be necessary but you might be better advised to consider more than one computer or even linking them in a local area network (LAN). A LAN enables personal computers to communicate with each other and share resources such as software and printers. Usually, a LAN is driven by a computer with a large hard disk which contains the software programs and the data input by the rest of the PCs or terminals on the system. It also provides access to shared printers on the network and to central back-up facilities.

Choosing the Right Equipment and Software

There are several ways in which you might set about choosing a computer, depending on the level of knowledge of staff in your organisation. Let us assume elementary or no knowledge of computers, in which case the options might be the following:

Dealers-may be OK if you know one who can be trusted or is recommended, but computer salesmen, like car salesmen, frequently try to get you to buy a machine of a higher specification than you need and blind you with technical jargon. Also limits your ability to shop around.

Consultants-usually expensive unless you can beg or borrow their services. Have a tendency to recommend separate software packages for each function or even writing a program to match needs exactly, rather than considering integrated packages which may require some compromise. OK if all the separate parts can work or are not required to work together.

Other organisations in the same sector who use computers in their work may be able to recommend equipment and software from their own experience. However, they may be reluctant to admit shortcomings and the equipment/ software they use may be out of date, given the speed of developments in this area.

Networks who support advice services or community and voluntary organisations, such as Community Computing Network (65 Litchfield Road, Cambridge. CBI 3SP) or the Computer Development Unit of LASA (2nd Floor, Universal House, Wentworth Street, London El 7SA Tel. 071-377 2798). Apart from giving direct advice themselves, they may be able to put you in touch with others who can help.

Within the space of this book, it is not possible to consider all the potential uses to which a community information service may wish to put a computer. I have therefore concentrated on those aspects which, though not peculiar to information services, are a major feature of their work, namely database management (including view data and access to external databases), welfare benefits programs and enquiry statistics.

Database Management

There are a number of reasons why a community information service would want to keep its information file on computer:

— beyond a certain point, a manual information file becomes unwieldy to use and update, whereas a computer database offers speed of access and ease of updating;

— complex searches involving, for example, subject, place and clientele, are extremely difficult and well nigh impossible to carry out with a manual system. A good computer database management system allows searching of the database by a combination of terms;

— individual records or subsets of the database can be printed out on demand in a variety of formats, e.g. as mailing labels, lists, standard letters;

— through networking, online or exchange of disks a computer database can be shared with other workers or organisations;

— customers can be given direct access to the database through public access terminals in your centre(s) or through a viewdata system;

— entries can be fed into a word-processing or desktop publishing package and printed out as leaflets, directories, handbooks, etc., either directly or, more usually, by preparing a master copy which is then printed by offset litho.

A computer database is made up of two parts—the software, sometimes referred to as DBMS (database management system) and the database file, which is very much like a manual card index system in that it contains individual records made-up of a number of fields (e.g. name, address, telephone number, hours of opening, etc.), into which bits of information (data) are entered (input). Most database systems store data in a highly structured form but there are some exceptions, known as free-form' or text-retrieval databases, which automatically index all significant words in a document.

Some database systems allow information to be viewed or edited either in tabular form, with each record represented by a row and fields as columns, or a record at a time. The tabular approach is fine if the fields in each record are short but is less convenient where you have lengthy fields. This can usually be circumvented by limiting the number of characters in each field, but then you cannot see the record as a whole. There are basically two types of database management systems: a simple DBMS produces a single file (sometimes called flat file) of records that have a predetermined number of fields and field lengths. This does not mean that field lengths cannot be varied or fields added or deleted, but any changes affect all the records in the file. A flat file may be adequate, say, for a database of community organisations, provided you have reasonably anticipated all the fields and field lengths that you are likely to require.

However, it will not be suitable for a situation where you want to create several different files and have the ability to link them. For example, in addition to the database of community organisations, you may want databases of leaflets, booklets and pamphlets; 'hard' information; individuals who can provide help; and enquiries plus the ability to search on all of them simultaneously. In those circumstances, you are more likely to need a relational DBMS, which would enable you to form links between the different files through the use of common denominators, such as subject or classification fields, client reference number, etc.

Examples of flat file systems are Cardbox Plus (used by Help for Health for its database), Q&A, and FileMaker Pro. Most database functions that come as part of a suite of integrated software, such as LotusWorks, will be flat file systems and may not be sophisticated enough to cope with anything but a simple database. Some of the leading relational systems are DataEase, Paradox, and dBase systems (dBase IV, FoxPro, Clipper). I have used for a number of years a system called Superfile which, at the time, had the advantage of not requiring any programming knowledge from the user. However, most of the database systems have become more user-friendly over the years.

The price you can expect to pay for a DBMS can range from £ 250 to £ 750 but some suppliers offer substantial discounts to voluntary organisations and charities.

Choosing the Software

Before choosing a database management system, you should first of all draw up a shopping list of requirements and questions that you need to ask, such as the following:

— Is there a maximum size of file that the system can handle? In most cases this will depend on the size of your computer's memory but some systems may operate more slowly as the database grows beyond a certain size.

— Is there any limit on the size of records, the number of fields in a record, or the size of a field? If so, is this likely to be a problem?
— Is it possible to add, expand or delete fields from a record without difficulty?
— Is it possible to print out all or a selection of records and fields within records in a structured format?
— Can the data be fed into other software packages to produce lists, standard letters, mailing labels, etc.?
— What features does the system provide for searching records? For example, wildcards (* ? =) are a means of truncating search terms to take account of varying forms or spelling, e.g. ORGAN will find any first word in a field beginning with the letters O R G A N, such as 'organize', 'organisation' etc.; 'sounds like' will find records with similar sounding words to the search term, e.g. FAIR and FARE; or Boolean logic will logically link or exclude terms to narrow the field of search, e.g. A and B, A or B, A-G but not C, etc.
— Is there a facility to check on the validity of entered data?
— Is there provision for constructing user-friendly menus?
— How much knowledge is required to set up a database? How long does it take to acquire that knowledge? Can it be self-taught?
— How good is the documentation that accompanies the system?
— Is there a multi-user version of the software?
— What support is offered by the supplier? Is there a user group?

Creating a Database

Having bought your software and installed it on the computer, the next step is to create the database(s). Each database management system will have its own methods for setting up the database and this section can give only a few tips from experience that should apply to whatever system you have chosen. First determine a structure for your database. If you are moving from a manual database, that structure may already exist.

The method of determining a structure for a database is similar to that for a manual system, as set down in Chapter Four. However, in addition to identifying what fields you require, you will also need to determine field lengths. Be generous, if your system will allow, so that you are not continually changing field lengths. Separate fields will need to be given field names or tags to which the computer attributes information that is being input.

For example, you might decide to call the 'organisation' field ORG and the 'address' field AD. Where a field consists of several lines, such as an address, it is advisable to split these into separate fields, particularly if you want to search on individual elements of an address field or to print them out as separate lines. An address field might thus be divided into two or three separate fields for building/house name, street, community/ neighbourhood /village and fields for town county and postcode. So your tags might look like AD1, AD2, COM, TOWN, COUNTY, PCODE. In some cases, a record might contain several repetitions of the same field types, e.g. names and addresses. Every time these occur you could give them the same tags, but this may cause problems when conducting searches or printing out information.

For example, if you wanted to see or print out a list of organisations in, for example, the 'Parkside' neighbourhood, the computer would also throw up records for, say, the secretary or second contact who lives in the 'Parkside' neighbourhood, if they all have the tag COM. The simplest way to get round this is to add to the root name for tags a separate number for each occurrence within your record structure, e.g. COM1, COM2, COM3, etc.

There are occasions when it is advantageous to use the same tag for several fields:

1. You might, for instance, want to have several fields for subject headings. By giving them the same tag name, this will ensure that, in whatever subject field a particular term is entered, the computer will throw up the record when a search is made under that term. It will also ensure that when records are printed out, the record will appear under each subject heading.
2. When you want to link files in a relational database system, this can be achieved by having a field with the same tag name appearing in each file, e.g. subject, name of organisation, client, etc. A search on that particular field will throw up records in all the files bearing the term entered.

It is advisable to have the database form in the same order as the form that you send out to organisations for their information, as this makes it easier for the person inputting the information. One of the advantages of having a computerized database is that it enables you to automate administrative processes.

For example, update letters can be customized and printed out as you determine. If you want to update continually throughout the year rather than in one session, then you will need to build into your record structure, field that enables you to do this, e.g. date record last updated, or added, AGM date. Dates must be entered in a consistent form. Alternatively, you could update a section of the file each month by the initial letter of each organisation, e.g. Month 1: A-B, Month 2: C-D, and so on. When sending out update letters or forms for information about new organisations, if people's names and addresses are going to be recorded, make sure that you conform to the requirements of the Data Protection legislation by including a statement to the effect that: 'This information is being stored on a computer. Please indicate any items of information that you do not wish to be recorded.' Details of the Data Protection legislation are available in a series of free booklets from the Office of the Data Protection Registrar, Springfield House, Water Lane, Wilmslow, Cheshire SK9 5AX.

When sending out an update letter, it is also helpful to send a printout of the organisation's current entry as a prompt for any changes that may have occurred and to save the secretary time in filling out information that has not changed.

All database systems will have some facility for sorting records according to predetermined parameters, e.g., alphabetical by subject, then alphabetical by name of organisation or by date, record number, place, etc. You can usually set up several sort codes for different purposes. The sorting facility is used mainly as a prerequisite when printing out records as a list. For this, you will use another element of a database management system, the report function. (name may vary according to DBMS used). This enables you to print out information in a predetermined format, using all or a selection of fields in a record. Again, more than one report can be set up to take account of different uses. In some integrated packages, the database records have to be fed into the word-processing package for printing out.

Uses of the Database

You will need to consider potential uses of the database. Is it a tool for the staff alone to use in answering enquiries? Are you going to allow public access to records either as printouts or through an online terminal? Are you going to make the database available to other organisations through exchange of disks or even sell it to interested parties? The answers to some of these questions will have implications for the kind of hardware and software you choose, and therefore need to be decided at an early stage.

For example, if you decide to allow public access via an online terminal, apart from the additional terminal and keyboard you will also need a multi-user version of the software which must be capable of building in safeguards to prevent the public either accidentally or deliberately erasing or changing records. Some database systems offer run-time versions of the software at a substantially lower price. These allow the system to be run but cannot be used to make changes to it or to develop something new. A good example of this is Help for Health's Cardbox database, where subscribers receive a run-time version of Cardbox and regular updated disks of the database. Obviously, a run-time database will not be as up to date as a true online system but could be a cheaper and safer alternative. You can also get video-text packages which run on PCs and enable you to marry attractive graphics with a database. A good example is the TAP (Training Access Point) terminals located in various public buildings to dispense information on training opportunities, both local and national. The national information was obtained through online links with remote databases. The local information database is maintained on a master PC and then transferred by disk at regular intervals to the PCs driving the TAP terminals. The advantages of this kind of system are:

— eye-catching graphics can be used to overcome the public's reluctance to use computers;
— protects your main database;
— less staff-intensive.

The disadvantages are:

— the information is not as up to date as a true online system;
— a PC is occupied just in running that function;
— it can become monopolized by young people who think it's the latest arcade game;
— if a printout facility is offered, this can easily be abused, requiring frequent staff attention;
— there is a need to make the equipment inaccessible or tamper-proof

PC-based video-text terminals have been used to good effect by some commercial organisations to provide local information and advertising. You find these most often in the windows of tourist information centres, the Computer display being operated by a touch-sensitive keypad attached to the window. Another variation is to get a graphics package, such as Animator, which has the facility of enabling the computer to rotate a series of frames of information at set intervals. Several years ago COIC (Careers & Occupational Information Centre) offered a video-text package called Rotaview which enabled users to create colourful graphics and link these with a database.

The system was used to prepare an off-line database, to access or download frames of information from Prestel, or to operate as a mini-viewdata service with users dialing into the database for information. Another attraction of Rotaview was its carousel feature that enabled up to 80 frames of information to be scrolled at intervals. A particular use for this type of feature in an information centre would be to display essential information to the public out of hours. Rotaview was marketed at schools, colleges and public libraries but, alas, was developed only for the BBC micro and is no longer available. However, with infinitely more powerful PCs around today, there must be similar software on the market, if you are interested in this way of displaying information.

Other Database Systems

Using a PC is not the only way of setting up a community information database. There are other options which may be open to services which are part of a much larger organisation, like public libraries, and are designed to cope with extremely large databases and/or allow wider access. In the early 1970s British Telecom (or The Post Office as it then was) developed a viewdata system called Prestel which enabled users to access thousands of pages of information on all sorts of topics via the medium of the telephone and specially adapted television sets (and later also PCs). The system has never quite lived up to its early promises but is still going strong. A number of public libraries were involved in the early trial of Prestel and recognized the potential of the system, if not of Prestel itself. Eventually, local authorities began to set up their own viewdata systems of information about council services and other local information. The public could access the system either directly from home or through terminals located in libraries and other public buildings. In some instances, the library service is responsible for maintaining the database, in others another department of the local authority. Viewdata systems offer the following:

- access to large amounts of continually updated information;
- use of graphics;
- public online access from home, workplace or publicly located terminals;
- centralized updating of information.

Some disadvantages are:

- they need dedicated teams to maintain the database;
- many clientele of community information services would not be able to afford equipment to access the database or would be afraid of using a computer;
- limitations on the amount of information that can reasonably be presented by the system.

Some public libraries, like Berkshire, are beginning to recognise the income-generating potential of viewdata, with electronic Yellow pages, teleshopping, rental pages, and closed user group facilities. Apart from running viewdata systems, mini or mainframe computers can also be the hosts for community information databases using such software as IBM's STAIRS (Storage and Information Retrieval System). Some libraries set up their databases on a mainframe computer from the start; others, like my own, have transferred from a PC-based system. Some advantages of using a mainframe computer are:

- almost limitless power;
- speed of operation;
- updating from multi-access points;

— online access from other council departments, the public, etc.;
— possible links into computer issue systems;
— feedback from other departments/ organisations using the system.

Disadvantages are:

— they are not particularly user-friendly - no use of graphics;
— it can be difficult to maintain consistency of input and updating with multi-access;
— they take second place to other council programmes, such as payroll.
— you need to 'log on' each time the system is used.

Reference was made above to possible links between mainframe-based community information files and library computerized issue systems. In fact, a number of these systems now offer add-on features for storing community information themselves. The advantages of these are:

— access to database through every library in the system;
— public access through OPACs (Online Public Access Terminals) or, in some cases, dial-in facilities;
— multi-access to inputting and updating database.

Disadvantages are:

— they can slow down issue system;
— they are not usually accessible by other council departments;
— with multi-access, it can be difficult to maintain consistency of inputting and updating;
— as with viewdata systems, information is usually limited.

Inevitably, with a technology that is developing so fast, there will be new systems for presenting information coming along. Glasgow has been experimenting with HyperCard, which allows the user to search through text, pictures, numbers and instructions to gain information about the city. Before long someone is bound to be using CD-ROM or interactive video.

Online Searching

Another use for your PC might be to link into other remote databases such, as a local authority's viewdata system. To do this you will need' a modem, to link the PC to the telephone network, and communications software, which could be a separate package or part of a suite of software such as LotusWorks or Windows 3.

Most online databases, except for local viewdata, are not appropriate for community information services, but there is one that is: Volnet UK, operated jointly by The Volunteer Centre UK and the Community Development Foundation. Volnet UK is an online information service containing thousands of references to press items, journal articles, reports, books and current research projects on community development, voluntary action, youth affairs, child care and social policy issues, stored on an easy-to-use computer database. For a modest annual fee (00 local community/ voluntary agency, £90 national/ regional voluntary agency, £150 central/local government department) you get unlimited access to the database (no charge per second or minute), a telephone help-line, plus a photocopy service for articles, for which a small charge is made. Details from Volnet UK, Community Development Foundation, 60 Highbury Grove, London N5 2AG.

Computer Programs

Providing welfare benefits information and advice can be notoriously difficult because of the complex legislation, which is why some advice centres and public library community information services are harnessing the power of PCs to help them. These are usually designed for use by advisers although, over the years, some attempt has been made to develop client-operated systems, largely without success.

Welfare benefits programs contain the rules of the benefits system and present on the computer screen a series of questions. The adviser, in consultation with the client, types in the details and the computer works out whether the client is entitled to benefit and, if so, how much. The results can usually be printed out and given to the client to take away. The main benefits covered include Income Support, Housing Benefit, Family Credit and Community Charge Benefit. The advantages of these programs are:

— consistency;
— they do not let you forget little-used regulations or passport benefits, where claiming one benefit, automatically entitles the client to others;
— a novice worker can tackle more complex cases;
— 'what-if' calculations can be done speedily;
— printouts of results can be obtained;
— some clients like the computerized approach and are reassured by it.

Some disadvantages are:

— they are only as accurate as the information input;
— the benefits, system is so complex, programs are bound to have quirks;
— an experienced adviser can usually work faster than the program;
— they ran lead to over-confidence;
— some clients can be intimidated by a computer and find it intrusive;
— they are no substitute for experience.

For more details, on the main welfare benefits programs available and their features, addresses etc., see the *Computanews* factsheet Welfare benefits programs (see Wow for details of all the fact-sheets in this series).

Other Programs

There are other PC-based programs and databases available on such areas as debt advice calculations and sources of funding from charitable trusts. Resource Information Services, in conjunction with LASA's Computer Development Unit, has developed a database of advice agencies in London which became functional in autumn 1992. A good way of keeping up with the latest developments in programs for advice work is to take out a subscription to the excellent bimonthly periodical *Computanews*.

Statistics

Collating and manipulating statistics is one area where computers have a distinct advantage. They can do in seconds or minutes what would take hours of work manually. However, if all you want the computer to do is add up and produce totals in broad categories, then you might as well stick to pencil and calculator. But, for detailed analysis of enquiries and graphic representation of information by means of pie or bar charts, graphs, etc., you will need a computer program. There are two options: either devise your own system using a spreadsheet package or a relational database system linked with a spreadsheet package, or buy a specially designed program. A lot of the pioneer work in this field has been carried out by LASA's Computer Development Unit which, some years ago, developed CRESS (Client Records and Enquiry Statistics Service) and is currently engaged on developing a model recording system called STATS. Finally, there are many more uses to which a community information service might wish to put its computer, e.g. word processing, accounts, graphics, desktop publishing, but these aspects apply to many other organisations and it is beyond the scope of this book to cover them.

SEVEN

Computers for Individuals

A major segment of the market involves entertainment: personal computers offer a growing array of games, music and art far more sophisticated, in terms of speed and capability, than video material based only on television. This advantage results from the ability of computers to store and quickly retrieve large amounts of information. The future will bring many more uses for computers in both business and home. Indeed, the possibilities are virtually endless. A great variety of computers, large and small, are now available. A personal computer is a machine or system meeting all the major qualifications.

The system is designed to accept secondary memory devices to supplement the primary, built-in memory, the user is expected to interact with the system continuously, not only at the beginning and end of a problem, at least one general language (Basic, Fortran, Cobol, Pascal, ADA, or Q is available for this interaction), the system is usable for a wide variety of problems and is not designed for any single application and the computer is distributed through mass-marketing channels, with the marketing emphasis on the first-time computer user.

A typical modern personal computer consists of a circuit board with a silicon chip, microprocessor and one or more memory chips attached. The microprocessor can perform hundreds of thousands of calculations every second, and the memory chips provide the primary storage for instructions and data.

External storage devices, such as cassette tape units or small recording disks (floppy disks), augment the memory capacity and provide a storage medium that can be physically transferred from one personal computer to another. (Typical users begin with cassette units but soon change to disks to gain the advantages of greater speed and capacity.) Input is through a typewriter-like keyboard unit. Output typically takes the form of words and numbers displayed either on a television screen or a similar specialized screen called a monitor. Most monitors are designed to display 24 lines of letters and figures, with each line containing a maximum of 80 characters. Adding a printer unit permits output in the form of a printed paper. A special device, called a "modem" (for modulator/demodulator), permits the computer to receive and transmit data over a conventional telephone line. But the significant factor in defining a personal computer is not its physical features but the characteristics of the operating system.

Designers of personal computers and software attempt to provide a friendly human-machine interface, even at the expense of brute computing power. Optional programs are also available so that the computer can be used for many different purposes. Although word processors and hobby computers have many characteristics of personal computers, they lack this flexibility.

Recent trends in microelectronics, memories, input-output mechanisms, and software suggest that the trend in microprocessors is toward larger "words" and higher circuit speeds. A computer capable of handling larger "words" is able to perform a complete operation in fewer machine cycles and to operate directly with larger memory. Both these assets enhance performance by increasing the speed of operations and the number that can be performed in a sequence without the operator's intervention. In addition to greater speed and memory access, these larger microprocessors have the advantage of greater accuracy.

The first wave of personal computers used 8-bit microprocessors—that is, microprocessors in which 8 binary digits (0 or 1) can be processed in parallel, giving the capacity to process in a single operation any number up to 256, or to address any file up to 256 points in the memory. Processing larger numbers with an 8-bit microprocessor requires multiple operations, which take more time. Newer systems use 16-bit microprocessors, and 32-bit microprocessors are now available.

The primary memory in personal computers is of two different types: read-only memory (ROM) and random-access memory (RAM). In the former information is fixed in the memory at the time of manufacture and is not lost when computer power is switched off. The role of such a ROM is to guide the computer through a fixed procedure, such as calculating square roots and translating a user's program into machine language. In a RAM, information such as special programs and data files can be "written in" or "read out" as frequently as desired, with any storage location directly accessible. RAMs are of two types: dynamic RAMs, which are cheaper but lose their store information unless they are "refreshed" often, and static RAMs, which are costlier but do not need to be "refreshed." If power is lost, both types of RAMs lose their stored information.

Over the past decade, the number of memory circuits per unit of area on a chip has increased by a factor of 64, and cost on a unit basis has been reduced by a factor of 50. Both these trends will continue. So-called 64K dynamic RAMs (each contains 65,536 bits of information [K = 210 = 1024; 64K = 64 x 10241) remained popular until about 1984, when the much larger memory capacity of 256K dynamic RAMs became standard.

Most manufacturers supply system programs on ROMs, and this practice will continue because programs are secure against power failures and users are less able to duplicate programs in this form. Computers' main memories will be supplemented by secondary storage devices that offer larger and relatively inexpensive, though slower, capacity for long-term storage of program and data files. The most popular of these is a so-called "floppy disk"-a disk of mylar coated with magnetic material on one or both sides. Data are stored in a series of spots-either magnetized or demagnetized-along the concentric tracks. Heads for reading or writing data can be moved radially across the disk to reach a specified segment of circular track.

Storage capacity depends on the format used for the stored data, the quality of magnetic surface, and the design of the reading-writing head. Floppy disks in current use typically have capacities of I to 4 million bits, sufficient to store 20,000 to 80,000 words of English text. During the next four years, higher-density floppy disks will be common. Indeed, disks offering capacities of 50 to 100 million bits are already becoming popular, but they are much costlier than floppy disks.

The primary display device used in all personal computer systems is a cathode-ray tube (CRT), either standing alone or as part of a television receiver. This system will continue through the foreseeable future. Output is typically presented in alphanumeric form-letters and numbers. Charts and game boards can be presented, but the memory and software required to display such graphical images is often complex and expensive. The letters and numbers in alphanumeric displays are patterns of dots programmed in special ROMs known as "character generators." The quality of the image depends on the number of points (or pixels) on the screen that can be addressed by the computer—that is, the number of points at which dots can be located. A typical low-resolution screen has a field of 6144 (128 times 48) pixels (picture elements-dots on the screen). High-resolution systems (100,000 pixels or more) allow sophisticated graphics for animation or detailed figures and may provide colour as well.

The Market and the Players

The personal computer industry has grown as a direct result of the evolution of the microprocessor. This evolution began when Intel in 1971 packaged a complete, if somewhat limited, processor with a 4-bit word size in a single integrated circuit. The company followed with the first 8-bit processor in 1972 and an improved version in 1974. One year later Micro Instrumentation and Telemetry Systems, Inc. (MITS), an Albuquerque firm, developed the first personal computer around Intel's 8-bit processor. The basic system sold for $395 in kit form and $621 in assembled form, not including accessories (peripherals).

Though the MITS system is no longer manufactured, its method for connecting peripherals and the main computer has become an industry standard. Within three years after it was introduced, Radio Shack, Apple, and Commodore had entered the market. Now makers of large computers such as IBM and Honeywell, as well as minicomputer leaders such as Digital Equipment Corp. and Data General Corp., are also making personal computers, having observed their traditional markets being eroded by the new low-priced products. The results are rapid development of the market and growing popular interest and faith in personal computers.

The pioneering manufacturers did not survive beyond the initial phase of personal computer development resulting from the preoccupation with the needs of hobbyists. New entrants such as Radio Shack, Commodore, and Apple captured a major share of the market in 1978 by promoting fully assembled, ready-to-operate systems that were easier to use. Though the success of the companies now in the competition will depend at least partly on their financial and technical resources, their products' ease of use (or "user-friendliness," in trade jargon) is a significant competitive feature.

The personal computer market can be divided into four segments in terms of the computer's intended use: business, home, research-technical, and educational. The business segment, by far the largest at present, accounted for 750,000 sales (retail value of $2 billion) in 1982, or 54 per cent of the units sold and 65 per cent of the dollar sales. This segment will continue to be a major factor, with more sales than all other segments put together. No wonder, then, that the leading makers of personal computers are concentrating their sales efforts on the business sector.

The most visible segment of the personal computer "revolution" is in homes, where computers are used for entertainment and education, sending messages, and home finances. The development of appropriate software will soon enable sophisticated users to pay bills, manage bank accounts, compute taxes, and even buy household items without written transactions.

The potential market for personal computers in education has barely been tapped. Computer manufacturers recognize that schools are the logical environment in which to develop computer skills, and they have reason to expect that students will purchase the machine on which they first learn computing. So manufacturers are offering many price incentives to the educational sector. Yet the educational sector will continue to account for only about 15 per cent of total sales of personal computers for the rest of this decade, as school funding will be limited.

Although they want to serve all four segments, the industry leaders (Apple, Radio Shack, Commodore, IBM, and Xerox) are focusing on the business market. Vector Graphics is also emphasizing this market by tailoring personal computer systems to the needs of particular industries. Atari (Warner), Intellivision (Mattel), and Texas Instruments are focusing on the home market. Hewlett-Packard is offering a highly specialized compact product with built-in printer, tape memory, and video monitor especially for the scientific community.

In general, both prices and manufacturers' profit margins are failing, and mass marketing is becoming the rule. The major competitors win each spend over $10 million in 1983 on advertising aimed at expanding the market and establishing their place in it. The dynamics of the entire personal computer market may be changed substantially by new entrants. Traditional barriers to entry into this field have crumbled. For example, manufacturing capability is not essential: the IBM personal computer is assembled almost entirely from pre-manufactured components not of IBM origin.

By contrast, strength in marketing and distribution is a significant advantage, and many organizations with this capability may seek to replicate IBM's strategy for rapid entry into the field. Organizations such as General Electric, Procter and Gamble, Phillip Morris, and Du Pont all have the ability to enter the personal computer market in this way.

Selling Personal Computers

Strategies for marketing mainframe computers are not appropriate for marketing personal computers—profit margins are not large enough to justify hiring internal sales forces to sell directly to end users. As a result, producers are now experimenting with a wide variety of distribution strategies.

Franchised retail chains constitute a major distribution channel. For example, the Computer-land chain sold $200 million worth of computers and related accessories in 1981. These stores distribute the products of many vendors, and their volume is large enough to support a technical and maintenance staff.

Manufacturer-owned retail stores have been used successfully by Radio Shack. Other computer manufacturers like IBM, Digital, and Xerox have also opened such stores but only to supplement existing distribution channels. Except for Xerox, each manufacturer's stores sell only that manufacturer's products, and a prospective buyer is thus obliged to visit several stores to compare equipment of different makes.

Department store outlets have generally been unsuccessful. Mass merchandisers depend for profits on fast-selling commodities. According to computer industry data, personal computer buyers make four shopping trips totalling as much as seven hours when selecting their machines. These buyers also expect sustained support and maintenance services that department stores are unaccustomed to providing.

Office equipment stores specializing in copiers, typewriters, word processors and other office equipment are well positioned to reach the most promising future market for personal computers and auxiliary equipment. If these stores can provide adequate servicing, they will become very popular.

Consumer electronics stores such as Tech Hi-Fi have been marketing personal computers with some success, but the lack of expertise at the store level has been a constraint. Japanese manufacturers have also established ties with such stores as distributors of other Japanese products. These channels will become a major factor as the Japanese increase their share of the personal computer market.

Independent retailers often lack the capital required to compete vigorously and are therefore not gaining in numbers and importance as quickly as other retail channels.

Catalog showrooms have been used by Texas Instruments, but personal computers require follow-up support that such showrooms have been unable to provide.

Mail order firms offering discounts ranging up to 30 per cent appeal to price-sensitive customers, but their total lack of continuing support disenchants users. Price-cutting penalizes the full-service dealers on which manufacturers want to rely as major sales outlets. Accordingly, major companies are trying to discourage sales through mail order firms.

Direct sales staffs are used for large-volume sales to government, educational institutions, and major corporations who prefer dealing directly with the manufacturer. But such direct sales tend to antagonize dealers by depriving them of some of their most profitable opportunities, and profit margins are inadequate to support direct sales to small-volume and individual buyers.

Value-added houses serve specialized users with coordinated hardware and software, such as printing companies that need word-processing and typesetting capabilities but have little or no internal computing expertise.

Though franchised retail outlets are the largest sellers of personal computers in the US, the one compelling characteristic of personal computer distribution today is diversity. No one form of vendor or market approach is dominant.

Selecting a Personal Computer

No one should select a personal computer system without taking into account the tasks for which it will be used and the environment in which it will function. Even within a single business, personal computer users come from different areas—personnel, accounting, management, manufacturing, research, sales—and have different computing needs.

Even within these departments there are many possible uses. The personnel department will run personnel files on a regular basis but also may want to study policy questions—the likely benefits in terms of increased revenue from a larger sales force, for example. Furthermore, different users have different strategies for using a computer for adhoc applications. Thus, the person who will use a personal computer should—almost by definition—choose it, partly on the basis of the services it will provide and partly on the basis of its apparent user-friendliness.

Indeed, the machine's responsiveness to its user's needs and style is the critical technological breakthrough. But it is not necessarily so simple. In many organizations, personal computers create and fill completely new needs that were not originally anticipated. Consider the example of a personal computer acquired to improve one lay worker's access to a main computer facility. While retrieving data from the main computer to run programs locally, this user finds that the personal computer can help evaluate alternative business strategies, answering, "what-if...?" questions instantaneously. The original application for which the personal computer was purchased becomes secondary. Given that the uses and benefits of personal computers are hard to predict, and that a wide range of systems and software are available, how should a potential buyer decide what and when to buy? And should such a buyer postpone action, expecting costs to decline in the future as they have in the past?

There are two arguments against postponing purchase in anticipation of a lower future price. One is the high cost of waiting, based on the computer's great (and often unexpected) uses. The other is that prices are expected to decline less dramatically in the future than in the past: future systems will offer technological improvements that increase performance at unchanged cost. Prices seem likely to stabilize at around $1500 for a full personal computer system.

Future personal computers are likely to offer increased memory size, increased processor power (16-bit and even 32-bit instead of 8-bit microprocessors), improved printers, more powerful programming languages that give the user more compact and natural communication with the computer, improved user access (better keyboards supplemented or replaced by touch screen and voice systems), and greater flexibility (the computer will be able to intermix and manipulate text, numbers, and graphics more easily). Yet today's personal computers may rapidly become obsolete, reflecting the rapid pace of change in semiconductor technology.

Changes in the microprocessor, the heart of any personal computer system, are most likely—and most serious. A computer with a 16-bit microprocessor cannot operate on hardware and software designed for an

8-bit microprocessor. Though many vendors provide add-ons that enable a new microprocessor to support old software, there is never total compatibility.

Software for a personal computer is expensive, and for most users new software development is impractical. So the amount and utility of the software available with a particular system is the key question in the buyer's choice of a personal computer. The variation can be large: 11,000 programs are available for the Apple II and Apple II Plus computers, whereas fewer than 100 widely distributed programs exist for the IBM personal computer, a more recent introduction. Of course, the IBM programs may be just the ones the buyer needs, and more software will be developed for the IBM.

Price is also an important variable, including not only the manufacturer's suggested price but the range of discounts offered by different retail computer sources. But discount alone does not tell the story. The critical factor is the ability of the distributor to provide both routine hardware maintenance, usually at the retail outlet, and software support and maintenance. A company store is likely to offer very good service on that company's product line but not necessarily on any other line, while a retail computer store is likely to offer moderately qualified service and support for many different manufacturers' systems. Mail order sales offer very little service but usually have the highest discounts.

The ultimate question for the purchaser is that of system configuration: what combination of hardware and software is appropriate for a buyer's current and projected needs? Perhaps the most important options are off-the-shelf software. To be sure they buy what they need, buyers should gain some experience with available programs before making the purchase; otherwise they risk having to modify their needs to fit the software they've purchased or embarking on the costly and time-consuming route of software development.

The programming tools on most personal computers are primitive compared with the equivalents available for mainframes, and the widespread dearth of good programmers is a complicating factor. Indeed, trying to develop special programming for a personal computer is inadvisable except under extraordinary circumstances. Memory capacity is another important variable. Most computers based on 8-bit microprocessors have no more than 48 or 64 kilobytes of built-in accessible memory (RAM), while the newer 16-bit processors can have 1, 4, or even 16 megabytes. (One byte is a string of 8 binary digits that represents one alphabetic character.)

Additional memory can be inserted into most personal computers by plugging cards into slots or inserting extra chips directly onto the microprocessor board. Storage capacity of the main memory for programs and instructions can also be supplemented by secondary devices, of which floppy disks are the most popular form. Disks with capacities ranging from 100,000 bytes to several megabytes now come in two standard sizes: 8 inches and 51/4 inches in diameter. But personal computer buyers should be aware that disks designed for one system will not generally operate in another maker's system because of different data-formatting conventions.

The critical performance characteristic of a floppy disk is the response time to a read-write request, which may range from 200 to 500 milliseconds. Though the response time for such a single file request is very short, a single compilation or word-processing command may involve 20 or more accesses to a floppy disk, and access time may account for 80 percent of the total command execution time. Where better performance-and/or capacity are required, "hard" disks in the same format as the 51/4-inch and 8-inch floppy disks are useful. They offer storage capacities ranging from 5 to 50 megabytes and response time averaging 20 to 40 milliseconds, but their high cost (around $2000, including interfaces) can double a buyer's total investment in hardware.

All personal computer systems use cathode-ray tubes as output devices. If "hard copy" is required, buyers have a choice of several printing devices. Thermal printers, costing under $500, create images by applying points of heat to special paper. A dot-matrix printer, costing from $500 to $1500, has a vertical array of 7 to 10 dot-printing elements. These are activated by the computer as the printing head passes across the paper, forming alphanumeric characters and also graphs and drawings.

The limited number of elements result in a somewhat stylized character, but the process is as fast as thermal printing—50 characters per second. If letter-quality printing is required, a printing mechanism with precisely machined character matrices that strike the paper through a ribbon is required. Such printers cost $750 or more and print between 30 and 90 characters per second. Various accessories such as floppy disks and printers or "peripherals" are attached to personal computers through special "slot" receptacles. The IBM personal computer offers five such slots for peripherals such as printer, colour display, dual floppy-disk controller, communications channel, and memory module.

All business-oriented personal computers offer text editing for word processing. Some permit use of a light pen to directly indicate on the screen the word to be edited, and a few check spelling against a mini-dictionary stored in the computer memory. Business data manipulation is facilitated by "spreadsheet" packages such as VISICALC—a very popular package that enables users to see on the CRT the impact of altering one figure (say, retail price) on all other figures (profit margin, corporate profit, return of investment).

Software is now becoming available for generating displays such as those accompanying this article with only a few minutes' effort. For example, ExecuVision permits displays combining information in numbers, text, and pictures. It also provides animation that can add emphasis to presentations. Though personal computers are today viewed primarily as stand-alone work stations for individual users, the ability to transmit and receive data and programs is an important attribute of most contemporary systems. Personal computers are already serving as intelligent terminals for retrieving and manipulating files from large computers. These networking and communication capabilities are important considerations in selecting a personal computer. Unfortunately, no industry standard for communication protocols and interfaces exists. Different vendors use different protocols, so it is difficult to link various personal computers in a single network.

Technological advances in communications are expanding the power of personal computers. For example, owners of personal computers can now receive current stock market quotations by telephone. Personal computers can be connected with nationwide electronic funds transfer systems, so users have access to full banking facilities. They may transfer funds from one account to another and buy and sell financial instruments.

In homes, personal computers are now primarily used for recreational activities such as computer games. But these applications should not be underrated. Computer games test and develop mental capabilities while generating familiarity with computer operations. Home computers are also used for office work at home, and many other applications are, or soon will be possible. For example, a tax return can be prepared with the aid of VISICALC. The user can analyse which forms and deductions to use for paying the minimum taxes. A personal computer can be used to balance check-books and plan investment strategies. And, of course, word-processing capabilities are useful in writing letters and professional papers.

As commercial databases such as Dow Jones News/ Retrieval and the Source are expanded and their subscription costs reduced, personal computers will be useful for locating stores and services, comparing prices, and placing orders. This would be especially appropriate for products or services whose prices fluctuate rapidly, such as those of the airline industry. Computers are being used increasingly for security applications. Traditional home security systems are limited, often failing to distinguish between a natural event and an intrusion by an unwanted visitor. Computer-controlled home security systems can include a variety of routines to analyse the signals from sensors before generating an alarm. Such an analysis can distinguish, for example, between the entry of a cat and the entry of a human being, or between the noise of a telephone ringing and a door opening.

Another home-oriented application is a medical information system that performs some of the diagnostics of a medical doctor. Perhaps the most useful application of this is a database for poison or first-aid: the user reports what substance the patient has ingested or the nature of an accident, and the computer responds with a step-by-step first aid plan. Children turn out to be heavy users of home computers. They enjoy computer games and seem able to learn computer languages faster than their elders. They find personal computers intriguing for recreation, school assignments, and self-paced courses that are now becoming available. Even though today's personal computers are equivalent in basic computing power to the mainframes of the 1960s and the minicomputers of the 1970s, they should not be used as substitutes for the earlier machines. Rather, users should capitalize on the new technical developments, including the diminishing cost of computing, and the English-like languages that replace the cumbersome digital languages of early machines, to make personal computers supplement, not displace, mainframe computers.

Personal computers should be used to analyse more issues in greater detail than ever before—not to do the more routine work typically assigned to mainframes. In the 1960s and 1970s, computers were used principally for well-structured, periodic jobs such as payroll accounting. Today's personal computers are especially useful for more casual, adhoc analyses and problem solving, such as studies of the impact of cost of living increases, new tax regulations, or alternative hiring and layoff policies.

Consider the stock analyst who uses data from company balance sheets to evaluate investment risks. The balance sheets contain detailed, current information on assets and liabilities. But interest rates fluctuate, markets rise and fall, taxes change, competition falters or gains. Each such change affects a company's current performance and future prospects.

Using a central computing facility is an expensive way to analyse such problems, resulting in time delays and partial answers. But using a personal computer in the interactive mode linked to the central computer and a financial database makes it possible to understand the impact of changes quickly and accurately.

Psychology of Personal Computers

Manufacturers, media, and computer scientists have presented the public to computers the way one might have presented Aladdin to his lamp. Here was the genie, the workhorse that we could mount today and ride onto the new millennium. By the end of the 1970s, with the mass manufacture of personal computers, the promises extended into the home. Home computers would teach us French, help us with financial planning, even do our taxes. For many Americans, the first thing that computers brought into their homes was not a more efficient medium for work but new worlds to conquer in play.

"Space Invaders" became a household word for people with home computers or the machines' junior siblings, computerized video games. The experience of the decade has resulted in a widely shared public rhetoric. When people talk about the computer and their futures, they tend to fall back on the two images they know best: the computer as tool and the computer as toy; objective instrumentality and engrossing play. But there is another dimension. What people do with computers weaves itself into the way they see the world. People use a discourse about computers for thinking and talking about other things, about politics, religion, education, and about themselves and other people.

The subjective dimension of the computer presence is not merely a matter of discourse, of using the computer intellectually as a metaphor or model. There is another, and more emotionally charged, aspect which does not engage ideas about computers as much as the immediate quality of an individual's experience when working with them. Working with computers can be a way of "working through" powerful feelings. In this essay we develop this idea by using as a case study a group of computer users for whom issues related to control are particularly salient. For a first generation of computer hobbyists, controlling the computer is a way to deal with frustrations and desires, both personal and political, that have nothing to do with the computer per se.

The subjects of my study are men and women who bought personal computer systems in the four years that followed the 1975 announcement of the "Altair"—the first computer small enough to sit on a desktop, powerful enough to support high-level language programming, and that you could build for only $420. My study began in 1978 with a questionnaire survey answered by 95 New England computer hobbyists (their names had been drawn from the roster of a home computer club and from the subscription list of a personal computer magazine), and continued during 1978 and 1979 with nearly 300 hours of conversation with 50 individuals who owned home computers. What we found can be read historically: a study of the pioneer users of an increasingly ubiquitous technology. But most central to the intent of this essay is to use the story of the early hobbyists as a window into the highly personal ways in which individuals appropriate technologies. It is a case study of the "subjective computer," the computer as a material for thinking, for feeling, for "working through."

It is acknowledged that the subjective is at odds with a widespread ideology that quickly grew up around the emergent computer hobbyist culture. The Altair, aimed at a strictly hobby market, was followed by other small systems—the Pet, the Sol, and, most successfully, the Radio Shack TRS-80 and the Apple, marketed to less specialized audiences of small businessmen and curious householders. With this explosion of hardware came a lot of rhetoric about a personal computer revolution. Most of the talk, both from the companies that marketed the machines and from those who claimed to be the most visionary spokesmen for the people who bought them, was about all of the things that a home computer could do for you.

The utilitarian, "genie in the bottle," ideology is expressed in the content of hobbyist conventions and magazines, filled with articles on how to make your home computer dim your lights, control your thermostat, run an inventory system for your kitchen or tool-room. And it is also found in writing on the personal computer from outside the hobbyist world. The view from the "outside" was well illustrated when, on May 14, 1979, the *Wall Street Journal* reported on its own little "evaluation experiment." The paper had drafted one of its staff reporters, Mitchell Lynch, into the ranks of the "home computer revolution." It presented him with a TRS-80 and asked him to take it home. His assignment was to report back six months later on what it had been like.

The terms of Lynch's evaluation are instrumental: what can the computer do? On these terms it fails, and when he goes to the experts for an explanation of what went wrong, he encounters an official "instrument" ideology from within: "Experts say that people like me have neither the technical training nor technical inclination to make a home computer strut its stuff." The point of having the computer, says the expert, is to make things happen. And, reassures the expert, they will happen with a better operator or a simpler computer. This instrumental view is an important ingredient of the hobbyist ideology but it is not the whole story. In the course of my work I found a very different answer from within.

Most hobbyists do make their computers "strut their stuff," but their sense of engagement and energy are found primarily in the non-instrumental uses of the technology. When asked in a questionnaire "What first attracted you to computers?" more than half the respondents gave reasons that were highly subjective. In response to an open-ended question, 26 per cent said that they were first attracted to computers by an appeal that was intellectual, aesthetic, involved with the fun of what I would call "cognitive play." They wrote of "puzzle solving," of "the elegance of using computer techniques to handle problems," of the "beauty of understanding a system at many levels of complexity." They described what they did with their home computers with metaphors like "mind stretching" and "using the computer's software to understand my wetware."

Another 26 per cent wrote of reasons for getting involved that seemed more emotional than intellectual. They wrote of the "ego boost" or "sense of power" that comes from knowing how to run a computer, of the "prestige of being a pioneer in a developing field," of the "feeling of control when I work in a safe environment of my own creation." The hobbyists who responded to my survey seemed familiar with Lynch's brand of scepticism, with people who ask them what they do with their computers and who won't take "cognitive play" for an answer. David, a 19 year-old undergraduate at a small engineering school, put it this way: "People come over and see my computer and they look at it, then they look at me, then they ask me what useful thing we do with it, like does it wash floors, clean laundry or do my income tax-when we respond no, they lose interest." David said that when he started out, he was attracted to the computer because "we liked the idea of making a pile of hardware do something useful, like doing real time data processing ... like picking up morse code with an amateur radio and transcribing it automatically into text," but in his list of things that he currently does with his computer.

Thirteen percent of those who responded to my questionnaire told a similar story. Like David, they began their relationship with personal computation for instrumental reasons (they had an image of a job to do, a specific task), but they became absorbed by the "holding power" of something else. A full two-thirds of my survey sample either began with or ended up with a primary interest in what I have called the "subjective computer," the computer seen in its relationship to personal meaning. Clearly, to understand what people are doing with their home computers we must go beyond the "performance criteria" shared by the hobbyist magazines and the *Wall Street Journal*. The simplest way of thinking about the subjective computer is through the metaphor of "the computer as Rorschach"-that is, seeing the computer as a projective screen for other concerns. In the Rorschach test, one is presented with a set of ink blots and asked to make some sense of them. Some people see the blots as threatening; others see them as benign. Some focus on small, complex details; others on global form. So too, the computer presents us with an ambiguous stimulus. For example, although most people think of the computer as an object without real intentionality, they accord computers enough autonomy in action to make "blaming the computer" a commonplace of daily life.

The very fact that the computer is a machine that touches on a sphere like intelligence that man has always considered uniquely his, is enough to make many people experience the computer as an object "betwixt and between," hard to classify, hard to pin down. Sometimes people deny the irreducibility of computation by asserting that, no matter how complex the computation "product," a move in a chess game for example, "all the computer really does is add." Of course in a certain sense this is correct. But saying that a computer "decided to move the Queen by adding" is a little bit like saying that Picasso "created Guernica by making brush-strokes."

Reducing things to this level of localness gives no satisfying way to grasp the whole. Just as in theoretical psychology there is a tension between the gestalt and the atomic, so too in computation there is a pervasive tension between the local simplicity of the individual acts that comprise a program or a computation, and what one might call the "global complexity" that can emerge when it is run. The elusiveness of computational processes and of simple descriptions of the computer's essential nature, the tension between local simplicity and global complexity, all contribute to making the computer an object of projective processes, and exemplary "constructed object." Different people apprehend it with very different descriptions and invest it with very different attributes. In views of the computer's internal process, individuals project their models of mind.

In descriptions of the computer's powers, people express feelings about their own intellectual, social, and political power, or their lack of it. Looking at the computer as Rorschach, has projected, puts the emphasis on aspects of the individual—from cognitive style to personal fears—that are revealed through behaviour with the machine. But of course the computer is more than a Rorschach. The Rorschach ink blots are evocative, revealing, but they stay on the page. They do not enter the life of the individual. The computer does. It is a constructive as well as a projective medium. For readers who have not had the experience of programming a computer, this idea may be sharpened by an analogy with another technology. In my own studies of people's emotional relationships with technologies, airplanes emerge as startlingly like computers in respect to the issues they raise for their hobbyists. Specifically, both are powerful media for working through the issue of control.

Local versus Global and the Question of Control

There is always a compelling tension between local simplicity and global complexity in the working of a computer and in the appreciation of a computer program. Locally, each step in a program is easy to understand; its effects are well defined. But the evolution of the global pattern is often not graspable. You are dealing with a system that surprises. This play between simplicity and complexity allows programmers, as pilots do with flying, to make of computation very different experiences that provide a context for working through different needs in relation to issues of personal control. Depending on how the programmer brings the computer's local simplicity and global complexity into focus, he or she will have a particular experience of the machine as completely understandable, under control, or as baffling, even as controlling. By focusing on the local, the line by line, you can feel in control. By focusing on the global, you can feel control slip away. In their style of programming a computer, people betray different levels of tolerance for temporary losses of control.

Some will avoid it at all costs; but others will seek it out, and enjoy "playing" with sensations of risk and danger. And so different people end up with very different relationships to controls and power in their programming work. To illustrate this point, it is useful to begin with an example at one extreme. We see a first style in Howard, an ex-programmer, now a university professor, who describes himself as "having been a computer hacker." Howard was not in my sample of hobbyists. He has a terminal at home which links to a large timesharing system, but when asked about home computers, he winced in distaste and said that he "wouldn't touch the stuff. It's too simple." His case, in which we see a love of programming for the feeling of "walking near the edge of a cliff," is meant as a contrast for what we found to be a prevalent and more conservative "hobbyist style." Howard described his longtime fantasy that he would walk up to any program, however complex, and "fix it, bend it to my will." As he described his intervention, he imitated the kind of hand gestures that a stage musician makes toward the hat before he pulls out the rabbit.

Wizards use spells, a powerful kind of local magic. Howard's magic was local too. He described his "hacker's approach" to any problem as a search for the "quick' and dirty fix." For Howard, what was most thrilling about the experience of programming was "walking down a narrow line," using the program's flexibility (for him defined as the possibility of making a local fix) in a struggle to keep the whole under control. Weekends at the terminal with little to eat and little or no rest were frequent, as was the experience of not being able to leave the terminal while debugging a program, even when the obvious need was for sleep and looking at the whole in the morning, instead of trying to "fix it" by looking at it line by line all night. For Howard, the urgency of these encounters was tied to his sense that through them he was grappling with a computational essence—the struggle to exert control over global complexity by mastery of local simplicity.

A second programmer, Bob, is a computer professional, a microprocessor engineer who works all day on the development of hardware for a large industrial data system. He has recently built a small computer system for his home and devotes much of his leisure time to programming it. Whereas, for Howard, the excitement of programming is that of a high-risk venture, Bob likes it as a chance to be in complete control. Although Bob works all day with computers, his building and programming them at home is not more of the same.

At work he sees himself as part of a process that he cannot see and over which he feels no mastery or ownership: "Like they say, we are just a cog." At home Bob works on well defined projects of his own choosing, projects whose beginning, middle, and end are all under his control. He describes the home projects as a compensation for the alienation of his job. He works most intensively on his home system when he feels furthest away from any understanding of "how the whole thing fits together at work."

Howard and Bob have very different opinions about what is most satisfying about programming. These translate into different choices of projects, into different choices of programming language and level to program at, and ultimately into what we might call different computational aesthetics. Howard likes to work on large, "almost out of control" projects. Bob likes to work on very precisely defined ones. Howard finds documentation a burdensome and unwelcome constraint; Bob enjoys documentation, he likes to have a clear, unambiguous record of what he has mastered. Indeed, much of his sense of power over the program derives from its precise specifications and from his continual attempts to enlarge the sphere of the program's local simplicity.

Hobbyist programmers mean different things when they say that machine language programming "puts them in control." For some, the reference seems objective. Given the primitive higher-level languages available on the first and second generation of hobby computers, machine-level programming seems to them the best instrumental solution. For others, the issue is more subjective. Many hobbyists who said, they felt Lineasy using systems programs for which they didn't have the source code also admitted never looking at the source code listing which they felt they had to have. Having access to the code was symbolic.

In other cases it was apparent that machine language programming was valued because the experience of doing it was pleasing in itself. It meant that the programmer was writing instructions that acted directly on the machine, no "building on top of" somebody else's interpreter. Most hobbyists have relationships with computation at work which involve sharing the machine with countless and nameless others. In personal computation they see a chance to be independent and alone.

The machine comes to them virgin. They have full possession. Finally, there is the issue of asserting control over an inferior. Bob, like many of the other hobbyists I spoke with, is a middle-level worker in the computer industry. He does not feel very good about the importance of his job. Proving that he is "better than any dumb compiler" gives a sense of importance. Using the computer to assert control was a central theme in my interviews with hobbyists. It was expressed directly, and also wove itself into four other issues that characterize the hobbyists' "subjective computer." These are using the computer to strengthen a sense of identity; to construct a completely intelligible piece of reality that is experienced as "transparent" and safe; to articulate a political ideology; and to experience a sense of wholeness that is absent in one's work life. It is these four issues to which I now turn.

Building Identity

In achieving a sense of mastery over the computer, in learning about the computer's "innards," people are learning to see themselves differently. Among other things, they are learning to see themselves as "the kind of people who can do science and maths." This was most striking among hobbyists who had no technical background. But it also came up among hobbyists who did see themselves as "technical people" but who for one reason or another had gotten "scared out of real science." Barry is 28 years old, an electronics technician at a large research laboratory. He went to college for two years, hoping to be an engineer, then dropped out and went to technical school. He has always loved to tinker with machines and to build things. His current job is to calibrate and repair complex instruments, and he is very happy with it because he gets a chance "to work on a lot of different equipment." But he came to his job with a feeling of having failed, of not being "analytic," "theoretical," of not being capable of "what is really important in science."

Five years ago, Barry bought a programmable calculator and started "fooling around with it and with numbers the way I have never been able to fool around before," and says that "it seemed natural to start working with computers as soon as I could." To hear him tell it, numbers stopped being theoretical, they became concrete, practical and playful, something he could tinker with. I'll pick up the calculator, and if we don't know how to do a problem I'll play with the calculator a few minutes, or a few hours and figure it out. It's not so much that the calculator does a particular calculation, but you do so many, have so much contact with the numbers and the results and how it all comes out that you start to see things differently ... The numbers are in your fingers.

When the calculator and the computer made numbers seem concrete, the numbers became "like him," and Barry felt an access to a kind of thinking that he had always felt "constitutionally" shut out of: "When I write in assembler I feel that mathematics is in my hands ... and I'm good with my hands." Barry claims to have "grown out of" his aspiration to be an engineer. He says he doesn't keep engineering as a pipe-dream or think of his computer skills as something that could make it real. In terms of his career he says that "nothing has changed." But a lot has changed. Barry has always thought of himself as a bundle of aptitudes and inaptitudes that define him as the kind of person who can do certain things and cannot do others.

Working with the computer has made him reconsider his categories. We really couldn't tell you what sort of thing I'm going to be doing with my computer in six months. It used to be that I could tell you exactly what we would be thinking about in six months. But the thing with this, with the computer, is that the deeper you get into it, there's no way an individual can say what he'll be thinking in six months, what we are going to be doing. But we honestly feel that it's going to be great. And that's one hell of a thing. For Barry, the world has always been divided between the people who think they know what they'll be thinking in six months and those who don't.

And in his mind, his home computer has gotten him across that line and "That's one hell of a thing." For Barry, part of what it means to have crossed the line is to start to call the line into question. When he was in school, his inability to do the kind of mathematics he had "respect" for made him lose respect for himself as a learner, The computer put mathematics in a form that he could participate in. Barry has three children, has bought them their own calculators, and encourages them to "mess around with the computer." He feels that they are going through the same problems with maths and science that he had and he wants them to have "a better start." For Barry, the computer holds the promise of a better start, not because it might teach his children a particular subject, but because it "might change their image of themselves. They might think of themselves as learners." Personal computers are certainly not the only hobby that people use to enhance their sense of identity. For my informants, "hobbies" have always been a way of life. Almost 90 per cent of them had been involved in a hobby other than computation, most usually in another "technical" hobby, such as photography, ham radio, or model railroading.

Fifteen percent of the hobbyists surveyed were using their computers to "augment" their participation in another hobby—for example, using the computer to keep an inventory of motorcycle parts, figure out ideal compression ratios for racing cars, interface with amateur radio equipment. For nearly a third of them, their home computer had completely replaced another hobby. People spoke of these abandoned hobbies as "fun" and as "good experiences," but their remarks about past hobbies underscored several ways in which in our day and time a computer hobby can be special. In particular, people spoke about their "switch to the computer" as making them part of something that was growing and that the society at large "really cared about."

Gregory is in his mid-forties, and has been in the electronics industry for all of his working life, as a technician, a programmer, and currently as a products designer. For two years, his computer shared space in his study with an elaborate model railroad system. A year and half before we met him he had bought a new hard copy printer and a graphics plotter. In the overcrowding that followed, the trains had finally found their way to storage in the basement.

We heard many echoes of Gregory's phrase, "it's part of the real world too." Hobbyists spoke about the computer offering them a connection with something beyond the hobby. For some, having a computer and "getting good at it" means crossing a frontier that separates "tinkering" from "real technology." They feel that the world sees their computer hobby as serious (several commented that friends and neighbours hardly even look at it as a hobby, as though the word were reserved for frivolities), and they start to see themselves that way too. Most first-generation hobbyists have technical educations, but many of them, like Barry, feel they have never been part of what is most exciting and important in the scientific and technical cultures. They see themselves as the low men on the totem pole.

Working with computers, even small computers, feels technologically "*avant garde*." A smaller group of hobbyists (but a group whose numbers are growing as new generations of personal computers become more accessible to the non-specialist) have always felt completely left out of the scientific and technical worlds. For them, owning a computer can mean crossing a "two cultures" divide. Alan, a 29-year-old high school French teacher who describes himself as "having a love affair with a TRS-80," has always felt he wasn't "smart enough to do science."

Alan majored in French ("It was easy for me ... my mother is from Montreal") and took up carpentry as a hobby. And although he was good at it, it only reinforced his sense of not being able to do intellectual things, which in his mind meant not being able to "do anything technical." When Barry began to do mathematics with his calculator, he felt that he started to cross a line to become the kind of person who could expect change and excitement in his intellectual life. For Alan, his TRS-80 led him across a line to become a member of a different culture, a scientific culture, a culture of "powerful people." For Barry, Alan, Gregory, relationships with computation enhanced self-image.

There is another way in which working with a computer can influence an individual's sense of identity. Ideas about computers, about how they work and what they can and cannot do, can be used to assert ideas about people. They can provide metaphors for thinking about oneself. Many hobbyists, fascinated by the idea of someday being able to trace out the complex relationships of electronic events, machine and assembly language instructions, and higher-level language commands within the computer, used the image of these many levels of intelligence to think about how people might work, about how people might, or might not, be like machines.

Some of this epistemology was implicit, for example when people made comments about "using the computer's software to think about my wetware." And sometimes, although less frequently, the issue became quite explicit. Conversations that began with descriptions of household "robotics" projects, such as a plan to build an energy monitoring system, led to reflections on how these projects would require programs that could represent the system's knowledge to itself, and from there into formal epistemological reveries: were these the kind of self-representing programs that ran inside people's heads? Do people have different kinds of self-representation programs for representing different kinds of knowledge, such as the knowledge of a dream and the knowledge of being awake? What kind of self-representation program might be running within people that allows them to remember and then forget their dreams?

Constructing Reality

Hobbyist's descriptions of what it is like to work with their own computers frequently referred to the idea that the computer provides a safe corner of reality. Other hobbies can give a similar sense of security but often exact a price. For example, people can feel safe but limited. Alan, the French major, now "in love with his TRS-80," felt secure in his carpentry hobby, but he experienced it as a safety that came from refusing challenge. "It was an 'arty' hobby. We couldn't see myself any other way." The computer is more likely than most other media to allow the experience of playing worlds (let us call them "microworlds") that are secure and also adventurous enough to allow for mind-stretching explorations. Almost all of the hobbyists we interviewed described some version of a limited, safe, and transparent microworld that was embodied in their personal computer.

For Alan just the fact of working with a computer created such a world. For others, it was more specific: a morse code microworld; a text editor microworld; and, most generally, the assembly language microworld. This use of the computer as a place to build a microworld is particularly salient for children when they are put in computational environments in which they have access to programming. Elsewhere I shall report on mv study of elementary schoolchildren and adolescents who have gone through an experience with a LOGO computer system—a System designed to allow children to work in visually represented microworlds that are sufficiently constrained for a child to be able to understand and control them, yet sufficiently open-ended to give a child a real experience of intellectual power. Here I will mention only briefly one child and what she made of the computer as a medium for world building. Deborah at 11 years old was the baby of her family, the youngest of three children. Her childhood, had been spotted with illnesses which further exaggerated her "baby" position.

The members of her family were always doing things for her, insisting that she was not old enough to do the things that she wanted to do most: take out the laundry, baby-sit, stay over at a friend's house, choose her own hair style and her own clothes. Dependent on others at home, very overweight, and with an image of herself as sick and weak, Deborah had little sense of her own boundaries, her ability to say no, to assert control. Even at 11, she had become involved with a crowd of older kids who were smoking, drinking, using drugs. Towards the end of her eleventh year a LOGO computer came into Deborah's classroom as part of an educational experiment. At first, she found the computer frightening and threatening: until one day she hit upon the idea of confining the designs she made with the computer to ones in which the lines always came together in multiples of 30 degrees. She called it her "30 degrees world."

This restriction defined for her a space in which she felt safe but in which she was able to produce designs of great ingenuity and complexity. When I interviewed her two years later I found that she had used her experience with the "30 degrees world" as a kind of model, an experience-to-think-with. In her mind it represented how you could take control and make things happen by the judicious use of constraint. In Deborah's words, it was the first time that she ever "laid down her own laws." It was a turning point in her ability to take control of other situations. She lost 20 pounds, has given up smoking and drugs, and says that she "only sometimes" has a drink. For Deborah and for many adult hobbyists the sense of safety with the computer derived from the feeling of working in a sphere of intelligibility and transparency, a sphere that is protected, much as the space of a psychotherapeutic or psychoanalytic relationship is set off, bracketed.'

People talked about feeling safe and secure in the world they had built with their home computers, a world where there were few surprises and "things didn't change unless you wanted them to." Of course, there was much talk of problems, of false starts, of frustrations. There are "bugs" in hardware and in programs. Things don't work; things go wrong. But bugs, with time, are either fixed or become "known" bugs. Joe is an insurance salesman in a small North California suburb who owns a second-hand Commodore Pet "with a lot of hardware problems." To Joe the bugs in his system "have become almost like friends": "We turn on the machine and we systematically check for my 'old friends,' and we swear, finding them there has a certain reassuring element."

Political Ideology

The use of the computer as a medium for building a transparent and intelligible world brings us to how it can be used to think through questions of political ideology. Fred sells components for a large electronics supply house. He narrowly escaped starvation in a prisoner of war camp during World War II, and from that experience he says that he took " a sense of optimism." "we mean, if there is something out there and you want to do it-do it, understand it, act." Fred has tried to live that way. He is active in local politics; he keeps up with the news; he writes letters to the editor of his town newspaper. He bought his TRS80 on an impulse because "it seemed that you wouldn't be able to understand American society any more if you didn't know about computers." When it comes to working with his computer, Fred wants to know "exactly how things work": "There is a big gap in my own mind between the fact that an electrical circuit can be on or off and the binary number system ... and again from there to the BASIC language. I've got to understand all of that."

When hobbyists like Fred spoke about "wanting to know exactly how things work" in their computers, they were usually talking about wanting to know how their systems were built up from level to level, from electrical circuit to high-level language command. Fred, for example, expressed sharp frustration at gaps in his ability to follow the system through: "we can't really follow the continuum. I'm going to a user's group meeting, I'm talking to people and reading books and some of it is helping, but I am really frustrated. We want to be able to follow the whole thing through."

Larry, 35, lives in a Boston suburb. His computer offered him the first way he ever found to challenge the school's judgments of his child's abilities. A year before we met him, Larry had bought an Apple computer for small business use and ended up bringing it home so that his children would be able to play with it. His 12-year-old son Todd had been judged "backward" by his teachers through six years of schooling. His maths scores were low; he could barely read. But Todd picked up the Apple manual, taught himself how to use the game packages, and then taught himself how to program in BASIC. As far as anyone knows, the Apple manual was the first book that Todd had ever read. In three weeks Todd was writing his own games, games that demanded an understanding of variables and a knowledge of geometry that his teachers claimed he didn't have. Larry is starting to demand more from his son's teachers. His experience with Todd has made him optimistic about what computers will mean for politics because "people will get used to understanding things, of being in control of things, and then, will demand more." This optimism is widely shared among hobbyists, part of a distinct style of talking about computers and politics.

Hobbyists, like Fred and Larry, take what is most characteristic about their relationships with the computer—using computers to build safe microworlds of transparent understanding-and turn it into a political metaphor. Hobbyists associate images of computational transparency and of "knowing how the in machine works" with a kind of politics where relations of power will be transparent, where people will control their destinies, where work will facilitate a rich and balanced cognitive life, and where decentralized power will follow from decentralized information resources. For many hobbyists a relationship with their home computer carries longings for a better and simpler life in a more transparent society. *Co-Evolution Quarterly*, *Mother Earth News*, *Runner's World*, and *Byte Magazine* lie together on hobbyists' coffee tables.

Small computers become the focus of hopes of building cottage industries that will allow people to work out of their homes, have more personal autonomy, not have to punch time cards, and be able to spend more time with their family and out of doors. Some see personal computers as a next step in the ecology movement: decentralized technology will mean less waste. Some see personal computers as a way for individuals to assert greater control over their children's education, believing that computerized curricula will soon offer children better education at home than can be offered in today's schools. Some see personal computers as a path to a new populism: personal computer networks will allow citizens to band together to send mail, run decentralized schools, information resources, and local governments. In sum, many of the computer hobbyists I have interviewed talk about the computers in their living-rooms as windows onto a future where relationships with technology will be more direct, where people will understand how things work, and where dependence on big government, big corporations, and big machines will end. They imagine the politics of this computer-rich future by generalizing from their special relationship to the technology, a relationship characterized by simplicity and a sense of control.

Alienation from Work

Over 40 per cent of those who responded to my survey worked, or had once worked, as computer programmers. The programmer is typically in a situation where he or she is in touch with only a very small part of the problem that is being worked on. Increasingly, programmers work in large teams where each individual has very little sense of the whole, of how it all fits together.

Programmers have watched their opportunities to exercise their skill as a whole activity being taken away (for those who are too young, the story of the process remains alive in the collective mythology of the shop). They have watched their work being routinized, being parcelled out into the well defined modules that make up the tasks of the structured programming team. They mythologize a golden age. This live experience at work made programmers particularly sensitive to the parcellization of knowledge and to the alienation from a sense of wholeness in work. And they bring this sensitivity to their home computer hobbies.

Hannah worked as a programming consultant for a large business system for ten years before starting her own consulting company through which she offers her services to other computer hobbyists. To her, nothing is more depressing than working on a tiny piece of a problem. In her old job, "most of the time we didn't even know what the whole problem was." She likes working with computers at home because she has more control of her time and can spend more time with her family. But she says that what is most important about working with a personal computer is that "we can finally think about a whole problem." Hannah's feelings were widely shared among the hobbyists interviewed, most notably among programmers, ex-programmers, and "team engineers." Images of lack of intellectual balance, of fragmentation, of not being connected, came up often, with the computer at home usually placed in the role of righting what had been wrong at work. As Hannah put it: "With my computer at home we do everything, we see my whole self, all my kinds of thinking."

For these people, having a computer at home meant "thinking experiences" where they could see what their real capacities were, where they had a chance to try things out. Karl, for example, who worked for a long time as a programmer and who is now an engineer in a microprocessor firm, had thought a lot about his mental "ecology."

In the first half of an engineering project there is literally nothing coming together and that's when we find that we need to go home and put something together. We used to make lamps like those in the living room out of glass bottles. But then towards the middle of an engineering project, things did start to come together and we would lose the need for cutting glass. But if you never or rarely get to finish things at work, if your job is basically to make little pieces and it's somebody else's job to make them fit into a whole, then working with the computer at home can give you an experience of getting it all together. You do the whole thing—building up from machine code to finished project. It makes you feel in balance.

When we spoke with Karl he was at a point where everything at work seemed "pretty disconnected." Or, as Karl put it, if he hadn't had his home computer, it would have been "glass cutting time". Karl saw his current work with the computer as a corrective to fragmentation on the job. During our interview, he spoke to me about his current plans for revamping his computer system. Computer systems like his which have several components (keyboard, display screen, printer) have to be connected by using devices, usually simple circuits, called "interfaces." Karl's plan called for using a separate microprocessor as an interface for each component. He had conflicts about the "rationality" of his project.

After all, he admitted, it was far from economical if one measures economy in terms of the cost of the hardware. Each interface circuit needed only a few specialized and inexpensive chips. There was no need for a general-purpose microprocessor at each node. Specialized chips could do the job more cheaply, but could not satisfy his desire to experience the system as maximally coherent.

By replacing the special-purpose circuits by separate, but identical, general-purpose microprocessors, the whole system, at least in Karl's eyes, became uniform, intelligible, and systematic. Karl could not help concluding: "We guess you could say that my choice of projects is not always rational." But from the perspective of using a relationship with computation as a way of "working through" personal concerns, Karl's rationale was clear. For Karl the "inexpensive" solution., using a collection of opaque, adhoc circuits, felt unintelligible. It felt to him like his work situation. His plans for his multiprocessor system were dictated by the logic of compensation rather than by the logic of material economy.

The Mind and Body of the Machine

In studying the hobbyist experience we have found people, largely people with technical backgrounds, in intense involvement with machines. They describe their work (or rather their leisure) with the computer as different from what they have done before with other hobbies. They describe it as an involvement with greater personal consequence. Some of the sense of consequence comes from an historical moment: the computer hobby is seen as signifying a place in the "*avant garde*". Although in some circles "computer person" is a term of derision, the hobbyist experiences it with pride.

Some of the sense of consequence comes from experiencing an individualistic and independent relationship with computation that can be mythologized as belonging to a now-past "golden age" of the programmer. But most of the sense of consequence comes from the holding power and intensity of the time spent with the computer. What is there about these people and these machines that makes possible relationships of such power and such intensity? For me, the relationships that hobbyists form with their home computers can be partially captured with a metaphor of the "mind" and the "body" of the machine.

The "mind" of the computer is that side of computation that involves thinking in terms of high-level programs. In this metaphor, relating to the "body" of the computer means not only working on hardware, but also, and indeed especially, working with programs in a way that is as close as possible to the machine code—that is to say, as close as possible to the core of the computer, its central processing unit (CPU). In terms of this metaphor I have found that the prototypical hobbyist is trying to get into a relationship with the body (rather than the mind) of the machine, to assert power and control in the relationship with the computer, and to create safe worlds of transparent understanding. In trying to find concepts for thinking more clearly about what draws the hobbyist to this kind of relationship with the CPU and about what its meaning might be, we find three issues particularly salient. We think, moreover, that, although we formulate them here in terms of computers, they are relevant to understanding relationships with other technologies as well. The first issue goes back to control.

The hobbyist complains of a work situation where everyone suffers from the constant presence of intermediaries. Bureaucracies stand between the programmer and the computer, a bureaucracy that schedules the computer, that decides its up and down time, that apportions the work for its software design and decides on priorities and procedures for access to it. At work, when something goes wrong with the system it is usually the fault of an intermediary person, one of the many "somebody elses" who deal with the machine. Or it may be the fault of a technical intermediary, one of the many elements in the computer system that mediate between the user and the bare machine: a compiler, an interpreter, an operating system, someone else's program. At home, the hobbyists feel themselves as working directly with the CPU, in complete and direct control of the machine's power.

And when something does blow up, the situation has a special immediacy. It is between them and the bare machine. When a FORTRAN program is run on a large IBM machine the events in the machine are far from being in one-to-one correspondence with the steps of code written by the programmer. Two factors contribute. First, it is in the nature of higher-level languages to work in a conceptual space different from that of the machine. FORTRAN works in a "formula" space, COBOL works in a "business" space, both very different from the space of bits and bytes. Second, the sense of indirect relationship is exacerbated when the compiled code is run by an operating system which allocates memory, mediates control of peripheral devices, and even interleaves the program with other programs. At home hobbyists can follow machine-language programs step by step as their instructions pass through the CPU. They can envision the changes in state of the whole system as being produced by specific actions of the CPU.

And if they suspect that there is a bug in the hardware they can pull out an oscilloscope and see whether the CPU is doing what it should in response to a given instruction. They can figure out where the signals should be going, they can collect their own evidence for what is going wrong, trap and fix the bug themselves. Again and again in my interviewing I heard about the pleasures of debugging-of "going in with meters and scopes and tracking it down." The procedure exhilarates. With every successfully tracked bug comes an affirmation of direct control over the machine. The issue of control was often explicitly recognized by the hobbyists we interviewed. But they lacked a language for naming a second issue which has to do with a notion referred to as "syntonicity" within the psychoanalytic tradition.

Syntonicity implies that we should look for "body-to-body" identification in every powerful relationship with a technology—the body of the person and the body of the machine. It implies that we should understand the appeal of machine language in terms of people's ability to identify with what is happening inside the machine. The CPU of the hobbyist computer lends itself to personal identification with its primary action: moving something that is conceptually almost a physical object (a byte of information) in and out of some thing (a register) that is almost a physical place. The metaphor is concrete and spatial. One can imagine finding the bytes, feeling them, doing something very simple to them, and passing them on. For many of the people that we met in the hobbyist culture, getting into this kind of identification feels safe. It makes the machine feel real. There is a third issue raised by the hobbyists' relationship to the CPU. It is an aesthetic one.

The generation of hobby computers that was born in the 1970s are very primitive machines. The hobbyist thinks of much about them as "klugey," a computerist's way of saying that one is dealing with a compromise, a collection of patches whose structure has been dictated by arbitrary corporate decisions, by economic necessities. The corner of the hobbyist machine that seems to them to have the greatest "intellectual integrity," that distills what they feel to be a tradition of some of the best ideas in computer science, that comes closest to being "clean," is the CPU. And so it is natural for the hobbyist to seek the closest possible contact with it. For a culture in which there is a widely shared aesthetic of simplicity, intelligibility, control, and transparency, getting into the "un-klugey" part of the machine and working in machine code seems the most aesthetically satisfying way to use the personal computer as an artistic medium.

EIGHT

Computer for Children

Instrumental uses of the computer to help people think have been dramatized in science fiction. Many cultural barriers impede children from making scientific knowledge their own. Among these barriers the most visible are the physically brutal effects of deprivation and isolation. Other barriers are more political. Many children who grow up in our cities are surrounded by the artifacts of science but have good reason to see them as belonging to "the others"; in many cases they are perceived as belonging to the social enemy. Still other obstacles are more abstract, though ultimately of the same nature.

Most branches of the most sophisticated modern culture of Europe and the US are so deeply "mathophobic" that many privileged children are as effectively (if more gently) kept from appropriating science as their own. Space-age objects, in the form of small computers, will cross these cultural barriers to enter the private worlds of children everywhere. They will do so not as mere physical objects. This chapter is about how computers can be carriers of powerful ideas and of the seeds of cultural change, how they can help people form new relationships with knowledge that cut across the traditional lines separating humanities from sciences and knowledge of the self from both of these. It is about using computers to challenge current beliefs about who can understand what and at what age. It is about using computers to question standard assumptions in developmental psychology and in the psychology of aptitudes and attitudes. It is about whether personal computers and the cultures in which they are used will continue to be the creatures of "engineers' alone or whether we can construct intellectual environments in which people who today think of themselves as "humanists" will feel part of, not alienated from, the process of constructing computational cultures.

But there is a world of difference between what computers can do and what society will choose to do with them. Society has many ways to resist fundamental and threatening change. Thus, this book is about facing choices that are ultimately political. It looks at some of the forces of change and of reaction to those forces that are called into play as the computer presence begins to enter the politically charged world of education.

In many schools today, the phrase "computer-aided instruction" means making the computer teach the child. One might say the computer is being used to program the child. The child programs the computer and, in doing so, both acquires a sense of mastery over a piece of the most modern and powerful technology and establishes an intimate contact with some of the deepest ideas from science, from mathematics, and from the art of intellectual model-building.

The learning paths have led hundreds of children to becoming quite sophisticated programmers. Once programming is seen in the proper perspective, there is nothing very surprising about the fact that this should happen. Programming a computer means nothing more or less than communicating to it in a language that it and the human user can both "understand." And learning languages is one of the things children do best. Every normal child learns to talk. Why then should a child not learn to "talk" to a computer? There are many reasons why someone might expect it to be difficult. For example, although babies learn to speak their native language with spectacular ease, most children have great difficulty learning foreign languages in schools and, indeed, often learn the written version of their own language none too successfully. Isn't learning a computer language more like the difficult process of learning a foreign written language than the easy one of learning to speak one's own language? And isn't the problem further compounded by all the difficulties most people encounter learning mathematics?

Two fundamental ideas run through this chapter. The first is that it is possible to design computers so that learning to communicate with them can be a natural process, more like learning French by living in France than like trying to learn it through the unnatural process of American foreign-language instruction in classrooms. Second, learning to communicate with a computer may change the way other learning takes place. The computer can be a mathematics-speaking and an alphabetic speaking entity. We are learning how to make computers with which children love to communicate. When this communication occurs, children learn mathematics as a living language.

Moreover, mathematical communication and alphabetic communication are thereby both transformed from the alien and therefore difficult things they are for most children into natural and therefore easy ones. The idea of "talking mathematics" to a computer can be generalized to a view of learning mathematics in "Mathland"; that is to say, in a context which is to learning mathematics what living in France is to learning French.

It is generally assumed that children cannot learn formal geometry until well into their school years and that most cannot learn it too well even then. But we can quickly see that these assumptions are based on extremely weak evidence by asking analogous questions about the ability of children to learn French. If we had to base our opinions on observation of how poorly children learned French in American schools, we would have to conclude that most people were incapable of mastering it. But we know that all normal children would learn it very easily if they lived in France. Much of what we now see as too "formal" or "too mathematical" will be learned just as easily when children grow up in the computer rich world of the very near future. We use the examination of our relationship with mathematics as a thematic example of how technological and social processes interact in the construction of ideas about human capacities. And mathematical examples will also help to describe a theory of how learning works and of how it goes wrong. Take from Jean Piaget a model of children as builders of their own intellectual structures. Children seem to be innately gifted learners, acquiring long before they go to school a vast quantity of knowledge by a process call "Piagetian learning", or "learning without being taught." For example, children learn to speak, learn the intuitive geometry needed to get around in space, and learn enough of logic and rhetorics to get around parents-all this without being "taught". We must ask why some learning takes place so early and spontaneously while some is delayed many years or does not happen at all without deliberately imposed formal instruction.

Child as Builder

If we really look at the "child as builder" we are on our way to an answer. All builders need materials to build with. Where we are at variance with Piaget is in the role we attribute to the surrounding cultures as a source of these materials. In some cases the culture supplies them in abundance, thus facilitating constructive Piagetian learning. For example, the fact that so many important things (knives and forks, mothers and fathers, shoes and socks) come in pairs is a "material" for the construction of an intuitive sense of number.

But in many cases where Piaget would explain the slower development of a particular concept by its greater complexity or formality, we see the critical factor as the relative poverty of the culture in those materials that would make the concept simple and concrete. In yet other cases the culture may provide materials but block their use. In the case of formal mathematics, there is both a shortage of' formal materials and a cultural block. The mathophobia endemic in contemporary culture blocks many people from learning anything they recognize as "math," although they may have no trouble with mathematical knowledge they do not perceive as such.

We shall see again and again that the consequences of mathophobia go far beyond obstructing the learning of mathematics and science. They interact with other endemic "cultural toxins," for example, with popular theories of aptitudes, to contaminate people's images of themselves as learners. Difficulty with school math is often the first step of an invasive intellectual process that leads us all to define ourselves as bundles of aptitudes and inaptitudes, as being "mathematical" or "not mathematical," "artistic" or "not artistic," "musical" or "not musical," "profound" or "superficial," "intelligent" or "dumb."

Deficiency becomes identity and learning is transformed from the early child's free exploration of the world to a chore beset by insecurities and self-imposed restrictions. Two major themes—that children can learn to use computers in a masterful way, and that learning to use computers can change the way they learn everything else—have shaped research agenda on computers and education. The metaphor of imitating the way the child learns to talk has been constantly with us in this work and has led to a vision of education and of education research very different from the traditional ones. For people in the teaching professions, the word "education" tends to evoke "teaching," particularly classroom teaching. The goal of education research tends therefore to be focused on how to improve classroom teaching. But if, the model of successful learning is the way a child learns to talk, a process that takes place without deliberate and organized teaching, the goal set is very different.

The classroom as an artificial and inefficient learning environment that society has been forced to invent because its informal environments fail in certain essential learning domains, such as writing or grammar or school math. The computer presence will enable us to so modify the learning environment outside the classrooms that much if not all the knowledge that schools presently try to teach with such pain and expense and such limited success will be learned, as the child learns to talk, painlessly, successfully and without organized instruction. This obviously implies that schools as we know them today will have no place in the future. But it is an open question whether they will adapt by transforming themselves into something new or will wither away and be replaced.

Although technology will play an essential role in the realization of our vision of the future of education, our central focus is not on the machine but on the mind, and particularly on the way in which intellectual movements and cultures define themselves and grow. Indeed, the role we give to the computer is that of a carrier of cultural "germs" or "seeds" whose intellectual products will not need technological support once they take root in an actively growing mind. Many if not all the children who grow up with a love and aptitude for mathematics owe this feeling, at least in part, to the fact that they happened to acquire "germs" of the "math culture" from adults, who, one might say, knew how to speak mathematics, even if only in the way that Moliere had M. Jourdain speak prose without knowing it. These "math-speaking" adults do not necessarily know how to solve equations; rather, they are marked by a turn of mind that shows up in the logic of their arguments and in the fact that for them to play is often to play with such things as puzzles, puns, and paradoxes. Those children who prove recalcitrant to math and science education include many whose environments happened to be relatively poor in math-speaking adults. Such children come to school lacking elements necessary for the easy learning of school math.

School has been unable to supply these missing elements, and, by forcing the children into learning situations doomed in advance, it generates powerful negative feelings about mathematics and perhaps about learning in general. Thus, is set up a vicious self-perpetuating cycle. For these same children will one day be parents and will not only fail to pass on mathematical germs but will almost certainly infect their children with the opposing and intellectually destructive germs of mathophobia. Fortunately, it is sufficient to break the self-perpetuating cycle at one point for it to remain broken forever.

The Turtle is a computer-controlled cybernetic animal. It exists within the cognitive minicultures of the "LOGO environment," LOGO being the computer language in which communication with the Turtle takes place. The Turtle serves no other purpose than of being good to program and good to think with. Some Turtles are abstract objects that live on computer screens. Others, like the floor Turtles shown in the frontispiece, are physical objects that can be picked up like any mechanical toy. A first encounter often begins by showing the child how a Turtle can be made to move by typing commands at a keyboard. FORWARD 100 makes the Turtle move in a straight line a distance of 100 Turtle steps of about a millimetre each. Typing RIGHT 90 causes the Turtle to pivot in place through 90 degrees. Typing PENDOWN causes the Turtle to lower a pen so as to leave a visible trace of its path, while PENUP instructs it to raise the pen. Of course, the child needs to explore a great deal before gaining mastery of what the numbers mean. But the task is engaging enough to carry most children through this learning process.

The idea of programming is introduced through the metaphor of teaching the Turtle a new word. This is simply done, and children often begin their programming experience by programming the Turtle to respond to new commands invented by the child such as SQUARE or TRIANGLE or SQ or TRI or whatever the child wishes, by drawing the appropriate shapes. New commands, once defined, can be used to define others. For example, just as the house is built out of a triangle and a square, the program for drawing it is built out of the commands for drawing a square and a triangle.

There are four steps in the evolution of this program. From these simple drawings the young programmer can go on in many different directions. Some work on more complex drawings, either figural or abstract. Some abandon the use of the Turtle as a drawing instrument and learn to use its touch sensors to program it to seek out or avoid objects. Later children learn that the computer can be programmed to make music as well as move Turtles and combine the two activities by programming Turtles to dance. Or they can move on from floor Turtles to "screen Turtles," which they program to draw moving pictures in bright colours. The examples are infinitely varied, but in each the child is learning how to exercise control over an exceptionally rich and sophisticated "micro-world".

Readers who have never seen an interactive computer display might find it hard to imagine where this can lead. As a mental exercise they might like to imagine an electronic sketchpad, a computer graphics display of the not-too-distant future. This is a television screen that can display moving pictures in colour. You can also "draw" on it, giving it instructions, perhaps by typing, perhaps by speaking, or perhaps by pointing with a wand.

On request, a palette of colours could appear on the screen. You can choose a colour by pointing at it with the wand. Until you change your choice, the wand draws in that colour. Up to this point the distinction from traditional art materials may seem slight, but the distinction becomes very real when you begin to think about editing the drawing. You can "talk to your drawing" in computer language. You can "tell" it to replace this colour with that. Or set a drawing in motion. Or make two copies and set them in counter-rotating motion. Or replace the colour palette with a sound palette and "draw" a piece of music.

You can file your work in computer memory and retrieve it at your pleasure, or have it delivered into the memory of any of the many millions of other computers linked to the central communication network for the pleasure of your friends. That all this would be fun needs no argument. But it is more than fun. Very powerful kinds of learning are taking place. Children working with an electronic sketchpad are learning a language for talking about shapes and fluxes of shapes, about velocities and rates of change, about processes and procedures. They are learning to speak mathematics, and are acquiring a new image of themselves as mathematicians.

Some of the children were highly successful in school, some were diagnosed as emotionally or cognitively disabled. Some of the children were so severely afflicted by cerebral palsy that they had never purposefully manipulated physical objects. Some of them had expressed their talents in "mathematical" forms, some in "verbal" forms, and some in artistically "visual" or in "musical" forms. Of course, these children did not achieve a fluency in programming that came close to matching their use of spoken language. If we take the Mathland metaphor seriously, their computer experience was more like learning French by spending a week or two on vacation in France than like living there. But like children who have spent a vacation with foreign-speaking cousins, they were clearly on their way to "speaking computer."

When we have thought about what these studies mean we are left with two clear impressions. First, that all children will, under the right conditions, acquire a proficiency with programming that will make it one of their more advanced intellectual accomplishments. Second, that the "right conditions" are very different from the kind of access to computers that is now becoming established as the norm in schools.

The conditions necessary for the kind of relationships with a computer that we will be writing about in this book require more and freer access to the computer than educational planners currently anticipate. And they require a kind of computer language and a learning environment around that language very different from those the schools are now providing. They even require a kind of computer rather different from those that the schools are currently buying. It will take most of this chapter to convey some sense of the choices among computers, computer languages, and, more generally, among computer cultures, that influence how well children will learn from working with computation and what benefits they will get from doing so. But the question of the economic feasibility of free access to computers for every child can be dealt with immediately.

Our vision of a new kind of learning environment demands free contact between children and computers. This could happen because the child's family buys one or a child's friends have one. For purposes of discussion here (and to extend our discussion to all social groups) let us assume that it happens because schools give every one of their students his or her own powerful personal computer. Most "practical" people (including parents, teachers, school principals, and foundation administrators) react to this idea in much the same way: "Even if computers could have all the effects, you talk about, it would still be impossible to put your ideas into action. Where would the money come from?" What these people are saying needs to be faced squarely. They are wrong. Let's consider the cohort of children who will enter kindergarten in the coming year, the "Class of 2010," and let's do some arithmetic. The direct public cost of schooling a child for thirteen years, from kindergarten through twelfth grade, is over $20,000 today (and for the class of 2000, it may be closer to $30,000).

A conservatively high estimate of the cost of supplying each of these children with a personal computer with enough power for it to serve the kinds of educational ends described in this book, and of upgrading, repairing, and replacing it when necessary would be about $1000 per student, distributed over 13 years in school. Thus, "computer costs" for the class of 2000 would represent only about 5 per cent of the total public expenditure on education, and this would be the case even if nothing else in the structure of educational costs changed because of the computer presence.

But in fact computers in education stand a good chance of making other aspects of education cheaper. Schools might be able to reduce their cycle from 13 years to 12 years; they might be able to take advantage of the greater autonomy the computer gives the students and increase the size of classes by one or two students without decreasing the personal attention each student is given. Either of these two moves would "recuperate" the computer cost. Our goal is not educational economies: it is not to use computation to save a year off the time a child spends in an otherwise unchanged school or to push an extra child into an elementary school classroom.

Impact on Students

George Mamunes, 14, a gangling ninth-grader dressed in flannel shirt, blue jeans and hiking boots, knits his thick, dark eyebrows while putting the finishing touches on a computer program, already nearly 300 lines long. For those uninitiated in the special languages of the computer age, it looks like a hopeless mess of numerical gibberish. But when completed, these instructions should produce a computer image of the heart detailed enough to show every major artery and vein, as well as valves and chambers.

The electronic heart is part of a teaching tool George is putting together for eighth-grade biology classes. A few feet away sits Pam Miller, 14, a ninth-grader with long, brown hair draped far down her back. She is operating a computer program or software that simulates the workings of a nuclear reactor. Today she is fine-tuning the section that governs the control rods, those regulators of the reactor's nuclear fires. Tapping away at the keyboard, Pam explains: "You have to maximize the power output without destroying the reactor." Suddenly, flashing numbers burst upon the screen. "There," says Pam, her face lighting up. "Reactor overheated. Power output low. Reactor core damaged. Melt-down!" A disaster that she has brought on intentionally, just to show how it could happen.

Other disciples, seated at terminals scattered around the room, are no less absorbed. Meilin Wong, 15, chic in blue velour blouse, jeans and Bass moccasins, is trying to figure out what went wrong with her business data management program. She is an old hand at such troubleshooting, having spent much of last semester "debugging" a program that, when printed-out, stretches over 30 ft. Jim McGuire, 13, is creating a video game called Spaceship, which will let electronic star warriors zap a boxy-looking orbital intruder.

A more mundane program is emerging from 15 year old Dave McCann's terminal: a verb test for seventh and eighth-grade Spanish classes. Off in a corner two youngsters are putting the impish face of *Mad* magazine's cartoon hero, Alfred E. Neuman, onto the computer screen. Says Muller, as he presides proudly over these after-hours computer converts: “No one told them they have to be here. They’re not usually doing assignments. They’re experimenting. They’re letting their imaginations run free.” Muller’s disciples are not all math whizzes. Or straight - A students. Or particularly precocious. They are reasonably normal youngsters who have grown up with computers. For them, in ways that few people over 30 can understand, manipulating these complex machines is as natural as riding a bike, playing baseball or even solving Rubik’s cube. Like thousands of others across the country, they are part of a revolutionary vanguard: the computer generation. Not only is this generation propelling traditional education down promising avenues, it is tugging at the entire social fabric, foreshadowing changes at least as startling and momentous as those ushered in by a new generation of automobile users more than a half-century ago.

In the classroom, where youngsters are being introduced to the machines as early as kindergarten, they astound-and often outpace-their teachers with their computer skills. After school they gather at the mushrooming number of computer stores (more than 1500 at last count) or join the computer clubs that are becoming a regular part of the youthful landscape. Huddling around any available machine, they argue over their programs, solve computer problems and swap software as intensely as kids once haggled over baseball cards.

In the summer, they may even go off to computer camps, another growth industry, and if they are Boy Scouts, they may try for a computer merit badge. During mischievous moments, they may tinker with one another’s programs, writing in steps that will flash an unexpected insult or obscenity across a buddy’s video screen. Some try to pick the encoded electronic locks on copyrighted software, taking glee in outwitting their elders, or spin fanciful plots to break into computer networks.

A few turn their skills to profit by showing baffled businessmen how to get idle, new computers to run, or by establishing Delaware based corporations to market their own software creations. To the bafflement of their parents, they talk in a jargon of their own (“Hey, Charlie, you should have POKED instead of PEEKED”). As with so many other changes in contemporary life, the spark for this revolution is technological: a bit of silicon sophistication variously known as the personal, home or microcomputer. No larger than an attache case, apart from its video screen, this mighty mite packs the computing power of machines that two decades ago occupied a full room. Yet the microsi as they are affectionately called, are a relative bargain to buy and are becoming steadily cheaper.

Many models cost under $1000, bringing them within reach of schools, parents or the children themselves. Last week, in the sharpest price break yet, Timex announced it will begin selling a small home computer for a suggested retail price of $99.95. But size and price cannot explain why computers have taken such a strong hold on so many youngsters. Certainly their interest has been stirred by a related rage, video games, whose computer-generated flashes, zaps, and pings have not only all the appeal pinball machines had for their elders but go a significant step further: they pique young minds to learn more about all that electronic prestidigitation. But many experts, and most of the young operatives, agree that the overwhelming attraction of the machines is the lure of control, the pleasure of being able to think out and then make something happen, a satisfaction all too often denied children.

Recognized by students and teachers, alike as his school's best computer programmer, Lewis works afternoons as an instructor for a computer consulting firm, introducing younger children to the machines. Last year his employers sent him to Chicago, where he displayed his special teaching gifts before a meeting of educators. As Lewis told *Time* correspondent Peter Stoler, "we love these machines. I've got all this power at my fingertips. Without computers, we don't know what do be. With them, I'm somebody."

Perhaps because of the faintly macho side of computers, the bug seems to strike many more boys than girls in the pre-adolescent years. Says Steve Siegelbaum, Lewis's teacher: "Maybe it's because boys are pushed more toward math and logic than girls are. Maybe it's because boys are just more aggressive." Paradoxically, the computer passion is often stirred in youngsters who seem least likely to be interested in high tech. Jay Harstad, 12, of Minnetonka, Minn., litters his house with poems and sketches but will do almost anything to avoid doing his math homework. Yet Jay is one of the Gatewood Elementary School's premier computerniks and regularly helps teachers introduce fourth-graders to the machines. At West High School in Wausau, Wis., Chris Schumann, 16, a junior, has made a name for himself by translating musical notes into digital form and getting a computer to play Bach and Vivaldi through its loudspeaker. Originally, Chris regarded computers as remote and forbidding, but that changed when he was introduced to his first micro. "It looked real friendly," he says. "It didn't overpower you. It wasn't this ominous thing but something you could get close to."

The closeness can be contagious. Explains Nick Newman, 15, Muller's chief disciple at Ridgewood: "The more you do on the machine, the more enjoyable it gets. It becomes habit-forming." In Alpena, Mich., youngsters who had learned computer skills in junior high were devastated when they got to senior high school and found too few machines to go around. Says Alpena Elementary School principal Burt Wright: "I've got high school kids begging to come in after school and use our machine." The truly addicted-known half scornfully, half admiringly as computer nerds-may drop out almost entirely from the everyday world. In Lexington, Mass., one legendary 16-year-old nerd got so deeply immersed in computers that he talked to no one, headed straight to his terminal after school and barely sat down for meals.

The only way his father could get him away from the terminal was to go down to the cellar and throw the house's main power switch, cutting off all electricity. Barry Porter, 14, of San Francisco, is a computer-age truant, so attached to the machine that he often skips school, rarely reads anything other than computer manuals and hangs out with his pals in the Market Street computer store, often plotting some new electronic scam. Barry (not his real name) currently boasts an illicit library of about 1000 pirated (i.e. illegally copied) programs worth about $50,000 at retail prices, including such software gems as VisiCalc, the popular business management and planning program. Before security was tightened up, he regularly plugged his computer into such distant databanks as the Source (which provides news bulletins, stock prices, etc.) via telephone without paying a cent.

Computer Generation

No one can say exactly when the computer generation began-certainly not earlier than the 1960s, when computers began appearing in schools. But even computer whizzes in their twenties are acutely aware of how soon they are likely to be outstripped by today's grade schoolers. Says Steven Jobs, the multimillionaire co-founder of Apple Computer Inc.: "These kids know more about the new software than we do." New York computer executive Charles Lecht goes further: "If you were born before 1965, boy, you're going to be out of it." Where their parents fear to tread, the microkids plunge right in, no more worried about pushing incorrect buttons or making errors than adults are about dialing a wrong telephone number. Says mathematician Louis Robinson, IBM's resident computer sage: "They know what computers can and cannot do, while adults still regard them as omnipotent." Hughes Aircraft Chairman Allen Puckett, who used to share an Apple with son Jim, 12 said: "A lot of adults grew up in a-slide-rule world and still reject computers. But computers are as natural to kids as milk and cookies."

More and more members of the computer generation are tasting the heady pleasure of teaching their own teachers how to use the machines and, if they are lucky enough to have computers at home, instructing their parents as well. Says Ridgewood's Newman, a regular teacher of teachers: "It's a sort of mutual doorway. The barriers between adult and child, between teacher and student, are broken, and it's person to person. Nobody's looking down on anyone; they're looking each other right in the eye." Often adults find it easier to ask a child how to do something than to ask another adult. Says University of Kansas education professor Mary Kay Corbitt: "One adult student of mine brought her son to computer class, and we discovered that he was doing her assignment while she watched. Two weeks later she overcame her anxieties and was participating fully." Confronted with the strange and unsettling world of the computer, teachers can get a useful perspective on what it is like to be a student again. After taking part in an elementary course in programming, Lois Brown, 54, a Wausau grade-school teacher, is thoroughly chastened. "Now we realize how little kids feel when there's a concept they don't understand. We sat in that course not wanting anyone to know all the things we didn't understand." Despite their obvious wariness of computers, parents are taking the lead in getting them into the schools. In Florida, communities have staged cake and candy sales, carnivals and tree plantings, weekend car washes, even a bike-athon to raise funds to buy computers. Says Marilyn Neff of Miami: "We feel computers will be the new paper and pencil." Of the 250 computers in the schools of Utica, Mich., more than two-thirds have been purchased by parent-sponsored fund drives. Says Utica Principal Paul Yelinsky: "Moms and dads are coming in and telling the counsellors they have to get their kids in computer classes because it's the wave of the future." So important is computer literacy that the Alfred P. Sloan Foundation is beginning a major program to get even such traditional liberal arts schools as St John's College in Maryland to begin giving courses in it. Though many schools began purchasing computers with federal aid budget cutbacks are drying up that well. Apple's jobs points out that other nations, especially Britain, France and the Soviet Union though surprisingly not the electronics-minded Japanese are paying far more attention to computer education than is the US. Earlier this year, Jobs persuaded California Congressman Pete Stark and Missouri Senator John Danforth to introduce bills in Congress that would allow computer manufacturers to take a hefty tax write-off for any machines they donate to elementary and high schools. Under the present law, full deductions for such scientific equipment are allowed only if it is given to colleges and universities.

Steve Jobs originally spoke of giving an Apple to every public elementary and secondary school in the country, more than 80,000 computers worth as much as $200 million retail. He thought private schools should be included and encouraged other manufacturers to join in the program as well. Meanwhile, Apple's archrival, the Tandy Corp., maker of the Radio Shack computer line, is taking a different tack: it has pledged $500,000 in equipment to spur development of educational programming, or courseware, for the classroom. Many of the approximately 100,000 computers now in US schools-roughly one for every 400 students-are in affluent suburbs like Ridgewood, a national leader in computer education. But the machines are also found in the unlikeliest of places. On a Chippewa Indian reservation in Wisconsin, computers are being used by young members of the tribe to learn their ancient and nearly forgotten language. Alaska's small rural schools have been ordering computers to meet a special need: they allow students of different ages and abilities in the same small classrooms to learn at their own pace. Dubuque, Iowa, the New Yorker founding editor Harold Ross disdainfully located his provincial old lady, has 13 machines and another 20 on order. Bill Holloway, a professor of, computer education at the University of Kansas, calls the spread of small computers in the classroom nothing less than an avalanche. According to various industry studies, there may be from 300,000 to 650,000 computers in the schools by 1985.

So far, the most common, and least interesting, way to use school computers is in direct drill and review. The machine simply quizzes, prods and grades the student, very much like a robot teacher. Hundreds of programs of this type are available for popular computers like the Apple II Plus, Radio Shack's TRS-80 and the Commodore PET. But many of these programs are little more than computerized rehashes of the old classroom flash cards that go back to the days of McGuffey's readers. One notable difference: today when the student answers correctly, the screen will light up with WOWS, HOORAYS or smiling animals. Wrong answers may produce sad or scowling faces, perhaps accompanied by a falling tear. Partly because of teachers' fears of the machines and for their jobs and partly because of the poor quality of software, the frequently heralded electronic revolution in the classroom has been slow to occur. Now, however, it is being pushed along by steady improvements in teaching programs, thanks to imaginative enterprises like the Minnesota Educational Computing Consortium.

One of its more refreshing drills: a program called Wrong Note, which helps teach sight reading of musical scores. As a simple tune emanates from the computer's loudspeaker, matching notes appear on the screen, but sometimes the quiz intentionally errs and obliges students to find the false note. In order to do so, they can order up a repetition of the tune as often as Bogie and company did in Casablanca. Says Kenneth Brumbaugh, director of the consortium's instructional services: "Imagine asking a teacher to play it again and again!" Even very young children can profit from such exercises. At the Nordstom Elementary School in Morgan Hill, Calif., a suburb of San Jose, Colin Devenish, 7, is working with a classmate on the arithmetic drill, honing his skills in addition and subtraction. Unlike youngsters doing such drilling in the past, Colin seems to be enjoying himself enormously. Why? "Because," he replies mischievously within earshot of his teacher, "the computer doesn't yell."

Computers, operated only by touching a few buttons, are also remarkably effective devices for educating the handicapped. At the California School for the Deaf in Fremont, Rhonda Revera, 16, has worked with computers for five years, studying every subject from fractions to spelling. Rhonda offers a paean to the machine in sign language: "Computer makes me remember. It is fast, easy and better than writing on paper." Still another important use of computers is as a remedial tool. One is a spelling drill with a special incentive built into it: if all the answers are correct, a video game pops onto the screen as a reward. When one youngster worked his way through the drill, even classroom hecklers were impressed. Said one: "Hey, Old Wentworth's getting better." More entertaining and demanding are think tank-type strategy games like Geography Search, which launches competing teams on a Columbus-like voyage of exploration. They must make their way across the Atlantic, taking into account currents and winds, finding their longitude and latitude by means of star patterns and the length of a shadow thrown by a stick at high noon (methods that worked for Columbus, after all), and coping with such unforeseen perils as an outbreak of scurvy, an attack by pirates and a tropical storm. Only shrewd planning, wise choices and cooperative action ensure survival. The simulated voyage becomes uncannily real to the participants. Says the game's creator, Thomas Snyder, 31, who heads Computer Learning Connection, Inc., of Cambridge, Mass.: "When they get near the end and the computer finally shows them another ship near by, they act as if they had actually spotted a ship at sea."

Until a few years ago, the few computers available in secondary schools were essentially "dumb" terminals linked by telephone lines to a large, centrally located machine that served a variety of users through an arrangement called time-sharing. All the courseware was stored in the big computer's powerful memory, which could be tapped at will by students and teachers. The most successful example of such a system-and the one still used by Wisconsin's Chippewa Indians-is PLATO (for Programmed Logic for Automatic Teaching Operations).

Developed in the 1960s by the University of Illinois and Control Data Corp., PLATO is an exemplary teacher containing more than 8000 hours of courseware, much of it in a continuous curriculum. Thus, if a youngster forgets a point from an earlier lesson, PLATO will search its prodigious memory and patiently recapitulate. But such time-sharing schemes are extremely expensive, since they require open lines to the central computer. They also can become backed up at peak hours, and do not always lend themselves readily to what is the most intellectually demanding use of the computer: learning how to program it. For this, the inexpensive, easy-to-operate personal computer, entirely self-contained and relying on equipment immediately at the student's side, is an ideal instrument-much more "user-friendly," as manufacturers like to say, than big machines. Yet even with a handy micro, programming can overwhelm the uninitiated. The programmer and computer must "speak" a common language. In the early days of the digital computer, this was, extremely difficult.

The machine reduces all the information it receives, whether it arrives as letters, numbers or graphic symbols, into the simplest possible electronic statements: either a yes or a no, represented by pulses of high or low voltage. To command the machine in its own internal language meant writing out endless strings of ones or zeroes, called bits and bytes, symbolizing those yes or no statements. But scientists soon began creating alternate languages for communicating with the machines that vaguely resemble everyday speech. The most popular of these computer tongues is BASIC (for Beginner's All-purpose Symbolic Instruction Code).

Developed at Dartmouth by mathematician John Kemeny and his colleague Thomas Kurtz to let even the least mathematically gifted student converse with the university's computers, it is "understood" by virtually all of today's personal computers. To show just how easy the language is, Kemeny offers this extremely simple lesson in programming: tell the computer to find the square roots (i.e., the numbers that, when multiplied by themselves, yield the original numbers) of eleven successive values, say 20 through 30.

The entire operation can be accomplished for a program of just four steps:

```
1 FOR N= 20 TO 30
2 PRINT N, SQR(N)
3 NEXT N
4 END
```

Translated into everyday language, the first line tells the computer to let N stand successively for 20 through 30. The second instructs the machine to print the first value of N (that is, the number 20), compute its square root (SQR) and print out the result. The third tells the computer to go on to each of the succeeding values all the way through 30. Finally, the program tells the computer to call it a day, its job having been done. Even the smallest machine can do such calculations in a flash, compared with the hours of work they might require of human computers. To preserve their creativity, the students can readily store their programs on magnetic tape or on a small, 45 rpm-size plastic record called a floppy disc—which is not, as some parents believe, a new form of back injury. Then when the occasion arises for using the program again, the computer operator merely loads the instructions back into the machine and punches in some new values for N.

The same broad principles apply to the creation of all software, even complex simulations like Geography Search. Literal-minded brutes that they are, computers do exactly what they are told. No more and no less. But youngsters of even the most tender age are surprising educators by showing they can master the beasts with startling ease. Computer software expert Leona Schauble of the Children's Television Workshop (producers of Sesame Street) recalls getting an eight-year-old boy at Manhattan's Little Red School House started on a simple computer game.

The game generated an image of a frog that would leap up and catch a butterfly, provided the right buttons were hit. After a few minutes, she checked back and found the frog jumping in slow motion. When she asked the youngster what happened, he replied, "Well, we wanted to make the frog catch more butterflies. So we got a listing of the variables and slowed him down." In other words, the youngster had broken into the game's program and changed it to suit himself.

To instruct very young children, even Kemeny's BASIC is much too mathematical. Instead, more and more schools are turning to an innovative computer language called LOGO (from the Greek word for reason), developed by Seymour Papert and his colleagues at MIT. A mathematician who studied with the Swiss psychologist Jean Piaget, Papert has become something of a guru of the computer generation, predicting that the machines will revolutionize learning by taking much of the mystery out of mathematics, science and technology. Says he: "The computer can make the most abstract things concrete."

With a deceptively simple set of commands, LOGO enables youngsters who know nothing of geometry and algebra, and barely know how to read, to manipulate a triangular figure, dubbed the Turtle, on a computer screen and trace all manner of shapes with it. At the Lamplighter School in Dallas, teachers using LOGO get youngsters of three or four to write simple computer instructions. In one game, they manoeuvre "cars" and "garages" on the computer screen in such a way that the cars are parked inside the garages. While playing with LOGO, the youngsters learn simple words, the difference between left and right, and geometric concepts that they would not ordinarily encounter until junior high.

The machines crop up in the lives of youngsters even before they enter school-and sometimes before they learn to walk or talk-in the guise of such siliconised gadgetry as Little Professor and Speak & Spell. With a few presses of the button, these computerized games produce flashing lights, squealing sounds and disembodied voices that inculcate the rudiments of spelling and calculating. A record of sorts may have been set by Corey Schou, a computer scientist at the University of Central Florida in Orlando: he rigged up a home computer so his five-month-old daughter could operate it by pressing buttons in her crib and changing the designs on a nearby screen. Says the proud papa: "Basically, it's an electronic kaleidoscope, another diversion, another learning device." Whatever it is, it prepares youngsters for all those buttons they will encounter soon enough in and out of school. Parents and teachers may shudder at the thought, but it is only a short hop from skilful operation of a video game to learning fundamentals of programming. Says MIT Sociologist Sherry Turkle, 33, who has been studying the youthful computer culture for five years: "The line between game playing and programming is very thin. Programming takes what is powerful about games this articulation of knowledge, this learning about strategy and carries it to a higher level of power." By the time the youthful programmers reach the eighth or ninth grade, their skills may reach a marketable level. In Chicago, Jonathan Dubman, 14, and Kay Borzsony, 13, have formed a company called Aristotle Software to sell their own computer games and graphics programs. Says Kay: "The nice thing about the computer business is that there is no real bias against children. In the computer magazines, you read articles by 12 and 13 year-olds." Laura Hyatt, 15, of Ridgewood, helps a stymied local insurance office figure out how to use its software. Says she: "It's better than babysitting." And, at $3.50 an hour, somewhat more profitable.

The prodigy of prodigies may be Eugene Volokh, 14, of Los Angeles. A Russian emigre, he earns $480 each week by doing 24 hours of programming for 20th Century-Fox, while carrying a full load of courses as a junior at UCLA. This year Greg Christensen, 18, of Anaheim, Calif., could make $100,000 in royalties from a video game he developed that was bought by Atari. Other youngsters are waiting at the sidelines in hopes of catching up with these young entrepreneurs. Every Tuesday night, Scott Whitfield, 13, and his brother Shawn, 11, appear at the Menlo Park, Calif., public library to get computer instruction. Says Scott: "We'll probably never get a job if we don't learn how to use a computer." Not all youngsters take equally to the machines. In a typical computer class, only about one in five students becomes seriously involved. Says Steven Scott, 16, of Wausau's West High: "Either you get the hang of it or you don't." Even so dedicated a computernik as Ridgewood's Nick Newman finds programming interesting only for a purpose. His own goal is to apply his computer knowledge to a career in science or medicine. Whatever these youngsters make of their computer experiences, they will surely confront the world differently from their parents.

The precise, orderly steps of logic required to use and program the machines promise to shape-and sharpen-the thought processes of the computer generation. Indeed, the youngsters playing all those strategy games are doing precisely what corporations do when they plan to launch a new product, or what military leaders do when they devise strategies to confront a potential foe.

Whether such abilities will change the world for the better is another matter. Princeton psychologist George Miller, for one, has doubts that "a few years of thinking like a computer can change patterns of irrational thought that have persisted throughout recorded history."

Other social critics ask if clear thinking is enough - if, in fact, there might not be a danger in raising a generation to believe that it has the analytical tools to contemplate any problem. Says MIT computer science professor Joseph Weizenbaum: "There's a whole world of real problems, of human problems, which is essentially being ignored." It is still impossible, after all, to reduce a human relationship to a printout or to solve a moral question by bits and bytes.

Some critics predict a future not unlike that portrayed in Isaac Asimov's I, *Robot,* a science fiction novel set in a society so thoroughly computer-dominated that the people cannot do arithmetic. Humanist critic George Steiner acerbically calls the computer generation the advance guard of a breed of "computer-mutants." Says Steiner: "They will be out of touch with certain springs of human identity and creativity, which belong to the full use of language rather than mathematical and symbolical codes."

Many others are much more sanguine. University of Chicago philosopher of science Stephen Toulmin predicts that computers will "re-intellectualize" the television generation. "TV relieved people of the necessity to do anything," says Toulmin. "Computers depend on what you do yourself." Catholic theologian David Tracy argues that "using computers sharpens the mind's ability to deal with our world: the world of technology." The final word may be simpler, and not pronounced by elders who find a devilish soul in the new machine. More so than adults, the young know the computer for what it is. Says a 10-year-old at Manhattan's Bank Street School: "It's dumb. We have to tell it everything." They also know something important is afoot. Says Shawn Whitfield: "When I grow up it's going to be the Computer Age.

NINE

Academic Aspects

Private data networks, academic and research networks have been a major influence on the development of information applications in wide area networking, since they offer a venue for experimentation and co-operative activity less constrained by strict commercial viability; the network is normally funded at a national level, by institutional subscription, or through co-operative arrangements, and thus access appears to be free at the point of use. The 1980s saw a steady broadening both of those able to access research networks and in the range of information services and facilities available, as the network infrastructure has developed.

In some cases, it appears that the research network may form the basis of a new generation of high capacity, fast broadband networks offering widespread public access.

Research networks, by their nature, are mostly composed of a collection of co-operating autonomous institutions: organizations involved in higher education and research institutes. The management of research networks tends to reflect this structure by the way in which control is distributed amongst the participants.

Technical control lies with the network 'hubs', those institutions which provide the basic backbone of the network and the major routing facilities, or with a management group specifically set up to control the development and standardization of facilities. This kind of group may exercise control over organizations which can attach to the network, but will not control the registration of individual users, and will not at present maintain any central directory of all users.

Similarly, services which are available over the network are the responsibility of the individual institutions which offer them, as is local control over which types of user are allowed access to these. Although the degree of decentralization will vary across different networks, in relation to the primary source of funding, there is generally no single central locus of control which determines all aspects of network use. There are, therefore, no centrally maintained directories or lists of all the services which might be available; it is up to the individual user to discover this information for themselves, perhaps aided by information or resource tools created through the initiative of other users, often on a rather *ad hoc* basis.

Navigating, through a research network is, at the moment, a matter of individual learning and initiative, presenting a steep learning curve to novice users. This barrier extends not only to knowledge of the specific facilities which might be available, but also to the way in which these facilities can be used on different computers and software systems, which may vary tremendously from one site to another.

The basic information services which are offered by research networks fall into three main categories. *Information exchange* services allow individuals and groups to communicate with each other across time zones and geographical space much more quickly than postal delivery would allow, and has proved the most popular and heavily used end-user application on networks by linking together people who would otherwise find it hard to meet face to face.

— *Electronic mail* facilities enable the exchange of messages and ideas on a one-to-one basis, whether more formally as in the sharing of notes and drafts of work-in-progress, or informally for general social contact.

— *Discussion groups* which use electronic mail in conjunction with mailing lists support the exchange of information and messages between all members of a group at the same time. This is popular for ongoing debates and the exchange of news items in discipline-specific areas, allowing members to share ideas and experiences in an informal setting, to 'meet' new contacts, and to follow and influence the development of new ideas in their field.

— *Bulletin boards* are used in discipline-specific settings, but require a positive action from the user in connecting to the service, and reading or downloading items of interest, rather than depending on the blanket distribution of material through mailing lists.

— *Computer conferencing* and interactive videoconferencing are still quite rare activities on such networks, but will become increasingly common as the technical infrastructure becomes capable of real-time interactive and multimedia services, and are likely to have substantial implications for distance learning education.

File and document distribution support the distribution of items such as text documents, image files, and software programs on an individual search and retrieve basis.

File transfer is permitted by the use of standard protocols so that data files can quickly be moved from one machine to another under the control of the end-user. As more electronic material is becoming available within the public domain, certain computer network sites offer an 'archive' service, in which data files can be stored on open access for retrieval by all network users.

Mail servers are commonly used as primitive database systems, returning files in response to commands received within electronic mail messages, as an automatic process with no manual intervention. This is often the method used for distribution of the composite articles within an electronic journal, for example—the user will request a particular file by name by placing the appropriate commands and filename within a mail message addressed to the mail server.

Information services are available through remote log-in protocols, which allow the user to link to a remote system and use it as though there were a direct connection between the two, within the constraints set by the host system. The service may be on a different network, as long as both end systems (user and host) are able to run the same protocols, and the networks and gateway used will support this.

Online public access catalogues available for remote searching from outside the library or campus site are a common use for this facility, as long as the OPAC allows for a remote connection. There is, of course, no common standard interface for the various online catalogues and so the user must know, or be able to learn, the functions and abilities of each different system.

Non-bibliographic databases are increasingly becoming available across research networks, particularly those which contain information within the public domain, or government-sponsored public information. Once again there is no standard searching approach, and the format and content of the databases may be very different.

The full range of academic and research networks worldwide includes both large and small-scale networks at differing stages in their development, and is too numerous to list here. Most research networks are organized on a national basis, or as a cooperative endeavour between smaller countries, as for example in the Gulf states. A comprehensive directory of networks is given in the books by Quarterman and LaQuey, together with a summary of protocols used, and a description of network interconnections.

RARE

The European academic and research networks are brought together in RARE (Reseaux Associes pour la Recherche Europeenne), the association for European research networking organizations and their users. This aims to overcome national boundaries in research networking by creating a harmonized computer communications infrastructure for the whole European research community. Based in Amsterdam, RARE has over 40 national and international research networking organizations as members, which are listed and described in the RARE 'State of the union' report.

The RARE association also sponsors a number of working groups which initiate and co-ordinate advances in network technology and services; these include, for example, work on user support issues, information retrieval across the networks, and developing multimedia services. The academic network for the United Kingdom is called *JANET* (Joint Academic Network). Similar networks exist for nearly all the other European countries, including Eastern Europe, although the range of facilities which each is able to offer can vary widely; in some former Communist countries, such as Romania, access is very difficult and services are unreliable. Others are far more highly developed, with extensive experience and expertise, and are moving towards a strategic position of greater independence from government support and self-sufficiency.

The Dutch academic network, *SURFnet,* is an example within this category. SURFnet networks together institutes of higher education, and scientific and industrial research organizations in the Netherlands, and has maintained a close involvement with the Dutch bibliographic utilities, taking a proactive approach to the development of information services.

In 1989 SURFnet became a private company, with a 49% share held by the Dutch telecommunications company PTT Telecom, and was required to become self-supporting in the near future. This has led it into substantial development in services and user support, working on-site with institutions and with individual discipline-based groups of users between institutions, to encourage and develop new users and uses for the network. There are several other major influences and initiatives which affect the shape of European academic networking at a transnational level, both outside and within the EC. These include:

Organizations which are co-operative efforts between countries grouped together by geography. *NORDUnet* is a collaborative effort between the Scandinavian group of national research networks: Denmark, Finland, Norway, Sweden, and Iceland aiming to enable an efficient connection to the rest of Europe. A similar arrangement has been proposed for the countries in the area of the former Yugoslavia.

EUnet is a large subscription-funded network, which serves European users from Scandinavia and the CIS (former Soviet Union) to North Africa. Founded by the European Forum for Open Systems, EUnet connects 6000 individual sites and networks together as Unix system users.

International treaty organizations which are established as European-wide enterprises, with funding from a number of countries. The most influential of these is CERN, the European Laboratory for Particle Physics, which has constructed *HEPnet* as a dedicated network serving research centre in high energy physics.

Ebone is the European backbone network, which provides a neutral interconnection between the backbone research networks which form part of the global TCP/IP based Internet described later in this chapter. The backbone provides a high capacity 'highway' between the major first-level nodes in Europe (Stockholm, Amsterdam, Geneva, Paris, and London) to which other sites connect.

EuropaNet has developed from an EC-funded project conducted by COSINE (Co-operation for Open Systems Interconnection Networking), intended to develop a coordinated approach to networking the European research community, and provide an international infrastructure accessible by academic organizations, research institutes, and organizations co-operating with EC research and development. A pilot International X.25 Infrastructure (IXI) network commenced in 1990, forming a backbone which linked research networks in many countries, both inside and outside the EC itself. This has now been replaced by EuropaNet as a high capacity multi-protocol backbone which will handle both OSI (open system) and TCP/IP Internet-style protocols.

The *EARN* (European Academic Research Networks) association provides networking services for researchers and academics in Europe, the Middle East and Africa, as a store and forward network based on IBM computers. EARN is closely integrated with the US-based BITNET network, and is in fact almost indistinguishable from this in use, effectively acting as its European arm. In some countries EARN forms the major research network, while in others there is a gateway connection between EARN and the national network.

There are three main international research networks with worldwide coverage, which are normally accessible from national-based networks, and which are based in or emanate from the United States.

The *Internet* described later in this chapter, is the largest of these; based in the US, it now has extensive worldwide coverage. The Internet is beginning to assume a quasi-public status, and is currently the subject of wide-ranging legislative attention.

BITNET is a major US-based co-operative computer network interconnecting over 3000 academic and research institutions in more than 35 countries, as far apart as Eastern Europe, Latin America and the Far East; the network encompasses a number of more local networks within its structure. BITNET was originally funded by the IBM corporation in the early 1980s, and was free to the academic community; this support is now terminated, and access is through institutional subscription. BITNET has close links with the Internet, and many sites are accessible through both networks, almost indistinguishably.

Usenet is a loose assemblage of Unix-based sites, including many thousands of computers, from large systems to personal workstations. It is strongly associated with the Unix world, and has an independent and co-operative spirit, with no central management or control, depending on allegiance to a code of good practice; the nature of Unix systems tends to limit the appeal of Usenet to a fairly specialist community.

The main feature of Usenet is the large number of discussion and newsgroups which exist within it, known as *Network News*, amounting to almost 30 Megabytes of 'news' daily. Essentially Usenet is not a technical data network, and does not have separate fixed lines between sites; it is more of an informal, agreement for the exchange and passing on of 'news' achieved by a variety of means, including dial-up over the public data network. Much of Usenet activity is now carried out via the Internet.

JANET

JANET, the Joint Academic Network, is an X.25 packet-switched service which provides inter-site access for and on behalf of the UK higher education and research community. Initially formed to support scientific data processing, JANET has since expanded to encompass a fundamental role as a communications medium for the academic and research communities as a whole, and now includes a broad range of information services and resources.

JANET was formed in 1984 through a rationalization of the existing academic and research networks which were established at that time in the UK. These included the two national research networks which belonged to the Science and Engineering Research Council and the Natural Environment Research Council, together with eight regional university-based networks.

Since these networks did not all work to the same standards, the introduction of an overall team for management and development of the network was required. The management and control of JANET is currently based at the SERC Rutherford Appleton laboratory in Oxfordshire: the joint network team deals with policy and development matters, and with management of the network infrastructure. Control of the local site facilities and their use of JANET, and administration of the physical connections to the JANET network, is the responsibility of the local representatives; however, these responsibilities must be fulfilled as a condition of connection to the network.

JANET provides the infrastructure which connects the facilities of its users together, rather than determining day to day use of the network. JANET is funded by the joint information systems committee of the UK Higher, Education Funding Councils on behalf of the research and educational institutions, by top-slicing a part of the overall national budget for higher education before it is devolved to institutions.

The network is thereby provided as a general educational and research facility, essentially funded through government agency. Organizations which have been permitted to connect to JANET but which do not fall within this category pay an annual subscription fee. It is expected that in the near future, the management of JANET will be taken over by UKERNA, the UK Educational and Research Networking Association; this represents a move away from direct central government funding. Restrictions on the type of organization which can connect to JANET have been slowly relaxed over time, as the needs of JANET's primary users demand a wider range of networking contacts.

Within the 'guidelines for acceptable use' which require that network usage is related to educational and research purposes, rather than to a directly profit-making enterprise, institutions with primary connections can now allow local secondary users to connect to JANET using their primary connection as an entry node. At the present time, the guidelines allow the following kinds of organizations to directly connect to the network:

- universities and other higher education and research institutes which are funded by a central government agency;
- government research establishments not directly funded;
- not-for-profit organizations engaged in collaborative research with higher education institutes;
- libraries and related organizations, including the national libraries and the UK bibliographic utilities;
- educational councils and research institutes;
- learned societies,
- commercial organizations involved in collaborative research, or providing a service or support to higher education (for example Blackwell's periodicals service).

The JANET network developed as an X.25 (84) packet switching technology as the underlying data transmission scheme. The recent upgrading of data lines has increased the bandwidth capacity to 2 Mbps on most sites, by the use of leased lines from Mercury Communications.

At the inception of a formal JANET service in 1984, X.25 was the only realistic, choice within the European market which was able to offer wide connectivity and reasonable performance, and which was compatible with public telecommunications, while avoiding the pitfalls of proprietary protocols. Since this time, however, the alternative approach of TCP/IP has emerged as a strongly supported protocol within the research community, affecting forward planning of JANET services.

The higher level network functions were originally provided by the proprietary *Coloured Book Software* protocols, designed as an interim measure. This set of protocols provide the basic network services; the Grey Book defines an electronic mail service, for example, and the Blue Book deals with file transfer. The network administration has had the long-term goal of making a transition to the use of OSI protocols; some services, such as X.400 mail, are already available. The development of a Domain Name Server provides a name to address translation to make network connections more transparent for the end user. This is made possible by a Name Registration Scheme (NRS) which maintains a directory of the individual systems which are JANET-accessible on all the 200 or so UK sites (and some others); these have both a network address (a string of numerals prefixed by 0000) and a more usable network name.

The naming scheme can be used by local nodes to distinguish between different machines at the same site by progressively structuring the address into smaller sub-domains. Thus, a full JANET address name will identify the *uk* geographical area, followed by an *ac* to denote the academic network, a single word name for the institution as a whole, and any necessary sub-units in the organization, until an actual machine name is identified (for example, the address *uk. ac. oxbridge, cc. vaxe).* All network node names are registered on the central database, which can be downloaded into local data storage and used by name servers at individual sites.

Individual users are registered on their local machines, and any others to which they have log-in rights, by the naming convention chosen by the individual site. Information objects addressed to a specific username (for example *user@uk. ac. oxbridge, cc. vaxe)* will be delivered to the site by the network; the local system is then responsible for delivering the data to the workspace of the local user.

The basic services of electronic mail, remote log-in and file transfer have given rise to a wide range of different specialized services and access to information resources which are available to all JANET users.

The academic community which JANET serves includes a wide variety of disciplines, many of which are not traditionally linked with computer and information usage. Electronic mail, as well as serving for one to one interpersonal messaging, is often used to support group communication and as a simple interface to online information services.

The Networked Information Services Project (NISP), based at Newcastle University, was set up to explore the possibilities for enhancing, group communication and information dissemination facilities on JANET, with a particular emphasis on drawing in those groups not yet familiar with computer networking. This has led to the development, from 1989 onwards, of a mail server service known as *Mailbase.* Working with groups of academics and researchers from areas as diverse as music, history and library management, Mailbase aims to provide groups with the ability to participate in focused discussions by joining e-mail based discussion lists, each limited to a narrow topic; and for groups to run these information services themselves. Each list has an owner, responsible for managing it, and may be 'open' to public access, or 'closed' to all except invited participants; the degree of access to archived discussion can also be controlled.

Contact with the service is through the sending and receipt of electronic mail messages to the mail server based at Newcastle, and can be done from any JANET site. Some of the lists are working on specific tasks—the development of network training materials, for example—the results of which are made publicly available as retrievable documents from the Mailbase service.

As of 1992, the project reported some 6,000 users of this service, with a total of 12,000 subscriptions to the 200 existing lists. Membership of lists varies from several hundred to single figures. One of the most enthusiastic users of this service has been the academic library community, which formed a number of lists on specific topics, most of which have now merged together under the name of *lis-link.* This generates some 5-10 messages per day, providing a forum for sharing ideas and experiences, seeking answers to specific reference questions, and testing opinion on library-related issues; with the advantage of fast response from other users.

The interchange on discussion lists is normally carried out from within the familiar mail package of local systems, with no requirement for the user to be aware of delivery and connection details. Remote log in is available by invoking the calling facilities of the JANET PAD (the local link into the X.25 packet switching system), and using the local workstation as though it were a terminal with direct connection to a distant system.

In many cases, the called computer will require the caller to have a registered username and password before making any facilities accessible. Some services, however, have been set up for general access by JANET users and do not require a caller to declare an identity.

Bulletin boards require the user to connect to the systems and choose items of interest typically organized under a multi-level menu structure; documents can be read on-screen, or dispatched to the user's home system. Examples of successful bulletin boards on JANET are the Joint Network Team news service, *JANET.news; HUMBUL,* a bulletin board for users in humanities disciplines; *History .news;* and the Bulletin Board for Libraries, BUBL. This latter is a library-oriented information service, containing news, events, announcements, current information on networking issues, and network training materials. The bulletin board also offers a current awareness service, by listing the contents pages of relevant periodicals, and acts as a forum for the sharing of ideas and expertise on a range of LIS topics. Although still used by a relatively low proportion of the library community, BUBL achieved a usage rate of 2400 per month in January 1993.

Online public access catalogues (OPACs) of the educational member institutions are normally available to anonymous remote users through a

JANET connection, which enables access to the system as though from within the library itself, though this may entail some degradation in the quality of the user interface. At least 60 UK library catalogues are available in this way, and these include those of the larger research libraries; details are listed in a directory by Peter Stone. Access to the new British Library online catalogue is likely to be available from late 1993.

A recent development is the mounting of the Institute for Scientific Information Citation Indexes for science, social sciences, and the humanities, and the Excerpta Medica database, as a JANET resource. This is accessible under the name of BIDS (Bath Information and Data Services). BIDS provides unlimited access to these indexes to all members of academic institutions which have paid an annual subscription fee. For the individual institution, the service is an interesting alternative to offering equivalent access through CD-ROM databases mounted on a local area network.

An added advantage of the JANET access method is that it will be able to extend the service to include document delivery as well as citation retrieval options as these become available. The British Library intends to offer a new table of contents database, and Inside Information, through the BIDS service containing article descriptions from the 10,000 most requested titles from the Document Supply Centre, linked to online ordering facilities. There are a number of other databases available through the network, containing various types of non-bibliographic information: for example, the Biron database contains information on datasets held by the ESRC (Economic and Social Science Research Council), and the HENSA archive gives access to public domain software.

Several fee-charging databases and electronic resources are also available, including access to the British Library *BLAISE* service and to major online bibliographic databases. This has the advantage to the user of saving on telecommunications' costs, since use of JANET as a carrier is effectively free to the end-user. The JANET User Group for Libraries (JUGL), one of the user groups which advises the network management, has overseen the development of a library-oriented guide describing many of the resources available over JANET through *Project Jupiter*, and has contributed to network training initiatives in this area. This has helped with the navigational problems in pointing to the existence and location of useful resources, and has helped to firmly establish JANET as a medium for reference of information resources within the library environment.

The JUGL user group has also initiated the *Plan for Library Action on NETworks* (PLANET) in an attempt to encourage library and information workers to become actively involved in developing and exploiting network resources and services through an awareness of and involvement in existing projects, and the encouragement of discussion and debate on future developments.

Interactive gateways allow a user to connect to a service on a different network, and to perform searches within these services interactively, by translating from JANET into the protocols running on the other network. This is particularly attractive for JANET users who wish to make use of Internet facilities, and are otherwise unable to do so. These gateways act as guest intermediate systems, which are connected to both networks; to search an Internet database, for example, a JANET user first contacts the appropriate gateway, and from there initiates a second link, which is controlled by the gateway computer, to the destination—the information exchange between the two end systems is conducted in two stages.

JANET gateways enable the use of remote log-in and file transfer with Internet systems in this way; the user is able to nominate the desired destination through the addressing scheme. The *NISS gateway* (National Information on Software and Services) is a more specific service offering links to specified information systems, both commercial and non-commercial, through hierarchical menus listed on sets of screens. Some of these services are freely available to all users, and others require previous registration and the use of a password and identifier.

The advantage of the gateway is that several popular services are brought together under one umbrella, without the need for special knowledge on methods of connection on the part of the user; some of the physical connections offered include services outside the UK which may not otherwise be available from all user sites. In the section on bibliographic services, the NISS gateway offers links to Datastar and Dialog bibliographic hosts; to the German DIMDI biosciences online and STN scientific and technical network; to the OCLC EPIC reference databases in Ohio, and the US Research Libraries Group Citadel system; to the ARTTel system of the British Library Document Supply Centre; to the MELVYL catalogue of the University of California, and the CARL database in Colorado; and to the Data Distribution Network of the European Space Agency. Neither research nor educational work is limited by national boundaries, and it is desirable to extend the communicative abilities networking offers beyond these geographic limits. JANET achieves this by the use of public gateways to other major networks, falling into two principal groups.

The *public data network* (known as the PSS) gateway enables stations on JANET to link into systems that are only available in this way, without the need to reconnect to a different network at the user end. This also offers a route into JANET for those who do not belong to one of the member institutions. As a part of a public telecommunications system, use of the gateway itself is charged for by British Telecom, although activity within JANET itself is not.

Relay gateways provide links with other major research networks in Europe and the United States. These do not offer a direct connection to services on the other network, but electronic messages and files can be passed from one network to the other, performing any necessary translation in the process, in a store and forward manner. This allows JANET users to exchange information with network users worldwide, and with relative case, although the addressing system is more complicated.

— The *Earn-Relay* gateway exchanges information packets with BITNET and EARN network services;
— The *Nsfnet-Relay* gateway performs a similar service for destinations on the Internet;
— *MHS-Relay* is an X.400 gateway giving access to Telecom Gold 400 for electronic mail messages in this format, and to an IBM Mail exchange with employees and customers of that corporation worldwide;

As pressure for direct access to the wealth of information and personal contacts on the Internet increased during the early 1990s, the JANET management team initiated a project to develop a JANET IP Service (JIPS), to enable JANET sites to behave as though directly connected to the Internet -network. This enabled sites to run the TCP/ IP protocols necessary for interworking with Internet sites, through a so-called IP 'tunnel', by carrying IP traffic as X.25 data—in effect running one protocol on top of the other, disguised as data packets. Implicitly, this development has recognized the importance of TCP/IP at the present time for the research community, and recognizes the value in allowing the existence of alternative standard protocols on the same network for the benefit of global connectivity.

A JANET site must have the appropriate hardware— an IP router—and software installed to use the service, and must develop an Internet compatible naming scheme for directly linked users. Parallel running of protocols requires a user to launch the correct application for a particular need (and to know which this is), and so is not a transparent process to the user. For those with networking skills, however, the JIPS development opens out an enormously wider range of information and communication possibilities, and is a step within the UK towards greater international connectivity, as well as opening up access to more advanced information applications. A UK researcher can, from an office desktop, log in to the catalogue of the New York Public Library, retrieve copies of archived files from Sydney, and communicate with colleagues on the Pacific Rim, unconstrained by distance and time differences. By September 1992, over half the data traffic carried on JANET was IP traffic.

The introduction of multiple protocols does, however, introduce network management problems in relation to the end-user: everything must have two names—a JANET name and an Internet name and there are often duplicate applications for the same task, confusing the user. There are pressures within JANET to move to the use of the TCP/IP protocol suite as the underlying native mode, and it seems likely that this will occur in some parts, so that X.25 traffic is carried on top of an IP-based network, reversing the current position.

From 1993, parts of the JANET network are to be upgraded to a substantially higher speed and capacity in the SuperJANET project, to be installed by British Telecom. This project is intended to make use of advanced communications technology to provide research and education within the UK with a high speed, high performance broadband network capable of transmitting voice, data, and image together. This new network, using optical fibre, will upgrade the bandwidth available to a maximum of 140 Mbps using ATM packet switching techniques.

Working at a speed some 1000 times faster than the present JANET, this will enable the development of a range of applications and information services previously unfeasible. It is anticipated that early applications will include the use of distance learning courseware, electronic publishing, library document distribution, medical imaging, and multimedia information services.

The SuperJANET project intends to link into the existing Livenet network of the University of London, at present used to transmit live video images of medical surgery in operating theatres to classes of medical students, for teaching purposes. There is a two-way audio link with the theatre, so that students may ask questions from the surgeon, and an image base of still and video images as background material is also available. SuperJANET will extend this link to a wider audience of medical schools. Within the library field, a number of projects have been proposed, indicating a recognition of the importance of networking in this area.

1 Seven university libraries will co-operate in a resource sharing project on the delivery of journal articles in electronic form, with a much faster response than normal interlibrary loan methods.
2 A need has been identified amongst researchers who require access to high quality images of rare documentary material and archives, such as illustrated manuscripts, at different locations; a project will use ancient Persian manuscripts held at Manchester University to investigate this area.
3 A number of publishers are co-operating in the creation of a body of material to build an experimental electronic journal, covering a broad subject area which tests a variety of information needs—colour maps, graphics, half-tones, and scientific notation, for example.

The history of the JANET network shows that in a few short years, the available facilities are being used for applications never imagined in the early stages, and it is likely that a similar turn of events will overtake these initial ideas for SuperJANET. While it is certain that this round of development will result in unanticipated services as needs and understanding develop, it is less certain that the future funding and access mechanisms will remain as at present.

The Internet is an extensive academic and research-based network linking a very large number of broadly educational, research, and other related organizations across the United States and worldwide. It has a very large number of users and makes accessible a wide range of information services and resources. The Internet proper, in technical terms, refers to the relatively small number of machines which constitute nodes on the main network in the United States. Most people use the term, however, to refer to the vast range of networks, which may emanate from other countries as far apart as Canada, Israel, Germany, Mexico, New Zealand, and Australia. Within the United States, the enormous success of the Internet over the last five years is moving the network away from its beginnings as a traditional academic orientation towards a new and radical direction as a major national educational and community resource for economic development.

The Internet originally developed as a military network from projects funded by the Advanced Research Projects Agency of the US Department of Defence from 1969 onwards, and was known originally as ARPANET. In 1989, ARPANET was finally completely decommissioned from military service, and devoted to educational and research purposes. Along with other emerging state and national network elements, this began to become interlinked through a major data communications backbone established by the National Science Foundation between six new supercomputer sites in 1985. In the late 1980s, the National Science Foundation developed a high speed network backbone called *NSFnet* from the existing infrastructure, linking-some 16 main sites around the United States with data rates of 1.5 Mbits per second (known as the T1 standard).

The NSFnet forms the main data highway for the Internet, providing links via the 16 main sites to regional networks, and thence onwards to local networks and individual organizations. These main nodes may also provide gateways through to many other national and international networks, such as Earn (the European counterpart of BITNET) and CA net in Canada. The regional networks cover broad geographical areas, although with some overlap as a result of recent growth, and are run by a variety of organizations from commercial enterprises to government funded groups. Parts of the NSFnet backbone are also provided by commercial enterprise, and this mix of public and private characterizes the Internet situation.

In principle, an organization which has an Internet connection, or ‘feed’, from a local or regional node, and which has the required equipment capacity, may itself provide any number of feeds to other organizations and individuals who are looking for an Internet account, relaying the network traffic back and forth on their behalf. In this way, the network is able to grow organically and, since one system user may have more than one feed available to it, can provide multiple communication pathways between sites, and therefore a more robust and reliable service.

By November 1992 nearly 8561 foreign, regional, state, and local networks constituted the Internet, connecting more than 700,000 host computers with an estimated 4 to 5 million users in 44 countries. In anticipation of serving a more widely expanded role, the NSFnet is currently being upgraded to support a much higher data transmission speed (45 Mbits per second, known as T3 standard), providing the bandwidth which will enable the development of a range of interactive services in multimedia telecommunications This strengthened backbone is now referred to as a 'cloud' to emphasize its non-linear, interconnected, nature. The glue which binds all these diverse systems and machines together is the use of the TCP/IP communication protocols suite.

Most of the machines which constitute the Internet proper, and many of those which communicate with it, use this rather old, but widely used and constantly developing set of protocols; their widespread use in the Internet is part of its inheritance from the ARPANET, and is a major reason why TCP/IP usage is seen as the main rival to the open-systems movement.

Control of the network is distributed through various levels, with (at the moment) no overall control imposed from a central point. The main backbone NSFnet is currently managed by a not-for-profit consortium consisting of IBM, MCI (an important commercial network provider), and the Merit organization, based in Ann Arbor, Michigan.

The network uses a Domain Name System, similar to that of JANET, so that IP address numbers (long strings of digits) can be replaced by more explanatory structured domain names; the main difference from the JANET method is that these names are structured in reverse order, progressing from the narrower to the broader domain: from organization on the left to country on the right. New system users register with regional co-ordinators, who have the responsibility for mapping connection points, in order to make their existence known to other network users. This gives the Internet a much more informal flavour compared with JANET, making it difficult to say exactly how many Internet users exist, since there is no one point at which they are all counted.

There are agreed conventions on the use of the Internet to which all users are expected to conform: that it is intended for educational use and should not be used for purposes of commercial gain. This restriction is commonly quite broadly interpreted, and by no means all of the traffic between educational and other institutions is strictly educational in content; although it does preclude the offering of straightforward commercial services across the network, even though many organizations with Internet connections are commercial in nature. Some believe that this restriction will be relaxed in the near future.

The main decision-making body for the Internet is the Internet Society, a voluntary organization which aims to promote Internet technology and developments. A sub-group of this, the Internet Architecture Board, is responsible, for approving new and upgraded standards and for allocating resources such as IP addresses.

The Internet Engineering Task Force (IETF) is another voluntary body with open membership, which represents the user voice in Internet debates. The, IETF holds regular meetings—in Europe in 1993—and has a number of working groups developing specific issues. The Internet, then, depends to a large extent for its success on willing co-operation among its participants, and conformity to the spirit of network culture and etiquette. For example, an unwritten rule is that users who are established on the network through the benevolence of a host organization supplying their feed are ethically bound to consider similar requests made to them in the future. The cost of these connections is normally born as a part of the network overhead; by this means the overall cost is spread among the participants, so that network use often appears to the individual user as free, or at least providing unlimited usage for minimal cost.

There is a stronger tradition of more widespread public access to the Internet in the United States, with the availability of low cost dial-in lines as an alternative to full peer connection. Within the UK, there are a number of private organizations which provide Internet access as a commercial enterprise, for those who do not have an entry through the JANET network. This approach has been very successful, with a phenomenal growth in network activity from 1989 onwards, since then network traffic has grown, on average, by 20% per month. In the year up to October 1992, the total amount of network traffic in bytes almost doubled to about 3000 gigabytes per month. This phenomenal growth has catalysed the discussion on the future development of the Internet, and its methods of funding and management may change dramatically in the future.

Internet provides electronic mail services based on the Simple Mail Transfer Protocol (SMTP) mail convention. Electronic mail has also given rise to a variety of alternative uses beyond the simple exchange of messages between individuals to a greater extent than in the UK. A large number of news groups and discussion groups have grown up over the Internet based on the use of mail servers; within BITNET, which has become identical to the

Internet for many practical purposes, these are known by the name of the Listserv software used on IBM machines.

These function as automatic distribution lists for subscribers interested in an individual topic, which rebroadcast messages from individuals to the entire group. The Usenet newsgroups are also distributed over the Internet, and are available in a similar way. Listserver-type systems are present on numerous sites, rather than being centralized as in the UK Mailbase project, and there are hundreds of interest groups, ranging from technical discussions on aspects of computer science, through topics based on subject disciplines and educational and research issues, to social and recreational groups.

An example of one of the more active, for library and information workers, is PACS-L, dedicated to the discussion of public access computer systems; this has about 2000 members in several countries, and generates some 10-15 messages per day. Within the LIS field, there are groups centred around specific information service types, library systems suppliers, and broader social issues. There are estimated to be over 3500 conferences active, in addition to the Usenet newsgroups. Electronic mail is also used for the distribution of several electronic journals which are available over the Internet, most with free subscriptions. Once edited and compiled by human intervention, the journal may be distributed either in its entirety as a long e-mail message, or in the more sophisticated cases only as a contents page; subscribers can then receive individual articles as an e-mail message by sending a command which is automatically processed. Over 100 electronic journals, newsletters, and digests exist in the Internet environment. In this case, the listserver can be seen to function as a primitive database system. This practice saves on the use of network bandwidth by not transmitting unwanted information in the shape of unused articles; but also allows users to gather the information in the format most appropriate to their needs. Journals may also include graphics files, and text files may be available as plain ASCII or Postscript (laser printable) versions.

For sites with a direct Internet connection, real-time interactive access with public online information systems is available to remote sites. As with JANET, a common use is for access to online library catalogues. Since most of the main educational sites in the US are connected, a wide variety of catalogues are available, some with major collections—for example the Melvyl catalogue at the University of California, and that of the New York Public Library. An estimated 337 library catalogues are available through use of the Telnet remote access protocol. The Library of Congress files are accessible over the Internet, indicating the extent to which it is being used for formal purposes by library institutions. These files offer access to 28 million catalogue records included in the LCMARC files, together with copyright files, public policy citations and federal bill status files.

The Library also allows Internet access to its technical processing and cataloguing system, and to its SCORPIO reference and retrieval system. A wide spectrum of non-library sources of information are also available. Many of the government agencies attached to Internet make information resources openly available; these include, for instance, databases mounted by NASA in space science and astronomy and the CIA World Factbook. This kind of access is likely to increase in the future as the Internet develops: recent legislation colloquially known as the GPO Access Bill, for example, to require the US Government Printing Office to provide online public access to the Federal Register and the Congressional Record, as well as other material. A range of resources are listed in *The whole Internet* by Ed Kroll.

Sites with direct Internet access have the ability to transfer files directly from one machine to another in their original format, using the FTP protocol. The practice of 'anonymous ftp' has arisen in the Internet community, by which ftp archives are established on a host machine at a particular network site, with named directories which contain archives of files available for transfer.

Although computer system users would normally have to be registered with the site system administration in order to connect to a particular machine, these hosts have allowed free access to these specified ftp archives without the necessity to register a username or password beforehand; although it is normal practice to enter the caller's network address as the password information. This practice of 'anonymous ftp' is widely used to transfer entire documents, saving on the time-consuming necessity to view and save on-screen copies, and retaining any original formatting included with the file. The file can then be downloaded to a user's own machine, and used in whatever way is appropriate; the archive is able to distribute image as well as text-based information. Many ftp sites allow remote users to deposit files in the archive for others to retrieve.

A recent OCLC report estimated that there are about 3 million such files at more than a thousand ftp sites on the network, although the 20 largest sites account for more than half of these. Much of the material is source code for public domain computer programs, although the method is increasingly being used for the distribution of text files between specific groups or more widely. A major problem is that, as the report states 'the ability for network users to share information surpasses by far their ability to discover information on the Internet'.

As with JANET, a main problem with resources on the Internet lies in the difficulty of finding out what information is available. Since there is no central administration of the network which controls all uses of it, there is no central directory of resources, although distributed network information centres are emerging.

The problem is of a much greater magnitude than that faced by JANET, however, due to the much larger number of individual connections, and consequently the larger number of services and resources potentially available. It is also the case that resources are often duplicated in different network sites, and may exist in varying and incomplete versions and formats.

A number of resource guides have been drawn up by enthusiasts, covering both library-based databases, and those offered by other organizations, as well as lists of discussion and news groups which are known to exist. These text documents may themselves be available in ftp archives, or retrievable as mail messages. None of these, however, are fully comprehensive, and all are subject to swift outdating in the climate of rapid change.

Long-term answers to the problem are the subject of debate, especially among some sections of the library community.

From 1992, the Internet is poised on the brink of a further phase of development. A concept sponsored by Vice-President Albert Gore for many years, and now enacted in the High Performance Computing Act of 1991, promises to continue national development in the existing interconnected array of networks to upgrade and expand the system, opening it out to more colleges, schools, and public libraries as a national research and educational resource. This 'information superhighway', to be known as the National Research and Education Network (NREN), was aimed at the development of a national 3 gigabit per second (a gigabit = one thousand million bits of data) network capacity by 1996. This represents an increase of approximately 2000 times over the current bandwidth, and as a rough guide enabled the network to carry 100,000 pages of text or 1000 bit-mapped photographic images each second.

As with SuperJANET, this supported a vast expansion in the number and range of services possible, and has made multimedia communications such as desktop videoconferencing feasible as an everyday reality. Another major goal for NREN was the standardization of protocols across the existing patchwork of networks, and the integration of TCP/IP with OSI protocols to enable both to be used in parallel. This is not an isolated development, but is indicative of a number of initiatives within this vein in the United States.

The Pacific Bell corporation, for example, is building a broadband network (CalREN) in California to link participating universities, research laboratories, major hospitals, and leading high technology firms in the Los Angeles and Bay Area to develop innovative applications for education, healthcare, and business. These include projects such as the use of tele-seminars for remote educational programmes, medical image processing, remote patient monitoring, and specialist diagnosis, and the development of collaborative precompetitive research with industry and commerce.

Proponents of the NREN argue that the United States is falling behind in economic competitiveness because government, business, industry, and education cannot access the information resources which it needs so as to maintain international competitiveness and economic strength. Another school of thought argues that such needs would be best met by. market forces in the shape of existing commercial networks such as MCI, funded by charging market prices for services. This argument has raised serious issues on future policies for the NREN on the issue of charging of services, which at the moment appear as though free to many end-users; and the use of NREN by commercial service providers.

Campus-Wide Information Systems

As the wider area networking scene has been undergoing substantial development, the local focus of higher educational networking has also seen significant developments. Concepts such as the 'electronic library', in which larger proportions of educational support material are delivered in electronic form to users who are more geographically dispersed, have raised the profile of local networking as a means of integrating access to the diverse range and types of information available from all sources within the campus environment. In recent years, such developments have been carried out under the name of Campus-Wide Information Systems (CWIS), which are intended to bring together this variety of local information within a single framework.

Access to this information should be simplified by an easy to use interface, which does not present barriers to the casual user, and is potentially available from all points of the local network. In many ways, this parallels some of the forms of community information systems which have emerged as civic facilities, but is more localized upon and targeted at a smaller and more easily identifiable community. The concept of a CWIS as a broad umbrella encompassing diverse information types precludes too strict a set of conditions on the structure and format of the information; in this sense, the CWIS forms an information base of both structured and less structured information, rather than a database in the precise sense of the term. A typical CWIS might include:

- staff and student directory information, including telephone, postal and e-mail addresses;
- a diary of local events and university calendars;
- course information, administrative details and timetables;
- a listing of research projects and interests;
- university rules and regulations;
- newsletters related to the local campus;
- community and travel information;
- library catalogues and library service information;

— locally produced indexes and databases;
— links to other CWIS, catalogues etc., located at different campuses.

By bringing this information together in one access structure, a CWIS can help to integrate a campus community which is split between different sites or spread over a wide area and develop a sense of community. There is also some interest in developing this approach as a method for inter-university communication at the campus level to encourage collaborative endeavours and cooperation, rather than through formal administrative structures. There are about 20 different campus wide information systems in the UK at the present time, although these can take very different forms, and a formal categorization is difficult; and there seems to be extensive interest in exploring the potential of the idea further.

The information system can be developed as an extension of a library online catalogue, and be included within the same framework as an available network resource, or be developed as a general computer-support facility. The direction from which the impetus for development comes is likely to influence the shape of the outcome, and the underlying objectives it is intended to achieve.

The CWIS may be centralized onto a single machine from which remote access is possible, or act as a gateway to other local systems, calling the information up as needed. Whether the CWIS takes a centralized or more distributed approach, the task of maintaining and updating the information is normally the responsibility of the information providers. This may be either by submitting updated files to the operators of a centralized system, or by directly maintaining the data in one of the nodes which constitute a distributed system, to which the CWIS offers gatewayed access. In either case, the CWIS is essentially acting as a framework for the presentation of data, rather than determining the content too precisely.

TEN

Community Service

Functions

There are a number of functions which can be performed by a community information service, of which the following are the main ones:

Self-help, as the name suggests, requires that users find the answers to their own problems. The community information service input is directed towards selecting appropriate materials, reprocessing information in a form that can be readily understood, packaging information, and arranging all these materials in a way that is easier for the customer to use. This kind of service is most suitable where there is insufficient staff or lack of trained staff to operate a personal enquiry service, such as a small branch or mobile library, or for deposit in an unmanned centre. It is usual for materials to be assembled, produced or packaged by a district or central office. Regarding support for other information services or for groups of professional workers, etc.—where there are adequate information services for the public, the greater need may be for an information service to support the work of other agencies and workers in the field. This could take one or several of the following forms:

— selective dissemination of information (SDI) which, in essence, involves channelling information to meet the expressed subject interests of groups or individuals;
— current awareness services;
— press cuttings service;
— provision of loan collections of reference books;
— provision of publicity and educational materials;
— provision of local information.

Information giving, which can range from simple directional information to the complex, such as eligibility for housing benefit, and may involve steering an enquirer to an agency from whom further help or advice can be obtained, without the service making contact with that source of help or advice.

Referral, on the other hand, is a more active form of steering in which contact or an appointment is made for the enquirer with an agency who can help. In some cases, it may be necessary to escort the client to the agency if it is felt that they do not have sufficient confidence to make contact, may be intimidated by officials, or are likely to have difficulties in explaining their case.

Advice is information tailored to individual need. Giving advice can be a fairly neutral activity, such as setting out options from which the enquirer must make his or her choice, or it can involve evaluation of options with recommendations as to the best course of action.

Practical help with writing letters, form filling or making telephone calls.

Advocacy is needed where a client is not capable of obtaining the information, services, benefits or justice to which he or she is entitled. A positive identification is made with the client's case, which is then argued in front of officials, tribunals or courts on the client's behalf.

Community Education, in the context of information and advice work, is a process of increasing the self-sufficiency of individuals or groups to manage their own affairs, obtain their rights, etc., or to improve their awareness and understanding of issues that affect them.

Community action involves the information service in playing an active role in precipitating change either by acting itself or by alerting other individuals and groups to campaign. Such action can arise out of an analysis of enquiries received, where it becomes apparent that there is a lack of a service or facility in the community or where a situation exists that is causing injustice or disadvantage to people.

Outreach is a means of providing information or a service to a clientele wider than that usually served by the community information service, in either geographical-or sociological terms. It covers the use of extension bureaux, mobiles and deposit collections as well as the use of media, e.g. newspaper or magazine articles, radio, television, advertising or viewdata.

Counselling requires a greater commitment of time and level of training to help individuals with problems. At one extreme it can involve simply lending a sympathetic ear to clients who, in externalizing their problems, may thus be better able to face them and arrive at a solution. On the other hand, counselling can lead on to diagnosis and analysis, with ultimate referral of clients to clinics for treatment. Counselling should not be undertaken without specialized training.

Not all these activities will be appropriate to your organisation or the community it is intended to serve. Some organisations, for example, have to be impartial, and therefore are precluded from activities which might be construed as 'political' or partisan or from taking sides in a dispute.

Although for the purposes of this book the different functions of an information service have been treated separately, in practice such clear distinctions can rarely be made. The process of information-advice-advocacy is often referred to as a 'continuum'. Boundaries are blurred and not generally recognized by users, even though, in setting up a service, limits have to be placed on its actions. Someone seeking information or help with a problem is not really interested in the philosophy of your organisation but in obtaining a satisfactory answer, or solution, irrespective of the means used to obtain it.

Consequently, some people have argued that it is wrong to set up an information service and attempt to limit its functions to just giving information or advice but not advocacy, since users' expectations may be raised but not completely satisfied. On the other hand, it can be argued that something is better than nothing. The decision must be yours but, in setting up a service, you must have a clear idea of what you expect it to achieve and, if it is necessary to set limits, how to deal satisfactorily with cases that need to be taken beyond those limits.

Structure

The next decision to be made concerns the structure of your service. Is it going to operate from one central point or from a number of outlets? From a static centre or a mobile unit? Obviously, the decision will depend to a great extent on the nature of your community. If it is a rural one, with small, scattered pockets of population, then you may want to consider either using a mobile service point or operating through a network of individuals, e.g. village contacts. If your community is a 'closed' one, such as a school or an organisation, then one static centre is likely to be more appropriate.

The structure of an information service whose community covers a large geographical area-nation-wide for instance may need to reflect a greater usage by telephone or post, rather than face-to-face contact with clientele. A general rule, already mentioned but worth repeating, is to start modestly and then expand from a secure base, rather than to over commit at the start and then have to retract.

Management

An information service is set up to meet the needs of a given community and one of the best ways to ensure that the service meets its objectives is for the community to be represented on its management. A fairly common pattern is for a management committee to be made up of representatives from statutory bodies, the service itself, and from individuals and groups in the community. No one group should dominate the management committee and its chairperson should preferably be neutral. An independent committee is particularly necessary where an information service works in the voluntary sector but receives most of its finding from a statutory body. This is in order to guarantee that undue pressure may not be exerted by the funding bodies on the way the service operates. It may be more difficult for a community information service belonging to a statutory organisation to have an independent management committee and, in those circumstances, an alternative would be to set up some form of advisory group representing users.

The responsibilities of a management committee or advisory group might include some or all of the following:

- — laying down rules for the use of the service;
- — monitoring use and recommending changes;
- — suggesting new strategies or areas of work;
- — ironing-out trouble spots;
- — taking community action where lack of provision or malpractice is identified by the information service;
- — lobbying for funds or improved facilities;
- — giving specialist advice and help on such matters as budgeting, publicity, law, etc.

What is needed to start an information service?

Establishing and operating a community information service requires certain resources:

- — finance: capital and revenue;
- — facilities: premises and equipment;
- — staff: paid/voluntary, training, health and safety, etc.
- — a system: for the collection, processing, storage, retrieval and dissemination of information.

Financial Resources

Whether you are setting up a community information service from scratch or simply adding another function to an existing service, there are bound to be financial implications. Capital expenditure will most probably be required to provide premises, furniture and equipment. Revenue will be required to pay for such considerations as:

- — staff salaries and-travelling expenses;
- — rent and rates;
- — heating, lighting, water and telephones;
- — insurances;
- — stationery;
- — publicity and promotion.

This is by no means a comprehensive list but it is essential that a budget or business plan is drawn up, setting out as accurately as possible the costs involved, so that expenses and sources of finance can be identified. A non-statutory organisation setting up an information service for the first time has a number of possible sources of funding:

Local authorities. Counties, district councils, parish councils are all empowered to make grants for this purpose. They will be the most likely source of substantial continuous funding and the most problematical. Local authority funding is usually given in the form of a grant which usually has to be applied for and agreed each year. Often the decision is made close to the end of the financial year and so an information service may find itself always existing on a financial knife-edge. Changes in the political control of a council can often spark off a reappraisal of support for services.

Moreover, there is sometimes covert pressure by local authorities to influence the policies of organisations they fund, particularly those policies of which the authority disapproves. Local authorities can also give grants for capital expenditure-premises, furniture and equipment. Some will use their lotteries fund for this kind of grant, but not for revenue funding. In recent years, the introduction of rate-capping and Standard Spending Assessments (SSAs) has severely curtailed the amount of money that local authorities are able to commit to the support of non-statutory services.

At the same time, other government legislation on, for example, compulsory competitive tendering is aimed at forcing local authorities to divest themselves of services that can, in the government's view, be better provided by the private or voluntary sectors. This may include local authority-run information and advice services. Where services are operated or taken over by voluntary organisations, there is a growing practice amongst local authorities to set down what is expected of the voluntary organisation in a contract which is signed by both parties and monitored.

If you cannot get funding from the local authority, it may be prepared to give support in kind, such as the provision of premises (short-term accommodation) or shared accommodation in one of its own buildings, e.g. the library, or community centre. The most likely local authority departments to approach are social services, leisure and amenities, libraries or education. Look out too for officers with titles like 'Community Development Officer', 'Youth and Community Officer', or 'Community Services Librarian', who may be able to advise you on what kinds of financial or other help are available. Your local Council for Voluntary Service (CVS) or Rural Community Council (RCC) may be able to advise you on the right approach to local authorities.

In the London area responsibility for grant-aiding voluntary organisations has been vested in a separate body known as the London Boroughs Grants Scheme. They can be contacted at 5th Floor, Regal House, London Road, TwickenhamTW1 3QS.

Government departments and agencies. There is no doubt that the voluntary sector's prospects of funding from central government are not rosy at this present time particularly for new projects. Funds available are contracting and many of the government departments are now channelling their grants through other agencies, such as Training and Enterprise Councils (see below). It would be advisable to consult your local Council for Voluntary Service or any of the books listed at the end of this section for the latest position regarding grants.

The following sources currently available may be of particular interest for a voluntary organisation wishing to set up an information or advice service.

The Public Library Development Incentive Scheme (PLDIS) funded by the Department of National Heritage and administered by the British Library Research & Development Department is also available to voluntary organisations.

Training and Enterprise Councils (TECs) are an important local source of funding, much of it coming from the Department for Employment. They will be particularly interested in projects that stimulate local employment, provide help for the unemployed. Or offer educational guidance. From 1992/3, TECs are administering an Ethnic Minority Grant which forms part of Section II funds but projects will almost certainly need to involve training, employment and/or enterprise. Unlike local government bids for Section II funding which can only cover revenue costs, TECs will be able to cover project costs.

Rural Development Commission is concerned to alleviate economic and social problems in rural England. Its Rural Social Partnership Fund supports voluntary organisations in partnership with other agencies which aim to tackle problems of rural disadvantage. This fund was set up for three years, starting in 1989, so you will need to contact them to see if it is continuing.

Local and national trusts are increasingly a useful source of funding for capital expenditure items but are less likely to provide revenue funding, except for innovative or experimental projects. Charities and trusts often have rules concerning what kind of project they can support. It is useful if you can find out what these are before applying since a little redrafting of your proposal could make it more acceptable. District Councils are supposed to keep a list of local charitable trusts and their terms of reference, but failing that, the local CVS or RCC will usually be able to advise on possible sources of local funding. National trusts and their areas of interest are listed in The directory of grant-making trusts, an annual publication from the Charities Aid Foundation. You should find a copy in the reference section of most major public libraries. A similar type of trust administration that has appeared in recent years is the 'community trust'. Community trusts are independent bodies who receive money from a variety of sources-trusts, industry, local government-and give grants to voluntary projects. Applications may have to be siphoned first through the local CVS or RCC.

Directory of Social Change also publishes a magazine, Trust monitor, which contains information on new trusts, changes in policies, facts and figures, feature articles, etc.; three issues a year, subscription £ 20.00.

Commercial organisations are unlikely to provide funding for the day-to-day running costs of an information service but may be a useful source of capital funding for equipment, furniture or special projects. They may also give help in kind-donations of redundant equipment, training, etc.- or offer sponsorship deals for fund-raising events and publicity.

Directory of Social Change also publishes a quarterly magazine, *Corporate citizen*, which gives up-to-date facts and help with understanding the corporate viewpoint; this costs £30.00 a year for voluntary organisations, £55.00 for others.

European Community, in certain circumstances makes grants to local authorities and to voluntary organisations for schemes which seek to alleviate deprivation or inequality. The European Social Fund is open to organisations that arrange training aimed at providing the skills needed for specific kinds of jobs. Applicants need to have agreement from central or local government to fund at least half (35% in Northern Ireland) of the costs of the training. The allocation of ESF grants has been decentralized and, as far as the voluntary sector is concerned, applications are handled and placed in priority order by the National Council for Voluntary Organisations. Final decisions are taken by the Department of the Environment. The European Community also administers many small schemes for making grants. These small schemes very rarely have more than a few thousand pounds, available for any one project and rarely make more than one grant for the same activity. Grants must also be used in the year they are made. They are thus not much use for core funding but may contribute to set-up costs. Areas covered by schemes include employment, families, older people, disabled people, human rights education, women and combating racism and xenophobia.

There are numerous guides to obtaining grants and fund raising which cover all the sources of funding listed above and many others. Some of these guides are general in extent, others cover particular areas of concern or geographical areas.

Directory of Social Change also produces a set of 8 and 12-page Fund-raising leaflets on Setting up; Planning a capital project; Drawing up a budget; Raising money locally; Earning money; Organising an appeal; Doing research; Writing an application; Social sponsorship; Developing a strategy; Organising an event; Fund-raising sources.

If a statutory organisation, such as a library, wants to set up a community information service then, in addition to whatever funding can be provided from within its own budget, the following are the most likely sources of funding:

The *Public Library Development Scheme (PLDIS)* was set up in 1988 to encourage new enterprises designed to extend or improve public library services in England. Activities likely to promote cooperation between public libraries and other libraries or organisations in the public or private sector are of particular interest. The Scheme is funded by the Department of National Heritage and administered by the British Library Research & Development Department. It was originally funded for three years, but in September 1990 the Minister for the Arts set aside up to £250,000 per year for the three financial years 1992/3-1994/5. The Scheme provides up to 40% of the funding for projects, the remainder to be found by the library authority or other organisation and its partners. Applications can be made by institutions/ organisations such as the following:

- English local library authorities or their co-operative agencies, such as regional library bureaux;
- any public or private organisation (including those in the voluntary sector) able to demonstrate that the proposed development will contribute to the quality and cost-effectiveness of public library services in England or in a particular region; applicants are normally expected to involve one or more public libraries in the project.

Types of project in the community information field that have been helped through this Scheme include a feasibility study and implementation of a Training Shop in Peterborough Central Library, and Camden's Free Leaflets Information Service.

At the time of writing, it is not known whether PLDIS funding will be available for a further three years. Details may be obtained from Christine Burden, Public Library Development Incentive Scheme, The British Library Research & Development Department, 2 Sheraton Street, London WIV 413H.

The British Library, Research & Development Department has given support over the years to innovative community information services such as South Hackney School 'Right to Know' project and South Molton Community Information project. If your proposed information service intends to break new ground then you may be able to interest the British Library (BL) in a research proposal. It would however, need to cover the research element of the work. Only in exceptional circumstances has the BL provided financial help towards the setting-up and running costs of projects.

The Department has a five-year plan for priorities in research of which the broad areas of concern are the applications and implications of information technology, information as a tradable commodity, and the effects of economic and social constraints on services. Further details of current priorities are contained in their booklet, 'Priorities for Research 1989-1994', available from The British Library Research and Development Department, 2 Sheraton Street, London WIV 4BH. Although most of the possible sources of funding listed above are for either capital items or relatively short-term running costs and do not solve the question of how to keep the service going in the long term, nevertheless they do make it possible to get services off the ground, to establish them on a sound footing and to demonstrate their worth to the community. Hopefully, this will then attract more secure and lasting financial support.

Many services provided by local authorities are encouraged to pursue income generation as a means of supplementing or making up a shortfall in budgets. There does not appear to be much scope in the field of community information for income-generation activities, other than marginal activities such as the sale of mailing lists to commercial organisations (a dubious practice at best). The one exception is in the provision of support activities, such as Camden's Free Leaflet Information Service, which is attracting support from libraries and other organisations nation-wide through a subscription system. Other organisations like the National Association of Citizens' Advice Bureaux successfully market their information service but this does not represent a major proportion of their funding. A community information service that is intended to be simply an extension of an existing statutory service will still need to be costed and the effect on other parts of the service assessed. It is highly unlikely that there will be spare capacity-either finance, staff or accommodation-lying around waiting to be used.

A sufficiently well thought out and argued scheme might attract additional funding from the authority or via other sources of finance, but the most likely scenario is that it will have to be provided out of the existing budget. This then means that some kind of re-assessment of current services must take place, in order that priority be given to the proposed service, and resources diverted to support it. It hardly needs to be said that for such a process to take place and be successful, both senior management and staff in the workplace need to be convinced of the value of the new service and to be committed to its development. The climate for developing such services in public libraries in the mid-1990s was not of the best.

Government legislation has focused attention on the 'core' service-lending books and providing reference and study facilities and other services are expected to be self-financing. Pressure on local authority budgets as a result of rate-capping or the threat of it has meant a retraction of staff involved in outreach services, such as community information. Nevertheless, public libraries are still managing to come forward with innovative schemes, but the watchwords now are 'partnership' and 'income generation'.

Facilities

The choice of a suitable location for your information service will depend to a certain extent on what can be afforded and on the community you will be serving. If you already have suitable premises with spare capacity which can be used, this will help keep costs to a minimum but, if not, you will need to look elsewhere. It is always a good principle to try to locate an information service where it will be most accessible for its users. Even in a 'closed' community, such as a school or organisation, there may be parts of the building that are natural congregating points, e.g. hall, refectory, social/ recreational rooms, by the drinks machine. Obviously, the location of an information service which does not deal face-to-face with clients, but mainly via the telephone and post, is less crucial.

Research into the use of advice centres indicates that most people use them because they are familiar and very local. They also tend to use those which are near their homes in preference to those which are near their workplaces. So the ideal location for an information centre that is to serve the public is likely to be in the neighbourhood shopping area. This will not only make it accessible for the purposive user but should also encourage casual use, especially if it has windows that can be used to promote the service, Such locations, however, can rarely be afforded by one agency alone and you might need to consider sharing with other agencies.

Currently, with the recession at its height, there are vast numbers of shop premises lying empty which landlords might consider letting on a short-term basis for a reduced rent. Sometimes district councils have empty property available for short-term lets. And short terms have a habit of being extended...

If obtaining your own prominently sited premises, either alone or shared, is out of the question, then consider the possibility of being housed by another organisation that already has premises in a suitable location, such as a library, community centre, council offices, nearly-new shop or even a commercial concern. Failing that, then the next best thing is to seek for premises where there is a significant traffic flow, even though it may not be with shops-for example near a popular bus stop, a school or clinic, on the way to a shopping centre, or in premises which regularly attract a large number of people, e.g. a community centre, library, health or day centre.

An information service that is intended to help those with special needs, such as people with disabilities or elderly people, might be best located in a place where they already meet for other purposes-a day centre or workshop for the handicapped, for example.

The following are some other features to look out for in locating a community information service:

— ground-floor accommodation for ease of access, especially for elderly and disabled people;
— prominent location-not tucked away in a back room or having rear access;
— facilities for display, such as a large picture-window;
— informal atmosphere-some buildings can look too official or have

unpleasant associations which may constitute a deterrent to users;
— drop-in facility so that potential users, especially the more timid, have a legitimate reason for going into the centre for the first time to 'suss it out' without being challenged.

Space

The space requirements for your community information service will depend very much on the type of activities it will be undertaking. A self-help service makes little demand on office accommodation and does not require an enquiry desk, but will need a prominent location near the major traffic flow in the building and space for noticeboards and leaflet/display racks, A support service or one that does not deal face-to-face with clientele may only need an office in any habitable location with space for processing materials.

A purely information-giving service dealing face-to-face with the public may need a waiting/ browsing area that is welcoming and relaxing, space for leaflet display and noticeboards, an enquiry desk prominently sited near to major traffic flow, and an office/workroom for the staff. An advice/counselling service, in addition to the above, will need an interview room(s) for those clients who need privacy. Some advice services put more emphasis on informality and seek collective solutions to problems through group discussion, rather than through a one-to-one approach. Here, the centre becomes more a social facility, encouraging people to drop in for a cup of tea or coffee and a chat, out of which information needs can arise incidentally. This style of operation calls for a kitchen or at least, space where drinks can be made, and a comfortable lounge-type area. An interview room may still be necessary for those clients who prefer privacy.

Furniture and Equipment

Desks, tables and chairs. The importance of the enquiry desk being located near the main traffic flow has already been stated. It should be immediately visible on entering the building/room and clearly signed. Where the location is unfavourable, this can be overcome to some extent by good directional signs. The desk should be low and form the least possible barrier between the information seeker and the staff. It will need to have a fairly large surface to accommodate such things as card files and/or computer, telephone, frequently used quick-reference books, etc.

Drawers will be needed for pens, stationery, frequently used information files and cash where the service has items for sale. A chair will be required for the use of staff, preferably one that is adjustable if more than one person is going to use it, and with a swivel where the staff need to turn round frequently to consult shelves or a noticeboard. There should also be a comfortable chair for clients at the enquiry desk and others nearby for those waiting to use the service, with possibly a low coffee table and some magazines for browsing while they wait.

An information service that handles a high volume of very quick enquiries, like a tourist information centre, might find it more convenient to have a higher bar-type counter with enquirers and staff either standing or using barstools.

Telephone. This is likely to be the most important item of equipment for the staff and the public in your information service. Some information services conduct most of their business with clients over the telephone. In such cases, where it is desirable to link a telephone enquirer immediately with another service, this can be achieved by means of a three-way or conference-line facility. The most useful location for the telephone is on the enquiry desk, although further lines or extensions may be required in interview rooms or offices. Opinions differ as to the desirability of having telephones in interview rooms. In some instances, when seeking information or trying to sort out a problem with another organisation on behalf of a client, it can be helpful to have the client by the phone so that points can be clarified as soon as they arise.

On the other hand, there will be occasions when you will want to talk frankly to another organisation without the client overhearing; in which case it is better if the telephone is located in a private office. Ideally, both options should be available but if this is not possible, then you will have to decide on the basis of which will cause the least inconvenience.

Where possible, your information service should have its own separate telephone line so that clients can make contact direct without having to go through a switchboard or waiting for lines to become available. If that is not possible, then an extension exclusively for the information service is the next best thing. It is helpful if you can obtain a telephone number that is easily remembered by users.

There are several instances where people living in more remote areas have been provided with free telephone access to a central information service through the installation of a direct line from a public building, e.g. a branch library. This has been used as an alternative to providing an extension bureau or to compensate for the withdrawal of local offices. The main space requirements would be for a soundproof booth or office.

A number of information and advice services have made it easier for clients to contact them by operating a Freefone service or using an 0800 number. Some information services, who need to recover the costs of providing information, are using pay lines, whereby the customer pays not only the cost of the call but also a fee for accessing the information. In these cases, the information is usually provided by means of recorded tape(s). Details of free telephone access services and pay lines can be obtained from your local British Telecom office. You may consider that it is important for your information service to be able to offer help outside its normal opening hours. In this case you will need to install a telephone answering machine, so that at any hour of the day or night when the service is closed, enquirers can leave messages. Some answering machines also enable you to record a message or information which is played back to callers when they ring the service. The kinds of information which might be appropriate to record in this way are contacts for emergency help, hotel vacancies, 'what's-on' events, bus or train times, etc.

British Telecom produce a useful free booklet, Telephone helplines: guidelines for good practice, which gives practical advice on setting up and running telephone helplines. Copies are available from your local BT office.

Computer. Now that microcomputers are within the reach of most information services, they are looked on as being a necessity rather than a luxury. However, if the computer is needed solely for the organisation of information, then it may only be of value where you have a particularly large database or need to access the information through a variety of search terms, Otherwise a card index system could be cheaper, quicker to consult and more user-friendly. Generally, a computer is needed for a variety of functions, e.g. maintaining client records, keeping statistics, correspondence, accessing online databases, designing publicity, etc.

Card files. If you are not using a computer for storing your information, it is most likely that your information file will be kept on cards, the most popular sizes being 5in x 3in (127mm. x 76mm), 6in x 4in (152mm. x 102mm) or 8in x 5in (203mm x 127mm). These cards will need to be housed either in drawers in a wooden or metal cabinet or in open boxes/drawers (wooden or cardboard). Open cardboard boxes designed for card filing are fairly cheap and can be found at most office equipment suppliers. If you are really desperate, shoe boxes make a good substitute. Open wooden boxes are much more expensive but anyone with a smattering of carpentry should be able to knock up some from offcuts of wood. They don't need to be very sophisticated or load-bearing. Open boxes are easy to access and, with appropriate guiding, display the information clearly, but they also have some disadvantages: it is not convenient to stack one on top of the other; they are open to dust, drips from coffee cups, discarded chewing-gum and other unmentionable objects; and, if knocked over, the cards spill out and you have the tiresome job of re-filing them.

Card drawers overcome these problems, even if the information is slightly less accessible. Again, you can get cheap drawers made of cardboard or expensive ones in wood but the most common type is the two and four drawer metal filing cabinet. Each drawer holds approximately 1,000 cards and, if you need greater capacity now or later, units can be stacked one on top of the other. Some drawers have a rod which holds the cards in place but this can be a nuisance if you need to take cards out frequently for updating or have information on the back of cards. For even greater storage capacity you will need to consider a free-standing metal or wooden cabinet, but these can be expensive items. With more and more public libraries going over to or computer catalogues, there may well be an opportunity to pick up such a cabinet at a knock-down price.

There are on the market other more sophisticated files which usually have their own specially shaped cards. If you decide to buy one of these, you will need to take account not only of the cost of the equipment but also of the continuing cost of the specialized stationery. The systems described previously use standard sizes, so that it is possible to use scrap cards or print your own with a predetermined grid to save costs. One card unit that does have some useful features is the Rondofile which, as the name suggests, works on a turntable system, enabling the user to consult a large number of cards from one sitting position. Cards are easily removed for amendment and are available in a variety of colours, so it is possible to introduce an elementary colour-coding of information, say by area, subject or type of organisation.

A Rondofile is compact enough to be housed on the enquiry desk and thus can also be accessible to the public, if desired. One file holds 1,000 cards and, if more capacity is required, a second unit can be stacked on top. For details and up-to-date prices of Rondofiles contact M Myers & Sons Ltd, PO Box 16, Vicarage Street, Langley Green, Oldbury, Warley, West Midlands B68 SHF. The same firm has other card filing systems, including the Rondex, a rotating desktop file that holds up to 1,000 cards, available in three versions to take any of the standard card sizes referred to above.

A similar card filing system using a vertical rotary action is made by Rotadex Systems Ltd, 3/5 Fortnum Close, Kitts Green, Birmingham B33 OJL. There are 24 models in the range suitable for five different card sizes, the largest housing 9,000 cards.

Many office equipment suppliers have these systems in stock.

Filing cabinets. Filing cabinets and dividers will be needed for keeping flimsy publications such as pamphlets, leaflets, broadsheets posters, etc., as well as correspondence. Cabinets come in one, two, three or four drawer units, usually in metal, although there are some cheap cardboard ones on the market. At a pinch you could always make do with cardboard boxes and used manila envelopes!

Bookshelves. Virtually every information service will have some reference books, varying from a mere handful to many hundreds of volumes. Small numbers can be kept on the enquiry desk itself using bookends (any heavy weight or 'posh' commercial ones), a book slope, or one of those lethal coiled-spring contraptions. Slightly more could be housed on a book trolley and sited conveniently near the desk for quick reference.

Larger numbers will require some form of book shelving, even if it is only cardboard or wooden cartons laid on their side or planks suspended across house bricks. There are an infinite number of commercially made metal and wooden shelving units, either free-standing or wall mounted. Your local library should be able to help you with addresses of library furniture and equipment suppliers. Their shelving tends to be more substantial and you may find that bookshop type shelving is both adequate, attractive and cheaper. Some of the major furniture retailers, like Habitat or MFI, have cheap shelving available in a range of bright colours.

Sloping shelves may be needed for displaying magazines, leaflets or information packs face outwards. Most modem library shelving systems have facilities to incorporate sloping shelves. Murrell's and Point Eight also do sloping shelf units, either free-standing or wall mounted.

Noticeboards. Noticeboards are needed for displaying posters and information sheets for your clientele. Rarely does a community information service have sufficient noticeboard space, so be generous in your estimate where funds allow. Nothing looks worse than seeing a noticeboard with posters all higgledy-piggledy and overlaping. For cheapness, buy a large sheet of pin-board and cut to size to fill empty wall spaces, pillars, doors, etc. Boards do not need too smart a finish since most of the time they will be hidden by posters.

A noticeboard standing or hanging in a window of your centre and visible to the public can be used to feature special events, information campaigns, surgeries, or out-of-hours information. Alternatively, consider using a mobile blackboard or external wall-mounted noticeboard for this purpose. External noticeboards are best covered, using vandalproof glass or perspex. There are many commercially produced free-standing noticeboards.

Leaflet racks and dispensers. Considering the vast number of leaflets available of all, shapes and sizes, the choice of ready-made equipment for displaying them is not wide and what does exist is often expensive. All the more reason for trying to make your own. The following are two cheap suggestions:

Cut stiff manila envelopes of the size required to make pockets, glue them in rows on a large piece of card and pin this onto a wall or noticeboard. An alternative material to use for pockets is adjustable book jacketing, which has an added advantage in allowing all the leaflet to be seen through the plastic front.

Screw cup hooks into a sheet of blockboard or similar material, attach to a wall, and hang leaflets from the hooks by punching a hole in the top of 'them. A copy of the cover of each leaflet could be pasted onto the board to act as an aide memoire when stocks have all been taken.

More up-market leaflet dispensers are available from Point Eight Ltd, Shaw Road, Dudley, West Midlands DY2 8TP, including the best one I have seen for holding flimsy A4 leaflets or single sheets. Point Eight also do a range of perspex holders for multiple copies of a single leaflet.

If money is no object, then you might be better off with a custom-designed and made leaflet dispenser.

Display boards. Display boards will be needed if you intend to put on free-standing exhibitions or displays. There are innumerable makes available in a variety of price ranges, many offering special features or accessories. It would pay you to look around first and get advice from someone in the know, such as the publicity officer of a local authority, large firm or other public body. Consider such matters as portability and stability-particularly if you are going to use the boards outside the centre or outdoors-and ease of assembly.

If any display boards can be said to be universal, then it must be those of Marler Haley Exposystems; Ltd, who do several different systems that are easy to erect, versatile and attractive. They also have a variety of free-standing noticeboards and a range of accessories such as shelves, lights, leaflet dispensers, etc., which go with both display boards and noticeboards. Prices vary according to range.

Strip indexes. Strip indexes may be used as an alternative to card files or as a way of presenting information for customers to use. They are limited as to the amount of information that can be contained on a strip, although a variety of widths is available. Strips are more awkward to interfile and update but they have the advantage that it is possible to photocopy quite easily the whole file or sections on demand. As personal computers (PCs) become more universally available, it is difficult to see a future for strip indexes. Kalamazoo do a variety of holders and strips but these are a little expensive.

Other furniture and equipment. It is possible to list a great many other items of furniture and equipment that a community information service might need, depending on the nature of the service, the resources available, and the existence or otherwise of back-up secretarial help. Simply note a few items that may be considered: (a) Pigeonholes for storing leaflets; (b) bench or flat surface for assembling packs or publications; (c) typewriter; (d) photocopier (be very careful about leasing a photocopier-preferably get legal advice before signing any contract); (e) machine for recording number of users/enquiries ('clicker'); (f) fax; (g) duplicator-stencil or offset.

Stationery

Pamphlet boxes may be needed for storing booklets, leaflets and other ephemeral materials for use by staff or clientele. They are available in several sizes and varying quality, from cardboard to plastic, closed or open at the top. Lawtons Ltd, 60 Vauxhall Road, Liverpool L69 3AU do a cheap A4-size cardboard pamphlet box, called the Lawco Box, which comes flat and is then folded into shape. A more up-market A4 fold-up file box in brightly coloured plastic-covered card is produced by Esselte Dymo Ltd, Spur Road, Feltham, Middlesex TW14 OTR. The same firm also has a more rigid (and more expensive) polystyrene plastic pamphlet/ magazine file box in bright colours, sizes A4 and A5. Library stationery suppliers also have ranges of pamphlet boxes on offer. Addresses can be found from your local library headquarters.

Manila folders and envelopes to file cuttings, leaflets, information sheets in vertical files. Plastic wallets are even better, if you can afford them, as they enable the contents to be immediately visible. They are particularly useful if you want to go in for compiling information packs for your clients. Celsur Plastics Ltd, Drake Avenue, Gresham Road, Staines, Middlesex TW18 2AW have them available in clear and opaque plastic to take A4 documents, and can produce other sizes to your specifications.

Cards: 5in X 3in for subject indexes or brief information, 6in x 4in, and 8in x 5in for recording more detailed information about organisations and services. Cards can be overprinted with a grid for entering information in a standard form. Use scrap cards for cheapness or 'greenness'. Cards are available from most office stationers.

Forms, headed paper, standard letters for collecting information, issuing press releases, correspondence, etc.

Notepads, printed with your logo, opening hours and services, on which to write information for enquirers. These can act as a form of publicity for the service. Statistics sheets for gathering information about the use made of the service.

Many organisations, from national campaigning bodies to local tenants' groups, from time to time identify a need to set up an information service to meet the requirements of the following:

— their own staff or membership;
— the general public, as a whole or belonging to a defined geographical area;
— sections of the public with particular needs, such as disabled or elderly people;
— workers and public associated with a particular purpose, e.g. a redevelopment scheme or antidampness campaign.

It may sound obvious to state that, before setting up an information service, you should ask yourself, first of all, some very basic questions-like those a journalist asks when writing a story: who, what, when, where, why, and how but not necessarily in that order. Surprisingly, there are many instances where these questions have not been asked and services have been set up which were either not wanted, the wrong type of service to meet clients' needs, a duplication of already existing services, too ambitious for available staffing or resources, or lacking in commitment from the organisation's management or staff. The moral is that careful thought and preparation at the outset can save a lot of wasted effort and be the sound foundation on which to build a successful service. It is always better to start modestly and grow, than to be too ambitious and have to retract or fail completely. This book tries to follow the same principle in the way its information is presented, moving from simple, cheap solutions to the more sophisticated. But first of all, let us consider those basic questions.

Everybody has information needs but not everybody is equally capable of satisfying those needs. This is increasingly so in what has been termed our 'information society', with its complexity and rapidly accelerating rate of political, social and economic change. Information is now also seen as a commodity, with a market value that may well put it beyond the reach of the 'have-nots'. More and more, it is being made available in electronic forms which may require the possession of special skills on the part of those wishing to access it. Therefore, a good reason for setting up an information service may well be to make it easier for individuals or groups of people who do not have their own adequate information networks to gain access to information that can be of benefit to them.

Some places or groups that warrant particular attention are the following:

- inner-city areas where multiple deprivation exists -bad housing, high unemployment, inadequate schools-or where there is a greater proportion of ethnic minorities and transient population;
- housing estates which lack basic social and other amenities-meeting places, recreational space, shops-where transport is inadequate or expensive, housing is bad, repairs deficient and vandalism rife;
- rural areas where facilities-shops, post offices, doctors, clergy, local government - services, transport-are fast disappearing;
- groups with special needs—low-income families, elderly and disabled people, ethnic minorities, young people, patients, community and social workers.

It is not a good idea to set up an information service because other agencies are doing so and you do not want to be left out. There has to be substantial evidence of need for such a service in the local community or amongst your clientele to justify provision. This leads in turn to the next question.

Who Needs an Information Service?

Rarely will the need for an information service be readily apparent or clearly voiced by your clientele. More often than not, that need will have to be identified and assessed. There are several ways of finding out who needs an information service, ranging from the simple to the sophisticated. Here are a few methods you may care to consider:

Talk to the 'gate-keepers' in your community—these are the people to whom, because of the nature of their work, their knowledge, approachability or status, others go for information and advice. A typical list of 'gate-keepers' might include community workers, trade union stewards, youth leaders, head teachers, Women's Institute secretaries, or GPs. A word of caution: you should not rely on the views of these people alone, as they may well be biased or blinkered. It is important to check out the views of community leaders or 'gate-keepers' against those of your potential clientele and by using some of the other methods below.

Contact groups and attend meetings. Find out if there are any groups which represent the people you wish to serve, as they will most likely be aware of the needs of their members or of the community. Particularly useful are 'umbrella' groups like Councils for Voluntary Service (CVSs), Rural Community Councils (RCCs), councils for the disabled, trades councils, and community associations. Look out also for less formal groups such as luncheon clubs for community workers, youth leaders or advice workers. Attending meetings of groups like these will give you a feel for the kind of problems that arise and for which provision is lacking. Such meetings might also be a useful forum to discuss the need for any proposed information service.

If your 'community' is a closed one, such as the staff of a particular organisation, then it is important for you as information officer or librarian to be thoroughly conversant with the issues of concern to that organisation. One way of achieving this more quickly is to have the right to sit in on committee meetings. This will help in identifying what the organisation's information needs are likely to be and how to meet them. However, you may need to convince committee members of the desirability of this. One of the best ways is to prove your worth by providing them with some information, possibly on subjects under discussion in the committee.

Arrange an informal meeting of representatives of groups in your community and other interested or useful people, at which you can discuss your ideas for an information service and then take things further if there is a positive response. Out of this meeting might well be set up a steering committee or working group which could canvass wider opinion, either informally by calling an open meeting, or by conducting a survey of the community. Such a group might also prepare a feasibility study.

Call an open meeting at which your 'community' can hear about the proposals and make comments. If the meeting is not being called by the kind of steering committee suggested above, in some circumstances it may be helpful for it to be held under the auspices of a neutral body; for example, where your organisation, rightly or wrongly, is perceived by the community as representing one particular viewpoint. Wide publicity for the meeting is essential, using whatever means are available-word of mouth, posters, leaflets, local press, periodicals, newsletters, etc. Try to get the publicity into those areas where large numbers of your 'community' arc likely to see it.

Any publicity should give details of what the meeting is about, when and where it is being held and at what time it starts. Choose a date and venue that is convenient to most of the people you want to reach, taking care to avoid other events or activities that are likely to prove a counter-attraction. This may even apply to what is on television at the time! It is wise to give some thought to the ambience of the venue for, in some cases, an imposing room or 'official' building may inhibit attendance or discussion. If at all possible, avoid holding the meeting in a room that is not accessible to disabled people. Refreshments, if only a cup of coffee, help to stimulate a feeling of informality and encourage discussion. Put yourself in the place of someone who might attend the meeting and ask what would attract you. Make careful preparations before the meeting takes place, including a plan of how you would like the meeting to be structured, so that discussion does not waffle on aimlessly and yet is flexible enough to encourage responses from the audience.

Graphic displays, audiovisual presentations and speakers from similar projects elsewhere can add another dimension and help to get your proposals across.

Some further tips about the meeting:

— you will need someone, preferably a respected figure in your community, to open the meeting, explain what it is all about, and chair the proceedings;

— keep items as brief as possible—people soon get bored with lengthy speeches; encourage questions and contributions from the floor even to the extent of planting a few 'stooges' to break the ice;

— end the meeting before people get restive or discussion degenerates into irrelevancies;

— make it clear what the next step is going to be. If there is support for your ideas, then you may want to form a committee from those present. Alternatively, this step may still be premature, in which case arrange a further meeting of interested parties. Take note of the names and addresses of those expressing an interest so that you can keep them informed of future meetings and developments.

So, by the end of the meeting, you should have some measure of the support for and/or interest in your proposed information service from the community. It is important that the community or users of the proposed service are given an opportunity to participate in both its planning and running, so that it matches their needs and is responsive to changing circumstances. If there is no support, it would be unwise to go ahead, at least in the form that was proposed.

Conduct a survey of potential users to establish their information needs. This can be done instead of or in addition to holding an open meeting. It is essential from the start that you have a clear idea of what you want to learn from the survey and at whom it is aimed, for this will determine the kind of questions you will want to ask and the form in which you present them. It is a good principle always to try to make the wording simple and unambiguous, with alternatives spelled out where appropriate, as in the example from South Molton shown here which was used to establish the need for a multi-agency information and advice service in a small market, own in Devon. When Terence Brake set up a community information service at South Hackney School in London, he rightly surmised that many of the pupils would not understand or use the term 'information' and some might even react with hostility to an emotive word like 'need'. Consequently, he devised a range of simple questionnaires related to the youngsters' interests, such as the one on use of leisure time illustrated here which elicit the information in a way which is less formal and with which the young people can identify.

The drafting of surveys and questionnaires requires a degree of skill not possible to examine here, but if you would like to know more, the following books should help. The method of distributing your survey questionnaire needs careful attention. Ideally, it helps if you can be present when the respondent fills in the form, so that any points or questions not understood can be clarified. Often valuable information is imparted verbally through hints and nuances of conversation that would not be transmitted on paper. However, the ideal is seldom attainable and you may have to be content with mailing out your questionnaires. This obviously saves time but you should not expect a very high response rate. It helps, of course, Where you can afford it, to include a reply-paid envelope for return of the form.

Surveys can be a help in pointing you in the right direction or by adding weight to already formative ideas; however, they should not be used on their own, but only in relation to evidence obtained by other methods. It is well to bear in mind that in the act of setting a framework of questions for your survey, you are already, to some extent, determining the response.

Analyse enquiry statistics to see if there is a demand from any particular group of user which requires special attention, such as a separate information or advice service. For example, recent years have witnessed the introduction of debt counselling services by Citizens' Advice Bureaux as a response to a Vast increase in enquiries in this area. Your organisation may already be operating a more traditional reference or information service. An analysis of its enquiries over a period of time, if these are kept in sufficient detail, may indicate areas of need not being adequately met. An alternative, if your statistics are not detailed enough, is to carry out a short-term monitoring exercise.

As well as your own statistics, you might get some help from those of other information and advice agencies operating in your area of interest. By using some or all of the above methods you should have arrived at an idea of the information needs of your community. However, before making a decision on the kind of service required to meet those needs, it is essential to have as complete a picture of your community as possible. The best way to achieve this is to compile a community profile.

Community Profiling

This is a systematic process developed initially by corporate planners and adopted by other agencies, including public librarians, to enable them to plan and provide services which meet the needs of a given community. A typical profile will include:

— Statistical data about the population: size, rate of growth, age ranges, sex, marital status, religion, ethnic groups, diseases and mortality, employment, education, housing, etc. Much of this information is obtainable from census documents, supplemented and updated by figures from local authority departments, health authorities, employment centres, police, etc.

— Socio-economic information: types of economic activity, community services-statutory, voluntary or private, recreational facilities, meeting places, clubs and societies, etc.

— Local issues: redevelopment, unemployment, poor housing, transport problems, racial prejudice, etc.

— Residents' viewpoints: letters to the press, community newspapers, local radio and TV, action groups, community noticeboards, graffiti, informal discussions with those who come into contact with a wide cross-section of the community, informal contacts in pubs, shops, public meetings, etc.

The important thing to remember when compiling a community profile is not to rely on one type of information alone as this will give you a biased view. Try to get a fully rounded picture by including both statistical data and subjective assessment, official opinion and residents' reactions, majority and minority viewpoints. Obviously You must tailor your approach to the kind of community you are analysing. Community profiles were initially designed for analysing geographical communities but the same principles will apply, although the techniques and sources of information may vary, if your techniques and community is 'a 'closed' one, such as a school or organisation, or a group with special needs, like elderly or disabled people.

There are also some spin-offs from compiling a community profile. If you do decide to go ahead and set up an information service, the information gathered for the profile could form the basis of your local information file and might also be made available to your clientele in printed form. The valuable contacts made with individuals and groups will form the basis for your information network.

The following conclusions, taken from the above issue of Library trends, sum up concisely the right way to approach the question of who needs an information service:

1 Formal and informal methods - are essential both for developing a clear and useful picture of 'who is out there' and for an understanding of what they need/want.
2 There is no substitute for direct, active information-seeking ... in order to avoid acting on unsubstantiated assumptions about groups of people.
3 Just as different communities will want or need different services, different approaches will have to be used ... There is no single approach which will always work, but respect and flexibility are always essential.

Remember also that some information needs may not become apparent until your community information service is up and running and it ought, therefore, to be flexible enough to adapt to them. Determining needs is not a once-and-for-all exercise but a continuous process.

What community information services already exist?

There is one more step that needs to be taken before you can get down to putting flesh on your ideas and that is to find out if there are any other agencies in your area operating identical or similar services to the one that you are planning. You may, in fact, want to take this step at an earlier stage or you might have found out the information in the course of conducting a community profile. Either way, it is an important stage that should not be omitted, since it is in nobody's interest to duplicate existing services, especially if there are gaps in provision elsewhere.

Information about existing services in a particular neighbourhood can be collected by means of a questionnaire. This will give you only the bare outlines and will need to be expanded by personal visits.

The kind of information you will need to collect about each agency consists of the following:

- — location, premises and equipment, opening hours, how to contact the agency;
- — staffing, information resources, training;
- — the agency's activities;
- — referrals, relationship with other agencies;
- — the agency's users;
- — feedback to policy-makers;
- — publicity;
- — management and funding.

This list was taken from Who knows?: guidelines for a review of local advice and information services, and how to publicise them (National Consumer Council, 1982) which is an excellent guide to conducting such reviews and provides more details under each of the headings advice services and preparing local development plans.

If it is local agencies that you want to identify and you are not a public library, then try your local library, for a start. Other sources might include the Citizens' Advice Bureau, Council for Voluntary Service, Rural Community Council, Social Services Department, civic information bureau and individuals, e.g. community workers, social workers.

Having gone through the above process, you should now have a clearer picture of which community information services already exist in your community or area of interest and where there are any gaps in provision. You should also have some idea of the strengths and weaknesses of existing services and thus be in a position to make a decision about your own proposed service. At this point there are *several choices open to you:*

Maintain and strengthen the existing provision. It may be that your community is adequately covered by information services, in which case your efforts might best be channelled towards helping to strengthen and maintain those services. The costs of providing an information service rise continually and finding for many organisations is frequently precarious. Material support or behind the scenes lobbying to help existing services will usually be more than welcome.

Extend an existing service. The solution to meeting the information needs of your community does not necessarily have to be your direct responsibility. It may not be appropriate, anyway, if your staff do not have the training, flexibility, time or independence to provide the kind of service that is required. There may well be an existing information and advice service which is prevented from providing an adequate service or extending its work through badly located premises or lack of funds for staffing, training, information resources, etc. Your organisation could make an appreciable difference by:

- — offering free use of suitable accommodation to relocate the service in a place more accessible for the public;

— offering similar accommodation for the location of extension services;
— offering the use of mobile service points, where feasible, so that remote areas of the community can be reached;
— providing help with staffing or resources, e.g. you may have staff who can offer a particular expertise on a surgery basis or you might be prepared to supply loan collections of more expensive reference materials;
— providing access points for the public in unserved areas on behalf of the information service, e.g. by having available postal enquiry forms or providing a direct-line telephone link to the service's central office.

Cooperate. Although your analysis may show that there are gaps or shortcomings in existing provision, that should not be a signal to leap in immediately and set up a new service. At a time of financial stringency, it makes sense to explore first of all the possibilities for cooperation and thus to get a more rational and coordinated use of resources. This could take the form of one or more of the following suggestions:

Sharing information. Public libraries, for example, could explore ways of putting their extensive information resources, both materials and staff, at the service of other information and advice agencies by agreeing to keep and update files of local information or by producing a directory of local information and advice services.

Joint collection, processing and dissemination of information. Often several agencies in a community will be collecting local information. This is wasted effort, as the information needs to be collected only once, and causes annoyance to those supplying the information. A joint approach is called for, either one organisation collecting the information on behalf of the others, or a consortium of organisations sharing the task.

Shared premises. There are distinct advantages in several agencies sharing premises, in addition to that of economy. It provides better access for the public who can see a number of agencies in one place without a time-consuming and often expensive run-around. Referral is made easier, which is particularly useful in multi-problem cases. There is the possibility of not only shared accommodation but also shared information resources, shared publicity, shared case-notes, etc. Sharing can make the service available for longer hours than one agency on its own could possibly attain.

Support groups. Cooperation of the kind listed in this section can be fostered by the formation of an 'umbrella' group which brings together all the community information agencies in a particular community or area of interest. Such a group could take on activities such as information sharing, joint collection of information, cooperative approach to funding, shared publicity, training and generally creating a greater awareness and understanding of each other's work. Sometimes a number of agencies in an area, who are heavily involved in handling problems of a particular kind and do not have the resources or time to cope adequately with them, can benefit from the creation of a specialist support group to provide training, research, practical assistance (e.g. tribunal work), information materials, current awareness services, or expertise in certain areas, such as law, debt counselling, housing, planning, tax, etc. Many public libraries have been involved in setting up and running 'umbrella' groups.

Face-to-face Contact with Clientele

This is the passive function, whereby a community information service is located in a particular building or room and the clientele call in person to seek information for themselves or from the staff by enquiry at an information desk. It has the advantages that the most up-to-date information is available, provided the service has carried out its information-gathering efficiently; the client can be questioned to reveal exactly what information is required; and a member of staff familiar with using the files or databases is on hand to search out information, to make contact with services, or to suggest further possible courses of action.

There is some evidence to suggest that many people prefer information to be transmitted verbally, especially if they have low literacy skills. Some disadvantages to having one defined location are that, unless it is prominently sited (in the High Street or shopping centre), it may be difficult to bring the service to the attention of potential users without a continual publicity campaign; moreover, it can be expensive to operate in terms of the resources needed to staff the enquiry desk continually and to stay open outside normal office hours.

Dealing with people face-to-face requires a certain amount of skill but, above all, is needed a pleasant and approachable personality, a characteristic that is not so easily learnt. *The following guidelines may be helpful in dealing with clients:*

— First of all, try to make the enquirer feel at ease, so that s/he is not afraid to ask questions or discuss a problem. It may seem a trivial enquiry to you but could be causing the client much anxiety.

— Do not exhibit feelings of shock, horror, amusement, disbelief or repugnance at anything the enquirer says. It is not part of your job to pass judgment on a client whose behaviour or ideas are contrary to the norms of society.

— Listen to the whole query and don't start jumping to conclusions before the enquirer has finished. Listen 'actively' by asking appropriate questions to clarify what is specifically being asked, since most people tend to phrase their questions in general terms. What the client seems to want may not necessarily be what she needs.

- — When you think you know what is being requested, state it in simple terms so that the enquirer can confirm that you are on the right track. I once spent ages looking through modern dance manuals to answer a request for information on 'how to jive' only to find, after proudly presenting a suitable book, that the enquirer had a speech impediment and really wanted to know 'how to drive a car'!
- — Try to identify with the enquirer's problem, at the same time remaining impartial and emotionally uninvolved. Such a counsel of perfection is probably unattainable but is only a warning that sometimes people ask the impossible or want confirmation or support for a course of action when they are in the wrong. Even where a service sets out to be impartial, it cannot support clients irrespective of whether they are right or wrong. It may be necessary at times to give information or advice that is contrary to what the client expects to hear.
- — Before any information is given it should, wherever possible, be double-checked to ensure that it is correct and, if reasonably brief, written down to give to the client.
- — Where there are alternative courses of action, explain these to the client in simple terms and leave her/him to make the choice unless you are specifically offering an advice service.
- — Never leave the client with a negative response even when you cannot answer the query. Suggest where else the client may go to find an answer, even making a telephone call, if necessary, to fix an appointment or make an introduction. If the enquiry needs further investigation and is not urgent, take down the details and offer to phone or write with the answer as soon as possible. Suggest that the enquirer calls back after a certain interval if they have not had a reply.

By Telephone

The facility to transmit information by telephone can be an advantage, especially where clientele have difficulty in reaching your service because of its location, hours of opening or their own physical disabilities. Some people actually prefer the greater anonymity of the telephone. The physical location of a community information service whose work is mainly or entirely conducted on the telephone is less important. Staff can be occupied with other work- in between calls and do not have to be tied to an enquiry desk; thus it may be possible to offer longer hours. With the additional use of an answer-phone, a service could offer a 24-hour, sevendays-a-week service. A conference-phone facility enables three-way link-ups to be made and clients can be put in direct contact with a service that can help them. If in order to answer enquiries you need to move away from the telephone or for periods of time it has to be left unstaffed, then you may need to consider the convenience of a cordless phone. However, telephones are still by no means universally available in Britain and public call-boxes are frequently inconvenient because of location, continual use or lack of privacy.

Some people are deterred from using telephones because they experience difficulties in expressing themselves via this medium or the cost of making calls is a financial burden. Some services have overcome the latter difficulty by offering a Freephone or 0800 number, or by a policy of ringing back the caller as soon as contact is made. Telephones are best suited for answering enquiries that are brief, clearly stated and uncomplicated. Time does not allow for lengthy interviewing of the client over the phone, nor is it possible to pick up those non-verbal signals which often help in assessing the true need of the client. Likewise, the client cannot see you-so make your voice welcoming. Here are some *tips on answering the telephone*, based on an English Tourist Board booklet supplied to Tourist Information Centres, called How can I help you?:

- — Make sure you always have a supply of pens and paper near the telephone and that the most used sources of information are readily at hand.
- — Answer the telephone promptly. You may think that letting it ring and ring is a smart way of getting over the message that you are busy, but it only annoys the caller, who imagines that you are either drinking tea or having an interesting discussion about last night's television programme or football match.
- — Never use 'Hello!'-it is too casual and, at this stage, the caller may be uncertain whether they have got through to the right service or not. State clearly the name of your service, your name, and offer to help, e.g. 'Exville Information Service. This is Joan Smith. Can I help you? If you don't want to give your name straight away, it can be left to the closing remarks, such as 'If you need to ring back, my name is Joan Smith. Thank you for calling.' If the caller gives their name, always use it.
- — Do not smoke, eat, or drink when speaking on the phone, and definitely not all at the same time, as it impedes your speech.
- — Jot down the details of the query as it is given and confirm with the caller the essence of their enquiry before answering.
- — If you have to ask the caller to wait or need to leave the telephone

to look up information, explain what is happening, how long you might be and, when returning, thank the caller for waiting.

— If it is obvious from the start or after an initial search that the query may take some time to answer, offer to ring the caller back-not forgetting to take down their telephone number and name (if you haven't got it already). If the enquirer is ringing from a public call-box, either ask them to phone back later (try to specify a time or length of time) or take down name and address and post the reply.

— When answering an enquiry by telephone, ask the caller to read back any crucial information such as amounts of money, telephone numbers, etc., as a check that they have been taken down correctly

— At the end of the enquiry, thank the person for calling. Be positive even if the reply is negative. It's better to say 'Sorry, we haven't been able to help you on this occasion but thank you for calling and please get in touch again if you have any other queries' than 'Sorry, can't help you'-click.

A very useful adjunct to the telephone is a fax machine which will enable you to transmit and receive documents to and from other information services.

By Post

Requests for information by post do not as a rule figure prominently in the work of a community information service, but they are just as important. There is a temptation to assume that postal enquiries are less urgent than the client who is breathing down your neck or hanging on the end of a telephone, and so can be dealt with whenever there is a spare moment—in other words, never! However, every enquiry is important to the person making it and should be dealt with expeditiously. A fairly common type of postal enquiry from individuals, organisations and commercial concerns is for lists of clubs and societies on a particular topic or in a particular area. It is obviously an advantage if you have a filing system or a computer database which enables you to print out a section on request.

An alternative for frequently requested information is to have it pre-printed as a list, broadsheet or leaflet. Where a postal request is likely to take a little time in being answered, send off a pre-printed acknowledgement slip on receipt of the enquiry, saying something like: 'Thank you for your enquiry, which is receiving our attention. We will send you a reply as soon as possible.'

By Display

People have an innate tendency to browse, picking up information in a casual fashion. An information service can take advantage of this by providing well-sited, attractive and interesting displays for putting over information to clients and passers-by. There are four main aspects of display: noticeboards; thematic displays; window displays; and outside display.

Noticeboards. A great deal of information is available in poster or broadsheet format. Some of it is free, some will come unsolicited and there are a few posters for which you have to pay. There are basically *four types of posters:*

— the event poster gives details of something that is to take place at a particular time;
— the service poster gives details of a service that is available-its opening hours, location, who is eligible, how to contact, etc.-and is not, as a rule, restricted to a finite date;
— the information poster presents hard information in an encapsulated form-benefit rates, changes in legislation, health education, eligibility for grants, etc.;
— the persuasive poster seeks to sell a product, change opinions, or enlist sympathies.

The displaying of posters is an activity that can frequently arouse great passion, particularly those posters concerned with moral, political or religious issues on which strong opposing views are held. This applies not only to campaigning posters but also to ones presenting straightforward information, such as an event or hours of opening of a service, e.g. family planning clinic. You need also to be sensitive to the location of a poster. A poster showing how to fit a condom may be appropriate for display in a youth information centre but not on a board just outside the children's library. Ideally, a community information service or its parent organisation should draw up guidelines on what is or is not allowed to be displayed. Try to avoid the overly simplistic approach of displaying either everything or nothing at all. An open door policy may, in theory, sound fine but in practice it can throw up just as many problems, such as your right to refuse to display a poster which may not be in the public interest but does not necessarily contravene the law. It has been usual to rule out commercial advertising posters, although some services now accept these but make a charge in order to raise money for the service.

There is more to displaying posters or other notices than simply sticking in a few pins. The following are *a few points you might find useful:*

— You can never (or rarely) have enough noticeboard space, so in planning your service- think BIG!
— If you have several boards, categorize them: coming events, sources of help, places to visit, official notices, etc.
— Don't overcrowd your noticeboards-posters should never be overlapping and obscuring each other. If you have only a limited space, change the posters regularly to give each one a chance.

- — Arrange posters to look attractive-contrast posters of different shapes or colours to add interest. Move long-term posters around from time to time to give the impression of change.
- — Check noticeboards regularly to ensure that out-of-date posters are not left on display and long-term posters have not become faded or tatty.

Thematic displays. Displays devoted to a particular topic or theme, if well produced, can be a useful way of drawing attention to information that is available. The theme may (a) concern a problem area identified from an analysis of enquiries; (b) relate to a current change in legislation or benefit rates; (c) be aimed at improving public awareness of an issue of current concern; (d) draw attention to benefits, grants, services or rights of which your clientele may not be aware; (e) canvass support for a cause; or (f) educate your clientele on personal matters, e.g. health, finance, etc.

Use noticeboards or, better still, free-standing boards to mount displays. There is some skill required in preparing materials and mounting an effective and attractive-display. Identify someone in your organisation who has a flair for this kind of work and is willing to help.

There are a number of organisations and government departments which have displays that can be borrowed free of charge or with carriage paid one way.

Window display. If your information centre is blessed with large windows looking onto a thoroughfare, then they can be used to good effect in giving out information not only to users of your service but also to casual passers-by. Displays such as those described in the section above, placed in the window, enable the information to reach beyond your own clientele and may even attract use of the service. An external window can also be used for displaying essential information when your service is closed, such as addresses and telephone numbers of emergency services, hotel or bed-and breakfast accommodation.

A number of commercial organisations have produced microcomputer-driven information points which usually stand in a window and feature local information and advertising. The computer either scrolls pages automatically or is controlled by means of touch terminals mounted on the outside of the window. Windows or glass entrance doors may be also be used to display small ads and job vacancies.

Small ads may seem trivial but they do fill an important information need in most communities and also have an added bonus of attracting people to look in your window and, perhaps, thereby become more aware of the service and of what it can offer. Small ads can also be a useful source of income for your service but, in setting up such a service, be sensitive to other small ads services in the immediate vicinity. A job vacancies board is a much appreciated service, particularly at times of high unemployment and where your service is located some distance from the nearest Job centre. Either contact local firms, businesses, shops, etc., and ask if they would be prepared to display their job vacancies in your window—you might even make a charge and earn a bit of income for the service — or contact your local Job centre which may be prepared to let you have job vacancy cards on a regular basis. If you can afford to do so, offer to contact the firm or Job centre on behalf of any client who is interested in a particular job advertised, to see if it is still vacant.

Outside display. As well as displaying information in your own centre, consider the possibilities of display in other buildings. It may be possible to circulate such displays around a number of centres and thus reach a much wider audience. Some commercial businesses, like banks and building societies, who rarely have anything interesting to show in their windows, encourage organisations to use these for displays.

With displays located on other people's premises, you will have to accept their right to reject all or part of your display if it is objectionable to them or not produced to a high enough standard. Find out first if the organisation or business has a policy on what can or cannot be displayed and whether they have their own display equipment which you can use. For outside displays, other than simply posters pinned to a wall, you will need spare display boards that can be released for a length of time and the means of transporting them.

The display should be checked regularly to maintain stocks of any give-away leaflets, etc., to ensure the currency of the information, and to keep it in good repair. Another possible outlet for information is community noticeboards, which are usually located in prominent positions such as market squares, shopping malls, etc. They are most likely to be owned by the local council but it is not unknown for councils to assign the running of them to a local information centre. The sort of information best suited for this type of board is directional (map of village, town or neighbourhood and street index), showing places of interest and sources of help or coming events.

Selective Dissemination of Information (SDI)

This is a rather forbidding term to describe a process for transmitting information which many information services adopt quite naturally without ever having heard of the phrase. It simply means taking note of the subject interests of particular individuals or groups who use your service and supplying them with new information on their areas of interest as it becomes available. 'Taking note' can be a fairly organized and formal system involving the sending out of a questionnaire initially to discover users' interests and maintaining a card or computer file (usually arranged by subject) of those interests.

However, more often than not, it is an informal system stored in the information worker's head (I know Joe Bloggs is interested in this, I'll send him a copy'). Good information workers can even anticipate the interests of their clientele. SDI is most often used by services whose clientele are a closed group a voluntary organisation, for example-rather than ones who serve the general public. Where generalist services operate SDI, it will usually be for other information and advice workers, community workers, professionals or groups.

An SDI service can include regular scanning of new publications, journals, newspapers, press releases, annual reports, product literature and other ephemeral material or it may be rather unsystematic and random. Either way, it is a useful service, particularly for those individuals, organisations and groups who need a supply of new information in their area of interest but do not have the time, resources or expertise to search it out. It is an excellent way to win friends for your information service who in turn may be able to help you on other occasions.

Deposit Files and Collections

One solution to the problem of having one fixed location from which to give out information is to duplicate your information file or database and make it available to other centres. This would make it more accessible to a wider clientele but has the *disadvantages that* (i) you need to have a mechanism for keeping these files or databases up to date, especially if you have to rely on someone else to do the updating, and (ii) staff, where they exist, need to be trained in using the file or database so that they can help clients. Files using loose-leaf binders or strip indexes are easier to duplicate than card files.

Computer databases can easily be copied onto floppy disk(s) but depend on the recipient having compatible equipment. Those information services whose databases are stored on a mainframe computer may be able to offer other centres online access. Some community information services have made available collections of books and other material for deposit in centres which have particular needs, such as youth clubs. This is quite a costly service to provide in terms of both resources and staff time in preparing the collections and keeping them up to date. With this type-of service, it is desirable to have a monitoring mechanism so that the value of the collections can be appraised.

A good example of an information service that is duplicated and made widely available is that operated by NACAB. The material consists of specially prepared loose-leaf information sheets; leaflets and booklets.; reference books; a comprehensive index; and other supplementary material, such as sources of leaflets and information about useful reference books.

The system is updated by means of monthly packs containing amendments to the information sheets and leaflets; new information sheets and leaflets; a booklist of new publications; Update-a magazine which provides a comprehensive summary of changes in the law, etc.; and new editions of reference books. As the system is so comprehensive, it provides the factual resource for a wide range of advice and counselling services. However, it does require a substantial amount of staff time to update the files each month and occupies several feet of shelving or a whole four-drawer filing cabinet.

For many organisations the costs may also be prohibitive: £644 (voluntary organisations), £900 statutory and commercial bodies) to set up the system. For those organisations who do not have this kind of money or do not need such detailed information, NACAB also produces a Basic Pack, contained in two loose-leaf binders, costing £29.50 (voluntary), £78.50 (statutory and commercial) to set up, with the annual updating subscription costing £37.00 (voluntary) and £61.50 (statutory and commercial). Further details can be obtained from NACAB, 115-123 Pentonville Road, London NI 9LZ. If your problem is lack of space or staff time to update the hard-copy version or you need an information system that is portable, NACAB also caters for these needs with a microfiche version of the hard-copy pack. Microfiche are 6in x 4in transparent cards which contain up to 60 pages of information.

Publications

These are one of the best means of making your information more accessible to a wider range of people than you could ever hope to reach through a static service point. Well-produced and attractive publications enhance the image of an information service and help to attract custom. The commitment of information to print inevitably attracts corrections, additions, etc., which enable your service to maintain an even more accurate and up-to-date database. Producing publications of reasonable quality can be costly in terms of both time and money, depending on the system used for storing your information base and access to in-house printing facilities.

If your database is stored on computer then there are a number of options, as follows:

If you have a word-processing or desktop publishing package on the same computer or access to these within your organisation, it may be possible to transfer the whole or sections of the database file into these packages and manipulate them to produce the required format and layout for publication; at which point there are again *several choices:*

- — print out the publication on demand or in small quantities (assuming your computer has a printer) not practical for lengthy documents;
- — print off one copy on your computer printer and photocopy best for

short documents with a limited run where quality is not critical;

— produce printed master copy for printing in-house or by a commercial printer using offset lithomaster. Copy needs to be produced on good quality laser or bubblejet printer; (Tip: You can improve the appearance of text produced on a dot matrix printer by reducing it slightly on a photocopier or, better still, asking the printer to do it photographically don't forget to allow for this by making your document format slightly larger.)

— copy the document to a floppy disk and take it to a printer who has equipment and software compatible with yours.

If you don't have access to word processing or desktop publishing, you may still be able to copy your database or sections of it to a floppy disk which can be read by a commercial printer with compatible equipment. The production costs will be higher as the printer will need to format, design or lay out the publication according to your specifications. When you are planning the computerisation of your information system, it is worth ensuring that all your software packages are compatible not only on your computer but also with systems that are to be found in most printing firms nowadays. It may be worth investing a little more to get an integrated package that offers a host of facilities—database, word processing, spreadsheet, graphics and communications service—that allow you to swap documents or bits of documents from one facility to another with ease. If you have a manual database then a lot more work is necessary to transcribe the data into a form suitable for printing.

Once you are at that stage, and if you do not have in-house printing facilities, it will pay you to shop around for printers. Get at least three estimates. You may be surprised at the variations in both price and the size of the publication (if this has not already been determined). Ask to see examples of the printer's work so that you can get an idea of quality. If you have a limited budget, look out for community printers or resource centres which offer services at virtually cost price of materials, though you may have to do much of the work yourself (with guidance).

Alternatively, explore the possibilities of: (a) joint publication with other organisations; (b) obtaining a grant from an appropriate body who may have an interest in your publication; (c) sponsorship; and (d) advertising, or a combination of these. Once information is printed, it is static and, therefore, liable to getting out of date in a short space of time. When producing a publication always consider your ability to sustain not only one issue but also any future new editions which may be necessary, to distribute copies, either as a one-off or regularly, or to handle sales. All these can be very time-consuming. It is an acceptable strategy, of course, to publish one issue in the hopes that its value will be recognized and future issues will attract funding.

The frequency with which you need to update the publication should determine the format you choose whether it should be loose-leaf or not, a leaflet or a booklet. Publications may be aimed at basically two kinds of users: those whose work or status consists wholly or in part in informing others and those whom we might term your general clientele. They may be the public, a section of the public, or members of a particular organisation or group. The major differences between the two in respect of publications will lie in the degree of detail and the way the information is presented.

How can you know what to publish?

— The kind of enquiries your community information receives will be one indicator. Frequent requests for the same type of information, e.g. addresses of playgroups, could be answered more quickly if a printed list or leaflet is available to give away.
— Identification of a gap in the coverage of a topic as a result of analysing enquiries or seeking information.
— Recent development or change which has not yet spawned any interpretative literature.
— Local issues.
— Discussions with other information and advice workers.
— Identification of specific groups who have defined information needs, e.g. elderly or disabled people, parents with children under five, school leavers, or the unemployed, which might be met with a suitable publication.

Types of Publications

These are the simplest forms of publication to compile, often arising out of repeated requests for the same kind of information. A list can be just a single side of paper or several sheets stapled together. The arrangement is usually straight alphabetical or alphabetical under subject or place headings, depending on the topic. Lists are given out in response to an enquiry, so do not require a great deal of creative design in their production.

The use of an attractive pre-printed headed paper with the logo of your service, if you have one, its address, telephone and fax numbers, can make a humble list a little more attractive and helps to publicize the service. Small numbers of lists can be run off on a photocopier or, where you have a computerized database, printed out on demand; for longer runs use a stencil duplicator or an offset litho printer. If you do not have print facilities yourself, most towns of any size nowadays usually have at least one instant-print shop where good-quality copies can be obtained at a fairly reasonable price and quickly.

Leaflets

A leaflet in one respect is a kind of extrovert list; it goes out of its way to attract attention-and thus encourage people to pick it up — through design and format. This is the information type of leaflet, the contents, of, which may be much the same as a list. The two other main, types are the educational leaflet, the objective of which is to increase understanding or awareness of a particular subject, and the campaigning leaflet, which sets out either to convert someone from one point of view to another, to gain sympathy and support for a cause, or to change a certain state of affairs for the better.

Thousands of leaflets are produced each year by all kinds of organisations. The quality of these leaflets varies from the extremely glossy and eye-catching to the cheap and nasty. If you are going to enter this arena, then it is important to give as much attention as possible to the design of your leaflets as to their contents, since they will have to compete for attention with commercially designed products. However, you do not have to be an artist or designer to be able to produce an attractive leaflet. Look around at examples of commercially produced leaflets for inspiration.

Advertising can also be a source of abundant 'free' or non-attributable art to illustrate and add impact to your own leaflets. Other sources of material might include old engravings (particularly of the Victorian era), drawings of the great masters (Leonardo da Vinci, Durer, Michelangelo, etc.) and instant art books. If you have access to a computer, the possibility is opened up to you of investing in a graphics package which enables you to draw and design your own artwork or to buy in copyright-free art files that can be manipulated for your own purposes. Lettering in a wide variety of styles which can be rubbed onto artwork is available in large and small sheets from most stationers, or your organisation may well have invested in a lettering machine. Use cartoons, diagrams, graphs and photographs to liven up the text. Keep the wording simple and short. Standard sizes of leaflets are A5 (149mm x 210mm) and 1/3A4, but experiment with unusual sizes based on standard sheets for added impact. Have enough copies printed to give a wide circulation and distribute them through as many outlets as possible, particularly those which attract large numbers of people. If the leaflet is aimed at a particular group in the community, identify specific ways of reaching them. Investigate other means of distribution that might give you an even wider coverage, such as insertion in a free community newspaper, with the council rates demand, or through organisations which regularly send out to a large mailing list.

Posters

The aim of a poster is to present a small amount of information with the maximum amount of impact so that 'he who runs may read'. Use posters to advertise events or your service, to draw attention to a few sources of information and advice, or to point out eligibility for a service or benefit. Much of the advice given above about the design of leaflets applies equally to designing posters. The design of a poster should emphasize and complement the information that you want to convey. Remember that some of the most successful posters contain no words at all. Let the image do the talking. Don't make the poster too-big as this may create problems for those whom you might wish to display it, since most organisations do not have unlimited display space. A3 (297mm x 420mm) is the most popular and useful size; A4 (210mm x 297mm) is a little small to make a great deal of impact but may be more 'acceptable if you want organisations to display the poster for a long time. Ideally, you should produce posters in a variety of sizes to suit varying locations. A one-off poster can be prepared by hand if you have someone artistically inclined, but for any quantity you will need to have access to printing facilities or a good photocopier capable of enlarging and reducing. Use coloured copier paper for effect if you are printing only in black.

Colour photocopiers are becoming more widely available but copies are considerably more expensive and may only work out cheaper than printing for very short runs. Another method of reproduction for posters is silk screen, which is a kind of stencil process. Some community arts centres make silk-screen facilities available to groups and will show you how to go about it.

Directories

A directory is one of the most useful publications that a community information service can produce. It may cover the whole of a community, be confined to the interests of a particular client group, or contain information about a certain type of organisation, e.g. social welfare agencies, local government services, etc.

A directory plays an important role in community development by informing people about the services that exist and thus encouraging them to take a more active part in community life. It helps newcomers fit into the community and is a tool for those people to whom others turn for help and advice. A directory will usually be based on an information service's database but reproduced in a format that is easy to understand and use. There are a number of *decisions to make before you embark on producing a directory*, such as the following:

Who is it aimed at? The general public (i.e. all households), sections of the public with special needs, community groups, community leaders, etc.

How many copies will be required? This will depend on whether you decide to distribute to everyone, in which case you need to obtain an estimate of your target group. If it is not feasible to give it such a wide distribution, then you must decide on how many copies it is realistically possible to produce within your budget.

You will need to work out as accurately as possible what size the directory will be so that estimates can be obtained from printers. Allow for advertising space if you are using this method to finance the directory. Newspapers work on a formula of two thirds advertising to one-third editorial (information) but they have greater overheads to, contend with. Half-and-half or even less may be an acceptable proportion, depending on your advertising rates and whether the advertising is intended to cover the whole cost.

How is it to be funded? From your own budget? By special allocation from a parent body? Joint funding with other agencies? Grants from outside bodies? Sponsorship? Advertising? Sales? Or a combination of several of these? Local newspapers often take a strong interest in producing, directories of their communities. These are generally paid for by advertising and are distributed free of charge either as a pull-out supplement or as an integral part of a particular issue of the newspaper. It may be worth pursuing this alternative at the outset as it could save your information service much time and expense, and ensure that the information is made widely available.

What format? This will depend to a certain extent on the number of copies to be produced and the target group. A directory aimed at the whole of a community and containing a lot of information will need to be produced as a book, rather like a telephone directory. An alternative is to publish it as a supplement to a newspaper, but this carries the danger that it will be discarded with the newspaper and not kept for reference.

If you intend to produce just a small number for use by other information centres and agencies who have contact with the public, then a loose-leaf format may be more appropriate, since it allows you to send out updated sheets as information changes. Consider any special needs of your target group, e.g. elderly people will require larger type, people with disabilities will require a format that can be handled easily.

How is it to be distributed? Delivered free to all households? Supplement in local newspaper? Insert in a community newspaper? On request or sold at certain specified centres? Through individuals who have contact with the public or target group, e.g. social workers, health visitors, teachers, etc.?

Content and arrangement. It is not possible here to go into detail about the scope and arrangement of directories because these will vary depending on the nature of the directory, but the following guidelines may be of help:

Scope. This should be clearly defined before you start. Decide the geographical area to be covered and the topics included. The majority of space will be taken up with entries for organisations, clubs, societies, agencies, services, etc., but some directories may include a certain amount of 'hard' information, even if this is simply a profile of the area. There is a thin line between what constitutes a directory and what a handbook.

Detail of entries. This will depend on the directory's target group. If you are producing it for use by other information agencies or community workers, then you may need to provide the full range of information collected for your database, as those using the directory will need as much information as possible in order to assess the usefulness of an organisation to their clients. For general use, it is best not to give too much information as this may result in confusion and render the directory too cumbersome for efficient use. The minimum amount of information for each entry should give the name of the organisation, persons to contact, address and/or telephone number, and a brief description of the aims of the organisation where this is not self-evident from the name. The examples listed here show a full and a short directory entry.

Arrangement. Directories can be arranged in a number of ways fairly common method is alphabetically by name of organisation under broad subject groups which in turn may be further subdivided. Where the directory covers a number of distinct and separate towns or neighbourhoods, you may want the primary arrangement to be geographical. Other possibilities include one alphabetical arrangement by name of organisation, with a subject index if required or an A-Z by subject. A contents page at the beginning may be necessary to locate sections quickly. Also include at the beginning any information that may be required, urgently, such as emergency services. Directories are used mainly to locate specific items of information quickly, so if the arrangement is not clear and well sign-posted, you may need to provide an alphabetical index. Certainly, in a directory aimed at other information services, there will be a need to index by the subject interests of agencies, groups, etc. as well as by their names. If there is more than one entry per page, then each entry should be given a running number to which the index should refer, rather than a page number.

Current Awareness Bulletins

Pleas for up-to-date information are often made by workers in many spheres of activity who do not have the time themselves to search the literature. Community workers, social workers, information and advice workers, doctors and health visitors are just a few of the groups which might benefit from a regular flow of information.

Current awareness bulletins are one of the best ways of supplying that information in an easily digestible medium. Awareness bulletin is a duplicated listing of items of information which would be of use to a certain group or groups of workers. Where a publication or periodical is available locally, it is helpful to state this at the end of each item using for brevity, if necessary, some sort of code to indicate location, with a key provided at the beginning or end of the bulletin. Each item should be given a distinctive number so that requests for further information, photocopies or loan of material can be identified clearly and simply.

The usual method is to use a combination of the year and a running number, e.g. 83/1, 83/2, etc. The most common way of arranging the entries in a current awareness bulletin is to group them under subject headings with a contents page at the beginning. Under each subject group arrange entries alphabetically. Before deciding to compile your own current awareness bulletin, find out what is already produced both locally and nationally, to avoid carrying out work that is already being done elsewhere. It may be possible, with permission, to adapt an existing bulletin to your own local use. Seek help from other agencies with compiling and finance.

The bulletin shown here was produced jointly by an Advice Workers' Group and the public library, with finance from the local CVS. If necessary, to cover costs or postage, the bulletin can be made available on payment of a subscription. You might be able to save on postage by using another organisation's distribution network, e.g. CVS, library, local government.

Handbooks often contain directory-type information but mainly the text comprises descriptive information, e.g. about services, rights of individuals or groups, advice on personal matters, courses of action, etc. Handbooks are usually aimed at groups of people who have particular needs—elderly people, people with disabilities, parents of children under five, young people—or they may cover a defined geographical area. You may have to get help with writing the text from people who have knowledge of the subject.

The arrangement of a handbook will depend to a certain extent on its subject but a simple method is to use broad headings arranged alphabetically. A contents page and index are essential for locating specific items of information quickly. Use illustrations to break up the text and add interest. Handbooks can be expensive to produce and you will probably need to get help with finance.

Packs are an increasingly popular form of making information available, especially to people who perhaps do not have a habit of reading books. They are also useful for community information services that have limited staff resources and where the service offered is mainly of the self-help type. Packs vary from a simple plastic bag to specially embossed or printed folders. Their chief attraction is that they enable a number of flimsy items on a related theme such as leaflets, information sheets, small booklets, flowcharts, sample forms and letters, booklists and address lists, etc., to be brought together in such a way that the whole is greater than the sum of the parts. Sometimes it is necessary to link items with a specially prepared text.

Packs are easily duplicated for distribution to outlying centres and, if well produced, can entice people to browse through them who otherwise would not use a more conventional book form. Some organisations even produce them to give away. Some, *disadvantages to packs* are that they need to be displayed face outwards for maximum impact; they need to be regularly checked and updated (less easy the more they are dispersed); and without careful and attractive presentation they can so easily appear to be just a rag-bag of assorted ephemera.

The following *guidelines* may be helpful in preparing packs:

— Choose the topic-this might well be (i) a subject of frequent enquiry, e.g. buying and selling a house, adoption, writing a CV; (ii) a subject on which there is very little information in book form or that which exists is not readily understandable to the lay-person, e.g. pensions, VAT registration; or (iii) a subject that meets the needs of a particular group e.g. one-parent families, those with debt problems.

— Identify suitable material and obtain in required quantities.

— Write or compile any linking material, lists of addresses and further sources of information. Get help from people who have special knowledge of the subject, if necessary.

— Number items and list them on a contents sheet.

— Design a title sheet which can be inserted into or stuck on the pack.

Manila wallets for packs can be obtained from most stationers, usually in A4 or foolscap sizes. You may also be able to find a stationer who stocks ready-made transparent plastic folders but, if not, Celsur Plastics Ltd stock these and will make them to your specifications, if necessary. If your funding will stretch to it, then you might like to go in for custom-designed plastic or manila folders.

The Media

Newspapers, local radio and, latterly, cable television can be important vehicles for getting information to the public, in addition to their usual fare of sensational stories. The media will often readily accept offers of free material from outside organisations, provided it is presented to them in the right way and would be of interest to their readers, listeners or viewers.

The great advantage in using the media is that your information will reach a much wider audience than it ever would by any other means of dissemination, even if said means is not a very permanent one. Of course, the decision as to how that information is presented to the public will be out of your control. However, we have found that in the case of regular articles, columns or scripts, provided you have clear guidelines on what is required by the particular media in terms of length, style and content and keep to them, the material will not be altered very much.

The kinds of topic in which the media might well be interested are (i) forthcoming events, either generally or of a particular kind (e.g. sports, arts, etc.); (ii) information on new services or groups; (iii) important changes in legislation affecting people's lives; (iv) reviews of new publications; (v) general rights education; (vi) general services knowledge; (vii) information for young people (jobs, education, rights, recreation, etc.). Think carefully about taking part in 'live' phone-in programmes on radio, if this is offered. You will not be able to bring your information service into the studio and you can be exposed to all kinds of trouble by giving instant advice or information over the air without being able to check its accuracy first.

ELEVEN

Community Networking

This chapter deals with the areas of information networking within the community context, ranging from voluntary community networks aiming to provide the citizen with general information for everyday life, to sophisticated systems for urban management and health care. Each of these areas has seen a growth in activity in recent years, as access to information becomes a more important factor in daily decision-making, whether at the individual or organizational level.

Recently, the ubiquitous spread of network access and microcomputers have encouraged the development of more proactive forms, under the aegis of local governments or local voluntary not-for-profit groups (particularly in the United States), which aim to provide the framework for community action which supports two-way networking, rather than the simple one-way passive consumption of display only systems.

Viewdata systems have been installed in public libraries since the early 1980s, following the emergence of the British Telecom Prestel public information system. This is now a rather outdated and unsophisticated system, and is not amenable towards tailoring for local community information. The late 1980s saw the growing use by public libraries of private viewdata systems which are developed and sometimes managed by commercial enterprises on behalf of their clients; these can be customized for individual applications and managed as a closed user group facility, offering specific information targeted at a well-defined user group. A recent survey showed that 40 public library authorities use a private viewdata system, and that this use is replacing that of Prestel (and broadcast teletext services such as Ceefax and Oracle) both in terms of numbers of systems and in overall usage by the public.

Added to the local focus in information content, there is also the advantage that some local authorities adopt private viewdata as an internal corporate information system, allowing the library to make use of any spare capacity. The Electronic Public Information Network (previously the Local Authority Videotex Association) acts as a forum for channelling ideas and advice in videotex, investigating hardware and software issues, and developing applications for community information needs. This has some 46 local authority members and a number of other public sector and private organizations.

Modern Community Information System

In the context of local government authorities rather than just public libraries, a survey by the Society of Information Technology Managers in 1992 revealed that 19 per cent of UK local authorities already have public access terminals installed, and that a further 14 per cent planned to do so in the near future. These were installed in venues as diverse as libraries, council buildings, and shopping precincts, containing information on:

- — local organizations and events;
- — details of general practitioners;
- — transport information;
- — tourist information, including accommodation listings;
- — information services aimed at local economic development, for example EC information on legislation, calls for tender, potential European business partners;
- — one-stop enquiry shops, staffed by assistants able to make use of the local authority computer network to answer information enquiries on council services and activities.

Since community information needs do not exactly match local authority boundaries, there is a growing tendency to combine together localized services to form geographically wider-ranging information systems. Doncaster Metropolitan Borough in South Yorkshire, for example, has developed a successful viewdata service known as DONFACTS, giving details of employment vacancies, and local events and activities; in 1989 the South Yorkshire Libraries Viewdata Information Service was formed to create a regional service from this and other participants in the area.

An alternative to Integrated library system module is the use of a purpose-designed module within an integrated library system such as GEAC. This has the advantage of allowing the information to be directly controlled by the library authority itself, and to integrate it with the other activities and information services offered by the public library.

Community information, for instance, may appear as one option on an Online Public Access Catalogue terminal, along with information relating to library holdings, updated and maintained in the same way. An example of this approach is the CINDEX system, the Camden Community Information Network Directory and Exchange, developed from the GEAC Local Information System. The use of an integrated system gave the possibility of direct access to the information, rather than requiring the user to traverse the 'tree' structure of menus typical of viewdata systems, with the use of various access points; and offered the functionality to manipulate and print subsets of the information base as specialized products. Once again, maintenance of this system requires co-operation with neighbouring areas, since urban information requirements are not constrained by administrative borders.

A recent development in the provision of community information has been the growth in co-operative host systems which, instead of building up a centralized bank of information for dissemination to users at terminal access points, concentrate rather on building a facilitating infrastructure which offers a framework for diverse groups to build and maintain their own information systems and services. Host systems offer general facilities for creating local information files, the exchange of information between specific small user groups, and the creation of developmental projects and workshops, as well as more traditional services such as access to external databases.

The Manchester Host System

The Manchester Host system uses an internationally available communications software system more widely known as Geo-net, which is installed in some 40 countries around the world; Manchester Host is directly linked to several of these other Hosts in Europe (including Eastern and Central Europe), Russia and the USA through the international telecommunications networks.

The Host is accessible to its subscribers through use of a phone line, PC and modem. Services include:

— *electronic mail* to other Host subscribers internationally, and to telex and facsimile services, and X.400 mail, as well as gateways to other e-mail systems, such as BT Gold, MCI Mail and Telemail;

— *data facilities* for the creation of bulletin boards, computer conferencing, and file exchange, and for the uploading of data to create new local online database services;

— *online database access* through a common searching language to external databases in both the commercial and voluntary sectors, for example to FT Profile, and the Amnesty International bibliographic database on human rights, together with an integrated billing mechanism which removes the need for each user to have individual subscription to specific database hosts;

— *network access* to other systems, such as Easynet, the Internet, GreenNet, and APC GreenNet.

The Host approach assumes that users are potential information providers as well as consumers of information, encouraging an active participation in creating new applications and making contributions to existing ones. This led the Host initiative to borrow the idea of electronic village halls from the telecottages movement, and establish these in localized areas of Manchester, offering new technology facilities and training for particular groups, organized and managed by these groups themselves: for example electronic village halls have been established by women's groups, for small local business, and for the local Bangladeshi community.

As well as supporting greater participation for the individual in the information age through a philosophy of low-cost access and distributed processing, the Host approach has been able to promote private and public sector partnership for the development of specific project initiatives for closed user groups. This involves cooperation between local government, employers and businesses, higher education institutes, trade unions, and non-governmental organizations, working on projects to extend the ISDN network and the consequent development of broadband information services, the extension of electronic data interchange (EDI) to smaller enterprises, and the development of a communications park for business. A number of *closed user groups* built around specific cooperative needs have evolved including:

— the Manchester Chinese Information Centre, supporting document exchange in Chinese script to the local ethnic community;
— the Manchester Business School, which uses the system to communicate with similar institutions on an international footing;
— the Manchester Asian Trading Network to improve trading links between local Asian business and suppliers in Bangladesh;
— the Manchester Design Community, in which smaller enterprises co-operate on bidding for larger contracts as a consortium.

Net Information System in various Nations

On a wider geographic level, links with other European cities allow for greater co-operation in trans-regional development initiatives. GeoNet systems are existing in Berlin, Athens, Moscow, Seville, Sofia, Siberia, Cyprus, Crete, Poland, and Finland among others. The European and international linkages are likely to develop more strongly in the longer term, encouraged by the recent bid by Manchester to host the Olympic Games.

The Host system brings together a diversity of community issues and related information activities focused on a locality. A range of networks have also sprung up to link together groups and organizations working in similar areas, though geographically dispersed. The 'networking' term is sometimes used rather freely in this context; some examples of these are closer to centralized database systems which offer remote access to users. Volnet, for example, is a central database host, providing a route to information for workers in the voluntary and community sectors at low-cost. The database offers online access to a bibliographic database, backed up by a photocopy-based document supply service, and to information on research projects and the details of UK parliamentary candidates.

GreenNet provides a similar service for non-governmental organizations and individuals working with ecological and environmental issues, giving access to a range of computer conferences and resources, with a common structure to the location of information files in the network. APC GreenNet links together progressive organizations with an interest in the development of information communication for social change, such as Amnesty International. A range of similar networking initiatives exist, some short- and some longer lived. These can operate at a purely national level, or work on an international scale, for example the *TUDIC* network which supports co-operation and the sharing of information between trade unions in Sweden, Denmark, and the UK.

While many of the existing community networks in the USA have evolved from local initiatives, stemming either from citizen groups or public authority organizations, recent years have seen the emergence of commercial organizations piloting or operating information networks directed at local users on a for-profit basis. These organizations make information produced by, others available whilst profiting from advertising through the network facilities. *American CitiNet* operates this service bureau arrangement, with particular customer bases in Omaha and Boston. The *regional Bell operating companies,* the US regional telecommunication suppliers; are beginning to offer gateway facilities to information providers who are then able to recover their costs from information users through the normal billing arrangements of the Bell companies. A number of local government agencies are providing information and services in this way.

As well as communicating with clients on their services, local government authorities have used information networking as a means of co-ordinating and delivering the services for which they are responsible. There is an increasing recognition for the need to integrate and consolidate the single-focus information sets that are incomplete and incompatible, or fragmented across departmental and administrative boundaries, in order to deliver services effectively and support local development.

The networking of databases and administrative systems increases access to information and forms a platform for the development of more powerful and sophisticated integrated systems. Local access to council services and information is effected by networking local branch offices and the creation of information points at which the public can find out about a range of council services, such as housing, employment, transport, and trading standards, together in one place; and can make specific enquiries about their individual progress, in areas such as payment of council tax or council house repairs.

Information access points such as these can be queried by telephone, and in some areas have taken the form of mobile services linked to a central computer system. This growth in the integration of information systems is linked to the development of new forms of service based on emerging information technology. Geographical information systems are a prime example of this, since much of the work of local authorities is spatially related.

Different information bases, for example, may contain mapping data on the location of gas, electricity, and water utility pipes, property ownership and land use, and demographic information for strategic planning. Although these may be distributed over different locations and agencies, their integration across a network offers the opportunity for more advanced services: for example, transport information management systems which link traffic flows with highway maintenance and emergency services; or demographic planning with library development and community analysis to determine information needs.

Local authorities are increasingly providing information services linked to local economic development. While the Manchester Host is one important example of a farreaching approach to this, more traditional information services include access to European databases for information on European Commission tenders and contracts, EC legislation, and searching for business partners in other EC countries. Staffordshire County Council, for example, is one of 25 areas connected to a European network for the communication of this type of information. A similar requirement encourages the integration of information access between different service delivery agencies. Recent changes in UK legislation have reinforced this trend, particularly with regard to the delivery of personal social services; this is one area in which the local authority is presented with a lead role as an enabler for their delivery.

Information System for various Social Causes

The Child Protection Register maintained by social service departments, for example, is in some cases available to accident and emergency departments of local hospitals in electronic form. The recent Care in the Community Act presupposes the exchange of information between social service departments and community health authorities.

The National Disability Information Project, formed under the auspices of the Department of Health, is a further influence developing this role for public service agencies. One of the pilot projects is the GUiDe system under development in Gloucestershire between the Gloucester and Cheltenham Health Authorities and the County Social Services Department. This deals with information for elderly and disabled people, and gives networked access to health professionals in hospitals, voluntary organizations, health centres, and social services offices. Another project is based at Gateshead Libraries and Arts Service, well known for its innovative viewdata-based home shopping experiment.

The Gateshead Disability Information project aims to make information from the council on disability available in various complementary forms—sound, text and video—of the user’s own choosing, and through a variety of delivery mechanisms.

There is a similar requirement for local government and local agencies to interwork with central government in an online networked environment, dealing with electronic information, often driven by central government directly. The Department for Education, for instance, uses electronic data interchange (EDI) to communicate with school educators on statistics and teachers' pensions matters, and similar arrangements for transport and trading standards are likely to emerge in the future.

Quasi-governmental organizations also have information exchange requirements with local areas, and as these bodies automate their information resources and archives they are developing online access arrangements to replace manual administration for remote users. Increasingly, these applications demand more sophisticated information exchange facilities and broadband network links to accommodate heavy data transfer loads and multimedia information types. A recent example is the MONARCH system developed by the Royal Commission on Historical Monuments of England, an electronic version of the National Monuments Record. The online version of this gives remote access for about 200 users to some 6 million items of information on ancient monuments, including graphic images with hypertext linkages as well as text-based searching.

The information is used by local government for planning purposes, by English Nature for Conservation Decisions, and by the National Trust for Property Management, and will in the future be linked to geographical information systems for map-based searching and retrieval.

Local authorities have traditionally been involved in the direct provision of services, including information services and information related to other activities. As information becomes more and more important as a basic commodity for daily life, many see local government as having a role in the planning and development of a local information infrastructure, aimed at facilitating the integration and sharing of information from many sources and encouraging a policy of common access and citizen participation as well as enabling local economic and community development.

The growing interest in teleworking and the establishment of service centres for telecommunications seem to call naturally on the support of local government, and this approach has found heavy investment in Japan, for example, where telecommuting has formed a major plank in the development of regional information systems; although in the UK, at the present time, there has not been the same display of interest.

As changes in society, and in the structure and nature of organizations create changes in work patterns and habits for those who work in them or use their products and services, the change towards client-focused rather than organization-focused services may offer opportunities for a new role for local government in this area.

Perhaps the first developments in this area can be seen in the recent creation of 'One stop shops' providing information, counselling, and advice support services for local businesses, established under the auspices of the Department for Trade and Industry. These are mostly based on joint arrangements between Training and Enterprise Councils, chambers of commerce and local authorities, taking different forms in different areas. Six main areas have been selected as the first pilots. These include urban areas such as Birmingham, where the Birmingham Libraries business information service is involved, to larger rural areas such as Hertfordshire, where the Hertis Information and Research consultancy and Hertfordshire Libraries are involved in the feasibility study.

Development Programmes through Networking in EC

Several of the development programmes started by the Commission for the European Community programmes are aimed at the development of networking as a supportive means for social and community development. The STAR programme, aimed at the less developed areas of the EC, has been used to develop the telecommunications infrastructure for rural areas to stimulate the use of information networking for employment and skills development.

The ACE-HI scheme, administered by the Association of Community Enterprises in the Highland and Islands of Scotland is an example of this type of activity (with additional funding from BT), resulting in a £16 million investment in ISDN networks, and their successful use for the delivering of training and the development of enterprises which conduct their business through such networks, in remote islands and isolated parts of the mainland. Other EC social programmes have made extensive use of networking and telecommunications. The COMETT programme has used networking to promote trans-national co-operation between commercial enterprises and universities for the delivery of training in technical skills.

The HORIZON programme is establishing a disabled people's network through Europe, while the New Opportunities for Women programme is funding the partnership of Athens, Manchester, and Seville in building a telematics development network for the exchange of skills, information, and experience on women's training, employment, and enterprise.

The European Community established the ROME (Reseau Observatoires Metropolitains en Europe) network in 1989 to advise the Commission on research and information development activities for urban areas. A three-year project known as the European Urban Observatory was initiated in 1992 linking ten European cities (Amsterdam, Athens, Barcelona, Birmingham, Berlin, Genoa, Glasgow, Lille, Lisbon, and Naples) in a European-wide network to support strategic decision-making and the common interpretation of information on a European-wide basis. The aim of this project is to monitor urban trends through the sharing of urban management information, the transfer of experience and comparative evaluation of policy.

Information is structured into three levels, dealing with basic indicators for policy decision-making, with quantitative comparative data, and with individual city data stored locally; these levels are linked together through the use of private networks, and the BT Tymnet wide area network employed to form a closed user group. This networking of public management information contributes to the development of pan-European urban policy issues, and gives a comparative background resource at the local level.

Information networking in the health service sector in the United Kingdom has become a vital activity in recent years, stimulated by a growing movement for the improvement of health information services, following the work of the Korner Group on information systems, and stimulated by the radical organizational changes creating an 'internal market' for purchasers and providers of health services. These tendencies have created a pressure for the harmonization and integration of all forms of information within the sector, and a demand for wider accessibility and sharing through computer networks and other applications of information technology.

The distinction between traditional bibliographic information services and other information sources has lessened with 'information' and 'intelligence' being taken to include, sometimes primarily, statistical and demographical information. Within this context, networks are used to link the office automation and service environment with information providers across many administrative levels.

Management information in this framework would include:

— literature in the form of written or audiovisual material: published books and journals, etc., as well as unpublished material in the form of correspondence, internal reports, and committee minutes;

— quantitative information, in the form of data relating to the activities and context of the organization; statistical returns, demographic analysis, performance indicators, trends analysis.

Within the health sector, an integrated delivery of information services more often requires sharing and exchange between fragmented units, now within a more competitive context. Library units, for example, may include within the same region headquarters libraries, departmental libraries, small service points in branch hospitals, postgraduate medical centre and research information points, and nursing libraries; each with varying degrees of autonomy.

With a growing amount of information available in electronic form, user requirements are likely to involve the retrieval of information from remote sources across a network, as much or more than from a local collection, and to varying forms and types of data. A medical researcher, for example, could require not only bibliographic references, and full text source material, but also images and the ability to process images, datasets of published social and epidemiological statistics, and so on. At the international level, medical information services are developing strengthened links through the use of networks.

The *HealthNet system*, for example, provides access for health professionals in developing countries to medical literature and databases and to colleagues elsewhere, through the use of low-cost packet satellite technology. HealthNet is initially targeted at Africa, and includes access to British and North American gateways which allow medical libraries and information units in these areas to 'twin' with similar organizations in Africa.

Generic facilities such as electronic mail and file transfer through computer networks open up the possibility of more extensive contact between health professionals in the context of an electronic community.

The increasing depth of special knowledge possessed by a widening spectrum of health professionals, located geographically separate from each other, and the interdisciplinary nature of health science, point towards a need for networked and integrated solutions to information needs. At the same time as information management has assumed a higher profile within the health sector, there is a growing interaction with external agencies fuelled by information interchange. Some instances of co-operative networking with local authorities have been mentioned earlier—a common interest in demographic analysis for strategic planning is another obvious example—and the same trend emerges in health service relations with other public service organizations.

The range and volume of data involved can be quite significant: a medium-sized hospital can send some 300-400 pathology and radiology reports per day to several dozen local general practitioner centres, and up to 100 discharge letters; a large general medical practice may make some 40 referrals per day to hospitals, and perhaps 7000 registrations per year to Family Health Service Authorities. Each of these activities involves the transmission of information and storage within the receiving systems, with the desirability of enabling the direct capture by the receiving application, and avoiding a need to re-key data which is already in electronic format.

The data needs to be exchanged between autonomous sites, loosely coupled in systems with a local focus. Externally, information flows will move from the National Health Service to product suppliers, service providers, government agencies, local authorities, European and international organizations, the insurance and banking sectors, private health care agencies, and other national health services. To achieve an electronic data integration across the whole health service delivery chain, and with organizations in its environment, the development and application of standard forms of information interchange is required.

The United Nations standard for electronic data interchange EDIFACT has been adopted by the Information Management Executive of the UK National Health Service as a widely used and comprehensible standard. The EDIFACT standard requires the use of structured information formats primarily intended for autonomous computer processing, and transferral across telecommunication links.

The standard provides for the definition of data elements, syntax rules, and the structure of messages built from these data elements in a way that can be adapted for various applications; EDIFACT is managed by the United Nations as a cross-sector global communications standard, and is also defined within GOSIP, the UK Government Open System Profile. The NHS Information Management Executive has initiated many projects intended to contribute towards the development of a common infrastructure through the use of information systems strategy.

A major plank within this strategy is the desire to create a unique NHS number for individual patients, enabling the creation of a population register designed for sharing across the sector, and forming a natural integration of basic administrative activities. This will underpin NHS-wide electronic networking, and together with agreed standards for information interchange, lead towards a shared language for health care throughout the sector.

A number of these projects are aimed at general practitioners, and their links with other health agencies, involving the identification of a basic information structure common to all GP practices, and the development of compatible systems. Some *examples of these development initiatives* are:

— the electronic linking of GPs to hospitals for the exchange of pathology and waiting list information via the RACAL Healthlink system and local regional networks, in the Oxford Region GP/Hospital Information Exchange Project;
— linking general practitioners with Family Health Service authorities (FHSA) to enable the direct registration and update of FHSA databases by local practices, and vice versa, and the linked updating of payments to doctors through FHSA systems;
— the update of FHSA databases to the NHS Central Register (NHSCR) as the most reliable measure of movement of the population between Health Authority areas;
— an information collection project to network GPs together to facilitate the continuous collection of morbidity data as the basis for a national agreed set of morbidity data.

More broadly, projects are developing standards which apply across the sector, for:

— the development of a common basic specification as an extension of the NHS Data Model, to cover data exchange standards, minimum datasets, clinical information for resource management; and the information demands of hospital and regional decision support systems;

— the development of an NHS Administrative Register to act as a shared register of basic administrative details meeting the requirements of District Health authorities, Family Health Service authorities, general practitioners, and hospitals;

— a Community Information Systems Project to facilitate the exchange of information to support contracting, and the development of datasets for ambulatory care in hospital and community settings, for use by community health services.

The National Health Service in the United Kingdom has seen a vast increase in the importance of information management, with information networking forming a vital part of this. It is certain that the importance of networking activities will continue to grow as a strategic element in the construction of integrated information and intelligence.

TWELVE

Service Monitoring

There are several ways of monitoring a service, each with varying degrees of complexity. The method(s) you choose will depend on the amount of staff time and resources you have available. They are the following:

(i) statistics;
(ii) feedback;
(iii) surveys;
(iv) research.

Statistics

There are very few people who would claim to like keeping statistics. At best, they are looked on as an irksome necessity, only tolerable when kept as unobtrusive as possible. Staff running a busy community information service will seldom have the time to keep detailed records of enquiries, so do try to restrict them to the absolute minimum. The barest minimum is simply to record the number of people using the centre by noting them down on a sheet of paper, using the time-honoured 'five barred gate' method or a number tally machine or 'clicker'. Both these methods rely on a human agent, however, and are therefore prone to inaccuracies— through forgetfulness, preoccupation, distraction, etc. There are now counting devices, available which are triggered off by customers treading on a pressure pad placed under a carpet or by breaking an infra-red beam. This statistic will only tell you the number of people who darkened your doorstep, however; they might only have been browsing or sheltering from the rain, so an alternative is to record manually the number of people making an enquiry.

Nonetheless, you will not know what subjects they enquired about, whether you were successful in finding the information, or how long it took. You can go some way towards collecting more detail and yet not create much more work by drawing up or pre-printing record sheets. These may be related to subject or length of time or both. Some centres, instead of using sheets, have a bank of 'clickers', but this is only practicable if you have a small number of categories, say no more than six; otherwise it becomes too complicated to remember which 'clicker' is which. All the above methods give only a broad indication of the number of enquiries and do not reveal any detail about their nature or the degree of success in answering them.

The latter aspect might be catered for by asking staff to record details of unsuccessful enquiries only, so that they can be analysed in order to detect any weaknesses in the information base that can be rectified or any sources that have been overlooked. A more detailed recording system, found mainly in advice and counselling centres, is the day book. In effect, this is a large diary in which details of enquiries are recorded, giving the name of the client, address, date and nature of the enquiry, action taken, further action being pursued, etc. A day book has the advantage of being easy to use and allows other information and advice workers to check on the details of an enquiry if the person calls back on another occasion. It is not so convenient for extracting statistics and, should you need to refer back to an enquiry, the chronological arrangement is not helpful— people can be notoriously vague about dates. A more satisfactory method is to enter the details on cards or pre-printed enquiry forms, one for each client. These can then be filed by the client's name, so that it is easy to refer back if some time has elapsed between visits.

The system of pre-printed enquiry forms is fairly common in information centres but is rarely used for all enquiries—usually only those that take some time or require further search after the enquirer has left the centre or hung up the phone. The *kinds of information that you may want to collect about each enquiry* can include the following:

- — *Client*—name, address, telephone number (if needed for follow-up), sex, Age range, ethnic origin.
- — *Enquiry— simple* precise of subject.
- — *Duration*—how long it took to answer the enquiry.
- — *Sources used or tried in answering the enquiry—useful* if a search is continued over a length of time to avoid duplication of effort.
- — *How enquiry was received—walk-in,* telephone, post, fax.
- — *Time—will* indicate spread of enquiries throughout the day.
- — *How client heard* of *service—may* help to measure success of publicity.
- — *Client's area* of *residence* (if address not taken)—will show from which areas most use comes and those not being reached.

Forms have the advantage that they can easily be stored for future reference and abstracting of statistics. With a subject heading or classification number added, they can become a subject file to be referred to when information on the same subject or statistics are required. In a very busy information centre, however, it is unlikely that staff will have time to fill in a form with this amount of detail for each enquiry, so you may need to consider a combination of (a) forms for lengthier enquiries and those needing to be followed up, and (b) a tally system under broad categories for quick enquiries.

The record sheet shown here is a kind of compromise hybrid that was used for recording enquiries in a public library community information centre shared with other agencies on a sessional basis. Each agency had its own pad of forms which were self-carbonated to produce an additional copy for the librarian to extract statistics. If you have invested in a computer, there is the temptation to use it for collecting statistics. After all, computers are eminently suited to number crunching, particularly if you need to ask detailed question of your enquiry statistics, such as 'How many women aged over sixty living in the Westward area of Exville enquired about bus passes in 1991?' Where you are simply keeping a tally of the number of enquiries in a number of broad categories, then it is debatable whether a computer will save you much time. Before you start using a computer to store statistics, make sure that you have an effective system for recording enquiries in the first place; otherwise the old computer adage of 'garbage in, garbage out' will apply.

London Advice Services Alliance is currently working on a model system for recording statistics for advice work, called STATS. There are four aspects to STATS: a Classification System, Sample Recording Forms, Guidance and Training. The project will take several years to complete and for more details and progress, contact LASA at 88-94 Wentworth Street, London El 7SA.

Feedback

Statistics can tell you how many people used your information service, how many enquiries they made, when and on what subjects, but they will not reveal whether users are satisfied with the service or why others do not use your service. To get some idea of the impact your service is making on its community, you need to have a means of getting feedback from that community. Two ways of doing this have already been referred to elsewhere in this book.

One of the functions of the management committee was identified as to 'monitor use and recommend changes'. If, as suggested, this committee includes representatives of the community, then their comments will be very important in assessing how well the service is: doing. It was suggested that a spin-off from publicizing a service by means of talks to community groups was that it also gave an opportunity for the speaker to ask for reactions to the service and discover any needs not being met. There is a practice, associated mainly with information and referral (I&R) centres in the United States, called 'follow-up', which is used to check on user satisfaction with the service or services to which they have been referred. At the time of the enquiry, details of the client's name and telephone number, the query and the service(s) to which referred are noted down. Then, at a later date, the client is rung back to find out if the information given was satisfactory or whether further help is required.

Any adverse comments about the I&R centre are noted, so that improvements can be made. Comments about other services are noted down on the back of the appropriate card for that service in the information file. In this way, clients are encouraged to feel that they are a part of the information service and can contribute to the continual evaluation of its information files. However, follow-up is a time-consuming and costly exercise which may well be beyond the means of most community information services except on an occasional and very selective basis.

You may be able to get some useful feedback by consulting other information and advice services or community leaders in your area. They may have had comments made to them about your service and be prepared to divulge them with frankness.

Surveys

A more formal and systematic way to obtain feedback is to conduct a user survey or a general survey of your community. You do not have to be an experienced researcher to carry out a survey, provided you exercise a degree of commonsense and care in drafting or asking questions. Before you conduct a survey, you must have a clear idea of what you want to find out and the use to which that information is going to be put. You may only want to gather information about the existing clientele of your service, in which case a user survey is called for.

On the other hand, it may be important to find out what impact your service is making on the community as a whole; therefore, the scope of the survey needs to be wider and to be conducted outside the information centre. There are basically two methods of conducting a survey, by questionnaire and by interview. A questionnaire survey of users is the easiest to carry out as, once the form has been drafted and printed, there is little staff involvement until the analysis stage. Questionnaires can be either left on the counter for clients to take or handed to clients when an enquiry is made. The following are points to look out for:

— Number forms so that you know how many have been issued and can calculate the percentage returned.
— Make sure there is an address for returning the form and a closing date—you do not necessarily have to keep to the closing date but it helps to encourage people to return them.
— Have an open question on the form to invite any other comments about the service—this often evokes the best replies.
— Arrange a system for return of forms that has a degree of anonymity (e.g. a posting box in the centre).

With this kind of survey, it is usual to find that responses tend to be biased towards the favourable, so you will have to allow for this when interpreting results. The other method of surveying users is by interview. This has the advantages that it is possible for the interviewer to probe for greater depth in responses, to record qualified replies to questions where the interviewee has difficulty in giving a straight yes/no response, and the interviewee can ask for clarification of any questions on which they are not sure.

A disadvantage of this method, of course, is that the more non-standardized the replies, the more difficult they are to analyse and represent statistically. You will need space in your centre to conduct interviews and staff time, although it may be possible to enlist outside help with this in the shape of volunteers or students on placement. It is unlikely that you will be able to interview every user, so an agreed proportion must be determined beforehand, say every other one or every tenth. Whatever the proportion, you will probably need two interviewers, as it is not possible for one person to interview and keep an eye on the numbers using the service in order to catch the next one. With more than one interviewer it becomes necessary to ensure that questions are posed in the same way, especially when there is a need to prompt interviewees, and that means some form of training or briefing beforehand. The remarks above about allowing for favourable bias in interpreting replies also applies here, possibly to an even greater extent, since clients are unlikely to be outwardly critical of a service when being interviewed on its premises.

A survey of the wider community needs to be drafted more carefully. Where possible, try to disguise the purpose of the survey and avoid asking leading questions ('Don't you think it's a good idea to . . .') as these may colour responses. It is unlikely with both questionnaires and interviews that you will be aiming for total coverage of the community, so decide on a representative sample. There are several possible ways of distributing questionnaires, the least satisfactory being to leave copies in public buildings (community centres, libraries, council offices, shops, clinics, etc.) for people to take. This is somewhat haphazard and may not result in a representative sample.

A better alternative is to deliver by hand to every tenth house (or whatever proportion you have decided), using your own staff or volunteers. There may, however, be a simpler way of distributing by, for example, inserting copies in the appropriate proportion of community newspapers (where they exist) that are delivered to each household. You might investigate the possibility of printing the survey as part of the community newspaper. Where you can afford it, include a post-paid, addressed envelope for reply; otherwise, arrange for them to be left at suitable locations in the community.

Interviewing of the wider community to determine use is unlikely to be a practical proposition for an information centre to conduct on its own. It is more likely to be used where an outside organization is conducting a research project on your service.

Research

Surveys are one form of research, of which there are a number of other types—operational research, action research, experimental research — most of which require a certain amount of expertise and, therefore, will need to be conducted by an outside agency, for example, The British Library, a library school, or a university department. Interest in your information service as a subject of research will depend on whether it has innovative features or whether the aspect to be investigated is symptomatic of other information centres and thus the research findings would have a wide applicability.

Continuous evaluation using one or several of the techniques outlined above is essential in order to keep the service relevant to its community. By indicating strengths and weaknesses, areas of unmet need and duplication of services, it can point the way to future development of the information service. An information service needs to be a growing organism because communities do not stand still but are constantly changing, some more rapidly than others. At the beginning of this book, we stressed the need to involve the community in the planning of your information service and that is my message at the end.

A community information service exists to serve its community and, therefore, the community needs to be involved at all stages. There is no guarantee, in these hard times, that your service will survive, but it stands a much better chance of succeeding and developing if it can be shown that it is respected, valued and used by the community and is performing a worthwhile function.

PART TWO

Application in Life and Business

THIRTEEN

Electronic Mail

The basic information services which are offered by research networks fall into three main categories: information exchange, file and document distribution and information services. Information exchange services allow individuals and groups to communicate with each other across time zones and geographical space much more quickly than postal delivery would allow, and has proved the most popular and heavily used end-user application on networks by linking together people who would otherwise find it hard to meet face to face.

Electronic mail facilities enable the exchange of messages and ideas on a one-to-one basis, whether more formally as in the sharing of notes and drafts of work-in-progress, or informally for general social contact. Discussion groups which use electronic mail in conjunction with mailing lists support the exchange of information and messages between all members of a group at the same time. This is popular for ongoing debates and the exchange of news items in discipline-specific areas, allowing members to share ideas and experiences in an informal setting, to 'meet' new contacts, and to follow and influence the development of new ideas in their field. Bulletin boards are used in discipline-specific settings, but require a positive action from the user in connecting to the service, and reading or downloading items of interest, rather than depending on the blanket distribution of material through mailing lists. Computer conferencing and interactive video conferencing are still quite rare activities on such networks, but will become increasingly common as the technical infrastructure becomes capable of real-time interactive and multimedia services, and are likely to have substantial implications for distance learning education.

File and document distribution support the distribution of items such as text documents, image files, and software programs on an individual search and retrieve basis. File transfer is permitted by the use of standard protocols so that data files can quickly be moved from one machine to another under the control of the end-user. As more electronic material is becoming available within the public domain, certain computer network sites offer an 'archive' service, in which data files can be stored on open access for retrieval by all network users.

Mail servers are commonly used as primitive database systems, returning files in response to commands received within electronic mail messages, as an automatic process with no manual intervention. This is often the method used for distribution of the composite articles within an electronic journal, for example, the user will request a particular file by name by placing the appropriate commands and filename within a mail message addressed to the mail server.

Information services are available through remote log-in protocols, which allow the user to link to a remote system and use it as though there were a direct connection between the two, within the constraints set by the host system. The service may be on a different network, as long as both end systems (user and host) are able to run the same protocols, and the networks and gateway used will support this. Online public access catalogues available for remote searching from outside the library or campus site are a common use for this facility, as long as the OPAC allows for a remote connection. There is, of course, no common standard interface for the various online catalogues and so the user must know, or be able to learn, the functions and abilities of each different system. Non-bibliographic databases are increasingly becoming available across research networks, particularly those which contain information within the public domain, or government-sponsored public information.

The full range of academic and research networks worldwide includes both large- and small-scale networks at differing stages in their development, and is too numerous to list here. Most research networks are organised on a national basis, or as a cooperative endeavour between smaller countries, as for example in the Gulf states.

RARE

The European academic and research networks are brought together in RARE (Reseaux Associes pour la Recherche Europeenne), the association for European research networking organisations and their users. This aims to overcome national boundaries in research networking by creating a harmonised computer communications infrastructure for the whole European research community. Based in Amsterdam, RARE has over 40 national and international research networking organisations as members, which are listed and described in the RARE 'State of the union' report. The RARE association also sponsors a number of working groups which initiate and co-ordinate advances in network technology and services; these include, for example, work on user support issues, information retrieval across the networks, and developing multimedia services.

The academic network for the United Kingdom is called JANET (joint Academic Network). Similar networks exist for nearly all the other European countries, including Eastern Europe, although the range of facilities which each is able to offer can vary widely. In some former Communist countries, such as Romania, access is very difficult and services are unreliable. Others are far more highly developed, with extensive experience and expertise, and are moving towards a strategic position of greater independence from government support and self-sufficiency.

The Dutch academic network, SURFnet, is an example within this category. SURFnet networks together the institutes of higher education, and scientific and industrial research organisations in the Netherlands, and has maintained a close involvement with the Dutch bibliographic utilities, taking a proactive approach to the development of information services. In 1989 SURFnet became a private company, with a 49 per cent share held by the Dutch telecommunications company PTT Telecom, and was required to become self-supporting in the near future. This has led it into substantial development in services and user support, working on-site with institutions and with individual discipline-based groups of users between institutions, to encourage and develop new users and uses for the network. There are several other major *influences and initiatives which affect the shape of European academic networking* at a transnational level, both outside and within the EC. These include:

Organisations which are co-operative efforts between countries grouped together by geography. NORDUnet is a collaborative effort between the Scandinavian group of national research networks: Denmark, Finland, Norway, Sweden, and Iceland aiming to enable an efficient connection to the rest of Europe. A similar arrangement has been proposed for the countries in the area of the former Yugoslavia.

EUnet is a large subscription-funded network, which serves European users from Scandinavia and the CIS (former Soviet Union) to North Africa. Founded by the European Forum for Open Systems, EUnet connects 6000 individual sites and networks together as Unix system users.

International treaty organisations which are established as European-wide enterprises, with funding from a number of countries. The most influential of these is CERN, the European Laboratory for Particle Physics, which has constructed HEPnet as a dedicated network serving research centres in high energy physics.

EBONE is the European backbone network, which provides a neutral interconnection between the backbone research networks which form part of the global TCP/IP based Internet. The backbone provides a high capacity 'highway' between the major first-level nodes in Europe (Stockholm, Amsterdam, Geneva, Paris, and London) to which other sites connect.

EuropaNet has developed from an EC-funded project conducted by COSINE (Co-operation for Open Systems Interconnection Networking), intended to develop a coordinated approach to networking the European research community, and provide an international infrastructure accessible by academic organisations, research institutes, and organisations co-operating with EC research and development. A pilot International X.25 Infrastructure (m) network commenced in 1990, forming a backbone which linked research networks in many countries, both inside and outside the EC itself. This has now been replaced by EuropaNet as a high capacity multi-protocol backbone which will handle both OSI (open system) and TCP/IP Internet-style protocols.

The EARN (European Academic Research Networks) association provides networking services for researchers and academics in Europe, the Middle East and Africa, as a store and forward network based on IBM computers. EARN is closely integrated with the us-based BITNET network, and is in fact almost indistinguishable from this in use, effectively acting as its European arm. In some countries EARN forms the major research network, while in others there is a gateway connection between EARN and the national network.

Global Internet Network

The Internet originally developed as a military network from projects funded by the Advanced Research Projects Agency of the US Department of Defence from 1969 onwards, and was known originally as ARPANET. In 1989, ARPANET was finally completely decommissioned from military service, and devoted to educational and research purposes.

Along with other emerging state and national network elements, this began to become interlinked through a major data communications backbone established by the National Science Foundation between six new supercomputer sites in 1985. In the late 1980s, the National Science Foundation developed a high speed network backbone called NSFnet from 'the existing infrastructure, linking-some 16 main sites around the United States with data rates of 1.5 Mbits per second (known as the TI standard). The NSFnet forms the main data highway for the Internet, providing links via the 16 main sites to regional networks, and thence onwards to local networks and individual organisations. These main nodes may also provide gateways through to many other national and international networks, such as Earn (the European counterpart of BITNET) and CA net in Canada. The regional networks cover broad geographical areas, although with some overlap as a result of recent growth, and are run by a variety of organisations from commercial enterprises to government funded groups.

Parts of the NSFnet backbone are also provided by commercial enterprise, and this mix of public and private characterises the Internet situation. In principle, an organisation which has an Internet connection, or 'feed', from a local or regional node, and which has the required equipment capacity, may itself provide any number of feeds to other organisations and individuals who are looking for an Internet account, relaying the network traffic back and forth on their behalf. In this way, the network is able to grow organically and, since one system user may have more than one feed available to it, can provide multiple communication pathways between sites, and therefore a more robust and reliable service.

In anticipation of serving a more widely expanded role, the NSFnet is currently being upgraded to support a much higher data transmission speed (45 Mbits per second, known as T3 standard), providing the bandwidth which will enable the development of a range of interactive services in multimedia telecommunications. This strengthened backbone is now referred to as a cloud' to emphasise its non-linear, interconnected, nature. The glue which binds all these diverse systems and machines together is the use of the TCP/IP communication protocols suite.

Most of the machines which constitute the Internet proper, and many of those which communicate with it, use this rather old, but widely used and constantly developing set of protocols. Their widespread use in the Internet is part of its inheritance from the ARPANET, and is a major reason why TCP/IP usage is seen as the main rival to the open-systems movement. Control of the network is distributed through various levels, with (at the moment) no overall control imposed from a central point.

The main backbone NSFnet is currently managed by a not-for-profit consortium consisting of IBM, MCI (an important commercial network provider), and the Merit organisation, based in Ann Arbor, Michigan. The network uses a Domain Name System, similar to that of JANET, so that IP address numbers (long strings of digits) can be replaced by more explanatory structured domain names; the main difference from the JANET method is that these names are structured in reverse order, progressing from the narrower to the broader domain: from organisation on the left to country on the right.

New system users register with regional co-ordinators, who have the responsibility for mapping connection points, in order to make their existence known to other network users. This gives the Internet a much more informal flavour compared with JANET, making it difficult to say exactly how many Internet users exist, since there is no one point at which they are all counted. There are agreed conventions on the use of the Internet to which all users are expected to conform: that it is intended for educational use and should not be used for purposes of commercial gain. This restriction is commonly quite broadly interpreted, and by no means all of the traffic between educational and other institutions is strictly educational in content; although it does preclude the offering of straightforward commercial services across the network, even though many organisations with Internet connections are commercial in nature. Some believe that this restriction will be relaxed in the near future.

Internet Society

The main decision-making body for the Internet is the Internet Society, a voluntary organisation which aims to promote Internet technology and developments. A sub-group of this, the Internet Architecture Board, is responsible for approving new and upgraded standards and for allocating resources such as IP addresses. The Internet Engineering Task Force (IETF) is another voluntary body with open membership, which represents the user voice in Internet debates.

The Internet depends to a large extent for its success on willing co-operation among its participants, and conformity to the spirit of network culture and etiquette. For example, an unwritten rule is that users who are established on the network through the benevolence of a host organisation supplying their feed are ethically bound to consider similar requests made to them in the future. The cost of these connections is normally born as a part of the network overhead; by this means the overall cost is spread among the participants, so that network use often appears to the individual user as free, or at least providing unlimited usage for minimal cost.

Internet provides electronic mail services based on the Simple Mail Transfer Protocol (SMTP) mail convention. Electronic mail has also given rise to a variety of alternative uses beyond the simple exchange of messages between individuals. A large number of news groups and discussion groups have grown up over the Internet based on the use of mail servers; within BITNET, which has become identical to the Internet for many practical purposes, these are known by the name of the listserv software used on IBM machines. These function as automatic distribution lists for subscribers interested in an individual topic, which rebroadcast messages from individuals to the entire group. The Usenet newsgroups are also distributed over the Internet, and are available in a similar way.

An example of one of the more active, for library and information workers, is PACS-L, dedicated to the discussion of public access computer systems; this has about 2000 members in several countries, and generates some 10-15 messages per day. Within the LIS field, there are groups centred around specific information service types, library systems suppliers, and broader social issues. There are estimated to be over 3500 conferences active, in addition to the Usenet newsgroups. Electronic mail is also used for the distribution of several electronic journals which are available over the Internet, most with free subscriptions.

Once edited and compiled by human intervention, the journal may be distributed either in its entirety as a long e-mail message, or in the more sophisticated cases only as a contents page. Subscribers can then receive individual articles as an e-mail message by sending a command which is automatically processed. Over 100 electronic journals, newsletters, and digests exist in the Internet environment. In this case, the listserver can be seen to function as a primitive database system. This practice saves on the use of network bandwidth by not transmitting unwanted information in the shape of unused articles; but also allows users to gather the information in the format most appropriate to their needs. Journals may also include graphics files, and text files may be available as plain ASCII or Postscript (laser printable) versions.

For sites with a direct Internet connection, real-time interactive access with public online information systems is available to remote sites. As with JANET, a common use is for access to online library catalogues. Since most of the main educational sites in the US are connected, a wide variety of catalogues are available, some with major collections, for example, the Melvyl catalogue at the University of California, and that of the New York Public Library. An estimated 337 library catalogues are available through use of the Telnet remote access protocol.

The Library of Congress files are accessible over the Internet, indicating the extent to which it is being used for formal purposes by library institutions. These files offer access to 28 million catalogue records included in the LCMARC files, together with copyright files, public policy citations and federal bill status files. The Library also allows Internet access to its technical processing and cataloguing system, and to its SCORPIO reference and retrieval system. A wide spectrum of non-library sources of information are also available. Many of the government agencies attached to Internet make information resources openly available; these include, for instance, databases mounted by NASA in space science and astronomy and the CIA World Factbook. This kind of access is likely to increase in the future as the Internet develops: recent legislation colloquially known as the GPO Access Bill.

Sites with direct Internet access have the ability to transfer files directly from one machine to another in their original format, using the FTP protocol. The practice of 'anonymous ftp' has arisen in the Internet community, by which ftp archives are established on a host machine at a particular network site, with named directories which contain archives of files available for transfer. Although computer system users would normally have to be registered with the site system administration in order to connect to a particular machine, these hosts have allowed free access to these specified ftp archives without the necessity to register a user name or password beforehand; although it is normal practice to enter the caller's network address as the password information.

This practice of 'anonymous ftp' is widely used to transfer entire documents, saving on the time-consuming necessity to view and save on-screen copies, and retaining any original formatting included with the file. The file can then be downloaded to a user's own machine, and used in whatever way is appropriate; the archive is able to distribute image as well as text-based information. Many ftp sites allow remote users to deposit files in the archive for others to retrieve.

A recent OCLC report estimated that there are about 3 million such files at more than a thousand ftp sites on the network, although the 20 largest sites account for more than half of these. Much of the material is source code for public domain computer programs, although the method is increasingly being used for the distribution of text files between specific groups or more widely. A major problem is that, as the report states 'the ability for network users to share information surpasses by far their ability to discover information on the Internet'. As with JANET, a main problem with resources on the Internet lies in the difficulty of finding out what information is available. Since there is no central administration of the network which controls all uses of it, there is no central directory of resources, although distributed network information centres are emerging. The problem is of a much greater magnitude than that faced by JANET, however, due to the much larger number of individual connections, and consequently the larger number of services and resources potentially available. It is also the case that resources are often duplicated in different network sites, and may exist in varying and incomplete versions and formats.

A number of resource guides have been drawn up by enthusiasts, covering both library-based databases, and those offered by other organisations, as well as lists of discussion and news groups which are known to exist. Long-term answers to the problem are the subject of debate, especially among some sections of the library community.

From 1992, the Internet is poised on the brink of a further phase of development. A concept sponsored by Vice-President Albert Gore for many years, and enacted in the High Performance Computing Act of 1991, promises to continue national development in the existing interconnected array of networks to upgrade and expand the system, opening it out to more colleges, schools, and public libraries as a national research and educational resource. This 'information superhighway', to be known as the National Research and Education Network (NIZEN), will aim for the development of a national 3 gigabit per second (a gigabit = one thousand million bits of data) network capacity. This represents an increase of approximately 2000 times over the current bandwidth, and as a rough guide would, make the network capable of carrying 100,000 pages of text or 1000 bit-mapped photographic images each second.

As with SuperJANET, this will support a vast expansion in the number and range of services possible, and make multimedia communications such as desktop video-conferencing feasible as an everyday reality. Another major goal for NREN is the standardisation of protocols across the existing patchwork of networks, and the integration of TCP/IP with OSI protocols to enable both to be used in parallel. This is not an isolated development, but is indicative of a number of initiatives within this vein in the United States. The Pacific Bell Corporation, for example, is building a broadband network (CalREN) in California to link participating universities, research laboratories, major hospitals, and leading high technology firms in the Los Angeles and Bay Area to develop innovative applications for education, healthcare, and business. These include projects such as the use of tele-seminars for remote educational programmes, medical image processing, remote patient monitoring, and specialist diagnosis, and the development of collaborative precompetitive research with industry and commerce.

Proponents of the NREN argue that the United States is falling behind in economic competitiveness because government, business, industry, and education cannot access the information resources which it needs to maintain international competitiveness and economic strength. Another school of thought argues that such needs would be best met by market forces in the shape of existing commercial networks such as MCI, funded by charging market prices for services. This argument has raised serious issues on future policies for the NREN on the issue of charging of services, which at the moment appear as though free to many end-users; and the use of NREN by commercial service providers.

Campus-wide Information Systems

Concepts such as the 'electronic library', in which larger proportions of educational support material are delivered in electronic form to users who are more geographically dispersed, have raised the profile of local networking as a means of integrating access to the diverse range and types of information available from all sources within the campus environment. In recent years, such developments have been carried out under the name of campus-wide information systems (CWIS), which are intended to bring together this variety of local information within a single framework. Access to this information should be simplified by an easy to use interface, which does not present barriers to the casual user, and is potentially available from all points of the local network. In many ways, this parallels some of the forms of community information systems which have emerged as civic facilities, but is more localised upon and targeted at a smaller and more easily identifiable community.

The concept of a CWIS as a broad umbrella encompassing diverse information types precludes too strict a set of conditions on the structure and format of the information. In this sense, the CWIS forms an information base of both structured and less structured information, rather than a database in the precise sense of the term. *A typical CWIS might include:*

— staff and student directory information, including telephone, postal and e-mail addresses;
— a diary of local events and university calendars;
— course information, administrative details and timetables;
— a listing of research projects and interests;
— university rules and regulations;
— newsletters related to the local campus;
— community and It-ravel information;
— library catalogues and library service information;
— locally produced indexes and databases;
— links to other CWIS, catalogues etc., located at different campuses.

By bringing this information together in one access structure, a CWIS can help to integrate a campus community which is split between different sites or spread over a wide area and develop a sense of community. There is also some interest in developing this approach as a method for inter-university communication at the campus level to encourage collaborative endeavours and cooperation, rather than through formal administrative structures. The information system can be developed as an extension of a library online catalogue, and be included within the same framework as an available network resource, or be developed as a general computer-support facility.

The direction from which the impetus for development comes is likely to influence the shape of the outcome, and the underlying objectives it is intended to achieve. The CWIS may be centralised onto a single machine from which remote access is possible, or act as a gateway to other local systems, calling the information up as needed. Whether the CWIS takes a centralised or more distributed approach, the task of maintaining and updating the information is normally the responsibility of the information providers. This may be either by submitting updated files to the operators of a centralised system, or by directly maintaining the data in one of the nodes which constitute a distributed system, to which the CWIS offers gatewayed access.

The ubiquitous spread of network access and microcomputers have encouraged the development of more proactive forms, under the aegis of local governments or local voluntary not-for-profit groups (particularly in the United States), which aim to provide the framework for community action which supports two-way networking, rather than the simple one-way passive consumption of display only systems. Viewdata systems have been installed in public libraries since the early 1980s, following the emergence of the British Telecom Prestel public information system. This is now a rather outdated and unsophisticated system, and is not amenable towards tailoring for local community information.

The late 1980s saw the growing use by public libraries of private viewdata systems which are developed and sometimes managed by commercial enterprises on behalf of their clients; these can be customised for individual applications and managed as a closed user group facility, offering specific information targeted at a well-defined user group. Added to the local focus in information content, there is also the advantage that some local authorities adopt private viewdata as an internal corporate information system, allowing the library to make use of any spare capacity.

The Electronic Public Information Network (previously the Local Authority Videotex Association) acts as a forum for channelling ideas and advice in videotex, investigating hardware and software issues, and developing applications for community information needs. This has some 46 local authority members and a number of other public sector and private organisations. In the context of local government authorities rather than just public libraries, a survey by the Society of Information Technology Managers in 1992 revealed that 19 per cent of UK local authorities already have public access terminals installed, and that a further 14 per cent planned to do so in the near future. These were installed in venues as diverse as libraries, council buildings, and shopping precincts, containing information on:

- local organisations and events;
- details of general practitioners;
- transport information;
- tourist information, including accommodation listings;
- information services aimed at local economic development, for example EC information on legislation, calls for tender, potential European business partners;
- one-stop enquiry shops, staffed by assistants able to make use of the local authority computer network to answer information enquiries on council services and activities.

An alternative to Integrated library system module is the use of a purpose-designed module within an integrated library system such as GEAC. This has the advantage of allowing the information to be directly controlled by the library authority itself, and to integrate it with the other activities and information services offered by the public library. Community information, for instance, may appear as one option on an Online Public Access Catalogue terminal, along with information relating to library holdings, updated and maintained in the same way.

The use of an integrated system gave the possibility of direct access to the information, rather than requiring the user to traverse the 'tree' structure of menus typical of viewdata systems, with the use of various access points; and offered the functionality to manipulate and print subsets of the information base as specialised products. Once again, maintenance of this system requires co-operation with neighbouring areas, since urban information requirements are not constrained by administrative borders. A recent development in the provision of community information has been the growth in co-operative host systems which, instead of building up a centralised bank of information for dissemination to users at terminal access points, concentrate rather on building a facilitating infrastructure which offers a framework for diverse groups to build and maintain their own information systems and services.

Host systems offer general facilities for creating local information files, the exchange of information between specific small user groups, and the creation of developmental projects and workshops, as well as more traditional services such as access to external databases. The Manchester Host system uses an internationally available communications software system more widely known as Geo-net, which is installed in some 40 countries around the world; Manchester Host is directly linked to several of these other Hosts in Europe (including Eastern and Central Europe), Russia and the USA through the international telecommunications networks.

The Host is accessible to its subscribers through use of a phone line, PC and modem. The Host approach assumes that users are potential information providers as well as consumers of information, encouraging an active participation in creating new applications and making contributions to existing ones. This led the Host initiative to borrow the idea of electronic village halls from the tele-cottages movement, and establish these in localised areas of Manchester, offering new technology facilities and training for particular groups, organised and managed by these groups themselves: for example electronic village halls have been established by women's groups, for small local business, and for the local Bangladeshi community.

As well as supporting greater participation for the individual in the information age through a philosophy of low-cost access and distributed processing, the Host approach has been to promote private and public sector partnership for the development of specific project initiatives for closed user groups. This involves cooperation between local government, employers and businesses, higher education institutes, trade unions, and non-governmental organisations, working on projects to extend the ISDN network and the consequent development of broadband information services, the extension of electronic data interchange (EDI) to smaller enterprises, and the development of a communications park for business. A number of closed user groups built around specific cooperative needs have evolved including:

— the Manchester Chinese Information Centre, supporting document exchange in Chinese script to the local ethnic community;

— the Manchester Business School, which uses the system to communicate with similar institutions on an international footing;
— the Manchester Asian Trading Network to improve trading links between local Asian business and suppliers in Bangladesh;
— the Manchester Design Community, in which smaller enterprises co-operate on bidding for larger contracts as a consortium.

On a wider geographic level, links with other European cities allow for greater co-operation in trans-regional development initiatives. The European and international linkages are likely to develop more strongly in the longer term, encouraged by the recent bid by Manchester to host the Olympic Games. The Host system brings together a diversity of community issues and related information activities focused on a locality. A range of networks have also sprung up to link together groups and organisations working in similar Areas, though geographically dispersed.

The 'networking' term is sometimes used rather freely in this context; some examples of these are closer to centralised database systems which offer remote access to users. Volnet, for example, is a central database host, providing a route to information for workers in the voluntary and community sectors at low-cost. The database offers online access to a bibliographic database, backed up by a photocopy-based document supply service, and to information on research projects and the details of UK parliamentary candidates. GreenNet provides a similar service for non-governmental organisations and individuals working with ecological and environmental issues, giving access to a range of computer conferences and resources, with a common structure to the location of information files in the network.

APC GreenNet links together progressive organisations with an interest in the development of information communication for social change, such as Amnesty International. A range of similar networking initiatives exist, some short and some longer lived. These can operate at a purely national level, or work on an international scale. While many of the existing community networks in the USA have evolved from local initiatives, stemming either from citizen groups or public authority organisations, recent years have seen the emergence of commercial organisations piloting or operating information networks directed at local users on a for-profit basis. These organisations make information produced by, others available whilst profiting from advertising through the network facilities. American CitiNet operates this service bureau arrangement, with particular customer bases in Omaha and Boston.

The regional Bell operating companies, the US regional telecommunication suppliers, are beginning to offer gateway facilities to information providers who are then able to recover their costs from information users through the normal billing arrangements of the Bell companies. A number of local government agencies are providing information and services in this way. As well as communicating with clients of their services, local government authorities have used information networking as a means of coordinating and delivering the services for which they are responsible.

There is an increasing recognition for the need to integrate and consolidate the single-focus information sets that are incomplete and incompatible, or fragmented across departmental and administrative boundaries, in order to deliver services effectively and support local development. The networking of databases and administrative systems increases access to information and forms a platform for the development of more powerful and sophisticated integrated systems. Local access to council services and information is effected by networking local branch offices and the creation of information points at which the public can find out about a range of council services, such as housing, employment, transport, and trading standards, together in one place; and can make specific enquiries about their individual progress, in areas such as payment of council tax or council house repairs.

Information access points such as these can be queried by telephone, and in some areas have taken the form of mobile services linked to a central computer system. This growth in the integration of information systems is linked to the development of new forms of service based on emerging information technology. Geographical information systems are a prime example of this, since much of the work of local authorities is spatially related. Different information bases, for example, may contain mapping data on the location of gas, electricity, and water utility pipes, property ownership and land use, and demographic information for strategic planning. Although these may be distributed over different locations and agencies, their integration across a network offers the opportunity for more advanced services: for example, transport information management systems which link traffic flows with highway maintenance and emergency services; or demographic planning with library development and community analysis to determine information needs.

Local authorities are increasingly providing information services linked to local economic development. While the Manchester Host is one important example of a farreaching approach to this, more traditional information services include access to European databases for information on European Commission tenders and contracts, EC legislation, and searching for business partners in other EC countries. A similar requirement encourages the integration of information access between different service delivery agencies. The Child Protection Register maintained by social service departments, for example, is in some cases available to accident and emergency departments of local hospitals in electronic form. The recent Care in the Community Act presupposes the exchange of information between social service departments and community health authorities.

The National Disability Information Project, formed under the auspices of the Department of Health, is a further influence developing this role for public service agencies. One of the pilot projects is the GUiDe system under development in Gloucestershire between the Gloucester and Cheltenham Health Authorities and the County Social Services Department. This deals with information for elderly and disabled people, and gives networked access to health professionals in hospitals, voluntary organisations, health centres, and social services offices. Another project is based at Gateshead Libraries and Arts Service, well known for its innovative viewdata based home shopping experiment. The Gateshead Disability Information project aims to make information from the council on disability available in various complementary forms—sound, text and video of the user's own choosing, and through a variety of delivery mechanisms.

There is a similar requirement for local government and local agencies to interwork with central government in an online networked environment, dealing with electronic information, often driven by central government directly. The Department for Education, for instance, uses electronic data interchange (EDI) to communicate with school educators on statistics and teachers' pensions matters, and similar arrangements for transport and trading standards are likely to emerge. Quasi-governmental organisations also have information exchange requirements with local areas, and as these bodies automate their information resources and archives, they are developing online access arrangements to replace manual administration for remote users. Increasingly, these applications demand more sophisticated information exchange facilities and broadband network links to accommodate heavy data transfer loads and multimedia information types. A recent example is the MONARCH system developed by the Royal Commission on Historical Monuments of England, an electronic version of the National Monuments Record. The online version of this gives remote access for about 200 users to some 6 million items of information on ancient monuments, including graphic images with hypertext linkages as well as text-based searching.

The information is used by local government for planning purposes, by English Nature for conservation decisions, and by the National Trust for property management, and will in the future be linked to geographical information systems for map-based searching and retrieval. Local authorities have traditionally been involved in the direct provision of services, including information services and information related to other activities.

As information becomes more and more important as a basic commodity for daily life, many see local government as having a role in the planning and development of a local information infrastructure, aimed at facilitating the integration and sharing of information from many sources and encouraging a policy of common access and citizen participation as well as enabling local economic and community development.

As changes in society, and in the structure and nature of organisations create changes in work patterns and habits for those who work in them or use their products and services, the change towards client-focused rather than organisation-focused services may offer opportunities for a new role for local government in this area. Perhaps the first developments in this area can be seen in the recent creation of 'One stop shops' providing information, counselling, and advice support services for local businesses, established under the auspices of the Department for Trade and Industry. These are mostly based on joint arrangements between Training and Enterprise Councils, chambers of commerce and local authorities, taking different forms in different areas.

Several of the development programmes started by the Commission for the European Community programmes are aimed at the development of networking as a supportive means for social and community development. The STAR programme, aimed at the less developed areas of the EC, has been used to develop the telecommunications infrastructure for rural areas to stimulate the use of information networking for employment and skills development. The ACE-HI scheme, administered by the Association of Community Enterprises in the Highland and Islands of Scotland is an example of this type of activity (with additional funding from BT), resulting in a E16 million investment in ISDN networks, and their successful use for the delivering of training and the development of enterprises which conduct their business through such networks, in remote islands and isolated parts of the mainland. Other EC social programmes have made extensive use of networking and telecommunications. The COMETT programme has used networking to promote transnational co-operation between commercial enterprises and universities for the delivery of training in technical skills. The HORIZON programme is establishing a disabled people's network through Europe, while the New Opportunities for Women programme is funding the partnership of Athens, Manchester, and Seville in building a telematics development network for the exchange of skills, information, and experience on women's training, employment, and enterprise. The European Community established the ROME (Reseau Observatoires Metropolitains en Europe) network in 1989 to advise the Commission on research and information development activities for urban areas.

European Urban Observatory

A three year project known as the European Urban Observatory was initiated in 1992 linking ten European cities in a European-wide network to support strategic decision-making and the common interpretation of information on a European-wide basis. The aim of this project is to monitor urban trends through the sharing of urban management information, the transfer of experience and comparative evaluation of policy. Information is structured into three levels, dealing with basic indicators for policy decision-making, with quantitative comparative data, and with individual city data stored locally; these levels are linked together through the use of X. 25 private networks, and the BT Tymnet wide area network employed to form a closed user group. This networking of public management information contributes to the development of pan-European urban policy issues, and gives a comparative background resource at the local level. Information networking in the health service sector in the United Kingdom has become a vital activity in recent years, stimulated by a growing movement for the improvement of health information services, following the work of the Korner Group on information systems, and stimulated by the radical organisational changes creating an 'internal market' for purchasers and providers of health services. These tendencies have created a pressure for the harmonisation and integration of all forms of information within the sector, and a demand for wider accessibility and sharing through computer networks and other applications of information technology.

The distinction between traditional bibliographic information services and other information sources has lessened with 'information' and 'intelligence' being taken to include, sometimes primarily, statistical and demographical information. Within this context, networks are used to link the office automation and service environment with information providers across many administrative levels. Management information in this framework would include:

— literature in the form of written or audiovisual material: published books and journals, etc., as well as unpublished material in the form of correspondence, internal reports, and committee minutes;
— quantitative information, in the form of data relating to the activities and context of the organisation; statistical returns, demographic analysis, performance indicators, trends analysis.

Within the health sector, an integrated delivery of information services more often requires sharing and exchange between fragmented units, now within a more competitive context. Library units, for example, may include within the same region headquarters libraries, departmental libraries, small service points in branch hospitals, postgraduate medical centre and research information points, and nursing libraries; each with varying degrees of autonomy. With a growing amount of information available in electronic form, user requirements are likely to involve the retrieval of information from remote sources across a network, as much or more than from a local collection, and to varying forms and types of data.

A medical researcher, for example, could require not only bibliographic references, and full text source material, but also images and the ability to process images, datasets of published social and epidemiological statistics, and so on. At the international level, medical information services are developing strengthened links through the use of networks. The HealthNet system, for example, provides access for health professionals in developing countries to medical literature and databases and to colleagues elsewhere, through the use of low-cost packet satellite technology. HealthNet is initially targeted at Africa, and includes access to British and North American gateways which allow medical libraries and information units in these areas to 'twin' with similar organisations in Africa.

Generic facilities such as electronic mail and file transfer through computer networks open up the possibility of more extensive contact between health professionals in the context of an electronic community. The increasing depth of special knowledge possessed by a widening spectrum of health professionals, located geographically separate from each other, and the interdisciplinary nature of health science, point towards a need for networked and integrated solutions to information needs. At the same time as information management has assumed a higher profile within the health sector, there is a growing interaction with external agencies fuelled by information interchange.

Some instances of co-operative networking with local authorities have been mentioned earlier—a common interest in demographic analysis for strategic planning is another obvious example—and the same trend emerges in health service relations with other public service organisations. The range and volume of data involved can be quite significant: a medium-sized hospital can send some 300-400 pathology and radiology reports per day to several dozen local general practitioner centres, and up to 100 discharge letters; a large general medical practice may make some 40 referrals per day to hospitals, and perhaps 7000 registrations per year to Family Health Service Authorities. Each of these activities involve the transmission of information and storage within the receiving systems, with the desirability of enabling the direct capture by the receiving application, and avoiding a need to re-key data which is already in electronic format.

The data needs to be exchanged between autonomous sites, loosely coupled in systems with a local focus. Externally, information flows will move from the National Health Service to product suppliers, service providers, government agencies, local authorities, European and international organisations, the insurance and banking sectors, private health care agencies, and other national health services. To achieve an electronic data integration across the whole health service delivery chain, and with organisations in its environment, the development and application of standard forms of information interchange is required.

The standard provides for the definition of data elements, syntax rules, and the structure of messages built from these data elements in a way that can be adapted for various applications; EDIFACT is managed by the United Nations as a cross-sector global communications standard, and is also defined within GOSIP. The NHS Information Management Executive has initiated many projects intended to contribute towards the development of a common infrastructure through the use of information systems strategy. A major plank within this strategy is the desire to create a unique NHS number for individual patients, enabling the creation of a population register designed for sharing across the sector, and forming a natural integration of basic administrative activities. This will underpin NHS-wide electronic networking, and together with agreed standards for information interchange, lead towards a shared language for health care throughout the sector.

A number of these projects are aimed at general practitioners, and their links with other health agencies, involving the identification of a basic information structure common to all GP practices, and the development of compatible systems.

More broadly, projects are developing standards which apply across the sector, for:

— the development of a common basic specification as an extension of the NHS Data Model, to cover data exchange standards, minimum datasets, clinical information for resource management; and the information demands of hospital and regional decision support systems;

— the development of an NHS Administrative Register to act as a shared register of basic administrative details meeting the requirements of District Health authorities, Family Health Service authorities, general practitioners, and hospitals;

— a Community Information Systems Project to facilitate the exchange of information to support contracting, and the development of datasets for ambulatory care in hospital and community settings, for use by community health services.

The National Health Service in the United Kingdom has seen a vast increase in the importance of information management, with information networking forming a vital part of this. It is certain that the importance of networking activities will continue to grow as a strategic element in the construction of integrated information and intelligence.

Open Systems

The concepts of open systems are now widely accepted in principle, with many believing that the CISI model forms the underlying solution to connectivity problems. In practice, however, the actual take-up of OSI in terms of working products has been slower than at first anticipated. The slow pace of standards development tends to prolong uncertainty amongst end-users, and discourages suppliers from manufacturing OSI products on a large scale due to the lack of a mass market. The lower network-dependent levels of the OSI model are well established in both local and wide area networks; the upper levels are much less well supported with few general or specialised products in most areas.

The second problem facing OSI products is that of conformance testing, to ensure that inter-operability standards are met. These are relatively new, and although normally carried out by an independent third party, do not necessarily guarantee full compatibility; variations within the paper standard mean that compliant products will not always work with each other, still further undermining user confidence and acceptance. The interchange technology should allow a user working on one system to use an application on another system, without concern with data formats; the long-term goal is to allow network users to co-operate seamlessly with each other, regardless of physical location or equipment.

A common framework for networked information interchange must incorporate generic application problems, such as remote database access, file transfer and interpersonal messaging, in a way that separates the local form of the application from the interchange aspects, enabling flexible variations in the former while standardising on the latter.

De facto and de jure standards may emerge from two main sources. De facto standards are those which derive from the use and popularity of easily available products, through a process in which the most successful become widely accepted as industry standards. Proprietary standards from single computer manufacturers fall into this group; while having the advantage of the backing of a major interest group, their future development is unpredictable, and they are 'closed' systems in that the full specification for the standard is not published but kept as a trade secret. De facto standards may also emerge through adoption by a number of different product developers, and have the general support of a wider range of users who may also participate in the standards making process itself. This type of process is likely to be founded on widespread availability, and may be stimulated by an initial development, or adoption, of the standard in government administration projects. This has been the case with the TCP/IP communications protocols, which are widely used in industry, government and education (particularly in the United States) and form one approach to the 'opening' of computer systems to each other. De jure standards are those which emerge from an organised, usually public, programme sponsored by governmental or international initiative, emerging from an open debate over several years; by their nature, these developments are more speculative in nature, formulating paper specifications rather than products.

The successful take-up of the standard may depend on the willingness of private manufacturers to build conformable products, or on the degree to which the standard can be imposed by regulation. The CCITT recommendations for telecommunications, for example, can be reinforced by regulations governing approval of equipment for connection to public networks. Standards for software in computer systems cannot be regulated in this way, but can be encouraged by making them mandatory in private tenders for substantial government information system projects.

The main international efforts towards the specification of a framework for open, intercommunicating, systems has been the work done under the aegis of the International Standards Organisation (ISO) and known as Open Systems Interconnection. Although this is a long-term programme intended to take effect slowly, at the present time the ISO Open Systems model has had more influence in Europe than elsewhere.

Open systems interconnection (OSI) aims to develop a framework for the development and approval of computer networking standards that provides communications-based user services which support inter-working with computer systems from different manufacturers, in different geographical locations. This framework defines a conceptual and functional model for supporting multi vendor, multimedia and multi-software networking by formalising and delineating the problems involved in sustaining it, and structuring them within an open architecture. The basic idea of OSI is not to invent a completely new set of protocols where adequate ones already exist, or to impose any one single approach or technology, but instead to acknowledge existing protocols by placing them within the framework of the model according to open architecture principles. In this way, OSI functions as a focus for coordinating standards development, and depends on achieving a consensus view from all interested parties.

The organisation progresses its work through a web of committees, subcommittees and working groups, developing and adapting standards through lengthy consultation processes. Before being formally adopted as an International Standard, a specification will pass through working paper, draft proposal and draft international standard phases, with public comment and feedback at each stage; ten years is a not untypical timespan for this process.

The basic OSI reference model was formally adopted in 1983, and has since been elaborated for particular task types and application areas. The mission of OSI is described as:

— to specify a universally applicable logical communications structure;
— to encourage different users to communicate with each other by compatible implementation of communication features;
— to serve as a reference point in the development of new communication services;
— to enable the steady evolution of applications through the flexible incorporation of technical advances and the expansion of user requirements.

The open systems interconnection model does not impose a single standard, but incorporates a collection of existing standards, placing them in the context of an abstract model of the communications process. Compliance with the model is intended to ensure functional equivalence between computer systems. The term system is here used to denote a collection of computers, software, peripherals, and human operators which have information transfer capabilities as an autonomous whole - that is, a tightly linked computer network as much as a single computer installation.

OSI is therefore a standard for network-to-network communications. Given the tremendous diversity of needs and options in this area, a single universal protocol is not possible. The OSI approach is to group functionally similar protocols into a set of layers, each of which represents a particular class of service available in open systems. The basic reference model defines each of these layers by specifying both the functionality of each layer and the interface between adjacent layers.

The specification of the functionality within any particular layer defines what is to be done rather than how it should be done. The way in which these functions are implemented in software programs therefore can vary between different products, although the outcomes should remain compatible. The specification of the interface between layers ensures that the information which is received and passed on between software modules is in a standard form. The functional layers are stacked in a pre-defined order to form organised and structured communication tasks. Each layer provides a set of services to the next highest layer in the stack on demand; to be able to provide these it receives a set of services from the next lowest layer in the stack. Since each layer is providing its own set of specialised tasks (specified as the functionality of the layer), it acts as a user in relation to the layer beneath, requesting services from it but knowing nothing about how they are achieved.

Data passes up and down the stack using the standard interfaces as a conduit, with each layer adding control information relevant to its own specialised tasks in the overall communication process, and passing the package on. An analogy might be a set of Russian dolls, in which the larger doll encapsulates a smaller one. In OSI terms the original nugget of data is passed down and encapsulated by larger control messages as it passes through each layer on the transmitting system; on receipt the process is reversed, and the larger containers are stripped off as the message passes back up the stack of equivalent layers on the receiving system. The model deals only with the logical structure of the communications process and says nothing about how this is to be physically implemented. Since each layer is acting as a self-contained module within its own boundaries, an individual implementation of layer tasks in any particular system can change without having an overall effect on the communication strategy, as long as it conforms to the logical requirements.

Software modules can then be changed in individual systems to accommodate changes in user needs and network technology, as long as the new version continues to provide the same set of services through a standard interface. In the OSI model the overall communications tasks are partitioned into seven layers, with no overlap between them. Each layer communicates with to equivalent 'peer' layer on the remote system by exchanging messages on matters for which it alone is responsible. This is actually achieved through using the services provided by subordinate layers, but each individual layer works as though it were speaking directly with its peer in the remote system. This peer- to-peer communication is made possible by the standard boundaries between layers which are the same on both linked systems. The seven layers in the model fall into two main sections. The three uppermost layers are primarily concerned with the applications communicating via. the linked systems, and with organising the provision of services to human users.

The three lowest layers are concerned with the technical transmission problems of sending data across networks; the fourth layer (the transport layer) acts as an interface between these two broad tasks. The three lower layers are supported in both the end systems and the network mediating between them; the upper layers are in the end systems only. The physical layer controls the actual physical connection between the two systems, the hardware standards used, and the method of data transmission, working at the level of the individual data bit.

The data link layer provides error handling and flow control, and works at the level of the data frame. The network layer controls both routing decisions and the relaying of data packets through intermediate systems, concealing from the upper layers the particular path through the network which is being used. The transport layer is the first to be primarily concerned with end-to-end operation, concerned with quality of service and ensuring a reliable connection for the upper layers. There are five different classes of service available from the transport layer, including connection-oriented and connectionless services.

The end systems must agree which class is in use for the session before exchange of data can commence.

The session layer is responsible for setting up an individual session between two systems, and managing the inter-system dialogue on behalf of the higher layers; for instance, synchronising the information flows between the applications, and allowing activities to be started, abandoned, paused, and restarted correctly. The presentation layer negotiates how data is to be presented to the application, dealing with character sets and codes, agreeing to exchange data in a shared common format. End systems which use different information coding sets (such as ASCII or EBCDIC) can use a special open transfer syntax notation called Abstract Syntax Notation I (ASN.1). The application layer, the topmost layer of the CISI stack, is responsible for interfacing with the local end system and ultimately for supplying direct network services to the user. In an open systems environment, all lower level OSI functions can only be accessed through the application layer, and not directly.

The application layer is not a user application in its own right, but rather handles the semantics of the communication process in the same way that the presentation layer handles the syntax. It has specialised elements for different generic applications, such as file transfer and electronic mail, which interface with proprietary packages for these applications installed on the local end system. For example, a person would send and receive electronic mail through the same package which they normally use on their own system. This local package translates the mail into terms acceptable to the OSI e-mail application layer element, which in turn initiates the setting up and management of the appropriate communication characteristics by the next layers, and so on down the stack. For the transfer to be successful, the receiving system must be able to support the same kinds of action, in reverse.

FOURTEEN

Communication Network

This chapter briefly indicates the fundamental problems of data communications, and gives an overview of the broad approaches which have been taken towards them. A computer network is intended to connect together a wide range of information technology equipment, which varies in the design and operation of the equipment itself, in the information types it handles, and in the internal processing which it undertakes. In order to be able to communicate successfully, these various systems must have a common 'language' for the exchange of information, while retaining their local independence.

As well as an agreed form for communication the network units must also agree on the *process* of communicating—the rules which govern the procedures for the exchange of information and the co-ordination of action between participants. In daily communication we use a language to express our ideas and intentions, and the voice as a means of transmitting this information to other people.

To ensure that the communication is successful, we must converse in an agreed language which the hearer is able to understand, and we must be able to respond to messages asking for clarification, repetition, and so on. In any conversation, there are well-understood rules which govern the process of communication—taking turns, allowing others a chance to speak, indicating that one has finished speaking, waiting until another has finished before speaking oneself, and so on. These signals are often given implicitly, through signs such as gestures or tone of voice.

The data communication process has to be much more structured and organized in style, spelled out in every detail, since these tacit factors are missing from a machine-based process. The rules for communication must be well defined and standardized; there is little room for 'working out' what a participant in the exchange really meant to say. This process underlies the development of sophisticated and 'user-friendly' applications which depend on network communication, even though this may not be immediately apparent.

A goal of information networking is to allow users to work in this medium as naturally as possible, finding the least obtrusive fit with the natural work environment. A dialogue which is mediated by computer communication places certain structural constraints on the process due to these underlying factors; the task of design is to find ways of incorporating the equivalent of the contextual factors which inform human communication while working within these constraints.

A basic model of the communication process identifies the entities involved, and their respective roles. There are *three main elements:*

— a sender, which originates the message, and is responsible for representing the content of the message in an agreed format;

— a channel, which carries the message from sender to recipient through some physical medium;

— a recipient, which receives the message, and is responsible for interpreting the format in order to extract the information content.

The communication is not simply one way, and the roles will at times be reversed: the recipient may wish to send a reply to the original message.

In network communications explicit signals and format rules are vital to the success of the communications process. For a computing machine to manage this role effectively, rules governing these issues must be highly detailed and agreed in advance, and encapsulated in software programs. Data and communications software must be able to deal with management of the exchange process, not only by sending and receiving messages for transmission but also by encoding and decoding between the internal format used by the computer and the format used for external communication. It is also necessary to monitor the quality of the exchange and to co-ordinate activity between participants at each end of the process for longer data exchanges. This is problematic because the channel which links the participants will have characteristics of its own which may affect the communication process.

The sender and recipient exchange the encoded message as a physical signal transmitted in some way through a common channel. The environment of this channel can potentially interfere with and distort the signals in the process, resulting in the loss of some information content and delivery of a corrupted message. This kind of interference is known as noise—the general effect of corruption and garbling of the signal which occurs after it has been transmitted, and before it is received. For example, noise may result in the appearance of unreadable characters on a computer screen in the midst of a normal text message, as the receiving system attempts to interpret corrupted signalling codes.

Although the use of more reliable technology for the transmission medium has reduced the problem of noise in recent years, the rules which govern the process of communication must still be capable of detecting and recovering from the effects of noise by returning a message to its original form. This requires the communications software to have some means of matching the two versions of the message—that sent and that received—and entails the addition of control information to the actual data content by one means or another. Thus the exchange will consist of more than the simple data content of the message.

Techniques for dealing with the effects of noise are concerned to preserve the quality of the message by ensuring that the integrity of the signs used to represent the information content have not been compromised. The communication process must also therefore be concerned with co-ordinating the system of codes which is used to express the semantic qualities of the data exchange— what might be loosely called the 'meaning' of the message; with the physical representation of the chosen coding system as transmittable signals; and with any control information which binds different elements of the message together.

There are various different coding systems for the task of mapping information generated by user applications into the binary codes used by the computer. For text, the ASCII (American Standard Character Set for Information Interchange) is widely used both within a computer and for data interchange purposes; however, completely different character sets are used in some circumstances, for instance the EBCDIC character set which dominates the IBM mainframe world.

There are several different representation formats for graphics, which may be either bit patterns or geometric coordinates, and for other forms such as video signals and multimedia information. For meaningful information to be exchanged there must be agreement between the end systems on the method by which the message is represented for transmission, even though this may not affect the actual exchange of data symbols through the channel. Since there is little standardization in the internal processes of computer equipment itself, the actual physical signals which the end systems generate must also be standardized for inter-operability, and may need translation between different physical forms in order to make use of different communication channels.

Synchronous Communication

The additional control information which is appended to the user data content carries information used in processing the message, and must also conform to some agreed system. A simple text message, for example, may contain an, overhead of information which is used to indicate boundaries between characters or chunks of text; this will appear as binary digits apparently identical to the data content unless there is some common method of 'reading' the stream of symbols.

A multimedia message which contains both sound and video elements requires control information to synchronize timing of the audio output with the screen display, which must be interpreted correctly.

There is a need to synchronize the flow of information between the two end systems, so that signals which will conflict with each other are not generated simultaneously, and so that the receiver is ready and capable of processing the incoming signals at the rate which they are sent—controlling the 'turn taking'.

Communication can take place in one of two alternative patterns, and in each of these synchronization is achieved in different ways. Asynchronous communication might be characterized as 'on demand' use of the communications channel. All network users have equal access to the channel, and use it to send messages at random intervals, characterized by short bursts of activity with quiet periods of varying length in between. The receiver waits passively once the connection is made, and is activated by the receipt of some message. A typical example of this type of activity would result from a user seated at a terminal screen, alternating interaction with a remote computer with periods of planning and thinking.

Synchronous communication depends on a fixed cycle, in which each user, and both sender and receiver, have access to the channel at defined time intervals for a fixed period. A typical use for this approach would be for large data file transfers in situations where the transfer is being managed completely by the computer systems as a background process, with no human intervention, or for the synchronization of composite sound and video signals in videoconferencing systems. This synchronicity characteristic delineates the time/ response dimension of the process, defining the appropriate activities in the end systems to control the process; there is also a characteristic of the communication channel which will affect the process.

Channel and Network Communication

Communication through a transmission channel can occur in one of three different modes determining the allowed direction of, and activity in, the transmission. Each of these requires a different level of sophistication in control of the communication process:

— in simplex mode, the communication is in one direction only;
— in half-duplex mode, the communication can be bi-directional, but at alternate times only;
— in full-duplex mode, the communication is bi-directional, and both end systems may transmit simultaneously.

The term *channel* is normally used to denote the more abstract concept of a logical channel rather than the actual physical media. One physical link may potentially offer several logical communication channels, which can be thought of as separate entities.

Since a network is a medium which is shared by many users, management of the communication process is a more complex affair than is suggested by the simplified one-to-one model. As well as agreeing on the terms by which the exchange will be conducted, the communication program must be able to supervise and regulate the dynamic elements of the process, making necessary adjustments as the exchange develops and operating conditions fluctuate. The communications program must be able to:

— establish and maintain a connection across the communications channel between the sender and recipient, agree the terms on which the exchange is to be conducted, and terminate the connection-once the transaction is completed;

— regulate the flow of information, between the two participants to prevent overloading of data by the recipient, if it is unable to process it quickly enough, causing a possible loss of information;

— control the quality of information received by checking for the presence of errors and taking corrective action as necessary, perhaps by requesting a re-transmission of corrupted information.

The term bandwidth is an indication of the information capacity of an individual element in the process: the higher the bandwidth of the channel, for example, the higher is its information carrying capacity. The overall communication process (including the end systems) can only utilize the information-carrying capacity of the element with the least available bandwidth in any particular situation; either the transmitter, the receiver, or the channel will have a more limited bandwidth and will define the overall information exchange rate. In many situations, the bandwidth of the communication channel will be the limiting factor.

The capacity of a channel is normally expressed as the number of distinct symbols per second which can be passed along it. This maximum *data transfer rate* is usually given as a figure measured in bits per second, since the number of bits which make up a useful 'unit' of information will differ according to the information type—between characters of text and scanned images in bit-map form, for instance. In this way, a channel which is described as having a particular data transfer rate can be matched to different information transfer processes through a common comparable measure.

In some circumstances, particularly in modem communications, the term *baud rate* is popularly used as an equivalent to this measure of bits per second; although this is not a strictly accurate use of the term, it can normally be taken to mean the same thing. Transmission channels can also be characterized by the method of signalling which is used. In *baseband* signalling the capacity of the channel is monopolised by a single message for the duration of the transmission, typically by varying the voltage along an electrical cable in a predetermined pattern.

Broadband signalling enables a physical medium to be shared by more than one transmission by varying the frequency of the signals to ensure separation of one from the other—as a radio can distinguish different frequency transmissions. Broadband offers a higher bandwidth, and is therefore more associated with information networking applications which require higher capacities, such as multimedia. A goal of advanced network technology is to support a flexible, use of bandwidth among multiple users according to fluctuating demand, changing the allocation dynamically according to the transmission needs of the moment.

A communication channel is a logical concept that can be implemented on a variety of physical media, each of which has its own particular characteristics, such as bandwidth. The choice of which to adopt depends on the purposes of the network and the kind of applications it must support, the predicted traffic and data rates together with required response times, and a cost/benefit calculation on the relative importance between these factors. *The most widely used physical transmission media* are mentioned here.

1. The *twisted pair* is the simplest form of transmission cable, two simple electrical cables carrying small voltage signals. This is a very cheap method of connecting individual computer devices to a network, and is easy to install, but is not suitable for carrying traffic from many devices due to the limited bandwidth which it possesses, nor for applications which require more than a 1 Megabit per second (Mbps) data rate.
2. *Co-axial cable* uses a more substantial copper cable to offer a higher bandwidth suitable for multiple access from a limited number of nodes on a local network; and for the transmission of information forms which require a higher throughput of data, such as video signals, at data rates of up to 10 Mbps. This higher capacity is bought, at the expense of a higher material cost and a more complex installation process.
3. *Optical fibre cable* provides a means of transcending the physical limits of electrical cable, by conveying signals on a fluctuating light beam enclosed within a very fine glass fibre. Since light waves have a much higher potential bandwidth than electrical waves, data rates up to and beyond 100 Mbps are possible, as well as providing a lower susceptibility to noisy or sensitive environments.

   A single glass fibre, smaller in dimension than a human hair, is provided for each signal, and an installed optical cable may consist of several hundred fibres bundled together, providing *a* high throughput for existing services and spare capacity for future growth. While optical fibre provides a means of building very high bandwidth networks which can support sophisticated multimedia services, it is expensive and more difficult to install, requiring special interface equipment between cable and computer.

4. Microwave transmission offers a non-terrestrial alternative to a physical point-to-point link. Data is conveyed on radio waves through free space, and is normally used in conjunction with geostationary satellites. The data is transmitted from a ground station, called an uplink, to an orbiting satellite which is in direct line-of-sight, and the signal is bounced back to a receiving station (the downlink) at some different geographical ground location. These re-transmissions can be broadcast over a wide area for receiving by many destinations simultaneously, or be highly focused onto a single site.

   Satellite relay is commonly used to leap across long distances, such as in establishing connections between the United Kingdom and North America or the Far East, or where the laying of a physical cable would be costly or impossible across difficult geographic terrain. Microwave transmission may also be used on a ground-to-ground basis, between stations situated on high ground within line of sight of each other.

5. Network communications have also been achieved by the use of infra-red and broadcast radio signals; there is a growing interest in the use of such 'wireless' communications. Some use is also made of radio data broadcasting at a national level.

Protocol

Essentially, a protocol is one attempt to address the basic problems of data communications which are relevant to networking within a particular arena. Since there is no universally applicable solution to these problems which will be effective in all situations, a range of protocols have arisen which have proved effective in specialized areas within the current feasible limits of technology. Since the development of communication protocols has been carried out within the context of perceived practical problems, there is a tendency for the available offerings to be predicated on the use of particular types and make of equipment.

Computer manufacturers have devised their own proprietary protocol suites intended to interconnect only their own Computer makes; and more generally applicable protocols have emerged for broader categories of computer, such as for local networks of microcomputers. Although a protocol is a logical specification rather than a particular implementation in software, the protocol is likely to make assumptions about available levels of functionality and processing capabilities.

As both computing technology and information applications have led to more complicated and sophisticated systems, so the communications process between systems has also developed a parallel level of complexity. The complexity of modern networking has led to the use of a layered approach to protocol specification. A *layered protocol* aims to analyse the overall communications process into well-defined and semi-autonomous functional modules, each of which provides a specialized service in the overall scheme and works in conjunction with the other modules.

Network communications can then be split into constituent activities, and these may be implemented on different parts of the system in different ways, not necessarily all together on the same piece of equipment. Any particular layer may potentially contain alternative specifications for reaching the same goals, each of which is more appropriate for a particular environment or system arrangement.

As long as the module provides the appropriate service to its co-operating modules in the process, the way in which it achieves this, need not concern the other elements involved. This kind of arrangement provides a greater flexibility in interconnecting different types of system which conform to the same layered framework. There is not, unfortunately, any universal or even dominant agreement on one scheme for dividing logical functions up into different layers.

The growing need for the capability of interconnecting computer systems and networks whatever their local characteristics has led to several initiatives aiming to develop a degree of standardization between the major alternatives, and to encourage the conformity of new products to these standards. These initiatives aim to develop a framework within which the most widely accepted and established 'industry standard' protocols can fit as alternative options. At the present time *networking protocols fall into two main groups:*

— those emanating from or supported by national and international standards making bodies, which are more 'open' in nature with published specifications;
— those emerging as proprietary standards from major computer manufacturers, which are 'closed' by virtue of their confidential specification and non-participative development.

A set of protocols will deal with aspects which are more closely related to the physical aspects of data and network communications and those which are more related to the information transfer aspects of user activity. At the lower layers, the data communications aspects can be seen as falling into three main areas: the physical connection, basic data management, and routing through the network.

The *physical level* is concerned with producing, receiving and monitoring the physical signals which are actually transmitted between systems, and with the standardization of methods for representing data in this form. A well-known example of this is the RS-232C standard, which defines the electrical interface on equipment as a multi-pin connection and assigns a particular function to each pin. The *data link level* is concerned with error control and recovery, and with flow control. In order to achieve this, most protocols encapsulate the user data in frames by dividing the stream of data into units of a pre-defined size and adding control information for use by the receiving system.

Since the network connection will be conveying data frames from other network users, the frames also contain header information ensuring delivery to the correct destination. Special frames may also be used to send only control information, for example to establish and terminate links with a remote computer. The frames produced by this layer are passed down to the physical layer for transmission. Error control is important at this level, since noise on the channel may result in the loss or corruption of frames; one purpose of splitting the whole message up into data frames is to retain control of the transmission as it develops, and take any remedial action necessary. Error control takes a number of simpler or more sophisticated forms, some of which require feedback from the receiver, for example:

— echo checking reflects characters received back to the sender for display on the user's screen, so that any corruption can be noticed and re-transmitted;

— acknowledgement schemes require the receiver to send an acknowledgement token back to the sender for each frame, before the next is transmitted;

— check digits can be added to the data frame by the sender, based on a computation carried out on the user data itself, and which the receiver can also use to check comparability.

Flow control is concerned to ensure that the data is not sent faster than it can be received; this function can make use of some of these devices, for instance acknowledgement schemes, or add features of its own to the exchange process. The network level is concerned with the management and control of routing through the network. The extent to which this is relevant will depend on the type of network in use and the mode in which it is operating. There are two broad *modes within which a network can operate*, providing two types of service:

— *connectionless mode,* in which the data is dispatched by the sender with the appropriate destination address, but no direct connection is made between them;

— *connection-oriented mode,* in which the two end systems behave as though directly connected by a physical medium and work in co-operation with each other.

A network may offer both these types of service as alternatives, but cannot normally support mixing between the two types. The network layer will typically organize the user data into packets, splitting up a long message into several units in this way, and adding source and destination information to each packet to control its path through the network. The user data is then made transparent to the network itself, which is only concerned with the 'external' address information. Different packets which constitute parts of the same message are then treated as independent units, and may travel through the network at different speeds and by different routes.

The receiving system takes delivery of the packets and reconstitutes the message in its original format using the control information added by the sender. An X.25 network is an example of this arrangement. The network layer passes the packets down to the data link layer for the addition of error and flow control information, once it has completed its own responsibilities. In this way, user data is passed between layers between both sender and receiver, with only one actual physical network connection among them. In wider area networks, these three layers are also contained within the intermediate stages internal to the network, as well as in the end systems.

The specific protocols which implement these broad functions within an individual context will depend on the particular elements of the system which is in use, and the equipment available to it. A range of protocols are implemented either as software in communications packages, or as hardware built into data communications equipment. Some have achieved a greater popularity and therefore a *de facto* standardization over time.

In order for the communication process to be successful, both systems which intend to exchange data must be capable of running the same protocol, and therefore require communications packages or hardware which are compatible with each other; larger systems may offer a range of communication alternatives which can sense the method used by a caller and select the appropriate option. The *types of device* which have a role to play in the entire networking process can be placed into two categories:

1. Data terminating equipment (DTE) describes the devices which originate, or provide the final destination for, a transmitted message; either that operated by an end-user, or a computer system working autonomously. Typical examples of DTE devices would be a desktop microcomputer, a mainframe computer system, a visual display

unit terminal, a printer, or a facsimile machine.

2. Data communicating equipment (DCE) describes the devices which are used to support onwards communication of a message received from some DTE device type—acting as 'steps in the chain' rather than as end-controllers of the process. This may be a device attached to the end-user's system, or some intermediate step within the network itself.

The data communicating equipment acts as a common interface for unlike data terminating equipment, they even out the differences through the common use of a standard protocol. In this way, dissimilar types of information technology equipment can be used for information networking without requiring an end user to be aware of the physical characteristics of the receiving system; the more widely used and open the standard protocol which is used, the wider range of equipment which can interface in this way. The kind of equipment which falls into this category fulfils specialized roles in the communications process: some of these are integral to the internal functioning of the network element itself while others are associated with user end systems. Some of the more commonly found types are described here, as indicative of the range.

The multiplexor is a device for concentrating a number of terminal devices onto a higher bandwidth line, to make the most effective use of this capacity and to remove the need for installing many duplicate lines (one for each device) in parallel with each other. This arrangement would typically be used to connect several terminals situated near each other to a central computer which is located at some distance.

The multiplexor is responsible for apportioning the line on an equitable basis between all the terminals connected. There are two main *methods of multiplexing:*

— time-division multiplexing, in which each terminal is allotted the full capacity of the channel for a fixed period of time, in a rotating sequence;

— frequency division multiplexing, in which the full channel capacity is divided between users by allowing simultaneous transmissions on different frequencies.

The most appropriate form to use depends on the applications which are using the system; commonly asynchronous enquiry-type applications, where activity is in unpredictable short bursts, use a time-division form of multiplexing. A multiplexor must be installed at the receiving end, to separate out the different terminal signals for processing.

A modem is another communicating device which allows one kind of DTE equipment, typically a personal computer, to establish a temporary connection with another remote computer system by using the normal telephone line. Since the telephone system was originally designed solely for voice traffic, which is analogue in nature, some conversion process must be used to translate digital computer data into a form suitable for analogue transmission. Digital signals are in a binary form, having two states—on and off—which can be represented as changes in voltage signals on a cable.

Since the normal telephone line is designed for transmission of sound waves, these binary signals must be transformed into an audible tone by the modem. This process is known as modulation; a modem (modulator—demodulator) must be placed at both ends of the telephone link, acting as the interface between the end systems and the line, encoding and decoding the digital data signals. A series of standard recommendations have been produced by the CCITT (Comite Consultatif Internationale de Telegraphie et Telephonic), a Geneva-based international organization concerned with issues in telecommunications development, for the use of modems.

The cost of modem technology falls steadily, faster transmission speeds are achieved at a higher expense due not only to the price of equipment but also to the additional communications processing which is necessary at higher speeds. Higher transmission speeds are achieved by a panoply of special techniques which require more sophisticated error checking capabilities to counter the increased influence of noise. While the CCITT recommendations are standard within Europe, the existence of a separate set of standards in North America, developed by the Bell Communications Corporation complicates the international situation. The most common of these are the Bell 203 (1200 bps) and Bell 212 (2400 bps) standards; although working at the same speeds they do not have an exact correspondence to V-series standards and are not compatible. A final example of data communicating equipment is the use of an independent computer processor to handle the communications loading of a large computer system. This would typically be used where a mainframe computer is accessed by a large number of remote callers; for example, a system which makes an online bibliographic database available to users on a global basis through the use of modems over a telecommunications network.

A smaller computer—such as a minicomputer—is dedicated to the task of dealing with the network communications, so that the mainframe does not become overloaded by this task, and is free to concentrate on database activity. The front end processor handles the external contact, passing queries back to the database machine and forwarding results to the remote caller. The remote caller is, of course, not aware of the separation of functions.

FIFTEEN

Telecommunications

In the past, telecommunications has largely been regarded as a natural monopoly in the same way as gas and electricity. The long-term nature of the investment needed to create a national infrastructure, and the massive amounts of capital involved, have discouraged private initiative and relied upon governmental support and backing.

During the 1980s, in many countries around the world, this arrangement began to be perceived as too parochial and inflexible; liberalization of telecommunications regulation was seen as a means of stimulating innovation, by encouraging the development of market forces through the introduction of new service suppliers.

The Impact of Liberalisation on Telecommunication

The liberalization of telecommunication has taken two main forms, they are the following:

— *Privatization:* In which the monopoly authority is established as a private corporation, which competes on the same terms as any other business;

— *Deregulation:* In which the new service operators are allowed to offer competitive services, by the removal of restrictive practices which legislate against this.

Many countries have opted for a mixture of these two forms, often with a transition period setting special conditions which alternative services must adhere to. The extent of deregulation has also differed considerably; since the telecommunications infrastructure is seen as so essential to economic development, the total abolition of all regulatory mechanisms is not considered desirable. For this reason, many people speak of re-regulation, stressing a change in the conditions of government control rather than their complete removal.

The United States has historically been different from most countries, as their telecommunications have always been controlled by a private corporation, licensed under government monopoly. The AT&T corporation which provided the telephone systems prior to liberalization was forbidden in a 1956 ruling from entering the computer or information business, in which it would be considered to have an unfair advantage. In this period, two major corporations, AT&T and IBM, dominated the telecommunications and computing industries respectively, but did not directly compete with each other. A Federal Court ruling of 1982 ended this arrangement by requiring the divestiture of AT&T under the corporate anti-trust laws. This resulted in the creation of seven independent regional Bell operating companies out of the former AT&T conglomerate, each of which provided the public telecommunications service in a separate geographical area of the United States.

AT&T itself remained to compete for the traffic between these regions with other long haul private network operators, such as MCI and US Sprint. AT&T was also permitted to enter the computer and information business in its own right which, as the original developer and licensor of the Unix operating system through Bell Laboratories, was an attractive new direction for its corporate development. This 1982 judgment did not, however, end all regulation over the sector. The divested AT&T was prevented from entering 'electronic publishing'—the provision of information services on their networks—for a period of seven years while competitor network suppliers developed their strengths. This restriction was removed in 1989. The regional Bell Operating Companies (RBOCs) are now permitted to offer:

- electronic messaging;
- audiotex;
- electronic gateways to other information services;
- integrated billing for online information through these gateways;

The AT&T corporation, however, acting as the carrier of inter-exchange traffic, is increasingly interested in developing information services through co-operation with the regional operating companies; for example by using national customer databases of names and addresses for credit card and target marketing schemes.

The exact relationship between monopoly and non-monopoly services has been the subject of debate and conflict between the Federal Court and the Federal Communications Commission throughout this period, with the former championing a more restrictive approach. Realizing that the computer industry which telecommunications serves has global interests, both AT&T and the RBOCs have also diversified into foreign markets, both as telecommunications carriers and as providers of online services, following the creation of new opportunities through liberalization abroad. Although the trend has had an uneven development, there has been a steady blurring of the distinction between information provider, information vendor, and telecommunications carrier which will continue in the present decade.

The Impact on Europe

The same global trends have affected the structure and development of the telecommunications and information industries in Europe, but here the influence of the large multinational corporations has been diluted by a more robust government policy concerned to afford some protection to national and regional interests.

Within the European Community (EC) public data networks evolved independently in separate countries, resulting in a smaller overall market for each and a relatively higher cost of equipment than in the United States. This concentration on a national scope has led to the development of somewhat different technology, in particular the installation of different digital switching systems, which reduces the potential for European connectivity as a whole. The European Commission has taken an active interest in the development of the telecommunications facilities within the member countries, where the role of the PTTs has been crucial in the growth from fledgling services to sophisticated international networking.

The EC commissioned the national PTTs to create a packet switching network capable of supporting access from all the member countries, and itself mounted a variety of information services on the network under its own auspices as well as encouraging the development of others. This initiative gave a boost to the developing of information industries in Europe, helping to create a market sector in which information vendors could establish themselves, while the PTTs were engaged in building up their own national packet switching networks to an adequate standard for a reliable international service. Once these had begun operation, Euronet was discontinued in 1984. This period saw the development of Packet SwitchStream in the United Kingdom, and similar ventures in France, the Netherlands, and Germany. These developments were driven through at the national level, with little detailed co-ordination between network suppliers who increasingly came to view each other as business competitors. This has resulted in the duplication of services and multiple international links to the United States and 'the Far East, with the network suppliers effectively competing to capture the market for data traffic entering and leaving Europe. The European Commission has been concerned to develop international research and development programmes, and has used the liberalization programme to stimulate innovation and corporate investment, in order to build a compatible European communication infrastructure capable of supporting the open internal market from 1993.

The creation of a European Telecommunications Standards Institute and the progressive reform of the PTT authorities have pursued explicit actions taken to support the *declared EC objectives* of:

— opening all network services, except the basic phone and data transmission services, to full competition;

— opening the terminal market to competition, within accepted operating standards;
— establishing uniform telecommunications charging across the EC;
— creating European standards for equipment to increase compatibility and remove trade barriers.

To support research and development, the Commission has embarked on programmes intended to develop the economic power of the region through advanced telecommunications technology. The RACE programme is a ten year effort, begun in 1985, to develop the standards and technology for integrated broadband communication through EC funded project teams drawn from industry and the academic world. The less favoured regions of the EC are supported by the STAR programme, which funds in part the development costs of upgrading communications infrastructures in the more rural areas of Greece, Ireland, Portugal, Spain, France and the United Kingdom.

The development of narrowband ISDN, however, has not progressed as fast as considered desirable by the EC, which set itself a target of 5 per cent market penetration by 1993 a decade earlier; in early 1992 there were only five commercial ISDN networks in operation in the EC and it was forecast that only 1 per cent would be achieved, failing to impact small and medium sized enterprises and the residential sector. This slow progress is due to a lack of standardization in ISDN technology, a consequent lack of cheap equipment available for the mass market, and an unwillingness on the part of network operators to take committed action in the face of this uncertainty. Within the EC, individual countries have taken a path of deregulation; rather than of complete privatization. Many former publicly owned PTTs have been transformed into private companies, but often with a majority shareholding still held by the State; and although these compete with private corporations in some areas, most still retain a degree of monopolistic privilege and control which in some cases can be quite substantial.

In the United Kingdom, changes in telecommunications occurred within a more general perspective of a Conservative government free-market philosophy committed to reversal of state monopoly industries in general. The successive Telecommunications Acts of 1981 and 1984 separated telecommunications from the postal authority, establishing British Telecom (now BT) as a private corporation.

At the same time, licences to offer similar services were granted to other companies, with the intention of creating a competitive market which would encourage innovation and efficiency. The principal competitor was Mercury Communications, founded for this purpose by Cable and Wireless; two other licences have also been granted for the establishment of cellular radiophone services. A curiosity of the British situation was the third licensed operator, the city of Kingston upon Hull, which for historical reasons already ran its own local phone system; so successfully, that it was able to provide free local phone calls. This was privatized in 1987 as Kingston Communications, eventually becoming part of a larger corporation. The 1984 Act also created OfTel, the Office of Telecommunications, charged with administering the regulatory structure and overseeing the implementation of the licensing conditions imposed on operators, to protect the public interest and ensure fair trading. These included, for instance, a requirement for Mercury to develop a commitment to geographical coverage of the United Kingdom as well as more specialist business services.

Mercury Communications was originally targeted at large corporations, but since obtaining the right to interconnect with the BT network, and thenceforth the ability to provide complete end-to-end services has extended beyond its limited geographic base to provide data services to medium and smaller businesses. As well as digital phone services, Mercury now provides a packet switched data network under the name of Mercury 5000, and an ISDN service called Premier 2100. In this way, it is now competing with BT in all major areas, and is actively co-operating with foreign companies such as AT&T to provide advanced satellite network services. While Mercury has built an all-digital network from the start but has limited geographic coverage, BT has a completely national service but is only just becoming fully digital in operation. At the same time as upgrading its national network, however, BT has developed international networking capabilities, for instance by the acquisition of the US-based value added network Tymnet, and has developed a range of value added services for delivery on these networks. There are two levels of integrated digital services, ISDN-2 and ISDN-30, with connections to ISDN networks in many other countries. In conjunction with this, BT provides videoconferencing and 'narrowcast' television services for intra-corporate communications. There is no restriction in the UK on the development of information services by the carrier operators, and BT in particular has developed many such services. The Prestel system was an early venture in videotex or viewdata, but has had little long lasting success despite being one of the first examples in this field.

Although successful in niche areas such as travel and insurance, after ten years of operation Prestel had only 95,000 terminals in use, and had failed to fulfil its original objectives of serving a mainly residential domestic market. Business oriented information services have proved more successful in areas such as financial, currency, market, and company information. Despite deriving some benefit from a loose regulatory structure which in practice has shielded BT from some of the effects of direct competition, the lack of a unifying structure under which services are offered, and a general philosophy of 'let the market decide' has retarded the development of a strong direction for telecommunications in the UK.

The relatively low cost of leased lines, for example, favours the building of private networks and encourages third-party suppliers of value added services, while discouraging the take-up of national services such as ISDN. BT must therefore compete on the same terms as these third party suppliers, and increasingly on an international basis. A further round of liberalization has opened up the possibility of other interested *organisations applying for a licence to offer telecommunications services*. These potentially may include:

— cable TV franchises, which have extensive links to customer premises, but no national infrastructure,
— cellular phone services, which are now allowed to build trunk networks, and could offer a packet switched radio-based network;
— recently de-regulated satellite services, able to both receive and transmit from small and cheap terminal devices;
— other national industries, such as British Rail, which have extensive rights of way for laying cable.

Telecommunications Network

The cost of entering the telecommunications market in a competitive age is likely to mean that any of these options would be likely to seek an alliance with a foreign network provider, with particular interest expressed from United States Bell companies.

The relationship between the telecommunications authority and the State has been markedly different in France. The national PTT was established as France Telecom in 1986, a private corporation owned by the State. The organisation largely retained control over equipment approvals policy and operating conditions, able, for instance, to levy a 30 per cent surcharge on leased lines used for value added services by competitors and to prevent certain kinds of services altogether. The Transpac X.25 public packet switching network is claimed to be the largest in the world, and has extended its value added services across Europe, including the recent acquisition of the London Regional Transport network in the United Kingdom. Transpac is already used by some United Kingdom companies for facilities such as electronic data interchange, in preference to leased lines from BT or Mercury, and is expected to be a major competitor in the European open market for telecommunications services.

The best known success of the French PTT has been the Minitel service. This videotex service, in contrast to the British Prestel, is now the dominant videotex system in Europe and has been sold to many other network operators, including NTT in Japan, the Italian PTT, and the Bell company US West.

From its first operation in 1983, Minitel grew to 4 million connected terminals in its first five years, and nearly 6 million in 1992, of which at least half are business users. Videotex is generally seen as providing access to large databases of information, presented through an easy to use interface which is accessible without special knowledge or instruction.

The system was originally developed by telecommunications interests as a suitable information service for linking the telephone with a domestic television set, but is now often supplemented by an intelligent terminal or personal computer.

A major reason for the initial success of Minitel lay in the policy of freely loaning the specially designed Minitel terminals to telephone subscribers, at the same time as replacing the printed telephone directory with an electronic version accessible only by videotex.

The Electronic Directory, provides access to a national database of phone subscribers through a simple enquiry language, as well as Yellow Pages information. The great success of this service very quickly created a substantial user base for Minitel, which encouraged other information suppliers to participate by offering *other information services*. These other services now cover a wide range, aimed at both the business and domestic markets, and including:

1. Electronic messaging; professional services; business and financial information; domestic shopping and banking services; and chat lines.
2. The Pancatalogue comprising the online catalogue of the Bibliotheque Nationale and of the SIBIL user group of 22 universities is accessible via Minitel; the new French national library, due to open in 1995, plans a Minitel link which will enable the user to reserve materials from home, and book a reading place for a particular arrival time.
3. Newspaper publishers offer a package of services including news bulletins; information on sport, entertainment and leisure; classified advertisements and games.
4. Professional level services include access to a range of online databases with sophisticated search interfaces. Services are organized in a kiosque system, in which the user has the available services displayed, but pays only for actual use with no fixed subscription fee. There are seven different levels of charging for information used, priced per hour of usage regardless of distance, ranging from almost nominal fees to typical online database searching rates. The vendor is responsible for mounting the information, and chooses the applicable charging rate. Teletel, the service umbrella under which Minitel is managed, provides access through the Transpac

network, and includes the user charge in the normal telephone bill on behalf of the vendor. An international service is being introduced which interlinks with other national videotex services, and with Teletel hosts abroad.

France Telecom opened its Numeris ISDN network in Brittany in 1987, and established a national network by 1990. The development of suitable applications was encouraged by a partnership scheme in which France Telecom sponsored pilot projects with commercial and client organisations, paying half of the research and development costs.

Many of these applications involved the transmission of high definition images, such as the transfer of X-ray and ultra-sound images between medical sites in emergency cases, the transmission of satellite weather images, and the distribution of high definition house pictures around a network-of estate agents.

In Germany the PTT authority was turned into a private corporation, owned by the State, in 1989. Prior to this, the Deutsches Bundespost had resisted the encroachment of competitor services by keeping prices high for leased lines, charging by volume rather than by fixed rental, and by refusing equipment approval.

The 1989 reform following re-unification created a Bundespost subsidiary Telekom, and allowed private companies to compete in establishing value added services, although it has been reluctant to change this policy. The Deutsches Bundespost established a basic rate ISDN in eight major German cities in 1989, and expanded this to a broader national network, using the promotion of applications as a development device in a more limited way than the French experience. *Telekom currently offers the following services:*

— a videotex system, called BTX, similar to Prestel;
— an electronic mail system, called Telebox;
— a facility for remote monitoring of industrial plant and processes.

The development of the telecommunications infrastructure has stimulated the growth of the information industry, creating a larger potential market for information suppliers by bringing many more people into the communications web, building a more reliable and higher quality service.

The European informatics organisation, Eusidic, organizes an annual monitoring exercise on the reliability of international packet switching services by asking its members to record a typical week of use. This has recorded a steady increase in reliability, showing a 25 per cent failure rate in calls in 1988 compared with a 14 per cent failure rate in 1991; a large proportion of these faults are believed to occur when a dial-up connection through the PSTN telephone line is used as a local link to the packet switching exchange. An increase in, the total number of packet, switching exchanges in many European countries has reduced the importance of this link, and simplified the installation of new network connections. More dependable communications together with a network infrastructure which is capable of offering a connection to all potential customers has led to a substantial growth in the number of databases produced in both Europe and the United States, and in the number of host systems which offer access to them.

By 1990, there were estimated to be 1048 databases produced in EC countries and 2214 produced in the United States, excluding those only available in videotex model; this represented an 11 per cent growth rate per annum in EC countries and a 7 per cent growth rate in the United States. Within the EC the same report counted 227 hosts offering access to a total of 1256 databases, and 299 online hosts giving access to 3348 databases in the United States, although this is counting some databases more than once when they are accessible from several different hosts.

Compared with earlier surveys, this figure indicates that the number of hosts in Europe is growing while those in the United States are remaining stable. A similar growth has taken place in the professional and business database usage sector of the videotex market in Europe, although the trends here are not so marked.

The French Minitel system dominates the area, with just over 85 per cent of the total number of installed terminals across Europe (6.6 million) in that country; some 50 per cent of the estimated 5 million or so Minitel customers are professional users compared to about 80 per cent of the 100,000 Prestel users in the United Kingdom.

Overall the European Commission estimated that there were 15,000 information providers on European videotex in 1991, although once again most of these are mounted on the Minitel system which provides the predominant means of access to professional information services in France.

The distribution of the online market is uneven, with the United Kingdom, France, and Germany accounting for about 70 per cent of the market share, the majority of which is held by the United Kingdom. The large number of potential databases, of which many can often be consulted through competing hosts at varying charging rates and using different command languages as the search interface, is a source of confusion and uncertainty for the end-user.

The diversity of choice has led to a growth in the use of gateways, links through from one system to another. *Gateways in telecommunications networks can take one of three main forms:*

1. Conversion facilities which link dissimilar systems and information types by translating back and forth between the two, typically used for the forwarding of electronic mail between systems.
2. Hosts which have a link through to other relevant hosts in their own

area of interest, so that the second can be contacted directly through the first: the user can drop out to the second host, consult the services available, and then return to the first host on completion.

3. Value added services, which provide a 'front end' to a range of databases on different hosts, to ease and enhance user access by masking detailed differences between hosts in a common presentation and interface.

The use of value added networks has been an established part of the online scene for many years, especially attractive in situations where the reliability of the public telecommunications network cannot be guaranteed. The Dialnet value added network used by Dialog Information Services is an example. This offers regular users a local network connection point with billing for telecommunications costs integrated with that for database usage, as an alternative to dealing separately with international telecommunications vendors.

A third-party gateway service takes this a step further by offering integrated access to many different hosts through one connection.

A well-known example is the Easynet service, which allows users to access more than 900 databases on various hosts; the user may select a database directly or enter the relevant criteria for the service to select automatically the most appropriate one itself. The service offers one command language (the Common Command Language) to all the host systems made available, so that the user is not required to learn new search languages and unfamiliar systems are made more accessible. The Common Command Language reduces the degree of sophistication possible in search statements by imposing commonality, so there are facilities for the experienced user to express the search in the native host language. The Easynet gateway also provides management facilities for monitoring the most effective database usage, and for controlling the costs of use throughout the process. This growth in the online information market in both provision and usage is predicted to continue at increasing rates, with market research forecasts continually revised upwards.

The expanding availability of information services and the almost ubiquitous presence of telecommunications networks is beginning to have a substantial impact not only on the way in which traditional information consumers such as libraries and information units are conducting their activities, but also on the development of new information-related occupations which can take advantage of these new communications features.

The practice of teleworking—the conduct of work-related business and interactions over a distance and through the telecommunications networks: with the employer, with colleagues, with customers and contractors—is one of these areas.

The development of information networking, in both the informal and personal sense as well as a formal substitution of electronically-mediated work relations, is fundamental to the growth of distance working. A major plank, in the telecommunications development programme of the European Community has been the encouragement of teleworking through funding of the necessary infrastructure, as a means of economic development in rural and less-favoured regions of Europe. An example of this is the development initiative conducted in conjunction with BT to provide digital exchanges and line plant for the installation of integrated digital services in the Highlands and Islands region of Scotland in the Hi-ISDN project.

A complementary project, called Hi-VANS, provides an easy to use. interface to national and international services and to the public data network, as well as local information services for the development of tourism and community enterprises. An example of one such enterprise is the use of the network to transmit scanned images of articles which require abstracting from the Elsevier journal publisher in the Netherlands to distance workers in the Highlands. The articles are scanned in the Netherlands from the print version and sent to abstracters in Scotland; the completed abstracts are then returned for publication to the Netherlands with no production delays arising from postal delivery times.

Similar pioneering efforts have seen the foundation of 'telecottages' in many rural areas of the United Kingdom. A 'telecottage' is a generic term for a rural information technology centre, which provides easy access to IT, usually in a public building and on an informal basis. The telecottage can provide formal and informal training, and can act as a catalyst for small enterprises by providing cost-effective access to equipment, as well as supporting the work of community groups.

Apart from public funding, this concept has found sponsorship in the United Kingdom from BT in the form of four Teleservice Centres, established as well-equipped telematic service centres; the first of these to begin activity was in the Shetland Islands in 1991. The willingness of a major telecommunications corporation to become so involved in stimulating the development of a networking infrastructure for information enterprises testifies to the potential possibilities for the future.

In Europe, the concept of teleworking was developed first in Scandinavia, spreading quickly to the United Kingdom, and is now strongly favoured by government and EC initiatives in France, Spain and Portugal. The Telecottage Association, formed to support the growth in teleworking through personal networks of groups and individuals by means of seminars and newsletters, estimates that there are about 590,000 self-employed teleworkers in the United Kingdom. Many of these people are professional workers to whom the production, use and exchange of information is vital to their work.

The recent acceptance of distance working indicates a trend towards more skilled information workers, and the replacement of traditional office work by advisory, and consultancy-based work which does not require full-time attendance at a particular geographical location to be carried out. This is demonstrated by the Telemedicine project in northern Norway, for example, which provides specialist medical services from a central hospital in Tromsoe to isolated local hospitals and health centres through the telecommunications network. This links together services such as pathology, radiology, and cardiology with access to medical databases and patient information systems for remote diagnosis, without requiring the patient to make a difficult journey to the central hospital facilities.

Generally, *teleworking can take a number of different organisational forms:*

- — home-working, in which the workers are based at their own home, making few visits to an employer premises;
- — telecentres, which offer new information technology facilities for workers from various different organisations;
- — satellite work centres, which are owned by a large company but are geographically dispersed;
- — flexible working, in which the work is carried out partly at the organisational premises and partly at home.

These new organisational forms imply changes in the overall structures for, organisational communications, and an increased emphasis on the development of networks which are capable of stretching beyond the physical confines of organisational premises alone.

Private Network

The backbone of the private network is often provided by leased lines rented on the basis of sole use from the local telecommunications authority. The costs and conditions which are imposed on the lease differ quite widely between countries. Where lines are cheaper, as in the United Kingdom for example, a large corporation may now be able to justify a leased line on the basis of voice transmission only, using one Megastream line to replace several analogue phone lines, and be able to transmit data 'for free' on top of this usage since a fixed annual rent is paid. High bandwidth circuits such as these supply point to point connections, and are appropriate between sites which have a high data traffic flow. Dial-up lines may be used where small amounts of data are sent at infrequent intervals, possibly using parts of a public packet switched network.

Since 1989 organisations operating private networks are allowed to resell spare capacity on their bandwidth, so that linking in with an existing private network already in use elsewhere may be one option; for example, banks may make network capacity available to some of their corporate customers. While terrestrial lines are normally controlled by a telecommunications authority, since the liberalization of services, nonterrestrial transmission offers a wider range of choices. Microwave links can be installed in line of sight between tall buildings, and satellite services can be employed to carry audio, video, and data services within a wide geographical area.

In the United States an 'open skies' policy, under which any qualified organisation can apply for a licence to operate a domestic satellite service, has been in operation since 1972; in the United Kingdom, however, the granting of full licences has been more reticent. Satellite capacity is available from a number of specialist data communications companies, and in Europe through the Eutelsat consortium of PTT authorities.

The pricing for Eutelsat usage is differentiated by member country with wide discrepancies in national rates; recently, clients have been able to book capacity through any member offering the best service. The development of very small aperture satellite transmission (VSAT) has lowered the relative costs of this medium, since a receive-only dish uses the same technology as satellite television, and can be installed with the same case without the need for formal planning permission. This can provide a low cost 'point-to-multipoint' data broadcasting, for fast updating of many sites simultaneously; and is also used as a stop-gap measure for establishing a network link quickly, or to link sites which are separated by difficult geographic terrain.

Cellular network radiophones can also be used for non-terrestrial data links, but are not yet ideal for this purpose. This method appears attractive for linking a portable personal computer fitted with integral modem into a corporate network from uncertain locations; however, the cellular channel is noisy, and the transition between cells introduces delays in which contact may easily be lost, preventing truly mobile data communications. Although the transmission of packet data over the VHF radio waveband is feasible, its widespread use has not yet developed in the United Kingdom; licences have only been granted to radio data companies for quite limited and specialized markets. The wide area networking elements also take on an international dimension if the organisation needs to link into sites in other countries.

Larger organisations have diverse information needs, and sometimes specialist but vital performance requirements, demanding the development of a networking strategy which is flexible and controllable. *For information networking to support corporate communications, the organisation has three main choices.*

1. To build a private network, using network elements from different sources.
2. To send data over one or other of the public data networks.

3. To hand over all network management aspects to a third party specializing in this field.

These choices have varying significance for different sizes and types of organisation operating in different fields; these disparate demands will lead to differing perceptions of the technology which is at the leading edge for a particular situation.

Smaller organisations, or those in which the information and technological infrastructure is less developed, may be well satisfied with options which are already considered obsolete by those in which information has a high strategic or mission importance. The use of public telecommunications in one organisation may have as farreaching effects within that context, as are sought by another which is planning frame relay and broadband networking as replacements for X.25 and ISDN.

Some significant *advantages of the use of public data services* are given below:

1. The development of private networks demand the use of scarce personnel with technical expertise in networking skills, and a high capital investment; while using public networks leaves the organisation free to concentrate on its main activities without draining off resources.
2. Once a private network is developed it is hard to change, and expensive to upgrade with technological advances. Information activities may demand a high capacity at critical times, but may not justify this permanently (for example, with seasonal peaks in retail transactions); variable use of the public data network will smooth out peaks and troughs of this nature.
3. The broader scope of public services can develop targeted services for particular needs in a cost effective way, and save on the duplication of effort.
4. In a similar way, organisations can take advantage of the emergence of international standards through participation in a more general service, without the cost of individual development; for example, electronic data interchange standards for basic business documents.

The use of *completely private networks, can have important attractions* especially to a larger organisation which is able to amortise costs over a longer period.

1. The strategic importance of information to the organisation may demand control over the data at all stages, without fear of compromising security and confidentiality.
2. Control of the network enables technological decisions on speed and efficiency, the use of proprietary protocols, and flexible changes in response to need.

3. Control also extends to economic factors, where the unpredictable nature of changes in public services are avoided while spare capacity can be resold to other clients.

The third choice available to organisations is the use of third-party services to take over the management of all network facilities on its behalf. This removes the need for specialist expertise and keeps the technology up to date, while retaining some control and security, by making this the responsibility of an expert networking company. This is an attractive option to a medium sized organisation where communications is important for its business-but not of vital strategic importance—particularly if this involves the use of international connections and the consequent complexity of dealing with different national authorities. Network management companies may offer project management in the development of a network for a specific problem within a defined budget, or full facilities management of all networking arrangements.

Liberalization of telecommunications has encouraged the growth of management services from both private companies and the telecommunications carriers, while the increasing complexity of networking has made their services attractive to the client. The operators have often evolved from organisations which have established their own extensive private networks.

SIXTEEN

Various Networks

Computer networks can be categorized along several dimensions and by differing criteria; however, there is no single set of exclusive categories into which all types can be fitted, and networks which are similar in some ways will display contrasting characteristics in others. A network consists essentially of two kinds of elements:

— the network nodes which are the actual computing end systems, for example larger computer systems, personal workstations, terminals etc;

— the connections between these along which data transmissions travel.

Network nodes can have a greater or lesser degree of autonomy in local processing power and ability to control the communication process. The connections can be implemented in various media, but in some way serve to switch data flows between these network nodes, to pass information from a source to a destination in the most appropriate way. As a first stage, networks can usefully be categorized by the type of switching which is used.

There are three basic types of switching:

— circuit switching,

— message switching;

— packet switching.

Circuit switching occurs when a physical link is made between the two network nodes for the duration of a session, along which the data can flow under the control of the end systems. The connection must be set up before the information exchange can begin, and released once the session is completed; in the interim, the circuit is only available to the two participants even though the channel may have spare capacity which is unused.

Message switching does not establish a physical connection, but relies on the network to deliver a message successfully once sent. The network is used as an intermediary: the source end system delivers the message to the network and then has no further involvement; the network will carry the message to its destination as a discrete process. The source end system is not aware of the details of this process, and does not work interactively with the destination end system. The network may use a broadcast technique to convey the message; or may route it through a series of intermediate staging posts, each of which receives the message, stores it, and forwards it on to the next step when possible.

Packet switching techniques take a more structured approach, dividing an individual message into separate 'packets', each of which can travel independently through the network. The network is responsible for the safe delivery of these packets, creating them from the message dispatched from the source end system at the point of entry to the network, and re-assembling them into the original message at the point of exit for the destination end system. Since the network is dealing with smaller standard size packets in a predetermined way, a faster transmission rate can be achieved. The network is able to make its own decisions on routing, and can change these dynamically as traffic conditions change, with no apparent difference to the end user.

A second characteristic of network types concerns the geographical extent of coverage, and the size in terms of the total number of network nodes which can be feasibly connected. A general classification is to describe networks as either:

— local area networks, which are limited to a single site or building, interconnecting nodes within a single organization;

— wide area networks, which are not geographically limited in scope, and are capable of national- and global-level interconnections.

This distinction arises from the very different technologies from which each of these approaches have grown; however, these categories are beginning to break down, as local area networks become more powerful and can themselves be interlinked to form larger entities. An intermediate category is now emerging, which can potentially provide a link between the first two:

— metropolitan area networks, which exploit the new transmission media such as fibre optic cable, to build city-wide capabilities.

Wide area networks can be usefully seen in two groups, distinguishing public and private systems:

— public telecommunications networks are international in scope, and accessible as a publicly available service;

— private data networks may also be on a national or international scale, but are intended for a distinct user group and not for general public access, although they are likely to make use of services offered by public telecommunications operators.

Each of these network types uses networking technologies which are most appropriate for their purposes. These may be quite unlike each other, employing different techniques, principles and protocols which are not easily compatible. Stemming from these different technologies, there are consequent differences in maximum data transmission rates, number of possible network connections, response times and overloading factors etc.

Network Control

The degree of control which is placed within the network, and the location of this control also arises from these differences in approach. This may be characterized along a continuum, which offers at the extremes:

- — a centralized network, with a single main switch to which all nodes are connected, and through which they communicate with each other;
- — a decentralised peer-to-peer network, in which each node possesses the same status as any other, and can communicate with any other of its peers on an equal basis.

For large-and small-scale networks, there are various compromise points along this continuum where they can be placed; an inter-network can also be constructed using various different topological arrangements which are interconnected at specific gateways-links from one type to the other type.

Terminal Networks

Terminal networks are perhaps the simplest form of computer network, consisting of a number of relatively unintelligent terminals, each of which is linked to the central computer system and not directly to other terminals. These two elements work in a masterslave relationship to each other, in that the central system has much more extensive control of the process than the terminal, even though this may in fact be a more powerful machine, such as a personal computer emulating a terminal. In such centralized systems, the connections between terminals and central computer are arranged in a star topology, with communications routed through the central switch; if this should fail the main communication channel is cut off.

Terminal networks are very commonly used in large organizations that wish to pursue a centralized information policy, holding data and information systems at a single main point and providing widespread access to this. There may be a high level of communication requirements for this arrangement, however, where speed of response is important; for example a lending desk terminal in a branch public library may be required to transmit loan information to a central database and check appropriate records in a very few seconds.

Micro-mainframe Connections

The growth in use of desktop and personal computers has had a significant influence on the development of centralized systems, swaying the balance of control more towards the end user. This has led to attempts at creating a greater integration and cohesion between information systems at an organisational level, linking the departmental/office information into the corporate database level.

Terminal emulation packages allow the end user to utilize a familiar desktop PC as a central terminal without the need to switch equipment, and increasingly to use the computing power of the PC to improve screen presentation rather than simply emulating a dumb terminal. This provides the user with a more familiar and 'friendly' workspace environment. Developments in software compatibility are integrating the standard software packages in use on both central and local systems for mainstream applications such as databases.

Local Area Networks

A local area network is a means of linking together a collection of computer-based devices, within a limited geographical area, on a more or less equal basis. This equality makes the local area network (LAN) radically different from a terminal-based network, because each of the computer systems connected is an independent self-governing entity working in a co-operative relationship. Local area networks fall principally into two types:

— proprietary networks which are used to interconnect a number of computers and their peripherals of the same type, for example a network of IBMPC or Macintosh personal computers;
— standards-based networks which can interlink computer systems and devices whatever their type and specification, through the use of common network technology

For both types of LAN, every device connected to the network must communicate through a uniform hardware interface and software protocol to achieve compatibility. Standardization of the communication process is therefore very important to the concept of local area networks. Since no one method of operation would be suitable for all possible needs, a number of alternative 'standards exist, reflected in the different types of LAN available. A local area network will typically be provided to support:

— resource sharing of expensive hardware resources, for example a colour printer, a high-speed modem, or a CDROM drive accessible to all nodes on the network;
— resource sharing of software and access to specialized facilities;
— file transfer and access to support the sharing of information between different users, and to provide wider access to central information while ensuring consistency by maintaining a single copy of the data;
— messaging and conferencing to enhance organizational communications and co-operative working;

The most commonly found use of proprietary local area networks is to interconnect IBM-PCs within a single site or building. There are broadly two types of PC-based local area network, distinguished by the operating system by which they are managed.

1. In peer-to-peer PC networks all the network nodes interrelate as equals, able to call on the services of other nodes and to offer their own at all times. This is an intuitively natural way of sharing resources, and is cheap and effective to implement for small workgroups.
2. In server-based PC networks, one or more network nodes act as a provider of services for the other client nodes. Clients can only call on the services of a server node, which is dedicated to this purpose; there are three main types of server: a file server dispenses data and program files across the network, either in whole or in part; a print server drives the printing devices and manages the print queues; and a communications server controls external links, such as facsimile machines, or modems.

Standards-based LANs

There are two industry wide standards for more broadly based local area networks, known as Ethernet and Token Ring.

Ethernet is the most widely installed local area network at the current time. This describes a specification for a LAN rather than an actual product; there are many implementations of Ethernet from a range of suppliers.

Ethernet aims to provide a high-speed network which gives all nodes stable and equal access to the network, in such a way as to prevent individual stations interfering with each other; and allowing a transmission to be addressed to individual nodes, a group of nodes, or the whole network community. The intention is that all Ethernet installations are compatible with each other, since they use the same set of protocols. The stations on an Ethernet are each linked directly to a main transmission channel in a bus topology. Stations may not transmit at the same time, but must wait for the channel to be free before sending. Transmissions are broadcast to all nodes on the network, but only utilized by those stations to which the message is addressed.

There is a mechanism for the detection of collisions when stations broadcast together, which requires a station to rebroadcast after a random interval. Thus, all stations are in contention for the channel and cannot be guaranteed a fixed rate of access. This method works well for small to medium size networks, where traffic is in irregular short bursts rather than at regular intervals.

Ethernet therefore suits the typical office environment where user activity follows this pattern. It has a relatively high data rate of up to 10 Megabits per second (Mbps) and is able to support large data file transfers without undue disruption to the user. It does not, however, support voice and video transmission, because the nature of the access medium cannot guarantee the continuity necessary for voice or moving image. There are technical limitations to the size of Ethernet networks, both in length of cabling and in the amount of data traffic which can be supported.

The Token Ring network takes a different approach and as the form chosen by IBM for its own local area network has become more popular. The nodes are connected in a ring configuration, with each node receiving information from its 'upstream' neighbour and re-transmitting it to its nearest 'downstream' neighhour Each node, either originates its own transmission, passes on the data exactly as is received, or removes the data which is addressed to itself.

To control transmission collisions, each node has the opportunity to transmit in turn, governed by the passing of an electronic 'token' around the ring. The Token Ring thus provides for both an egalitarian 'round robin' access, and multiple priority levels for urgent messages. The method is self-governing to quite a high degree, and able to react dynamically to changing conditions.

IEEE Standards

International standards development for local area networks has been co-ordinated by the Institute of Electrical and Electronic Engineers (IEEE) in the United States. This work is intended to develop a framework within which the major types of LAN can be encompassed rather than to specify one single standard for all requirements.

The standards are known by numerical, rather than product names, and cover both Ethernet and Token Ring style networks.

1. The 802.3 standard defines a broadcast bus-based network, which is effectively the same as the Ethernet type.
2. The 802.5 standard defines a token passing ring network similar to the IBM Token Ring local area network.
3. The 802.4 standard describes a token passing bus network, which combines features of the first two types.
4. The 802.6 standard defines a high-speed broadband network, better known as a metropolitan area network.

LANs Interconnections

As the local area network grows in size and complexity the amount of traffic increases, overloading the network and degrading response time and reliability. A number of devices have emerged to support both the partitioning of over-large LANs and the desirable interlinking of others.

1. A repeater carries messages between two sections of a single network,

allowing it to be extended beyond the maximum cable limit.

2. A bridge is a station which is connected to two networks, through which the two LANs can exchange data frames. The networks must be using the same protocols so that they look like one logical network, but only traffic destined for nodes on the other network will pass over.
3. A router is a device for sending data frames between two dissimilar LANs, since the router can deal with both network types.
4. A gateway links a local area network to a different networking environment altogether, such as a wide area network.

The development of standards for local area networks is continuing with the emergence of newer technology and physical media. In particular, the use of fibre optics is having a substantial impact, potentially opening up the geographic and bandwidth limitations in this field to such an extent that distinctions between local and wide area networks are becoming less meaningful.

Wide Area Networks

In the wide area environment, the communications process is seen as a multi-stage process. Networks are constructed from point-to-point links between different nodes on a store and forward basis. Data transmission are passed from one node to the next, *en route*, with each node making its own routing decision based on matching the information which it holds in a routing table with the destination address on the data packet, and in the context of current network traffic conditions.

Control and management of the network is therefore distributed throughout, and not located in any single central point. The network can be extended by adding new connections between nodes and making appropriate updates to routing tables. There are likely to be alternative routes between nodes which can be used at different times. The relaying of data between nodes in this way requires the use of protocols to ensure correct delivery and protect against the corruption of transmissions *en route.*

The protocols normally operate both within the end systems and within the network relay nodes themselves. There are several alternative protocols suites which may be in use; however, all nodes on the network must be running the same protocols to communicate successfully, since this is the glue which binds the disparate systems together.

As an indication, *some common wide area protocols are:*

— the X.25 protocols which have been developed for public telecommunications, and ensure international compatibility in this area; these have also been adopted in some private networks;

— the TCP/IP protocol suite, widely used in private networks and in the academic and research world.

Private Data Networks

Private data networks can be constructed from a number of different elements, as appropriate to requirements and conditions. They, may be national or international in scope, and involve the co-ordination of several different network operators and administrations, using elements from:

— public telecommunications operators, in the form of leased data lines or use of the public data network;
— private network operators, including commercial long-haul network operators and satellite communications suppliers with independent local or wider scale data lines directly owned.

The Leasing Telecommunication Lines

The leasing of dedicated digital data lines is an alternative to use of the public data network. These are private circuits, provided as a carrier by the telecommunications authority, but under the control of the organisation which is renting the line.

Leased lines are point to point connections, permanently linking one site with another, rather than dynamically switchable channels. Since they are leased at a fixed annual rate, the greater the use made of the line the more efficient it becomes.

Multiple rented private circuits are used by large organizations to create private networks, linking remote sites together on a point-to-point basis, or providing a high-speed network backbone for a high volume of traffic between main sites. In the United Kingdom BT provide a range of leased circuits with varying data carrying capacities:

— *Kilostream* and Kilostream N lines provide bulk data transmission at rates from 64 Kbps up to about 1 Megabit per second (Mbps);
— *Megastream* lines provide higher speeds of up to 2 Mbps;
— *Satstream* utilizes satellite capacity to offer point-to-multi-point connections over a wide geographical area.

The use of leased lines is a more secure and flexible option for the larger organization, which is able to deploy the systems and protocols which it finds most suitable, rather than having to conform to those used in public systems.

Telecommunications Networks

There are three types of telecommunications networks presently available for use for computer-based information networking; each of these represents a phase of the developmental efforts to improve services over several decades, using quite different technology and approaches.

Public Switched Telephone Networks

This is an extensive network of cabling reaching into almost all inhabited parts of developed countries, and providing international connections worldwide. This made it attractive as an immediately available option for computer networking in its early stages. The PSTN network is designed for voice communication, and therefore requires the use of an interfacing data communication device, usually a modem, to be used.

The network makes use of circuit switching technology for the same reason, and is a rather noisy and unreliable channel, prone to errors and mis-connection. It is poorly suited to fast and reliable data transfer and depends on the end systems alone to maintain quality through the use of suitable protocols.

The PSTN makes inefficient use of the available carrying capacity because the circuit is idle when neither end system is transmitting but cannot be used by waiting nodes. Since a switched phone line is charged on a time and distance basis, regardless of the amount of traffic, it is not ideal for inter-computer communication over long distances.

Public Data Networks

The public data networks have been constructed specifically for the transmission of computer information, and cannot be used for voice conversations. They provide fast and reliable data transfer, more flexibility and international communications at a reasonable cost, since they are charged for by the amount of data transferred.

The *advantages of packet switching* are the following:

— there is no need for a direct physical connection between sender and receiver;

— information from different calls can then be interleaved on the network channels to make maximum use of the available bandwidth;

— the user, therefore, pays only for information actually transmitted, rather than for elapsed time regardless of the transfer activity;

— standard use of communication protocols facilitates the inter-connection of equipment from different manufacturers, operating at different data rates.

The X.25 protocol set is the standard for public packet switched networks. This is an open standard which is available for end systems as well as internally to the network.

Public data networks are mostly managed on a national basis with international connections between them, and so X.25 forms something approaching a global standard in this area. There are, however, two versions of the X.25 protocol; the latest 1988 version is not used by all public telecommunications operators. The public data network is essentially a store and forward network which operates at speeds fast enough to give the impression of a direct connection to the end user.

The network can be considered as an opaque entity—a 'cloud'—by the end systems which need only to specify an address for the remote end system which is being called, having no concern with the details of routing. The network 'cloud'- consists of a collection of packet switching exchanges (PSEs) which forward packets through the high-speed network from the local exchange to the destination PSE.

Integrated Services Digital Networks

The development of public data networks for digital communications led to the existence of different networks for different purposes-one for voice, one for data-and requiring organizations using these systems to connect to both. A duplication of networks is undesirable, and also works against the integration of information types, for instance the co-ordination of voice, data, and image, which is increasingly needed to support multimedia information networking. Two *development programmes in the telecommunications infrastructure* provided the platform for the development of integrated services digital networks:

1. The digitization of the trunk network lines, the main arteries used by the network operator.
2. The replacement of electro-mechanical local telephone exchanges by digital switches, such as System X and System Y.

The ISDN aims to build a public network providing end-to-end digital connectivity which is capable of supporting a range of digital devices and services, both voice and non-voice, on the same basis. This allows users to send and receive data, text and image from computer systems as well as offering an enhanced telephone service. This is achieved over a single ISDN connection, so that one network may be used for all applications, at the same time as providing higher transmission speeds. The ISDN approach is to enable the existing local telephone circuitry-the 'local loop-to carry digital signals through the use of specialized equipment and adapters, rather than a wholesale and long-term replacement programme.

This requires the installation of suitable network connections in the customer premises, which acts as the interface to the digital network. Despite this fully digital transmission system it is still not possible to connect all computer equipment directly to the network, unless the equipment is designed or adapted to work-with the ISDN system. Access to the ISDN network is available at two different capacity rates, known as basic rate and primary rate.

The access rate refers to the number of logical channels which are available on a single ISDN connection. Each of these channels constitutes a circuit which is independently usable by different information media, and can be switched to different destination points.

The Basic Access Rate provides two *bearer* channels and one *data* channel, and for this reason is otherwise known as 213 + D. Each of the bearer channels can simultaneously carry a separate information stream, at a transfer rate of 64 Kbps—for example, one for a voice call and the other for data.

The Primary Access Rate gives a higher number of bearer channels, although the exact number is variable between countries—the United Kingdom, Europe and Australia supply 30 bearer channels to the user in one ISDN connection (known as SOB + D); in the United States, Japan and South East Asia 23 channels are used. SOB + D gives an overall data rate of some 2 Mbps (the equivalent to a BT Megastream leased line) and necessitates the use of extra equipment at the subscriber premises to manage the multiple channels.

Potential benefits from the use of ISDN are the following:

— all 30 ISDN connections can be switched to different locations;
— clearer calls and more reliable data transfers, with less line noise;
— the ability to send two information streams to the same locations, to support remote co-operative working; for example, two users might discuss over the phone the contents of a computer screen which can be transferred between their terminals;
— enabling domestic services, such as home shopping and banking, online database access, security, and control services for the home;
— feasible desktop videoconferencing;
— providing a better infrastructure for teleworking, for example by providing a gateway to an office LAN for home-based workers.

Frame Relay Networks

Frame relay networks technologies are in development which may threaten to overwhelm ISDN before its full potential is reached. These involve the use of high-speed packet switching on high bandwidth networks which will offer much greater information transmission capabilities than the present public data networks. Frame relay networks are made possible by use of less noisy transmission media, such as fibre optic cable, in the network infrastructure, exploiting this greater reliability to simplify the protocols necessary and thus speed up the transmission process.

Integrated Broadband Networks

Integrated broadband networks are sometimes referred to as broadband ISDN (B-ISDN) systems and will also exploit the advantages of fibre optic to offer a dramatically enhanced communication facility, far more powerful than exists at present and able to deliver a wider range of services more effectively.

The intention of developing integrated broadband communications systems is to provide a single network for all services, which can be flexibly and dynamically configured according to individual circumstances and priorities. Digital networks allow both data, voice and image to be sent along the same connection, but these information types have different desirable transmission characteristics.

Data and text require a high degree of accuracy, since the quality of the received information is more easily affected by corruption, but are more tolerant of short transmission delays. Voice and the moving image, on the other hand, are more tolerant of errors since these can be interpreted by human interpretation, but are more susceptible to variations in transmission delays. Existing or narrowband ISDN circuits do not have a high degree of flexibility in adjusting to these varying demands: for instance, a voice call does not use the full capacity of a 64 Kbps channel, but the spare cannot be freed up for use by others during a session.

Broadband networks aim to overcome such problems by using a universal fast packet switching technique for all information streams, and allocating 'bandwidth on demand' rather than relying on fixed bandwidth channels. Two rates have been proposed for B-ISDN: a basic rate service with a digital access rate of 155 Mbps, and a primary rate service at 622 Mbps, although the full specifications for these have yet to be defined.

A technique known as the Asynchronous Transfer Mode (ATM) is used to transfer the data. This divides the data stream into cells, which are smaller than X.25 packets in size, and can travel through the network at a faster rate. Since the channel is more reliable, quality control is performed at each end of the link and not by each intermediate network switch as well. This ATM approach is particularly suitable for information activity which is inherently 'bursty' in nature.

For example, a service would be able to ensure fast delivery of high definition images, and then free the channel for other use during the reading and thinking time in which the user is not generating any network activity. This will allow the development of true multimedia information services.

Developments in the wider networking field are increasing aiming to deliver faster data transfer rates, on integrated networks which make no distinction between information types.

As the amount of electronic information available across networks grows exponentially, and the sophistication of end-user computing equipment encourages the sharing of information in more intensive forms, network infrastructures respond with the ability to communicate sophisticated forms fast enough to make actual use of the network transparent to the user. At the organizational level, local area networks are also undergoing technical developments which increase their speed and potential for the delivery of more sophisticated and high-powered information applications.

A situation in which these applications can make use of the network 'on demand' in a flexible manner which can respond to changes in circumstance is foreseeable. There is, however, much effort required to make this a reality; the current situation reflects a confusing picture of many alternative networking options which conflict and compete with each other, and which are poorly suited to widespread interconnection between different types of network. The concept of connectivity stresses the importance of this ability to interconnect networks and information networking applications in a reliable and effective manner, and without the need for equipment which is difficult to obtain, install, and use, and soon becomes redundant.

SEVENTEEN

Model Open Systems

The concepts of open systems are now widely accepted in principle, with many believing that the OSI model forms the underlying solution to connectivity problems. In practice, however, the actual take-up of OSI in terms of working products has been slower than at first anticipated. The slow pace of standards development tends to prolong uncertainty amongst end-users, and discourages suppliers from manufacturing OSI products on a large scale due to the lack of a mass market.

The lower network-dependent levels of the OSI model are well established in both local and wide area networks; the upper levels are much less well supported with few general or specialized products in most areas. The second problem facing OSI products is that of conformance testing, to ensure that inter-operability standards are met. These are relatively new, and although normally carried out by an independent third party, do not necessarily guarantee full compatibility. Variations within the paper standard mean that compliant products will not always work with each other, still further undermining user confidence and acceptance.

The interchange technology should allow a user working on one system to use an application on another system, without concern with data formats; the long-term goal is to allow network users to co-operate seamlessly with each other, regardless of physical location or equipment. A common framework for networked information interchange must incorporate generic application problems, such as remote database access, file transfer and interpersonal messaging, in a way that separates the local form of the application from the interchange aspects, enabling flexible variations in the former while standardizing on the latter.

De facto and de jure standards may emerge from two main sources. De facto standards are those which derive from the use and popularity of easily available products, through a process in which the most successful become widely accepted as industry standards. Proprietary standards from single computer manufacturers (such as IBM's Systems Network Architecture) fall into this group; while having the advantage of the backing of a major interest group, their future development is unpredictable, and they are 'closed' systems in that the full specification for the standard is not published but kept as a trade secret. De facto standards may also emerge through adoption by a number of different product developers, and have the general support of a wider range of users who may also participate in the standards making process itself. This type of process is likely to be founded on widespread availability, and may be stimulated by an initial development, or adoption, of the standard in government administration projects. This has been the case with the TCP/IP communications protocols, which are widely used in industry, government and education (particularly in the United States) and form one approach to the 'opening' of computer systems to each other.

De jure standards are those which emerge from an organized, usually public, programme sponsored by governmental or international initiative, emerging from an open debate over several years; by their nature, these developments are more speculative in nature, formulating paper specifications rather than products. The successful take-up of the standard may depend on the willingness of private manufacturers to build conformable products, or on the degree to which the standard can be imposed by regulation.

Various Interconnection Systems

The CCITT recommendations for telecommunications, for example, can be reinforced by regulations governing approval of equipment for connection to public networks. Standards for software in computer systems cannot be regulated in this way, but can be encouraged by making them mandatory in private tenders for substantial government information system projects. The main international efforts towards the specification of a framework for open, intercommunicating, systems has been the work done under the aegis of the International Standards Organization (ISO) and known as Open Systems Interconnection. Although this is a long-term programme intended to take effect slowly, at the present time the ISO Open Systems model has had more influence in Europe than elsewhere.

Open systems interconnection (OSI) aims to develop a framework for the development and approval of computer networking standards that provides communications-based user services which support inter-working with computer systems from different manufacturers, in different geographical locations. This framework defines a conceptual and functional model for supporting multi vendor, multimedia and multi-software networking by formalizing and delineating the problems involved in sustaining it, and structuring them within an open architecture.

The basic idea of OSI is not to invent a completely new set of protocols where adequate ones already exist, or to impose any one single approach or technology, but instead to acknowledge existing protocols by placing them within the framework of the model according to open architecture principles.

In this way, OSI functions as a focus for co-ordinating standards development, and depends on achieving a consensus view from all interested parties.

The International Standards Organization is a worldwide standards agency, based in Geneva, for the establishment of scientific and technical standards. It is composed of representatives from a wide range of other organizations, especially the national standards institutes—for example the British Standards Institute and the US equivalent, ANSI (the American National Standards Institute)— and works in collaboration with other standards bodies, such as the CCITT International Telecommunications Authority, the US-based Institute of Electrical and Electronic Engineers, the European Computer Manufacturers Association, and professional groupings representing the user population.

The organization progresses its work through a web of committees, subcommittees and working groups, developing and adapting standards through lengthy consultation processes. Before being formally adopted as an International Standard, a specification will pass through working paper, draft proposal and draft international standard phases, with public comment and feedback at each stage; ten years is a not untypical timespan for this process. The basic OSI reference model was formally adopted in 1983, and has since been elaborated for particular task types and application areas. *The mission of OSI is described as:*

- to specify a universally applicable logical communications structure;
- to encourage different users to communicate with each other by compatible implementation of communication features;
- to serve as a reference point in the development of new communication services;
- to enable the steady evolution of applications through the flexible incorporation of technical advances and the expansion of user requirements.

The open systems interconnection model does not impose a single standard, but incorporates a collection of existing standards, placing them in the context of an abstract model of the communications process. Compliance with the model is intended to ensure functional equivalence between computer systems.

Functioning of Interconnected System

The term system is here used to denote a collection of computers, software, peripherals, and human operators which have information transfer capabilities as an autonomous whole—that is, a tightly linked computer network as much as a single computer installation. OSI is therefore a standard for network-to-network communications. Given the tremendous diversity of needs and options in this area, a single universal protocol is not possible. The OSI approach is to group functionally similar protocols into a set of layers, each of which represents a particular class of service available in open systems. The basic reference model defines each of these layers by specifying both the functionality of each layer and the interface between adjacent layers. The specification of the functionality within any particular layer defines what is to be done rather than how it should be done. The way in which these functions are implemented in software programs therefore can vary between different products, although the outcomes should remain compatible.

The specification of the interface between layers ensures that the information which is received and passed on between software modules is in a standard form. The functional layers are stacked in a pre-defined order to form organized and structured communication tasks. Each layer provides a set of services to the next highest layer in the stack on demand; to be able to provide these it receives a set of services from the next lowest layer in the stack. Since each layer is providing its own set of specialized tasks (specified as the functionality of the layer), it acts as a user in relation to the layer beneath, requesting services from it but knowing nothing about how they are achieved. Data passes up and down the stack using the standard interfaces as a conduit, with each layer adding control information relevant to its own specialized tasks in the overall communication process, and passing the package on.

An analogy might be a set of Russian dolls, in which the larger doll encapsulates a smaller one. In OSI terms the original nugget of data is passed down and encapsulated by larger control messages as it passes through each layer on the transmitting system; on receipt the process is reversed, and the larger containers are stripped off as the message passes back up the stack of equivalent layers on the receiving system.

The model deals only with the logical structure of the communications process and says nothing about how this is to be physically implemented. Since each layer is acting as a self-contained module within its own boundaries, an individual implementation of layer tasks in any particular system can change without having an overall effect on the communication strategy, as long as it conforms to the logical requirements.

Software modules can then be changed in individual systems to accommodate changes in user needs and network technology, as long as the new version continues to provide the same set of services through a standard interface. In the OSI model the overall communications tasks are partitioned into seven layers, with no overlap between them. Each layer communicates with an equivalent 'peer' layer on the remote system by exchanging messages on matters for which it alone is responsible. This is actually achieved through using the services provided by subordinate layers, but each individual layer works as though it were speaking directly with its peer in the remote system. This peer-to-peer communication is made possible by the standard boundaries between layers which are the same on both linked systems. The seven layers in the model fall into two main sections. The three uppermost layers are primarily concerned with the applications communicating via the linked systems, and with organizing the provision of services to human users.

The three lowest layers are concerned with the technical transmission problems of sending data across networks; the fourth layer (the transport layer) acts as an interface between these two broad tasks. The three lower layers are supported in both the end systems and the network mediating between them; the upper layers are in the end systems only.

The physical layer controls the actual physical connection between the two systems, the hardware standards used, and the method of data transmission, working at the level of the individual data bit. The data link layer provides error handling and flow control, and works at the level of the data frame. The network layer controls both routing decisions and the relaying of data packets through intermediate systems, concealing from the upper layers the particular path through the network which is being used.

The transport layer is the first to be primarily concerned with end-to-end operation, concerned with quality of service and ensuring a reliable connection for the upper layers. There are five different classes of service available from the transport layer, including connection-oriented and connectionless services. The end systems must agree which class is in use for the session before exchange of data can commence. The session layer is responsible for setting up an individual session between two systems, and managing the inter-system dialogue on behalf of the higher layers; for instance, synchronizing the information flows between the applications, and allowing activities to be started, abandoned, paused, and restarted correctly.

The presentation layer negotiates how data is to be presented to the application, dealing with character sets and codes, agreeing to exchange data in a shared common format. End systems which use different information coding sets (such as ASCII or EBCDIC) can use a special open transfer syntax notation called Abstract Syntax Notation 1 (ASN. 1). The application layer, the topmost layer of the OSI stack, is responsible for interfacing with the local end system and ultimately for supplying direct network services to the user. In an open systems environment, all lower level OSI functions can only be accessed through the application layer, and not directly.

The application layer is not a user application in its own right, but rather handles the semantics of the communication process in the same way that the presentation layer handles the syntax. It has specialized elements for different generic applications, such as file transfer and electronic mail, which interface with proprietary packages for these applications installed on the local end system. For example, a person would send and receive electronic mail through the same package which they normally use on their own system; this local package translates the mail into terms acceptable to the OSI e-mail application layer element, which in turn initiates the setting up and management of the appropriate communication characteristics by the next layers, and so on down the stack. For the transfer to be successful, the receiving system must be able to support the same kinds of action, in reverse.

As a logical model of the whole communication process, each layer is effectively split across the two end systems which are exchanging information; the interaction between pairs of layers determines the particular characteristics of an individual connection during the time it takes place. Each layer is able to 'speak' directly to its peer about its own concerns, even though there is only one actual physical connection, controlled by the lowest layer; and even though each layer may in fact be implemented by different products from different manufacturers. There is a basic level of functionality specified for each layer which ensures an essential core of services across all implementations, known as the kernel functional unit. Each layer may also offer extra functions in the shape of additional activities, also grouped together as functional units.

The basic level of service for each layer is mandatory for working in an open-systems environment; additional functional units are named after the type of service they offer, and can be optional, so that individual OSI products may contain only some of those that are possible.

Application layer standards request service from lower layers in the form of functional units. One of the responsibilities of the session layer is to establish whether an equivalent set of functional units to those first requested (including the optional ones) is available on the system which is being called, to negotiate an acceptable common environment within which the exchange can take place. Not all combinations of options within the OSI framework will be compatible with each other. For example, the OSI model includes both a connectionless and a connection-oriented type of service, enabling a local area network to interconnect with an X.25 service for example. However, an upper layer working in connection-oriented mode cannot co-operate with a peer which works only in connectionless mode.

The application layer in the OSI strategy is the one most closely associated with the movement of information, rather than data, between users on systems. The layer supports particular types of applications by defining an application agent for each type. The range is growing as more types are specified, and includes both generic and specialised applications. Generic applications are those which are widely used in many different contexts, such as file transfer, terminal access, and electronic mail, and were the earliest to be defined. Specialized applications are those which are relevant to quite narrow areas of activity; for instance in the library field, these include inter-system communications for information retrieval and interlibrary loan.

File transfer is one of the more difficult problems in data networking, since the way in which files are stored and handled varies considerably between different manufacturers. The File Transfer and Access Method provides a reliable transfer of all or parts of data files from one system to another, even if the file structures and access methods are different, and is guaranteed to result in a correct copy, or a message that the activity has failed with an automatic restart. FTAM also provides the ability to change the characteristics and contents of a file remotely from a different computer system.

In a networked environment, there are many different terminal types from different manufacturers which may potentially wish to connect to an information system, so that the software running on a host system has the problem of recognising and responding to these different terminal types. This can be quite a significant problem, especially with IBM mainframe hosts, and may require the user to operate different terminal emulation programs for different hosts

For example, the ability to use online catalogues across a network can be severely limited by the idiosyncrasies of terminal access and often leads to a 'lowest common denominator' approach to screen presentation design. Ideally, the user should not need to be aware of terminal characteristics in using networked systems.

Virtual Terminal

The Virtual Terminal (VT) standard is designed to allow host computer systems to communicate successfully with a variety of terminals by separating the user application from the terminal display. A virtual terminal software product would therefore be in two parts, one part installed on the host system with the user application, and the other part on the particular terminal or workstation in use, co-operating together through a standard protocol. This allows a user to retain the same 'look and feet' of the interface across different host applications

Job Transfer and Manipulation allows a user to send processing tasks to a different computer without having to know the details of its control language. Although this is not widely used at present, in the long term it is hoped that JTM will support automatic scheduling and sharing of tasks between all available resources on the network, so that this is not apparent to the user—in effect, the user would be to able to use the full combined power of the network as though it were a single computer system.

Remote Database Access supports the ability to query relational databases across a network in a standard manner. Since the relational data model is well defined, there is an underlying consistency to different packages available and a widely used database query language, the Standard Query Language (SQL), which is an ISO standard in its own right. Remote Database Access is broadly compatible with the SOL language.

The concept of the OSI model is to design a framework for the inclusion of many possible standards (for example, the major local area network standards and the X.25 packet switching service) within a common functional structure. Within this structure there are different possible options for building up a communications process from different combinations of elements within and between layers; not all possible combinations are internally feasible or compatible. To simplify the model for use with a particular user context, specialized versions or subsets have been identified by interested user groups for use with their own applications. These are known as functional profiles, a thread of standards running down through all the layers of the model, but not encompassing all facets across the breadth of all levels.

A functional profile may include a limited set of options within itself, and can act as a detailed specification to which computer suppliers can build specialized products with confidence of user acceptance in that particular field. The specification of functional profiles is also a means by which large powerful bodies can influence the rate of take up of OSI standards. One of the first user groups to work out such a functional profile was the UK government, which now makes use of the Government OSI Profile (GOSIP) mandatory in government information system projects wherever it is appropriate.

Equivalent work is being undertaken by the European standards organizations to define compatible profiles in an EC context, and to harmonize these with the GOSIP work. The United States government has taken similar steps by making OSI a feature of federal information projects. The MAP/TOP profile has been defined by General Motors and Boeing for manufacturing and technical office automation.

A similar project is in progress in the library communications field, with close involvement from national libraries and representatives of major national library and information interest groups. The particular nature of bibliographic information, the primary requirements of those who would be earliest users of open systems, and the context of existing bibliographic information systems into which OSI would fit, determine the greater appropriateness of some elements of the model and highlight the need to develop more specialized elements as yet unrepresented.

The National Library of Canada has taken a strong lead in much of this work, which has been progressing slowly over the last decade—Application standards for information retrieval between bibliographic information systems which use different command languages, and for the handling of automated interlibrary loan requests.

The OSI approach has been to analyse, incorporate, and specify the entire networking process in all its aspects, with the hope of defining a *de jure* standard to which future developments will conform. The slow and meticulous pace of this work is not well suited to keeping in step with new developments in technology or new demands from network users.

An alternative to the OSI scheme which has achieved more dominance in recent years are the two protocols known as Transmission Control Protocol and Internet Protocol—which together form the more easily referred to TCP/IP. TCP/IP was originally developed for the United States military computer network ARPANET under government auspices, and has been recommended for government use in information systems projects for many years. It is not a rigorously defined specification in the same way as the OSI model, and does not cover the same breadth of ground; in fact, it is now rather old and unsophisticated as an approach to networking problems.

However, since the specification for the standard is published and in the public domain, it has been widely taken up by commercial manufacturers and others, and in this sense is considered as an 'open standard', the details of which are freely available. Due to this widespread availability and comparatively straightforward approach TCP/IP is in wide use in both wide area and local area networks, and a number of user-oriented applications have been developed and adapted in response to perceived user need. New applications are not required to go through such a formal process of approval as OSI equivalents, so that the time scales between development and installation are much reduced. In this sense, TCP/IP forms a *de facto* standard with well-tested and well-understood applications, which can be put into practice immediately, with a strong appeal to the more pragmatically minded as a flexible working solution to internetworking problems which has evolved heuristically from real needs rather than from paper specifications.

Although it is not without its own problems, which become more obvious as the protocol becomes more widely used, TCP/IP has been strongly championed by some as an alternative to the OSI approach in the so-called 'protocol wars'. The recent astonishing growth of the US-based Internet network, which uses TCP/IP to bind together a very large number of subnetworks, has given some credence to this notion at least in the short term.

The protocols fall into two parts. The Internet Protocol (IP) provides routing services broadly equivalent to (but not matching) the network layer 3 in the OSI model. The Transmission Control Protocol (TCP) deals with the equivalent of OSI transport layer 4 and session layer 5 functions, providing both a connectionless and a connection-oriented service. Neither of these, however, is a complete implementation of the functionality defined in these OSI layers, and does not divide the functions between layers in the same way; and so is not compatible.

There is no OSI-style application layer to provide the interface linking with user applications, and so these functions must be produced separately for individual systems and will vary accordingly. There are, however, a few generic application processes available which perform some equivalent functions to OSI application layer standards for the main user activities which take place on a network.

The *file transfer protocol (ftp)* supports access to and transfer of files between a user and a host system across *a* network, for both text and binary (program) type files, and in some cases bit-mapped image files. From a user's point of view, the protocol consists of a set of commands which reflect a Unix system environment. These allow a user to open a connection to a remote host on a TCP/IP network by giving the IP network address to identify the called system; to read directories and filenames once logged-in to the host; to transfer files from the remote host to the home system; to transfer files in the opposite direction if the host system allows this; and to close the connection once the transaction is completed.

The *Telnet* protocol supports interactive access to information applications on a remote computer by emulating a standard terminal. The telnet application runs on the home system, and disguises some terminal-specific features, allowing a user to log-in to the remote system and use it as though there were a direct connection. The *Simple Mail Transfer Protocol (SMTP)* provides interpersonal messaging between systems running different electronic mail packages, providing a standard format through which messages are generated and by which they can be received.

The OSI model, as a portfolio of standards and an architecture that binds them together, is often judged a more comprehensive and long-term approach, capable of handling a large number of interlinked systems with seamless access to all points, as well as providing better security, back-up, and recovery mechanisms.

The experience gained from design and use of TCP/ IP has fed into the open-systems development process, and the longer time-scales have enabled more considered and participative decision-making. There is some uncertainty about the future of both approaches, with TCP/IP being strongly supported in North America (partly due to its use on the Internet) and OSI most strongly favoured by efforts towards European network integration. There is no straightforward compatibility between the two, and in many situations both sets of protocols can coexist together on the same network. This can happen at three levels:

A dual-stack configuration, where both sets are present together, but do not interconnect. This implies two parallel sets of applications, requiring the user to know which one to select for certain conditions, and complicating network administration and maintenance.

A gateway which links two systems using different protocols. This arrangement can make transparent conversions (e.g. from SMTP to X.400 mail), but some of the power deriving from specialized applications can be lost in the process. The building of hybrid stacks, with OSI top layers over TCP/IP lower layers, or vice versa; such a network can only directly communicate with those who have the same arrangement.

The OSI approach clearly separates user interface issues from network communication issues, so that standardization on the networking component does not prevent the design of different user interfaces for different user groups on the local end system. This movement is founded on the Unix operating system, which provides the one example of a non-proprietary computer operating system that is not bound to any particular make of machine, but is available and used across the range of hardware platforms.

Although not noted for its ease of use, Unix is flexible and adaptable to circumstances, and the core facilities are able to be presented through a range of different user interface shells. Unfortunately, there are a number of different competing versions of Unix, with various interests vying to establish one in particular as the standards version. The Open Software Foundation is an association of computer industry bodies, including IBM, and favouring a version known as Motif; the alternative, Unix International, includes AT&T and champions the Unix System V version. Many of these manufacturers are also shareholders in the X/Open corporation, an international organization attempting to group the confusing variety of open-systems standards under a common Umbrella through the development of a common applications environment.

EIGHTEEN

Future Prospects

A major development in the short- to medium-term future will be a substantial growth in multimedia telecommunications in one form or another. In this way, multimedia information networks will correspond more closely to the experience of communication and information interchange in everyday life, with a flexible mix of sound, picture, and words. The exact nature of this experience, and its effect on information networking in general, remains to be determined.

A major preoccupation of national governments and regional associations is to build the infrastructure for these developments as quickly as possible (and preferably faster than their competitors). In the United States, for example, in an attempt to accelerate the extension of fibre optic cabling to the home and workplace on a massive scale, the Federal Communications Commission (FCC) decided in 1992 to allow the regional Bell holding companies to transmit video programmes over their telecommunications lines, in competition with the cable television companies. This attempt to create a communications infrastructure for the next decade is seen by many Americans as crucial to the continued economic dominance of the United States.

The FCC, for example, places much weight on multimedia devices such as high definition television (HDTV), a digital two-way interactive television set, as the vehicle for this information expansion, aiming at a 15-year deadline for complete HDTV conversion. The technological futurism is the subject of much political controversy and manoeuvring at the current time; it seems likely though, that the new presidential administration will lead out more liberalization to support competition-led innovation and technological advances in communications media, and there is a clear emphasis on the use of technology for economic development in the shaping of public policy statements.

One more immediate result is likely to be the more rapid development of the National Research and Education Network, sponsored for many years by Vice-President Al Gore. The building of NREN services will bring into focus many of the issues on information access and price, and of the emerging adjustments in role of all those involved in the information cycle. In Europe, a similar urgency has fuelled a number of projects sponsored by the European Commission to work towards the creation of integrated broadband communications systems by the latter part of this decade.

The RACE programme, funding research into advanced communications technology and services, indicates by its name some of its underlying objectives, and the ESPRIT programme has initiated the IDOMENEUS project to co-ordinate and improve European efforts in the development of new information environments in the area of interacting open media.

Adopting lessons learned from the French experience in the establishment of the Minitel system, the EC has initiated a number of pilot projects in fields ranging from banking and finance, manufacturing, transport, distribution, and health care; other projects include distributed multimedia publishing, and a broadband network based library for the petroleum and chemical industries, with a 1995 target date for the introduction of commercial services.

Picture Phone

The existence of a technology, however, does not always guarantee its success. It is worth remembering that the Picture Phone was launched by AT&T at the World Fair of 1964. This was a first example of the videophone; in 1992, almost 30 years later, British Telecom launched a low-cost consumer videophone, which can be connected directly to the normal telephone socket; this has yet to make any substantial impact in the market-place.

A development which is likely to be of more significance, though, is the planned availability of audio-visual workstations based on the standard PC. Such a device will make the full integration of data and voice networks within the workplace a reality, and will open up a wider market for multimedia information services. Many of these exist in prototype form within the existing infrastructure, Multimedia electronic mail, for example, is now possible even in a limited text-based medium such as the Internet.

The multimedia information mail extensions (MIME) to the standard Internet mail protocol permit the inclusion of still images, group III fax images, compressed full motion video, and multipart playback sequences of mixed information types. Desktop videoconferencing, the ability to conduct meetings with a geographically remote person each at their own desk, has been established as a feasible communications technology, although so far finding only a niche rather than global market. The experience of this videoconferencing raises some key issues since it is clear that, as with most other information technologies, the technical issues are much less important than the social and organisational ones. Technical advances which have made the idea more feasible include the emergence of high bandwidth ISDN networks, video compression techniques, accepted standards, and powerful graphics-based workstations.

The degree of success in implementing this technology is determined rather by the extent to which it supports the way people work. This involves an understanding of the communication patterns between individuals and groups, and of their information needs and the uses to which information is put; of the way in which contacts with individuals and formal information sources assists the innovative process, and an appreciation of the interdependent balance between formal and informal communication.

Technology has often been used to support formal communication patterns, rather than the informal, and sometimes is intended to replace rather than support either. Multimedia communications systems cannot replace face-to-face contact any more than the telephone (although a similar argument was made against the telephone at its introduction). Thoughtful design may, however, be used to appropriately support existing human communication patterns. The Cruiser Videoconferencing System designed by Bellcore, for example, uses the metaphor of a 'virtual hallway'—floor plan representing the groups of participants—along which a user may navigate, 'cruise' and interact with others in the process. For multimedia multi-user systems such as this, the interface must include some form of social protocol which enables decisions about the nature and style of desired interaction to be taken.

Many issues which will effectively delineate the shape of multimedia networking in the future remain as yet poorly understood and little explored. In some areas, however, there is more immediacy on concerns which are being raised by the substantial growth in network traffic and users, and the potential changes which can be perceived in the near future.

Digital Information

The growing importance of digital information and its availability over networks which demand physical and technical access create barriers, which are increasingly insurmountable for many groups of potential users. While there is no space here to discuss the social role of information systems, there are a few examples of attempts to bypass these barriers in recent developments. The Cleveland FreeNet experiment, for example, is an attempt to increase access to Internet information to a wider audience. The vast amounts of information, and the worldwide nature of the Internet, also bring to the fore issues of intellectual property and copyright surrounding electronic information. Much of the information held on this network is freely available, as long as it is used for non-commercial purposes.

The subsequent use of information taken from computer networks cannot be controlled, however, once it has been downloaded from the original source, it may appear in many diverse guises, be re-partitioned, amalgamated, and reused, taking on a life of its own. Print-based copyright legislation does not translate easily to the electronic environment. A feature of the legislation to found the advanced NREN network requires the provision of charging mechanisms to be put in place, suggesting a trend towards more control over use and users.

The characteristics which make electronic information difficult to define and 'freeze' at one point in time also raise questions on the nature of bibliographic control and mapping of information resources. Electronic documents and journals which can be updated immediately, replacing the original copy, are difficult to define as a bibliographic entity which has a declared and fixed content, which can be cited for later use. The identification of various versions of a logical item becomes a complex (or perhaps impossible) task, as does preservation and archiving as a permanent record of publication. These problems also rebound upon ethical issues involving the quality and accuracy of the document content. Electronic documents may easily be changed, or 'falsified' with changing intellectual fashions in a way that print being fixed at one point in time cannot.

The unstructured and fluid process which allows a network user to make information available without recourse to formal publisher or refereeing process has great advantages in stimulating an innovative and creative dialectic, but at the cost of a large proportion of false, misleading, or incorrect information which must also be sorted through and rejected. Whilst the research community, in which much of this type of activity takes place at present, has evolved a fairly sophisticated code of practical ethics for network co-operation and exchange, it is doubtful that this will survive a broadening of access and the introduction of commercial interests.

The increasing ubiquity of the networking environment is raising questions about the relative responsibilities of those involved in the information transfer process. Traditionally, this process could be described in terms of five distinct roles: the information provider, who provides the original source material; the information host, who makes it available and accessible, 'publishing' in a print or electronic sense; the carrier or distributor, providing a channel for the information; the intermediary, traditional role for the information service professional; and the information user, the final consumer. This model has long ago begun to break down when applied to the networking environment.

The information provider is often the same as the information host; the common carrier, the telecommunications authority, provides information services in its own right; and the information user more often wishes to find and retrieve the information directly. Many people have begun to question the role of the publisher in the electronic environment: since the author can mount a document in its electronic form on a local computer (quite probably the format it was originally composed in) and the network facilities act as a distribution mechanism, it becomes harder to gauge the function of a traditional publisher. Although this is clearly not appropriate for many types of publishing, in some areas it is possible to see advantages.

For instance, the traditional venue for reporting research has been through a journal article in some summarized form; network publishing would allow a reader to optionally download original datasets and intermediate analysis, as well as authorial commentary, directly and without a time delay. The development of network publishing is likely to involve identifiable changing characteristics compared to the traditional activities:

- a growth in publishing, with one source file producing a variety of complementary versions in different formats;
- multimedia electronic publishing, integrating database, text, image, sound, and video;
- the use of information tools to skim, navigate, browse, and utilize these new media products;
- selective and customized publishing targeted at particular groups;
- a linking with the concept of a 'virtual library', where access and materials are not constrained by physical barriers.

The information user and the information provider potentially have a much more direct and interactive relationship, possibly to the extent that the information exchange cannot be completely separated from the communication process itself. This raises a question mark against the received role of the intermediary, as a filter between originator, distributor, and recipient.

As the growth in electronic information begins to demand a new philosophy of the nature and function of information, so also the activities of the information professional must, in this setting at least, change and adapt. No longer able to adopt a custodial role in managing and organizing the materials and artefacts of information, the librarian or information scientist must adopt a facilitating strategy, aimed at creating the conditions for a user to negotiate the information sphere successfully. The librarian no longer 'owns' the information since it is not contained within any physical institution; and it cannot therefore be 'managed' in the same physical sense. The information user, however, must still negotiate an information space, now across a digital network as much as before within the confines of a library, in order to achieve a successful resolution to the information transfer process.

NINETEEN

Production of Computers

In Japan manufacturing productivity is currently growing at the rate of 4.1 per cent a year. France and Germany have manufacturing productivity growth rates of 4.9 and 5.0 per cent a year, respectively. Meanwhile, in the US the rate of growth in manufacturing productivity has fallen sharply. From 1969 to 1973, output per man-hour increased at a compound rate of 2.9 per cent. From 1973 to 1979, the gains dwindled to 1.6 per cent a year. In the US the program to reverse this productivity pattern significantly must rely on the continued development of advanced technology and its application.

Perhaps the most important element in this reliance on innovation is increased factory automation and a growing use of computers and microprocessor technology in manufacturing. Today, we are on the technological and sociological edge of a dramatic increase in the use of computers in our factories. This will have a profound impact on the nation's productivity growth in the next decade. Within the next 10 to 15 years, four evolutionary trends will meet on the factory floor: (1) the increasing power and simplification of computers, (2) a widespread appreciation of the practicality of computerized manufacturing and robotic applications, (3) a new realization of the impact of computers on people and of people on computers, and (4) a growing awareness of the urgent need for manufacturing innovation in our society.

Many computerized factory systems exist today as islands of automation. The immediate task of the scientific and technical communities is to use the increased power and simplicity of computers to link these elements into an integrated system. Making use of low-cost computer hardware to perform more and more jobs will make such an integrated factory system economically viable. Some difficult technical problems remain. We must develop generally accepted, standardized interfaces between computerized design engineering and computerized manufacturing, between individual machines and machining centers, and between computers and the people using them.

We must also refine the present state of application technology and reduce the cost of the factory automated system through the increased use of computers. Advances in factory automation are dependent on advances in computer technology.

Computer Evolution

The technological changes in the computer field during the past several years clearly equal any technological change that has occurred in our society over the past 100 years. The introduction of high-performance, low-cost microprocessor and storage technology has dramatically improved and enhanced the functions and capabilities of computer software and hardware. With today's increased power, manufacturing computer systems can be made more adaptable to the manufacturing environment, thus cutting systems engineering costs and time per installation. The cost of computer hardware itself has been steadily declining. This trend will continue with the introduction of new technologies such as very high-performance microprocessor chips based on very large scale integration (VLSI). On the other hand, the cost of the human and software resources for systems engineering and programming has gone up. In 1955, 85 per cent of the total cost for processing information was hardware; it is estimated that by 1985 hardware will account for only 15 per cent of this total cost. To improve manufacturing productivity we must reverse this trend by optimizing the use of our human resources and taking advantage of the increased computer power and reduced hardware costs. The need to make increased use of available computer power is heightened by the decline in available technical manpower.

The National Science Foundation found the annual growth rate of scientific and R&D personnel between 1954 and 1969 to be 5.9 per cent. Some 5,56,000 employees were involved in technical work in 1969, but the number fell to 5,17,000 in 1973 and then grew to only 6,10,000 by 1980s—a rate of only 2.8 per cent annually. The increased capability of computer technology and broad availability of application software packages, including the expanded use of problem-oriented languages and database software, can reduce the cost of program development and maintenance by a factor of 10.

The graphics capabilities of engineering computers can link integrated design and drafting systems to manufacturing systems. With these "user-friendly" approaches to provide a bridge between man and computer, the user can interact directly with the manufacturing system without traditional interfaces and jargon-heavy manuals. Computers and the applications they support can now be considered for functions or activities that only recently were impractical. The opportunity for the application of this technology, for all practical purposes, is unlimited.

Manufacturing Automation

Advances in the two primary elements of factory computerization-computer-aided design (CAD) and computer-aided manufacturing (CAM) will create a new industrial revolution. By integrating design with manufacturing, we can not only turn out new product designs much faster, but also program the computer to make sure the designs provide quality and reliability as well as the lowest possible manufacturing costs. CAD/ CAM is the integrated use of advanced computer technology in engineering and manufacturing. It is a common database of part and product geometry and related information which makes it easier to translate a creative idea into a final product at a reduced cost. With CAD, a user can define a part shape, analyse stresses and other factors, check mechanical actions, and automatically produce engineering drawings from a graphics terminal. When CAD is combined with the CAM system, the user can also manipulate nongraphic data such as bills of material, shop information, and cost factors.

The end result is greater design flexibility and what is referred to as designing to cost. The CAD functions can be grouped in four categories: design and geometric modelling, engineering analysis, kinematics, and drafting.

— In *design and geometric* modelling, the designer describes the shape of a structure with a geometric model constructed graphically on a cathode-ray tube. The computer converts this picture into a mathematical model, which is stored in the computer database for later use. Many other design functions depend heavily on the model. It can, for example, be used to create a finite-element model for stress analysis, serve as input for automated drafting to make a drawing, or be used to create numerical control tapes for the factory.

— After the geometric model has been created, the engineer can easily calculate such things as weight, volume, surface area, moment of inertia, or center of gravity of a part. But the most powerful method of analysing a structure is probably finite-element analysis. In this technique, the structure is broken down into a network of simple elements that the computer uses to determine stresses, deflections, and other structural characteristics. The designer can see how the structure will behave before it is built and can modify it without building costly physical models and prototypes. This procedure can be expanded to a complete systems model, and the operation of a product can be simulated.

— With *computer kinematics*, the user can examine the effects of moving parts on other parts of the structure or design and can analyse more complex mechanisms.

— Finally, the CAD system automatically drafts drawings for use in manufacturing.

Computer-aided design is a good example of the transition of expensive, state-of-the-art computer technology to a commercial, economically justifiable system. Recent advances in CAD technology have increased the productivity and effectiveness of design engineering groups. Such systems will be even more common in the next five to ten years. Manufacturing groups can draw on the geometric and numerically coded description produced by CAD to create numerical control tapes, which allow direct computer control of shop machines, determine process plans and scheduling, instruct robots, computerize testing, and in general improve the management of plant operations.

Computer-aided manufacturing has five main functions: tool design, machine control, process and materials planning, robotics, and factory management.

— Manufacturing engineering and *tool design* deals with the machines and fixtures needed to make a new product. In effect, the set of machines, tooling, and fixtures is a new product, and all the techniques of CAD are used in fashioning it. The CAD techniques are then used to simulate plant operation and the integration of machines and materials handling.

— *Machine automation* consists of a chain of increasingly sophisticated control techniques. At the lower end of the spectrum are fixed automation with relays or cams and programmable controllers, where relays have been replaced by electronics. Moving up the spectrum, numerical control (NC) refers to controlling a machine with prerecorded, numerically coded information to fabricate a part. In this case, the machine is hardwired and not readily reprogrammed. In computer numerical control (CNC) the machine is directly controlled by a mini-computer, which stores the machining instructions as software that is relatively easy to reprogram. Because of the computer control, CNC has the advantages of much higher storage capability and increased flexibility. Virtually all numerical control is computer-based, yet only ten years ago CNC was an expensive exception.

— *Process planning* considers the detailed sequence of production steps from start to finish. The process plan describes the state of the workpiece at each work station. An important element in process planning is group technology, in which similar parts are organized into families to allow standardized fabrication steps; this permits significant savings by avoiding duplicate tooling and systems engineering. Most automated process-planning systems use a retrieval technique based on part families and existing databases for standard tool in and fabrication processes. Materials planning or manufacturing resource planning is concerned with the precise flow and timing of manpower, materials, and processes; it is a detailed look at how everything comes together. The ultimate goal is to have continuous use of all production equipment, no bottlenecks, and a minimum inventory.

— Because they are widely applicable, *robots* have a distinct advantage over specialized, highly engineered manufacturing systems. The economic advantage of a mass-produced, readily adaptable robot over a one-of-a-kind system with a great deal of engineering content is obvious. Robots are now being used to perform materials -handling functions in CAM systems. They can select and position tools and workpieces for NC or CNC tools, operate tools such as drills and welders, or perform test and inspection functions. Through visual or tactile sensors, the robot can manipulate objects. Through its computer intelligence, it can inspect the object and provide the machine with corrective feedback or actually reprogram the machine or change the tooling.

— *Factory management* coordinates the operations of an entire plant. Factory management systems tie together individual machine tools, test stations, robots, and materials-handling systems into manufacturing cells and the cells into an integrated whole. An integrated CAM system of this sort is usually hierarchical, with Microprocessors handling specific machining functions or robot operation, middle-level computers controlling the operation and work scheduling of one or more manufacturing cells, and a large central computer controlling the overall system.

Reliability is greatly improved by structuring the control system correctly. Local, distributed control (with defined responsibilities) reports up to a supervisory control that, in turn, is linked to a managerial computer. This parallels the structure of the typical industrial organization.

Ultimately, the digital output from the CAD computer will be simply plugged into the CAM system to reprogram the plant's manufacturing computers. In such an integrated system, the databases will be organized in a way that avoids redundancy and reformatting of information. And any change in one part of the system will automatically revise dependent or related information in other parts of the system. Bridging the CAM and CAD systems will be one of our major jobs in the future. A fundamental difference that has to be reconciled is that CAD makes use of a pictorial, graphics oriented computer database, while CAM involves a great deal of text-oriented information. In other words, we need to find a way for the computer doing the drawing to speak the same language as the computer directing the manufacturing plant. Layering is one way to link these systems.

Layering is a particular technique for structuring the CAD and CAM databases. It enables various people to input data without losing control of the overall design and manufacturing process. Equally important, it enables shop people to see information that is meaningful to them without having to sort through and understand the rest of the information that is normally included in a drawing. To do this, all information is organized in an arrangement resembling layers, or slices, inside the database. The engineer or users in other departments of an organization can provide pertinent information or examine any or all layers of information according to their particular needs. As an example, a printed circuit board may have 250 to 300 layers of information.

A manufacturing engineer inputs layers of information that deal with fabrication and assembly. In turn, machine operators concerned with the details of the drilling and cutting configuration may access layers dealing with this part of the drawing. Other layers provide information pertinent to the needs of the purchasing department or component assemblers. Another major effort to integrate computer systems is an Air Force program called ICAM (integrated computer-aided manufacturing). This is a practical attempt to greatly shorten the time span for the implementation of compatible and standardized computer-manufacturing techniques and to provide a unified direction for industry.

The ICAM program provides seed money for the establishment, within private industry, of modular subsystems designed to computerize and tie together various phases of design, fabrication, and distribution processes and their associated management hierarchy. As appropriate, these mutually compatible modules will be combined to demonstrate a comprehensive control and management package capable of continual adjustment as production needs and the state of the art change.

The *ICAM program is divided into five major parts.*

- — Defining the manufacturing architecture. This permits a concentration on problems of generic scope and wide applicability as the basis for later projects in integration, support, and application systems and demonstrations.
- — Developing integration methodology. This activity provides a bridge between industry and ICAM for the transfer of ICAM technology for the integrated factory of tomorrow. The projects addressed include establishing factory simulation techniques, ICAM implementation techniques, configuration management, modelling tools, software integration simulation, automated systems engineering methodology, and various system analysis and design capabilities.
- — Establishing support systems. This is concerned with the portion of the ICAM system involving computer operations, including both software and hardware and both operational and managerial aspects

of computerized manufacturing.

- Establishing application systems. This includes such items as manufacturing cost and design guides, the design manufacturing interface, manufacturing standards, group technology concepts, and scheduling and process planning. Under an ICAM contract, the National Bureau of Standards considered standards in computer communications, languages, and networks to identify potential conflicts within an integrated manufacturing environment. Other areas of concern include robotics, prototype integrated production cells, integrated materials-handling and storage systems, and integrated manufacturing control and material management.
- Demonstrating the ICAM program. The ultimate goal in ICAM is the use of totally integrated manufacturing systems by industry in the completely automated factory.

Robotics

The problems encountered in trying to integrate advanced computer concepts, new manufacturing technologies, and the shop floor are clearly evident in the evolution of robotics. Robots are classified according to the way we provide them with information and the amount of self-adaptability they possess. The most *comprehensive categorization of robots* is provided by Japan's Industrial Robot Association:

- Manual manipulators are worked by an operator.
- A fixed-sequence robot has a manipulator that repetitively performs successive steps of an operation according to a predetermined sequence which cannot be easily changed.
- A variable-sequence robot is similar to the fixed sequence robot except that the set information can be easily changed.
- Playback robot reproduces, from its computer memory, operations that were originally executed under human control.
- An NC robot is a manipulator, whose tasks are programmed by using numerical control tapes or cards.
- An intelligent robot, using sensory perception, detects changes in the work environment and proceeds accordingly, using its decision-making capability.

Industry today is focusing on the development of NC and intelligent robots. There has been a growth in the number of firms that manufacture and sell such robots which is reminiscent of the proliferation of minicomputer companies in the 1960s.

Basically, robots are microprocessor-controlled mechanical devices that perform a function or provide an intelligent interface between machines and processes. They can be intelligent enough to make on-the-spot manufacturing "decisions." But for robots to become practical, we must reduce their size, mechanical complexity, and installed cost—primarily through the expanded use of computer and control technology.

Robots can duplicate human manipulative skills with accuracy and precision. Their flexibility and versatility, as opposed to hard automation, make robots ideally suited to the kinds of small batch jobs that constitute the bulk of industry's manufacturing activity. Today, robots, are freeing people from jobs that present serious health hazards, are mundane, or are highly repetitive. In most cases their use is justified for non-economic reasons. In the US, industry has been slow to adopt robotics. This reluctance appears to be due primarily to the large initial investment and the general availability of relatively inexpensive manual labour. Why install a $100,000 or $150,000 robot to perform a $25,000-a-year job? Also, the majority of today's robots are monsters: bulky, unwieldy mixtures of hydraulic and mechanical contraptions with a machine tool heritage.

This situation is changing. Robots are becoming more streamlined, and, when they are manufactured in large quantities, will rapidly decline in cost. Many technologically innovative firms are entering the business. Equally important, system engineering, which represents as much as two-thirds-of the cost of a robotic application, is being greatly reduced. It is not difficult to imagine that in a short time the cost of a typical robotic system will be paid back in one or two years.

In the next decade the cost of a robot is likely to be down to $10,000 to $20,000, while skilled labour costs might easily be $25 or $30 an hour. When this economic threshold is reached, there will be a virtual flood of robotic applications. When this happens, robots will play an important part in the totally integrated factory of the future. Most of our plants will have a direct numerical control supervisory computer that coordinates the activities of several NC and CNC machines or hardwired machining centers and robots and connects all the machines into a system. The robot interface will handle the transfer of material and, with newly developed sensory capabilities, will also perform the in-line inspection of parts. For the near future, however, our factories will be some particular mix of machines, robots, and people that makes the most economic sense. Eventually in a decade or so-robots will fill a void in the supply of skilled labour. There has been a shift in the labour force from blue-collar to white-collar workers and from production jobs to service jobs. Currently, about two-thirds of our work force and 85 per cent of all college graduates are employed in a service-related activity.

The total service-oriented labour force is expected to increase by 20 per cent in the 1980s—to about 85 million people. This shift to a service economy, coupled with the slowdown in the growth of the US population, suggests that many businesses will find factory labour in short supply: this is already the case in Sweden today. At a leading university in Japan, there is a robot with human-type hands and legs, TV-camera eyes, artificial ears and mouth, and touch and joint sensing. These technologies are combined to provide the robot with some of the capabilities of a two or three year old child.

For example, when ordered to fetch an item in the room, the robot looks around the room and finds the article, walks to it, picks it up, and brings it back. If the robot does not understand a command, it speaks up. Such robots are essentially showcase examples: they are not appropriate for the majority of industrial applications. In fact, a universal person-like robot would make little sense except for very limited, specialized applications. At present, the major applications for robotics are in arc welding and material transfer. In the future, the major application will be in assembly. Artificial intelligence is a worthwhile goal,, but for the moment industry has more than enough applications for "dumb" robots. Employment of robots in these applications is actually limited by the extensive engineering required to put them to use. It is estimated that by 1990 two-thirds of the robots sold to industry will be off-the-shelf, modular units rather than specially designed systems. Looking at automated systems generally, by the end of the decade, 30 per cent of our systems will consist of hard automation, about 20 per cent will be adaptive control, and the remaining 50 per cent will be systems of a universal programmable nature. Respondents to a study conducted by the Society of Manufacturing Engineers ranked the technical and performance barriers constraining the rapid utilization of robots in US industry.

The leading technical barriers were mechanical manipulation, vision systems, tactile systems, sensory systems, programming, and control systems. The primary performance barriers included accuracy, speed, and the ratio of capacity to size. Robots will have a much greater impact in manufacturing when their total installed cost is reduced and they are more easily programmed. As with CAD and CAM systems, the most difficult thing about putting in a robot is interfacing it with the factory—both the machines and the people.

Research today is focused on the development of (1) equipment that will make greater use of computer technology to cut the cost of systems engineering and power electronics, improve servo-motor technology, and rapidly move from hydraulics to electric; and (2) sensors that will enable robots to perform more reliably and with greater precision and adaptability.

A good *overview of sensor types* has been provided by Bejczy:

> The nonvisual sensor information is used in controlling the physical contact or near contact of the mechanical arm/hand with objects in the environment. It is obtained from proximity, force-torque, and

> touch-slip sensors integrated with the mechanical hand. These sensors provide the information needed to perform terminal orientation and dynamic compliance control with fine manipulator motions ...
>
> Terminal orientation and dynamic compliance control are essential and intricate elements of manipulation. Soft and adaptive grasp of objects, gentle load transfer in emplacing objects, assembling or disassembling parts with narrow tolerances, and performing geometrically and dynamically constrained motions (like opening or closing a latch or fitting two parts together) are typical examples of manipulator control problems that challenge both sensor and control engineering ...

Vision systems close the control loop and allow the robot to interact in a dynamic, changing environment. A second use of vision will be for the critical inspection of the batch-produced parts. The position and orientation of the part can be used in advanced automation systems to direct a robot manipulator to pick up the part for an assembly or transfer operation.

The Robotics Laboratory at Westinghouse is working on state-of-the-art applications in many of these areas. Systems and development engineers are working on the integration of controls, tooling, processes, computer, and other elements of the automated factory. Specialists in robotics are concentrating, on developing and applying high-speed vision systems, tactile and force feedback sensors, high-performance electric servo systems, adaptable programmable assembly techniques and computer control, and artificial intelligence.

One of these projects is called APAS (automated programmable assembly system). Funded in part by the National Science Foundation, APAS is a pilot program in which robots are used to assemble components into the end bells of the Westinghouse line of fractional horsepower motors. It is a development project intended to transfer newly developed technology to the factory. The fractional-horse power motors are currently assembled in batches averaging 600, units at a time. There is a 20-second assembly time per motor to put together 30 different parts, and there are 13 changeovers a day to handle 450 different motor styles. The first section of the line puts parts on the motor end bells. To start the 15-second sub-assembly operation, a vision system in conjunction with a five-axis PUMA (programmable universal machine for assembly) robot inspects the end bells to make sure they are the style currently being assembled.

The end bell is then oriented and placed on a pallet. The next step in the assembly is the insertion of the uppermost components: a thrust washer, a bearing cap, and a felt washer. An auto-place robot picks up the parts and loads them onto an anvil. The end bell is moved into the station, and the parts are passed on. At the same time a semisolid lubricant is injected. At the next station, four screws, a plastic plug, and a contact point are inserted. Following this there is a complicated assembly procedure. In order to assemble all the different styles of end bells, several styles of certain parts are required. More precisely, six styles of mounting rings and three styles of dust caps are needed. Programmable feeders are used to accomplish the feeding and orienting of all these parts. At this station, a PUMA robot picks up a mounting ring, dust cap, and felt washer on an oil finger from the programmable feeder and places them on an insertion device, where they are fitted onto the end bell.

At the final station, a vision system and a PUMA robot are used to perform the final inspection of the end bell, pick it up from the pallet, and remove it from the system. The computer control arid sensory parts of APAS are the most revolutionary elements of the application. A distributed microprocessor system is needed to handle many simultaneous tasks in the short 15-second assembly time, and a visual sensory system is required to provide orientation and feedback.

One master microprocessor controls the entire system. Under its control are three types of smaller microprocessors for Vision control, local process control, and robot path control. All these controllers work in conjunction with the master computer to coordinate the inspection and assembly procedures. The vision system on this project recognizes randomly oriented parts on an assembly line after a multipass learning cycle controlled by an operator. It can also be used to rotate the part to any given angle.

Fostering Innovation

The technology exists in many parts of the world to achieve significant advances in many areas of factory automation. Many European countries are working on numerical control, process planning, and group technology approaches rather than the graphics approach emphasized in the US. They are well advanced in integrating CAD and CAM systems. In Japan, major advances are achieved through the efforts of the Ministry of International Trade and Industry (MITI). For example, hundreds of millions of dollars and some of the finest minds in business and universities are being applied to the task of developing a fifth-generation computer.

Another example is a $60 million government funded project to develop a flexible machine system. This system will use high-energy lasers to manufacture small batches of machined parts with assembly line efficiency. The project involves more than 500 engineers from 20 Japanese companies, and it could revolutionize much of manufacturing. MITI makes use of Japan's homogeneity and organizational milieu. It would be inappropriate for the US to adopt the same methods to foster manufacturing innovation.

However, we must recognize the urgency of the problem. The US is one of the few major industrialized nations in the world without a significant coordinated industry – government - university program directed at improving manufacturing technology. We need a national strategy for productivity improvement that brings together government, business, labour, and academia in a cooperative, rather than adversary, relationship. We will have to remove many of the disincentives to innovation and find new ways to capitalize on our diversity and our proven creative and inventive abilities. Just as MITI capitalizes on Japan's homogeneity, we must find new ways to foster, encourage, and channel our innovative diversity.

At present, Westinghouse is working with the National Science Foundation and the universities of Rhode Island, Florida, and Wisconsin on technology development programs. Along with the Robotics Institute of Carnegie-Mellon University, we are developing "seeing," "feeling," and "thinking" robotic systems for several of our factories. We are also very interested in the Air Force's ICAM program to coordinate sophisticated design and manufacturing techniques now used by industry on a piecemeal basis. This program attempts to integrate design, analysis, fabrication, materials handling, and inspection and to develop hardware and software demonstration manufacturing cells in selected aerospace plants.

Sociological Impact of Automation

The easier it is for people to use computers, the broader will be their applications in manufacturing. We are moving away from the airplane cockpit approach, with rows of complex devices, to create simple computer tools. Compare the ease with which we use personal computers today with the way we approached computers in the early 1970s.

The same changes will occur in manufacturing. The computer has three language levels: machine language, programming language such as COBOL or FORTRAN, and user interface or problem-oriented language. Designing the computer so that it can be quickly used by someone familiar with a problem-but not with computers-is the most difficult of all programming tasks. Today, computers "converse" with users in pictures, in ladder diagrams, or in the secretarial language of word processing.

Voice recognition systems will free workers' hands to perform other tasks and make it even easier for them to use computers. In the future, when research efforts begin to pay off in systems with some understanding of natural language and with "common sense," communications with computers may become as simple as talking to a three-year-old. An important area of concern in manufacturing is worker safety. Automated systems and robots must be able to work side by side with humans.

Major developments are taking place in the sensor area, particularly with proximity detectors and ultrasonic sensors, to make robots more suitable for inclusion in existing factories. The application of automated systems in manufacturing will have several major affects on the people involved in production. It will make our jobs more interesting and challenging; it will enhance job security; and it will multiply the productivity increases. Workers today are looking for greater job satisfaction through greater involvement and increased sophistication. New technologies provide this added dimension to the workplace. For instance, draftsmen use CAD today to perform work that was normally performed by engineers just five or ten years ago. Engineers, in turn, are freed to delve into even more technically sophisticated areas. As a peripheral advantage, the critical need for technical manpower is partially satisfied.

To manage technological change, we must manage our human resources better. For example, we must commit ourselves to ensuring that none of our workers is laid off because of technological changes, as long as they are willing to be retrained and accept new job assignments. Our experience at Westinghouse has been that employees displaced by robots normally move up to better, more challenging work. We should also rethink who can do what job.

For instance, there is a tremendous potential for productivity improvement if the person who knows the machine better than anyone else also has the skills to program the machine while it is working on another job and the skills to debug the programs on the spot. In Westinghouse, by putting the programs for people first, we expect to multiply the productivity improvements that are gained through technology and capital investments. With participative management, for instance, employees welcome advanced technology because they feel in charge of it. Only when these programs are in place will we emerge from the showcase and token automation phase that manufacturing is presently in.

Today, scientific work in the application of computers to factory automation is in the embryonic stage. We are on the verge of seeing the cost of NC, CNC, and robotics become low enough for these systems to be economically justifiable for many more applications. The cost will continue to decline as application problems are resolved and the computer becomes an understood and respected partner in the manufacturing environment.

First came numerically controlled machine tools, which performed their operations automatically according to coded instructions on paper or Mylar tape. Then came computer-aided manufacturing, or CAD/ CAM, which replaced the drafting board with the CRT screen and the numerical control tape with the computer. The new systems integrate all these elements. They consist of computer-controlled machining centers that sculpt complicated metal parts at high speed and with great reliability, robots that handle the parts, and remotely guided carts that deliver materials.

The components are linked by electronic controls that dictate what will happen at each stage of the manufacturing sequence, even automatically replacing worn-out or broken drill bits and other implements. Measured against some of the machinery they replace, flexible manufacturing systems seem expensive. A full-scale system, encompassing computer controls, five or more machining centers, and the accompanying transfer robots, can cost $25 million. Even a rudimentary system built around a single machine tool-say, a computer-controlled turning center-might cost about $3,25,000, while a conventional numerically-controlled turning tool would cost only about $175,000.

But the direct comparison is a poor guide to the economies flexible automation offers, even taking into account the phenomenal productivity gains and asset utilization rates that come with virtually unmanned round-the-clock operation. Because an FMS can be instantly reprogrammed to make new, parts or products, a single system can replace several different conventional machining lines, yielding huge savings in capital investment and plant size. Flexible automation's greatest potential for radical change lies in its capacity to manufacture goods cheaply in small volumes.

Since the era, of Henry Ford, the unchallenged low cost production system has been Detroitstyle "hard" automation that stamps out look alike parts in huge volume. There is little flexibility in hard automation's transfer lines, which get their name from the transfer of the product being worked on via a conveyor from one metalworking machine to another. But such mass, production is shrinking in importance compared with "batch production" in lots of anywhere from several thousand to one. Seventy-five percent of all machined parts today are produced in batches of 50 or fewer.

Many assembled products too, ranging from airplanes and tractors to office desks and large computers, are made in batches. Even such stalwarts of inflexible mass I production as the automakers are developing systems to produce more low-volume models for small market segments. In the past, batch manufacturing required machines dedicated to a single task. These machines had to be either rebuilt or replaced at the time of product change. Flexible manufacturing brings a degree of diversity to manufacturing never before available. Different products can be made on the same line at will. General Electric, for instance, uses flexible automation to make 2000 different versions of its basic electric meter at its Somersworth, New Hampshire, plant with total output of more than I million meters a year. The strategic implications for the manufacturer are truly staggering. Under hard automation the greatest economies were realized only at the most massive scales. But flexible automation makes similar economies available at a wide range of scales.

A flexible automation system can turn out a small batch or even a single copy of a product as efficiently as a production line designed to turn out a million identical items. Enthusiasts of flexible automation refer to this capability as "economy of scope." Economy of scope shatters the tenets of conventional manufacturing. There is no long trip down the learning curve on the factory floor, thanks to the unprecedented precision the system brings to each step of the manufacturing process, from machining to inspection.

The manufacturer will be able to meet a far greater array of market needs, including quick-changing ones seeing the needs of markets the company is not in now. He can keep up with changing fashions in the marketplace-or set them himself by updating his product or launching a new one. He has many more options for building a new plant: FMS frees manufacturers from the tyranny of large-scale investments in hard automation, allowing construction of smaller plants closer to markets. Flexible manufacturing is the ultimate entrepreneurial system: it will allow fast-thinking manufacturers to move swiftly into brand-new fields and to leave them just as swiftly if need be-at the expense of less agile older producers. As the new tools come increasingly into use, "some companies will find themselves blind-sided by competitors they never imagined existed," says Joseph D. Romano, a vice-president at A. T. Kearney Inc., management consultants. Flexible manufacturing systems were developed in the US more than ten years ago by Cincinnati Milacron, Kearney and Trecker, and White Consolidated. The US remains a world leader in the technology: the major machine tool builders are being joined by new suppliers with great financial resources and technical abilities, such as GE, Westinghouse, and Bendix.

The joint venture will bring together GM's considerable capabilities in design and software and Fanuc's expertise in building and applying robot systems. GMF plans to start building products next year. However, most of the action in flexible automation is now in Japan, and both American and European manufacturers will soon start feeling the pressure. Like many other manufacturing technologies conceived in the US – among them numerically-controlled machine tools and industrial robots-the FMS was greeted with a yawn by US manufacturers.

The Japanese have become the implementers par excellence of this new type of factory automation, not because they are great technical innovators, which they admit they are not, but because they have moved fast in putting the new systems into their factories. Once again, the path to success in a new manufacturing method leads through those Japanese factories set up as spotless little towns with flower beds and tree-lined streets. A visitor to Japan these days finds the new manufacturing system turning out parts for machine tools in Nagoya, electric motors near Mount Fuji, diesel cylinder blocks in Niigata, and many other products elsewhere. In most cases these plants run on three shifts. During the day skeleton crews work with the machines. At night the robots and the machines work alone.

In Fanuc Ltd's cavernous, bumblebee-yellow buildings in a pipe forest near Mount Fuji, automatic machining centers and robots typically toil unattended through the night, with only subdued blue warning lights flashing as unmanned delivery carts move like ghostly messengers through the eerie semidarkness. This plant, one of two in the Fuji complex, makes parts for robots and machine tools (which are assembled manually, however). The machining operation, occupying 54,000 square feet, is supervised at night by a single controller, who watches the machines on close-circuit TV. If something goes wrong, he can shut down that particular part of the operation and reroute the work around it. Some Americans think that Fanuc's Fuji complex is just a showcase. Some showcase.

The total cost of the plant was about $32 million, including the cost of 30 machining cells, which consist of computer-controlled machine tools loaded and unloaded by robots, along with materials-handling robots, monitors, and a programmable controller to orchestrate the operation. Fanuc estimates that it probably would have needed ten times the capital investment for the same output with conventional manufacturing. It also would have needed ten times its labour force of about 100. In this plant one employee supervised ten machining cells; the others act as maintenance men and perform assembly. All in all, the plant is about five times as productive as its conventional counterpart would be.

Across the street, 60 machining cells and 101 robots toil in a big two-storey facility automatically machining parts and assembling them into 10,000 electric motors a month. There is nothing else like it in the world. Men perform maintenance functions here in the daytime. The robots work through the night, in silence marred only by hydraulic sighs and the sibilance of those automatic carts. The first floor of the plant contains the machining cells and 52 robots. Machining is carried out on about 900 types and sizes of motor parts, in lots ranging from 20 to 1000 units.

Machined parts are temporarily stored in an automatic warehouse; they are automatically retrieved when they are scheduled for assembly on the second floor. Yamazaki Machinery Works Ltd operates a flexible automation plant near Nagoya that makes parts of computerized numerically-controlled lathes and machining centers; the latter combine several metalworking machines and incorporate automatic tool changers. In the daytime 12 workers man the $20 million plant. At night only a lone watchman with a flashlight is on duty while the machines keep on working. A conventional machining system with similar production volume, according to Yamazaki, would require 215 workers and nearly four times as many machines, and would take three months to turn out the parts the new plant makes in three days.

The company estimates that over five years of operation its plant will produce after-tax profits of $12 million, compared with $800,000 for a conventional plant that size. Yamazaki is now transferring this technology to its machine-tool-making plant in Florence, Kentucky, bad news for Yamazaki's American competitors. But the most astonishing Japanese automated factory will be started up by Yamazaki about 20 miles from its headquarters near Nagoya. This will be what Tsunehiko "Tony" Yamazaki, the personable senior executive managing director, describes as his company's twenty- first-century factory. The new plant's 65 computer-controlled machine tools and 34 robots will be linked via a fiber-optic cable with the computerized design center back in headquarters. From there the flexible factory can be directed to manufacture the required types of parts-as well as to make the tools and fixtures to produce the parts-by entering into the computer's memory names of various machine tool models scheduled to be produced and pressing a few buttons to get production going.

The Yamazaki plant will be the world's first automated factory to be run by telephone from corporate headquarters. The plant will have workmen, to be sure: 215 men helping produce what would take 2500 in a conventional factory. At maximum capacity the plant will be able to turn out about $230 million of machine tools a year. But production is so organized that sales can be reduced to $80 million a year, if need be, without laying off workers.

The Yamazaki plant illustrates yet another aspect of economy of scope: with flexible automation, a manufacturer can economically shrink production capacity to match lower market demand. Though Japanese machine tool makers are the most ambitious installer's of flexible automation, they are by no means alone. FMS is spreading throughout Japanese manufacturing, with Panasonic, Mitsubishi, and other consumer and industrial goods producers installing the new systems. So far, nothing even remotely comparable is happening in manufacturing in the US or anywhere else in the world. Disturbingly, all of US industry can boast only about 30 flexible manufacturing systems in place; in Japan one large industrial company, Toyoda Machine Tool Co., has more than 30. Frets David Nitzan, director of industrial robotics at the research and consulting firm SRI International, "We are facing another sputnik-a Japanese sputnik."

The growing Japanese lead underscores frequently heard charges that US managers are too remote from technical disciplines to appreciate the potential of such new technologies, and too engrossed with short-run financial results to invest in them. More often than not, machine tool makers report, executives of US manufacturing companies look at something like a flexible machining system only in relation to the narrowly defined functions of the conventional tools it might replace-not for its potential to provide a different, and far more efficient, organization of the manufacturing process. "Cost accounting is a very poor language to communicate new ideas in," observes Paul R. Haas, vice president of Kearney and Trecker's special-products division. One consequence of this myopia is apparent in the ageing of the US machine tool stock. "Many American factories are barrier reefs—one old, tired technology piled on top of another," says James A. Baker, the executive vice president in charge of GE's drive to develop automated systems. "Even when they build new factories, Americans tend to use the same old machines, shipping them from the old plant to the new plant."

Uncle Sam is the Methuselah of machine tools: more than 34 per cent of US tools are 20 or more years old, the highest proportion in any major industrialized nation. Even England is better off; only 24 per cent of its machine tools are similarly ancient. In Japan, only 18 per cent of machine tools are 20 or more years old; 61 per cent are less than 10 years old, against 31 per cent in the US. In fact, fewer than 4 per cent of machine tools installed in the US are numerically controlled—though the concept has been commercially available for a quarter of a century.

The art of managing the factory is no less antiquated. Fixated by the short run, managers have pursued all sorts of piecemeal efforts to hold down costs without stopping to map out systematic ways of organizing the factory floor for more efficiency. Too often production procedures in US factories appear to be little more than accretions of adhoc solutions to problems ranging from space shortages to union-dictated work rules.

Managers focus obsessively on chipping away at direct labour costs rather than exploring better ways to organize the work-force—or investigating the extent to which new technologies are making direct labour costs less important. Japanese corporate leaders, by contrast, tend to be sympathetic with technical disciplines—a far greater proportion of them are engineers—and they have more freedom to incur short-run costs in pursuit of long-run strategic objectives. Perhaps most important of all, a technology like flexible automation is a logical extension of a manufacturing philosophy that views the production of goods as a seamless activity that starts with product design and ends with support in the field a philosophy, as the Japanese put it, of "making the goods flow like water."

"Japanese management takes a holistic view of manufacturing," says James F. Lardner, a Deere and Co. vice-president who supervised a major restructuring of the farm equipment company's manufacturing operation. "They apply logic and common sense to their problems rather than laboratory investigations and discounted cash-flow calculations."

Even before they began to adopt flexible automation, Japanese plants typically employed far fewer people for a given output than American and European plants. The 'Japanese were able to do that by reorganizing production-including the placement of machine tools on factory floors. In the US and Europe, machine tools are usually grouped by type, and parts are directed to them as required.

The Japanese instead place different types of machines together so that each given part can be processed in one place. Much quicker than anyone else, too, the Japanese have taken to such important concepts as "group technology "the grouping of similar parts into families for easier manufacture and better inventory control. Like so many other Japanese manufacturing methods, group technology isn't new; the idea evolved in the 1920s in Germany. A central element of the Japanese manufacturing philosophy is the famous just-in-time concept, the system in which materials and components are delivered as required on the shop floor, not accumulated and stored for future.

Since one-third of factory space is usually employed for storage, the savings are substantial. But there is much more. By reducing inventories to the lowest level at which operations can be sustained, the Japanese force their manufacturing organizations to deal with problems previously hidden. For example, the Japanese already practised preventive maintenance on their machine tools to a degree unknown in the US. Just-in-time has forced them to do even better, because the flow through of products requires every single machine to function perfectly all the time. Yet the Japanese do nothing that Westerners can't—if they only decide to do it. Marvellously efficient factories using the latest automated equipment exist in the US and Europe, towering like islands of excellence in a sea of stagnation and yielding remarkable benefits to their owners. Take Deere and Co's giant new tractor assembly plant in Waterloo, Iowa, which need not take a back seat to any plant in any industry anywhere. In the past few years, Deere has restructured the Waterloo complex from a gigantic, somewhat chaotic job shop into a world-class producer. Deere has poured $500 million into the complex, $150 million of it into the 2.1 million square foot assembly plant. Chassis and engines received from sister plants are joined with tractor cabs and bodies made at Waterloo into gleaming mechanical behemoths. Almost all the materials handling at Waterloo is under computer control. Each part or sub-assembly-engine, transmission, wheels, and so on is automatically assigned to a specific customized tractor ordered by a dealer; it is retrieved from storage and delivered automatically to the assembly line just when it is needed. Putting just-in-time to work, Deere has cut inventory in some areas by as much as 50 per cent, saving millions of dollars.

Flexible automation allows Deere to build a tractor at least twice as fast as before. And it has given the company a new agility: Deere can now successfully compete not only against other big manufacturers but also against "short-liners" that make only one farm implement in higher volumes. What's more, the company is right now bidding on a defence contract worth hundreds of millions of dollars. If it wins the contract, it will start making bulldozers, graders, and other heavy construction equipment on its new flexible automation lines-putting new pressure on such established manufacturers as Caterpillar and Case.

But if manufacturers think that they have to pour hundreds of millions of dollars into flexible automation to reap its rewards, they are mistaken. They can begin by acquiring smaller machining centers to modernize portions of their operations. Clifford R. Meyer, president and chief operating officer of Cincinnati Milacron, recalls his delight at recently visiting one of his client companies, which occupies a garage at the end of an alley in a Los Angeles suburb. Inside, a father-son entrepreneurial team mans $450,000 worth of the latest computerized machine tools—successfully competing as a parts supplier to aerospace companies against much larger firms with older production equipment. Furthermore, not all plants that install flexible automation have to be started from scratch. GE, Ford, and GM are among the manufacturers that have successfully revitalized old plants by installing new machinery.

Like many other American companies, GE had for years lagged in automating its own factories. Some old GE plants, James Baker notes, "make Santa's workshop look like the factory of the future." Today, however, GE can boast notable successes in converting some of its old plants into what it calls "factories with a future"—a marketing slogan the company uses to impress upon potential customers that they need not build entirely new factories of the future. Its meter plant in New Hampshire, modernized at a cost of $25 million, is the epitome of an antiquated multi-storey mill building.

Another ancient GE plant, the Erie, Pennsylvania, locomotive facility, is being transformed with a $300 million investment into an ultramodern automated factory—inside if not on the outside. Building a batch of locomotive frames formerly took about 70 skilled machine operators 16 days; the newly automated factory will turn out these frames in a day—untouched by human hands. The displaced workers are being retrained for other, more sophisticated jobs. As a general matter, in fact, flexible automation threatens employment less than might be supposed.

The US faces a shortage of skilled machinists for the rest of the decade, and automation of assembly, where semiskilled jobs predominate, will proceed much more slowly than automation of machining. GM is also advancing. Last October it installed its first flexible automation system, an Italian-built Comau system with three machining centers, at the Chevrolet Gear and Axle Division in Detroit. Almost immediately GM discovered just how valuable the new manufacturing flexibility can be. When an outside supplier failed to deliver a front-axle component up to quality standards, GM brought the job in-house. It designed and built the tooling for the component on the FMS in ten weeks—a job that normally would have taken up to a year. As yet, such examples are rare exceptions.

The majority of US manufacturers either do not yet grasp the significance of the new technologies, or do not want to invest while money is still expensive and future markets uncertain. Moreover, there are still no turnkey flexible manufacturing systems available. Installation of big systems requires skilled people not all companies may have on hand. But as GE's Baker argues: "We're running out of time and excuses." The price of delay may be disaster, as the experience of some American machine tool manufacturers show.

During the 1970s the Japanese became the world's first mass producers of computerized machine tools. While the majority of US machine tool makers—small companies, in the main—plodded along using old technologies and methods, the Japanese redesigned their tools for easier manufacturing with flexible automation and equipped these tools with advanced yet simple electronic controls that the humblest job shop could understand. Then, in the late 1970s, the Japanese caught the US tool-makers napping in the midst of the capacity crunch, with delivery times stretched up to two years.

The Japanese, to be sure, played an additional trump. Their machine tool industry had been mobilized by Japan's Ministry of International Trade and Industry (MITI) into a cartel and bolstered with millions of dollars in government funds, as shown in documents obtained in Japan by lawyers for Houdaille Industries, the US machine tool maker that has asked President Reagan to deny investment tax credits for some Japanese tools sold in the US. This double-barrelled assault left US makers of numerically controlled machining centers and lathes in shambles. By the end of last year the agile Japanese had captured more than 50 per cent of the US market for those machines.

The US machine tool industry has lost its erstwhile position as the world's leading producer to competitors not only in Japan but also in Europe and even Taiwan. The National Academy of Engineering, usually not given to alarmist statements, in a report soon to be issued calls the Japanese and other foreign inroads "a very threatening development that could seriously endanger the future economic security of the American industry."

The reason for concern is simple: the machine tool industry is central to the growth of all manufacturing. There will be more Japanese surprises in the years to come. This October, for instance, the Japanese government's mechanical engineering laboratory near Tokyo will unveil a small prototype factory of the future where novel, laser-equipped machine tools will take production automation a big step beyond where it is now. The new machines will perform metalworking processes now done separately—such as turning, drilling, and milling—all at once, cutting batch production time of metal parts in half and reducing the number of production processes by 60 per cent.

TWENTY

The Set of Instructions

Software—the set of instructions that tell a computer what to do—is becoming more important than the computer itself. The booming software industry is therefore setting the pace of the information technology, revolution and the pace is frantic. This article first appeared in *Business Week*, February 27, 1984. Hardware is no longer where the action is. Computers are becoming remarkably similar—in many cases they are turning into off-the-shelf commodity products. Now the computer wars are being fought on a new battleground: software-the instructions that tell computers how to do everything from processing payrolls to playing video games.

A key reason for the change in emphasis is an overwhelming demand from customers for packaged software that will let them apply computer power to a broad range of new tasks. Increasingly, corporations are finding they do not have the resources to write the programs they need. As a result, most companies have stopped writing their own software and are instead buying standard software packages. So far, producers have been unable to keep up with the need, and there has been a severe shortage of software able to take advantage of the power of the latest machines. "We're getting the hardware we need from the industry, but software has not moved out at the same rate," says Robert J. Metzler, vice-president of First Computer Services, the data processing arm of First Union National Bank in Charlotte, NC. "There's a tremendous cry in the industry [for software]." The flood of personal computers pouring into small businesses and homes has created an even faster-growing new market.

The booming demand for new and better programs has quickly turned software into big business. During the 1970s software was stiff a cottage industry, 'with sales totaling just $2.7 billion annually. This year-according to estimates by Input, a California market research firm-sales are expected to top $10 billion.

With the industry growing that fast, several once-small software companies have become sizable corporations. Management Science America Inc. (MSA), for example, the largest independent software supplier, quadrupled in size during the past three years, topping $145 million in sales in 1983. Trying to capture some of the soaring market, thousands of new companies have entered the software business in recent years; by one count there are more than 3000 software companies now.

While over the long term a shakeout is probably inevitable, the near-term outlook for most of these companies appears extremely bright. "The upside potential has barely been tapped," says Robert M. Freeman, senior analyst at Input. He expects the market to keep growing by a dizzying 32 per cent a year, topping $30 billion in 1988. Software sales of that magnitude would amount to half of the $60 billion hardware business expected for the same year; today, revenue from software is equal to only 27 per cent of the value of all the computer hardware sold.

Sales of software for personal computers-the fastest-growing part of the software industry-should rise by an astounding 44 per cent annually over the same five years. "We're finally at the point where software applications are going to be a big moneymaker," says Jack M. Scanlon, vice-president of the Computer Systems Division of AT&T Technologies Inc. But with the promise of such rapid market growth, has come a feverish competition that is beginning to restructure the entire industry.

Until recently, nearly all software companies concentrated on a particular niche and fit neatly into one of three distinct market segments that were divided along the same lines for both mainframe and personal computers.

Systems software, which handles basic housekeeping operations such as controlling the printer and memory, was supplied primarily by the computer makers. They were joined by independent software companies in providing utility software, which, among other things, helps programmers write programs. And a host of other independent software companies competed in the third market; applications packages, which tell a computer how to carry out specific tasks such as accounting payroll or word processing.

But those distinctions are blurring:

Consolidation: International Business Machines Corp., to move faster into applications software, placed all its software efforts into a single, entrepreneurial unit—the same type of organization it used to launch the highly successful Personal Computer.

Acquisitions: To provide a full complement of software for their computers, makers such as Hewlett-Packard, Burroughs, and Prime Computer are rushing to buy applications software companies. Other manufacturers, including Honey-well, Sperry, and Digital Equipment, are setting up joint ventures with software suppliers. Says Roger T. Hobbs, vice-president for software products and services at Burroughs Corp.: "The demands for software are increasing so rapidly that it is impossible for the manufacturer to keep up."

Expanding Product Lines: Software companies are expanding product lines to maintain their competitive edge. Those that supply programs mostly for large mainframes—MSA, Cullinet Software, and Computer Associates International, for example—are snapping up personal computer software houses. And companies that specialize in systems software for personal computers-among them, AshtonTate, Digital Research, and Microsoft—are adding general-purpose applications software to their product lines. "Today you have to have a broad line," says Terry L. Opdendyk, president of VisiCorp. Otherwise, "the vendor is increasingly vulnerable to competition."

New Players: To grab a piece of the action, publishers and other communications companies outside the computer business—including CBS, Dow Jones, Dun & Bradstreet, McGraw-Hill, and Simon and Schuster-are licensing programs they then sell through their own distribution channels. "It's a very fragmented industry, but the potential is absolutely huge," says Richard W. Young, president of Houghton Mifflin Co.

Japan's Drive: To catch up with their US rivals, Japanese computer makers are launching major software development efforts. Perhaps the best evidence of just how important software has become is the drive by Japan's Ministry of International Trade & Industry to change the law to boost Japan's fledgling software industry. One result of many of these trends is that mergers and acquisitions in the industry are at an all-time high. In 1983 there were 146 acquisitions, valued at more than $1 billion-up 130 per cent from 1982—reports Broadview Associates, a New Jersey company handling such transactions. And this frantic pace of activity is turning the software business into a hot investment area.

To complete successfully amid all the turmoil, software suppliers are struggling to formulate new strategies. One common focus is advanced design-especially the development of software aimed at making computers easier to use. But just as important to the success of new software as advanced design is marketing. Nowhere is the new attention to marketing more noticeable than in the hotly competitive personal computer business. To reach the millions of personal computer users, software companies are spending huge amounts to introduce and advertise their products. Industry watchers have dubbed this obsession with splashy promotion "the Lotus syndrome—a reference to the more than $1 million that Lotus Development Corp. spent over a three-month period in 1983 to launch its first product, the highly successful 1-2-3 package. "Lotus advertised so much that companies are going to be forced to step up their advertising just to be heard above the noise," says David E. Gold, a San Jose (Calif.) computer consultant.

Advertising, in fact, has grown so important to software success that it has become one of the biggest barriers keeping out new companies. "The cost of technology development is dwarfed by the marketing cost; the ante has really been upped," says Rodney N. Turner, vice-president for sales at Ashton-Tate, a Culver City (Calif.) software company best known for its dBase II package. It takes as much as $8 million to launch a new software product today, he estimates. Ashton-Tate, by contrast, was founded on shoestring-$7,500. Any personal computer program must be carefully packaged and promoted. “It has to look good and has to be well supported by national advertising,” maintains retailer Gregg E. Olson, a salesman at Mr Software in Boulder, Colo. “A company that wants to sell to us has to have all their marketing elements in place to establish credibility with us and our customers.”

Even companies that write software for the large minicomputers and mainframes are plowing more money into marketing. While programs for the larger systems, unlike those for personal computers, do not require heavy consumer advertising, says Martin A. Goetz, senior vice-president at Applied Data Research Inc., his company still boosted its 1984 advertising budget by 60 per cent over 1983, to about $2 million.

This emphasis on marketing comes chiefly from the need to reach a different, far broader group of potential customers. “Ten years ago, [our customer] was the data processing department, and if you had a product, you sold it to the technicians,” says Robert D. Baskerville, group vice-president for product management at Computer Sciences Corp. in El Segundo, Calif. Today, he notes, a software company has “to sell to end users, and you have to emphasize more than the technical capability—you have to really sell the benefit.”

Software marketing now means beginning with detailed product planning, so that a program will better meet the needs of customers. Before Wyly Corp. even began a $20 million program to diversify into applications software, for example, it brought in all its salesmen from the field to tell programmers what customers wanted-something most software companies had never done. “With modern software development techniques, you can produce almost perfect code, but if you don’t understand the market, it could still be useless,” says Ron W. Brittian, Wyly’s vice-president for research and development.

Critical, too, is a reputation for good service and customer support. “To grow from a startup into a large software company is not so easy anymore; it’s a question of distribution and of educating the customer in how to use the product,” says Anthony W. Wang, executive vice-president at Computer Associates International Inc. Toward this end, his company recently installed a $250,000 television studio—costing $200,000 or more a year to run—to make training videotapes for its software customers.

As competition heightens and customers become more demanding, companies selling personal computer software are providing a level of customer support far greater than anything they offered before. MicroPro International Corp., for example, has plowed $2.5 million into developing computer-aided instruction and retail-support aids for its software, which includes the best-selling WordStar. The company is also hiring journalists and other non-technical writers to produce more readable instructions for today’s less sophisticated customers. "A product is not a product until you have computer-aided instructions, it's not a product until you have, video instructions, it's not a product until you have honest-to-God understandable language [in the manual]," says H. Glen Haney, president of MicroPro.

Establishing a marketing presence is especially critical to companies that provide personal computer software, because they must win space on already crowded retail shelves. Computer retailers are reluctant to take on new software without feeling confident it is a winner. "Unless we're absolutely convinced that it's a great program, we'll wait to see how other channels of distribution do with it," says John H. Rollins, national manager for Sears Business Systems Centers, the computer retailing arm of Sears, Roebuck & Co. One group trying hard to enter the software business may have an edge: the traditional book publishers. Because they already sell books to computer stores, the publishers "provide ready-made distribution channels," says Input's Freeman. Agrees William M. Graves, president of MSA: "There is a tremendous similarity in the microcomputer business and publishing."

No matter how they distribute their products, software companies are finding that they must offer a broader range of products than ever before. "The strategy is to offer a complete solution," asserts Robert N. Goldman, president of Cullinet. "If a customer ends up with eight different vendors, none of the software works together." Cullinet, for instance, is augmenting its data base management software with such applications as general ledger accounting and manufacturing control programs.

Broadening a software line also helps a supplier leverage the large amounts of money spent on establishing a reputation and brand recognition. VisiCorp, which gained market fame with its hit program VisiCalc, is trying to expand its applications software with its family of Visi On integrated software. Similarly, Microsoft is plunging into such applications software as word processing and financial analysis to build on its reputation as a supplier of personal computer operating systems.

The trend toward comprehensive offerings will force small companies to target well-defined market niches. There will continue to be demand for specialized packages designed for the needs of a particular industry or profession. "There will be several very large software companies," contends John P. Imlay Jr., chairman of MSA, "but there will be literally hundreds of small companies, with under $100 million [in sales], that have specialized market niches."

Perhaps the most fundamental change in the software industry is the blurring distinctions among the suppliers. No longer can they be neatly divided into companies that make basic systems software and those that write programs for specific applications.

Moreover, the top mainframe software companies are rushing to market with software for personal computers. And vendors serving only the personal computer market are joining with suppliers of mainframe software. VisiCorp, for example, recently teamed up with Informatics General Corp. to offer Visi-Answer, a program that allows a user of an IBM Personal Computer to retrieve information from an IBM mainframe data base. These changes resulted in large part from the growing number of customers linking personal computers to large mainframe systems.

As the emphasis in the data processing industry shifts to software and, as software companies strengthen their sales, service, and distribution, the big-system makers, too, are scrambling to do more to provide their customers with software. "In the old days, our customer wrote his own application [software]," notes Jon Tempas, vice-president for software products at Sperry Corp.'s Computer Systems operation.

Today, "there's an increased expectation for hardware suppliers to provide the complete solution." That means the equipment makers will need to provide more of their own software. Sperry, for example, now writes 95 per cent of the software it sells for its computer line.

Most of the big-system companies, however, are turning to software specialists for help, since much of the demand is for applications software finely tuned to specific industries. "I don't think there's any hardware manufacturer that can-or should-provide all the software," says John E. Steuri, general manager of Information Services Business Unit, IBM's new independent software group. "We'll be more and more dependent over time on software developed outside the company." The pressure to team up with independent software companies is especially strong among makers of personal computers.

"Without an adequate software base, a microcomputer dies," says Eugene W. Helms, vice-president for business development at Texas Instruments Inc.'s Data Systems Group. So TI is recruiting software suppliers to adapt their best-selling programs for its Professional Computer. Perhaps the most extensive effort to sign up independent software companies was made by Apple Computer Inc., which courted more than 100 companies to write software for its new Macintosh computer. The California company was successful in signing up more than 80 of them.

Companies that do not move quickly to develop a broad line of software will have trouble keeping up. Consider Tymshare Inc., a computer timesharing company. As computer prices began to drop precipitously, more and more of Tymshare's customers stopped renting computer time and purchased their own machines. Tymshare did not have any software packages to sell to its former clients. "We used to look at software as simply an in-house tool that we needed to offer time-sharing," says one Tymshare executive. That attitude, he admits, "came back and bit us in the rear end." In 1983 the California company lost $1.7

million while sales fell 3 per cent to $288 million. Now it is engaged in a major effort to expand into applications software, through internal development and by licensing packages developed by other companies.

With most software companies trying hard to move in the same direction, everyone will face stiffer competition. Already, prices for personal computer software are sliding wherever similar programs have proliferated-in electronic spreadsheets and word processors, for example. And a big battle is under way among suppliers of the various "windowing" software packages—software "environments" that permit the users of personal computers to display several tasks at once. VisiCorp has been forced to slash the price of its Visi On environment package from $495 to $95. "The consumer is going to force the prices down simply by demand," asserts Alvin B. Reuben, executive vice-president at Simon & Schuster Inc.'s electronic publishing division.

As prices fall, the opportunity for, newcomers to jump in and grab quick and easy profits will all but disappear. The cost of developing and marketing new programs is mushrooming just when margins are shrinking. VisiCorp successfully launched VisiCalc in 1978 with a $500 budget. But the California company has spent more than $10 million developing its latest product, the Visi On environment. Product life is also getting shorter as new products come out faster. As a result, says Softsel's Wagman, "the stakes have gotten higher and much riskier from a development point of view."

For many companies, the risks will ultimately prove fatal. "There are companies out there that are already feeling the pinch of increased costs, new product announcements, and a changing marketplace," warns consultant Gold. Industry watchers predict that some big-name failures will occur within the next 18 months. By the end of the decade, many experts expect the software industry to have consolidated its thousands of suppliers into a few major players.

For those that make it, the future is promising. The demand by the growing army of computer users for easier-to-use software to handle an exploding variety of tasks will drive the industry to create software with far more capabilities than anything available today. "They only limit," says Stuart A. Walker, vice-president of marketing at Konware Inc., "is the limit of new ideas."

Software in the Office: One of the main reasons software has become so highly visible is that the computer has moved from the back room into the office.

Programs of the past, buried inside corporate data processing centers, primarily handled clerical and accounting tasks, such as turning out payrolls and keeping track of accounts receivable. Today's computer handles a rapidly increasing variety of management tasks— employing information as a competitive marketing weapon, for example.

That trend, coupled with the explosion in the use of personal computers by executives, has given new prominence and importance to software in us corporations. "The Programs that we install to do our company's business are analogous to the tools an auto maker puts in place to manufacture auto parts," says Jeffry A Alperin, an assistant vice-president in the Information Systems Support Dept at Aetna Life & Casualty Co. "Our plant is our data Processing system." Since the same computer hardware is available to everyone it is the software that often gives a company a competitive edge. At oil companies, software is now helping drive the search for oil and gas," says Michael C. Balay, general manager of information technology at Gulf Oil Corp. "You can have the same [seismic] information, but a unique software package may help you better interpret the data. "

Nowhere is this increasing importance more obvious than in corporate data processing budgets. Belay estimates that Gulf spends about half of its annual data processing budget to write its own software and to buy packages from independent suppliers. At Aetna, half of the 4000 people on the data processing staff are programmers. And Alperin says that 4 per cent of the $237 million that Aetna spent on information processing last year went to purchase packaged software.

Moreover, individual owners of personal computers are no different from large corporations: over the life of a computer, its owner will spend $2 on software for every $1 spent on hardware, according to industry estimates. But users are finding it increasingly difficult to produce their own programs. Writing software is still a time-consuming process, and users find that, while they want to employ their computers to handle more applications they do not have enough programmers, time, or money to write all of the necessary software in-house. In fact, at many companies users must now wait as long as 18 months for new programs to be written. "The end users are demanding faster development of solutions," declares Richard R. Douglas, group vice-president for the US Marketing and Services Group of Honeywell Inc.

To come up with this software faster, computer users are compromising-buying more ready-to-run software packages than ever before. "There has been a turnaround from five or six years ago," says James T. Manion, sales vice-president at ASK Computer Systems Inc. Few computer users write all their own programs anymore, he says. But while users sacrifice some uniqueness by buying a software package, that disadvantage is often outweighed by the time and money saved. A package can cut the time it takes to get a system up and running by more than half, figures Mayford L. Roark, executive director of systems at Ford Motor Co. He also estimates that buying a program from an outside vendor can cut development costs by 30 to 75 per cent. "It's almost impossible to think of a function you might want to do on a computer that you can't find several software packages for," says Roark.

Where a custom solution to a problem is required, companies often find they can save time by modifying a standard software package for their own use. That is what Aetna did when it needed software to offer a new type of life insurance called universal life, which features flexible premiums. That allowed the company to match its competitors' universal life offerings quickly. With such a variety of software available, companies are finding that writing their own programs from scratch makes about as much sense as drawing their own road maps for their salesmen.

US corporations brought more software packages last year than in the previous year, says International Data Corp., a market researcher. Such purchases, along with software services, now account for more than 8 per cent of the average data processing budget, up from 6 per cent in 1980, IDC says. Similarly, purchases of personal computer software packages soared 74 per cent last year to top $1 billion, according to market researchers at California's Input.

The customer's new attitude toward software is beginning to alter fundamentally the way that computer systems are sold. "We're selling to a much more sophisticated consumer base than in the past," points out Elizabeth M. R. Hall, product manager at Information Science Inc., a Montvale (NJ) software company. From now on, adds Honeywell's Douglas, "Software looms as a much more important factor in [a user's purchase] decision than hardware."

Computing used to be the solitary domain of a priesthood of programmers. The arcane languages that they employed to command their giant mainframe systems were shrouded in such 'complexity that few laymen could understand, much less control, these behemoths. But the inexpensive yet powerful personal computer is changing all that. Millions of nontechnical users are now running computers. These new operators, however, refuse even to read an instruction manual, let alone memorize the cryptic commands of the priesthood. Their aversion, coupled with the dramatic slide in the cost of computer power, has created a revolution in the way software is written.

Today, making a program simple to control and easy to use is just as important as what the program actually does. Says Jeanne M. Baccash, a software engineer at American Telephone & Telegraph Co.'s Bell Laboratories: "There's a whole new thrust to reach a market that doesn't know or care what it means to 'boot a system' [start a computer].

These efforts to make software simpler are crucial if the information processing industry is to continue its fast growth. "Rapidly advancing technology has left consumers, trainers, and computer salesmen behind," says Terry L. Opdendyk, president of VisiCorp. "How to spread the word of how to use the products is the key limiting factor [to the industry's growth]."

Man-machine Communications: To reduce the number of commands that a user must memorize and type into a computer to get it to work, the latest software enables users to communicate with the machine in new ways. Some of these techniques are relatively simple, such as establishing menus to show users what commands are available to choose from. The most elaborate methods, used on such machines as Apple Computer Inc.'s new Macintosh, replace commands with an array of tiny icons or pictures—for example, a file folder to indicate filing. To tell the computer what to do, users point to the appropriate picture on the screen by moving a pointing device called a "mouse" over a desktop.

Data Exchange: Early programs were built to accomplish tasks, such as calculating financial forecasts, with only the data located within the user's own computer. But users have grown more sophisticated and now want to get data stored in other machines—say, the corporation's central mainframe system.

Artificial Intelligence: This embryonic method of programming, which enables software to mimic human thought more closely, might one day be the easiest approach to use. It is starting to show up in programs such as Artificial Intelligence Corp.'s Intellect and Microrim Corp.'s CLIO, which let a user ask a computer for data with English sentences rather than esoteric commands.

Most programmers are ill-equipped to figure out how office workers and other nontechnical users best handle computers. So companies are bringing in professionals from disciplines as far afield as education and psychology to help in the design of man-to-machine communications. AT&T, for example, had a staff of psychologists survey about 400 computer users to help it decide how to add commands to its Unix operating system. Designing easier-to-use software also requires new development techniques. For instance, International Business Machines Corp. now tests its new software in "usability labs." In these labs, volunteers try to use IBM software while researchers with video cameras watch from behind one-way mirrors. "In the long haul, those that have easy-to-understand software will be the successful companies," says John E. Steuri, general manager of IBM's Information Services Business Unit.

A lot of discussion is going on in research circles over how best to use menus, pictograms, and other techniques to help both computer novices and experts. The type of coaching that a neophyte computer user requires becomes annoying once the user has learned to use the software. So designers want "to make sure the user is not presented with extraneous information when he has to decide what to do next," says Brian K. Reid, a Stanford University electrical engineering professor. Micropro International Corp., for one, is trying to solve this problem by using a time-delay activator.

If a user types in a command less than two seconds after being asked for it, the program assumes the user is proficient and skips over the menu listing of options. Users who take longer are given the benefit of directions from a menu. Researchers are also struggling with the problem of how to standardize commands. Each computer program now uses different commands to accomplish the same function. An early attempt to solve this problem is the "environment" or "windowing" software packages for personal computers that are now coming to market. These packages-products such as Microsoft Corp's Windows and VisiCorp's Visi On-divide the computer screen into segments, each of which shows a different task. Common functions are performed the same way in each window. The large amount of this kind of innovative software now being developed to make personal computers easier to use is also forcing the mainframe software companies to follow suit. "Mainframe vendors have been forced to make software easier' because users have been spoiled by micros," says David Ferris, an industry consultant.

Information Builders Inc. has developed a version of its Focus data base management software that can run on an IBM Personal Computer. But the PC version, called Pc-Focus, "has more features than the mainframe product because after using a personal computer customers have come to expect more," says Gerald D. Cohen, president of the New York company. One of the Pc-Focus features is called Table Talk. Each time a user moves from one step to another in the course of retrieving information from a data base, the computer automatically presents him with a menu of the appropriate choices.

Another software challenge is to find simpler ways to exchange data between programs. In most cases, swapping data between a large mainframe and a personal computer is difficult because the two units use different formats for storing their data and different commands for retrieving them.

The problem is similar to that of a person who is trying to communicate with 10 people, each of whom speaks a different language, explains Robert J. Spinrad, director of systems technology at Xerox Corp. "I could either learn all nine other languages," he says, "or we could all learn a common one-say, Latin." But software vendors are not waiting until everyone in the industry agrees on what common language all computers should learn. In the past year a host of vendors, including Cullinet Software, Informatics General, and Cincom Systems, have announced products that link personal computers to large mainframes. Such capabilities will be a requisite for any software in the marketplace, says President Frank H. Dodge of McCormack & Dodge Corp. "If mainframe software doesn't allow that link," he adds, "it's going to be pushed aside pretty rapidly."

Some experts think the ultimate mechanism for exchanging data is the use of sophisticated file-management software, often called data base management systems. These large programs index information and then store it in such a way that it can be retrieved using a variety of names-much the way a library card catalog lists the same book by author, subject, and title. By providing this uniform filing structure, a data base system simplifies the exchange of data.

What may soon make computers even easier to use is artificial intelligence (AI) software. Some companies are already working on AI programs that will eventually be able to remember an individual's habits in using the computer. The first AI applications from Microsoft-expected within the next year-will be likely to use these rudimentary pattern-recognition techniques in tutorial programs that teach novices how to operate software. These programs will adjust the level of tutorial difficulty by determining the proficiency of the student running the program. Says Microsoft Chairman William H. Gates: "Just as humans take actions based on past experience without having to be told again, and again, so will software.

As the second-largest group of computer companies in the world-after the US—they exported $2.7 billion worth of computers in 1983. But in their single-minded determination to build and export hardware that is faster and less expensive, the Japanese have given short shrift to software development. "Hardware manufacturers have been lazy about developing software," acknowledges Hisao Ishihara, managing director of the Japan Software Industry Assn. But now, he points out, "that is suddenly changing." Faced with the twin problems of rapidly falling hardware prices and the growing percentage of computer budgets being spent on software, Japanese hardware makers are scrambling to shift their resources into software development.

In 1981, for example, Hitachi Ltd's Computer Division spent just 10 per cent of its research and development budget on software. In 1984, even though software accounted for about 10 per cent of its computer revenues, Hitachi spent 30 per cent of its R&D money on software. "These days, more and more of the value added in information processing comes from the software," says Toshimitsu Kaihatsu, head of Toshiba Corp.'s Software Management Dept.

The Japanese are also keenly aware that to continue to be competitive internationally in computers they must now start developing world-class software. "We have to sell more overseas if we want to recover our software development costs," comments Shoichi Ninomiya, general manager of Fujitsu Ltd.'s Information Processing Group. His company spent close to $100 million a year-more than one-third Of its R&D budget-to develop software for its mainframe computers. But almost all of that software is now sold to Japanese customers. Until now, software has been the biggest handicap the Japanese have had in selling their equipment abroad-especially in the office and personal computer markets. This point was driven home in 1982 when International Business Machines Corp. sued Hitachi, one of the largest Japanese exporters of computers, charging the company with copying IBM software and reselling it with Hitachi's machines.

Hitachi agreed in an out-of-court settlement to pay IBM between $2 million and $4 million a month in software license fees and agreed to let the

US computer giant inspect all new Hitachi products before they go on the market to ensure that they do not infringe any IBM copyrights. To avoid a similar lawsuit, Fujitsu has also agreed to pay IBM many millions of dollars and promised not to copy the US company's software, although it would not disclose the details of its agreement. Because of language and cultural differences, software written in Japan is obviously difficult to export. Not only must instruction manuals be translated, but often the programs have to be completely rewritten. Japanese accounting rules, for example, are different from those in the US, so accounting software written in Japan is useless in the US. This shortage of programs has severely limited the export potential of Japanese home and personal computers.

Japan's software industry is also held back by the preference of Japanese buyers for custom software rather than the standard software products that US customers are increasingly buying.

As a result, an independent Japanese software industry has been slow to develop. And the computer manufacturers' efforts to develop all the software themselves have been crimped by the large sums they must invest to improve methods of getting information into the computer and processing it, using Japanese *kanji* characters. But it would be dangerous to write off the Japanese as competitors in world-class software. "The Japanese are as dedicated to computers as they were to autos and shipbuilding," says John P. Imlay Jr, chairman of Management Science America Inc. "If US companies do not innovate and supply quality products and customer support, the Japanese could make inroads."

Already, Japan has a big chunk of the US market for the simplest type of software, that written for video games. Cultural differences are not a barrier to the export of such hit games as Pac-Man. The fast-growing software markets for engineering and scientific applications-which manipulate schematic drawings and numbers rather than words-are other areas where Japanese software could be sold overseas with little modification. The Japanese have also demonstrated their technical ability by creating powerful supercomputer software and sophisticated banking and airline reservation systems. "People are misreading the capability of the Japanese when they say Japanese can't build good software," maintains Joseph C. Berston, president of Comstute Inc., a software consulting firm based in Japan.

Indeed, Japanese companies may actually have some advantages over their US rivals. Because of the legendary thoroughness of Japanese workers, "the finished product here is better, more reliable, and easier to maintain," says consultant Berston. Labor costs are also lower for Japanese software makers. Well-educated, highly disciplined Japanese programmers are paid an average of about $10,000 a year. Their US counterparts can expect a starting salary twice that—from $19,000 to $25,000 according to a survey by Robert Half International Inc., a personnel agency.

In addition to that salary differential, the Japanese claim their programmers are 10 to 15 per cent more productive than their US counterparts because of Japanese investments in program development aids. To widen that margin, they are now building software factories that give their programmers access to even more sophisticated tools.

Toshiba recently completed a factory that employs 3000 software engineers to develop industrial software. Now the company is building a second software factory that will employ 2000 more programmers. "To overcome Japan's language problem and compete with the US," says Toshiba's Kaihatsu, "we have to have productivity double that of the US." NEC Corp., which already spends a significant portion of its annual $400 million software budget on productivity tools, says it will use its productivity and quality advantages to crack the US market. It has hired US software engineers to analyze the needs of US computer users. The resulting lists of requirements are then fed into computers at NEC's Japanese software factories, where programmers write the software. In this fashion, NEC has already started developing commonly used business applications, says Yukio Mizuno, an NEC vice-president. Perhaps the best illustration of the importance the Japanese now attach to software is the move by the Ministry of International Trade & Industry (MITI) to back the industry. MITI has set up several research laboratories to work on software, including a lab that is developing the so-called fifth-generation computer and software. That lab is budgeted to receive $23 million in MITI support this year. The ministry is also giving low-interest loans and tax breaks to software developers.

MITI hopes to boost its fledgling software industry further with a proposed revision in the copyright law.

The ministry is pushing to allow Japanese companies to save money and programming time by making it legal for them to copy portions of existing software products without the permission of the original developers.

Industry observers say the law, if passed, would help the Japanese leapfrog US software companies by enabling them to copy popular US programs and incorporate them into Japanese software products. But even with MITI's help, it will be a long time before Japan's software industry catches up with its US competitors-if it ever does. In 1982, software sales in Japan were only $1.4 billion-just one-quarter of the total US software sales that year.

TWENTY ONE

Business through Internet

Electronic Commerce is a rapidly growing research and development area of very high practical relevance. A major challenge in successfully designing EC applications is to identify existing building-block technologies and integrate them into a common application framework. The last decade has witnessed major efforts towards generalizing the functionality of transaction-based applications, and this has been one of the roots of today's workflow technology. Thus, an important objective and, to some extent, achievement of workflow technology is to reconcile rich specification languages for the control flow (and data flow) of complex, long-lived, often widely distributed, organizationally decentralized applications with the industrial-strength dependability features of transaction processing systems. To this end, workflow technology does itself make intensive use of database systems and transaction services. Thus, we believe that workflow technology today provides the most mature and promising basis for becoming the backbone of EC applications.

A second overriding requirement of EC is the need for interoperability among largely autonomous, often heterogeneous subsystems. This can span a spectrum from widely established standard protocol's between information servers all the way to "intelligent agents" that have capabilities for negotiation, sophisticated exception handling, etc. Again, we believe that the database and transaction technology has a fairly good success story in dealing with heterogeneity and interoperability issues. Thus, workflow technology inherits a number of these salient properties, and this is why it can integrate almost arbitrary application programs into a unifying execution framework. However, there are still major deficiencies in the interoperability between different workflow management systems. These need to be overcome to provide full-fledged support for enterprise-spanning applications including EC applications.

Initial ideas along these lines do exist, but more research is needed. It should be noted, in any case, that the highly structured, rigorous yet dynamically adjustable nature of workflows make them an excellent candidate for addressing the interoperability problems, as opposed to say a general-purpose, arbitrarily complex, distributed programming language such as Java.

We discuss the state of the art in workflow research and technology with regard to its applicability to and benefits for EC. We identify necessary extensions to workflow management systems to support EC, and sketch the research issues towards these extensions. The chapter is organized as follows: the next section discusses properties and requirements of workflows in the EC context. Then, we discuss contributions of workflow technology to address these requirements. For the sake of concreteness, we will occasionally refer to our own prototype system, Mentor, as an example for what state-of-the-art workflow management systems can do for EC.

Workflow technology aims to provide as much computer support as possible to the modelling, execution, supervision, and possibly reengineering of business processes. Such computer-supported processes are then referred to as workflows. It seems natural to consider EC as the glue between previously independently modelled business processes. Hence, it is beneficial to model major parts of EC applications also as workflows. Following this approach, we can identify several specific requirements that EC applications pose on workflow technology.

The probably most important requirement is dependability, which subsumes a reliable, fault tolerant, secure, and provably correct execution on a highly available platform. While these requirements are ubiquitous in mission-critical computer applications, there are a few key technologies that have achieved substantial progress towards these goals over the last three decades. In particular, database and transaction systems are the backbone of today's mission-critical commercial applications such as banking and insurance businesses, travel planning and booking, or financial trading.

Workflows that capture (major parts of) EC applications have specific properties. In the following, we denote such workflows as EC workflows for simplicity. In this section we discuss these specific properties of EC workflows and derive from them specific requirements on workflow management systems. The discussion is broken down into build-time issues that arise in the modelling and analysis of an EC workflow, and run-time issues that are relevant during the execution of EC workflows.

We can identify a number of specific properties of EC workflow specifications:

— The core part of EC workflows, which is typically centered around the payment scheme (i.e., the sales phase), is rather static. So there is no need for supporting dynamic modifications of running workflows. On the contrary, this core part can be viewed as a specific protocol that has been agreed upon by all EC partners at one point, and thus must not be changed further on. In fact, it is highly desirable that certain properties of the payment protocol can be verified in a formal manner. This is only feasible for static workflows.

— Both the pre and post-sales phases of EC workflows should, however, allow for flexibility at run-time. So a workflow model for these phases is more likely to be a skeleton that is dynamically enriched by introducing additional activities along with their control and data flow, and also possibly skipping parts of the pre-specified workflow skeleton. The need for such improvisation is particularly evident in the pre-sales phase that includes filtering ads, bargaining, etc.

— In all phases, EC workflows involve multiple parties. In the pre-sales phase, there are rather loose interactions between customers and merchant. In the sales and post-sales phase, customer and merchant and banks or credit card companies are more tightly coupled. The coupling is based on standard interfaces, enabling the establishment of the sales and post-sales phases without complex negotiations.

The core part of EC workflows can be fully automated. So none of the underlying activities should require human intervention, other than simple typing of initial input parameters. In contrast to certain high-value workflows in banking (e.g., approval of company loans), for example, EC payment schemes rarely involve intellectual decision processes. Such automation is indeed an important factor in promoting EC as a mass opportunity.

From these properties we can directly derive the following requirements that EC workflows pose on the underlying workflow management systems:

Verifiability (Provable Correctness): For the core part of an EC workflow, guaranteeing the correctness of the specification is of utmost importance. This involves proving both safety and liveness properties, examples being "the customer's credit card number is not shown to anybody other than the merchant's bank," and "money is eventually credited on the merchant's bank account", respectively. So, unless such critical properties are formally verified, an error-prone specification bears the risk of exposing confidential information or even losing money. Superficially, it may seem that the required correctness guarantees can also be given intellectually by "carefully looking at the specification", but it often turns out that such distributed payment protocols can raise a fairly large number of different exceptions all of which need to be handled correctly as well. Once the effect of multiple of these exception handling methods can interfere with each other or with the "mainstream" of the workflow, it is often extremely hard to capture all the resulting effects intellectually. Thus, automatic verification methods or at least support tools are called for. These should be able to deal with the complete control flow specification, including all exception handling, and also the data flow specification.

Composability of Building Blocks: As we discussed above, the payment scheme itself, albeit constituting the core part, is only one phase of an entire EC workflow. Both the pre- and post-sales phases require much more flexibility, and may involve sub-workflows that are provided and have been designed by different parties. Thus, EC workflows are inherently heterogeneous in that independently designed workflows have to be composed into more comprehensive higher-level processes. The specification method must support this kind of composability in an easy-to-manage manner.

Interoperability: Interaction between different parties involved in an EC workflow is required in all phases. However, the type of interaction differs, which requires different interfaces for interoperability. The loose coupling of customer and merchants in the pre-sales phase can be supported by a web-like interface, posing little requirements on specifications. As the customer has to face incomplete and outdated information anyway, the focus is on a fast access to information with low overhead, not on dependability. In the sales phase, the picture changes completely. As correctness and dependability are of utmost importance here, it must be possible to specify the corresponding interoperability protocols.

As for the run-time issues of EC workflows, we can again highlight certain properties. These are not exclusively specific for EC workflows, but they are of absolutely crucial importance in the EC area and play less of a role in most other application areas of workflow technology.

— EC workflows in-the-large usually involve several autonomous parties and would therefore often run in a truly distributed manner on several, possibly heterogeneous workflow management systems.

— The actual EC workflow will be accompanied or surrounded by state-of-the-art security protocols for encryption, authentication, digital signatures, etc. Without confidentiality or even anonymity and, especially, the guaranteed value of "cybercash" currencies, EC applications cannot be successfully established.

— EC workflows will ideally become mass operations: millions of customers initiating billions of workflows every month or even week, with heavy load peaks in certain popular time periods.

— Frequent failures and unavailability of EC workflow servers would immediately weaken the market position of the merchant (or the EC provider that acts on behalf of one or more merchants) and eventually result in significant financial losses.

From these properties we can infer the following execution-related requirements that EC workflows pose on the underlying workflow management systems:

Scalability: The mass-business characteristics of EC workflows requires a high-throughput execution engine. Thus, load distribution across multiple workflow servers is necessary to ensure this kind of scalability.

Dependability: This subsumes a number of technical requirements—fault tolerance, reliability, availability, and also security. The execution environment for EC workflows must be fault tolerant, in that, transient failures of system components (e.g., workflow servers) do not lead to incorrect behaviour of the workflow execution. Further, failures that hamper the progress of workflow executions should be very infrequent; that is, the overall execution should be highly reliable. This may already call for redundant components within the system, so that many failures can be completely masked. In addition, the few failures that cannot be completely masked should be repaired sufficiently fast, for example, by restarting system components based on logged data. The net effect should then be very high availability of the EC services: customers are satisfactorily served without noticeable delays with extremely high probability. Finally, security measures such as encryption must be well coordinated with the other system components to ensure the desired overall dependability.

Trackability: Comprehensive traces of EC workflow instances must be available for documentation purposes and possibly also for data mining towards better marketing. This calls for an elaborate workflow history management component with powerful temporal querying. On the technical side, extremely reliable and long-lived archiving must be provided, too.

Workflow specifications serve as a basis for the largely automated execution of business processes. They are often derived from business process specifications by refining the business process specification into a more detailed and executable form. Automated and computer assisted execution means that a workflow management system (WJMS) controls the processing of work steps, denoted as activities, within a workflow. Some activities may have a manual or intellectual part, to be performed by a human. But the workflow management system is in charge of determining the (partial) invocation order of these activities. In contrast to business process specifications, this requires a formal specification of control flow and data flow.

Build-Time Issues

The use of workflow management systems in different application areas has led to different approaches for specifying workflows. There exists a wide range of languages including script-based, net-based, rule-based, and logic-based languages. They differ widely in their expressiveness, formal foundations, and ease of use. Beyond the "pure" approaches that fit into the given categories, hybrid methods have also been proposed. Because of its designated role as an industry standard, the Workflow Process Definition Language (WPDL) of the Workflow Management Coalition is especially notable among these hybrid methods.

Provable Correctness

For specifying the sales phase of EC workflows, a high-level specification language with mathematically rigorous semantics is required. Its semantics must be sufficiently expressive for the complex protocols executed during the sales phase, Control flow as well as data flow must be captured in order to model all aspects which are relevant for the correctness of the EC workflow. On the other hand, the specification language should be simple enough to support automatic verification of workflow properties by means of verification tools.

The formalism of state charts and activity charts meets these demands. In the Mentor project, we have explored its use for workflows, and it seems that state and activity charts are gradually finding wider use and acceptance. They have been adopted for the behavioural dimension of the UML industry standard.

State and activity charts comprise two dual views of a specification. Activities reflect the functional decomposition of a system and denote the "active" components of a specification; they correspond directly to the activities of a workflow.

An activity chart specifies the data flow between activities, in the form of a directed graph with data items as arc annotations. State charts capture the behaviour of a system by specifying the control flow between activities. A state chart is essentially a finite state machine with a distinguished initial state and transitions driven by Event-Condition-Action rules (ECA rules). In the Mentor project, we have investigated the feasibility of verifying certain properties for workflows specified as state and activity charts. A prerequisite for the verification is the modelling of these properties in a formal language. For this purpose, variants of temporal logic are well established.

Temporal logic extends predicate logic by temporal operators. With these operators, invariants (both safety and liveness properties) can be modelled, such as: 'If p was true in the past, q will be true at some point in the future'. A temporal logic with an expressive power suitable for properties of workflows (according to our experience) is CTL (Computation Tree Logic).

Once the critical workflow properties are formally stated, they can be verified against the formal specification of the workflow. Two different approaches are possible here. The first approach is using theorem provers, which automatically verify that a given specification has the desired properties. However, in many cases this approach is computationally infeasible. As a less powerful, but much more efficient approach, model checking can be used. In essence, model checking verifies whether a given finite state automation (the workflow specification) is a model of a given temporal logic formula. The most efficient variant of model checking is known as symbolic model checking. Symbolic model checking is based on a compact, symbolic representation of a finite state automaton in terms of an ordered binary decision diagram (OBDD). Because state charts are closely related to finite state automation, model checking can be applied to state charts. Tools for symbolic model checking already exist (at least as research prototypes) and are used mainly in hardware design and for reactive embedded systems. In the Mentor project we have used a tool for symbolic model checking described in. Although the resource requirements were high, we found symbolic model checking to be suitable even for the interactive verification of workflow specifications of non-trivial size.

As EC workflows involve multiple parties with already existing business processes, the specification language has to provide means for combining existing workflows into higher-level specifications that provide more comprehensive EC support.

Starting with a skeleton workflow, e.g., only specifying the existence of the sequence of pre-sales, sales, and post-sales phase, workflow composition can be used to dynamically refine the workflow, even during execution. The mechanism of nested states, i.e., substituting states by entire state charts provides this composability for state and activity charts. The semantics is that upon entering the higher-level state, the initial state of the embedded lower-level state chart is automatically entered, and upon leaving the higher-level state all embedded lower-level states are left. In addition to the composability of workflow specifications, EC workflows need a flexible exception handling. Exception handling code should not be sprinkled throughout the specification.

Otherwise, the specification can easily become unmanageable. In state and activity charts, orthogonal components provide a convenient means for incorporating exception handling in an easy and clear (i.e., truly orthogonal) manner. Orthogonal components denote the parallel execution of two state charts that are embedded in the same higher-level state.

Both components enter their initial states simultaneously, and the transitions in the two components proceed in parallel, subject to the preconditions for a transition to fire. Using this concept, an execution handling for a "watchdog" component would simply stay in its initial state until a certain condition becomes true or a certain event is raised.

Run-time Issues

The run-time requirements for EC workflows stated previously are quite similar to those for current business workflows in large enterprises. Workflow management systems supporting such workflows or even workflows that span several enterprises also have to face scalability, dependability and trackability problems.

At the merchant's and at the bank's or credit card company's site, a large number of EC sub-workflows have to be executed concurrently. This is usually not the case at the customer's sites, as a customer will typically be involved only in the purchase of a small number of items at a time. However, this does not mean that the issue of scalability is limited to the merchant and the bank or credit card company. In fact, *the need for scalability may arise in three different forms:*

1. Upward scalability is needed mostly from a throughput perspective, that is, to cope with larger numbers of workflow instances simultaneously while still guaranteeing acceptable response times. This is the kind of scalability required for the merchant and the bank or credit card company. Upward scalability is usually achieved by using distributed, multi-server workflow management systems.
2. Downward scalability is desirable towards the lower end of performance demands in environments with scarce resources, the goal being a light-weight workflow management system with a small footprint. As a customer is typically involved only in a limited number of EC workflows at a time, and rarely uses advanced functionality like analysing the history of a workflow, light-weight and tailorable workflow management systems are the systems of choice at the customer's site.
3. Sideward (or horizontal) scalability aims at a high degree of interoperability with other workflow management systems. This is of crucial importance for EC workflows.

In all phases, EC workflows involve multiple parties, which will typically use different workflow management systems. The Workflow Management Coalition (WFMC) has proposed a standard interoperability interface for workflow management systems. However, this interface is currently in a proposal stage. It remains open how suitable it is for EC workflows.

Dependability

The sales phase of EC workflows requires a highly dependable execution platform. This subsumes the following issues:

1. *Fault Tolerance:* In a fault tolerant execution, system failures do not harm the correctness of workflow execution. When a failure occurs, the system is rolled back to the most recent consistent state. Workflow execution can then be resumed from this state. In data systems, fault tolerance is achieved by atomic transactions. Workflows are typically more complex, consist of more activities and are of longer duration than database transactions. Workflows are usually executed as sequences of short transactions. In case of distributed workflows, distributed transactions are required. Middleware services like TQ

monitors or OTS transaction services are used for controlling the execution of distributed transactions, also in case multiple different database systems are involved. Reliable message queues provide the "glue" between transactions implementing the steps of a workflow. Their usage guarantees that after a transaction has successfully finished, the next transaction is eventually started.

2. *Reliability:* In reliable systems, failures must be infrequent. Obviously, this requires well designed and rigorously tested software, running on reliable hardware platforms. Workflow management systems rely on database management systems as one of their core components, which have proven to be among the most reliable software products. In addition, redundancy can be used to increase reliability. For example, a failed server process can be dynamically replaced by a "stand-by" process that has access to the underlying persistent data (i.e., a database and/or a log file) which may again be replicated if necessary. This way, even less reliable components can be used in a system with high reliability demands.
3. *Availability:* Highly available systems ensure that users find the system operational with very high probability. This can be achieved by ensuring that complete system failures are sufficiently infrequent and by providing very fast recovery and restart after such failures. In addition to very low system downtimes, a workflow should almost never be delayed in a user-noticeable manner-Redundant workflow servers are again an appropriate means to increase availability. For less demanding workflows, for example in the pre-sales phase of EC workflows, simpler modes with lower or no redundancy should be supported to improve the cost/performance ratio of the overall system. Such a configurable approach has been studied in the Exotica project as an extension to IBM's Flowmark.
4. *Security:* EC workflows probably have the highest demands on security among all classes of workflow applications. Up to now, this issue has only received limited attention in the workflow community. Hiding confidential workflow data from unauthorized third parties can be achieved by using encryption protocols. In addition, specific secure payment and negotiation protocols have to be implemented, for example, as a special high-security class of workflows.

In all environments where non-trivial amounts of money are handled, trackability is a major requirement. The need for archiving the history of workflows is already considered a major requirement in many workflow management systems, because the history can also be used to improve the underlying business process model, e.g. by determining workflow steps that frequently miss their deadlines, involve more resources as expected, etc. Improving the overall workflow is typically based on aggregated, statistical data, whereas tracking EC workflows requires detailed data on a per instance basis. Workflow management systems can provide the necessary archiving service for this purpose as well.

Research Demands on Workflow Technology

The previous sections have shown that using workflow technology in EC applications is promising and already solves many of the encountered problems. However, substantial further research is needed to solve problems that are not yet addressed by workflow technology and arise in the new application environment of EC.

Dynamics

The need for dynamic adaptation of workflows has already been recognized independently of EC. For example, in medical applications, the treatment of patients can usually not be determined in advance. Instead, the corresponding workflow has to be constantly adapted to new diagnostic results. The current generation of workflow management systems still provides only primitive support for such dynamics. Systems supporting modifications of control flow and data flow at run-time are often weak in other respects such as administration, correctness of execution, reliability, scalability, and security.

EC workflows involve multiple, autonomous parties. Thus interoperability is crucial. A specific kind of interoperability has already been addressed by most workflow management systems, as they have to provide access to arbitrarily heterogeneous applications. This is achieved through the wrapper/ mediator paradigm, that is, encapsulating the actual application, surrounding it with a standard-compliant interface, and invoking applications through a broker service (the mediator in this case). Usually appropriate middleware such as CORBA or DCOM is employed. Only little work has been done, however, on the interoperability between different workflow management systems. In this regard, EC workflows have higher demands than most other workflow applications, since EC involves multiple, autonomous parties by definition.

The most promising approach to address the interoperability between different workflow management systems seems to be again a wrapper/ mediator architecture, where each system is encapsulated and enriched by a common language for process definition., state introspection, and manipulation. Once such 4 common language were established, different systems could communicate in a straightforward manner, possibly again through some “broker” or “mediator, services. More research is needed along these lines.

Security issues have to be addressed jointly with the security research community. The challenge is to integrate security protocols into workflow management systems in a seamless yet easily maintainable manner. This is not a straightforward task, as multiple components of workflow management systems such as communication managers and history managers are involved. Modifying these components accordingly may conflict with existing solutions for scalability and availability. In addition, different levels of security have to be supported in order to execute, for example, information gathering workflows in the pre-sales phase as inexpensively as possible. If different security levels are used in different phases of an EC workflow, we have to make sure that the security of the critical workflow parts is not compromised by parts executed at low security levels or by transition problems between security levels.

Electronic commerce (EC) can be regarded as a specific kind of business process involving several enterprises. As workflow management systems have been developed to implement such business processes, EC can largely benefit from current workflow technology. In particular, the core of electronic commerce, sales and money-handling transactions, can be modelled as special kinds of workflows. We have identified the main requirements that EC poses on workflow management systems. Pre-sales and post-sales phases of EC workflows require highly flexible, light-weight workflow management, whereas the core sales phase including money transfer has to be executed in a highly fault tolerant, reliable, and secure manner on a highly available system.

TWENTY TWO

E-commerce in India

The Marketing activities, concepts and strategies are bound to change due to changing situations and environment. Since 1985 a competitive, customer driven economy has been emerging in India. The abolition of Licensing and liberalisation of market entry regulations since July 1991 has further added impetus to this process. It has been given the final shape by the formation of WTO at Markesh. Most of the people and the industrialists have felt that whether a company wants to go for global competition or not, in the present liberalised situation the global competition will arrive at their door step. By endorsing the final act of the Uruguay Round the developed world has in effect endorsed free trade in a multilateral setting, a phenomenon that will work to the advantage of developing countries that need new markets. Countries like India will get an opportunity through WTO to negotiate better deal for itself. This can be possible only when the quality and standard of the products is improved. The policy of liberalisation has created a situation in which companies like Uniliver, Glaxo, Pepsi, Whirlpool, Phillips have increased their stake of equity. New products, new brands and new formulations are being announced virtually every week, mergers and acquisitions like Brook-Bond-Lipton designed to build competitive fortress have become more common. "Down Sizing" restructuring and engineering are the words that have potent force for emboldened executives. Thus, the whole market environment has undergone a severe change. India has become one of the largest emerging markets attracting substantial flows of equity investment and also consumer goods. It focuses that there are challenges of improved products of advanced nations, on the one hand and on the other hand, if the executives have vision, foresight and creativity there can be enough opportunity for domestic products both in India and in other emerging markets.

The first challenge by opening up the Indian market is of survival of Indian companies and Indian brands. Brand here understand should be considered as a brand created for the Indian market and owned by a company of Indian origin. Brand building is directly related to the scale of business. Indian companies like Lakme or Nirma do not have financial boost or technology like Hindustan Lever or Colgate Palmolive. It has been the experience in the past that whichever country has permitted multinational companies to enter their market their local brands have been wiped away. Indian brands cannot have a different situation. Though Swaminathan, S. Anklesaria Aiyar editor of the *Economic Times* feels 'A very large number of low end consumers have never heard of foreign brands and are not hungering for them". The magic of the foreign brands cannot be denied. A section of new middle class definitely runs after them. If the foreign brands in the long run become Indian or perceived as Indian like Lux, Surf, then the matter is different but a Coke will always remain foreign. What attracts the consumer towards the foreign brands produced by multinational companies (MNC) is firstly the quality and secondly the status symbol and a new way of life.

The survival of the Indian brands will depend not only on brand, but on product, the company backing it, and the people who run that company. Each of these factors determines the state and future of the brand. For example Amul has been promoted so consistently that no MNC will easily shake its grip on this Rs. 160 crores butter market. However, majority of Indian companies do not invest consistently and also adequate amount, they want quick harvest. Those Indian brands will survive which have a very clear positioning. Brands which depend not only on the advertising buck but on the entire gamut of brand building activities and what the consumer perceives as the value delivery. The technology difference won't be as important as the speed of action, the right package at the right price, the right upgradation at the right time. The products in which aesthetic function is more valuable than the physical attributes will face severe problems of survival. One may recall that in the past American consumers forum had protested on the import of Indian Shiffon Skirts and the USA government had issued a notification to stop their imports.

Another important issue related to survival of Indian brands is that there is lack of commitment to product quality in the Indian producers. What is important is that the product should meet the consumer need. You may recall that when American Cars were best, those who could afford bought them. Today when Japanese cars are the best every one wants to buy them. This is the function of quality not of origin or of brand alone. Indian brands have been isolated since 40-50 years, therefore the product quality does not stack up to international standards. In fact, branding is a process for adding distinctiveness to products or service which offer the consumer quality, value and satisfaction. Apart from Superior Quality what foreignness also signals is status and a new way of life. The producers will have to take care of this need. If they do not look after their consumers need somebody else will take care of their need. If the Indian producers do not realise these facts soon, they will face serious threat of loosing the market. In the liberalised economy a market once won can not remain for ever for a particular company or product unless the company is continuously engaged in innovation of technology and products as per need of the consumers.

Another important challenge is from the technology side because people look for a product that has a good image and is backed by good technology. It is believed that the multinationals have better technology to back their products. Indian producers have always imported technology and have never made any attempt to Indianise them or improve upon them and have a better technology of their own. Technology imports have been costlier and will definitely increase the cost of production. To meet the challenges posed by the multinationals it is essential that the Indian producers should also use better technology. Bajaj is one example in spite of the fact that Bajaj entered in business of two wheelers with the imported technology, they have Indianised it and they have also continuously made effort to improve the technology and improve the product. The Bajaj Scooter is one of the quality product backed by a better technology. No multinationals can make Bajaj to loose its market. If a company has invested in technology and in quality of its product it will survive like Bajaj & VIP Luggage because they have ensured that the product lives up to the promise that they have made. Hence, the Indian producers have to upgrade, update, renovate and improve their technology to have quality product which can match the foreign products and meet the consumers need both in domestic and foreign markets. This will definitely depend upon a company's access to the technology, if it can buy the technology it is better, if not they will be compelled to enter into joint ventures or alliance with foreign companies to have access to their technology. Only having the technology is not enough. There is need to invest in technology improvement and continuous upgradation of technology. Technology and its development has become one of the essentials of economic development. We are aware that only one company of Japan invests an amount in R&D of technology which has been equal to the total amount of investment in R&D in India.

The opening of Indian market for world producers and the market of other countries for Indian producers will have any meaning only when the Indian producers have any product backed by a good technology and superior in quality to match the competitors product. Hence Indian producers need continuous innovation and investment in innovation to survive.

Everyone, who is a student of marketing might have found that in Indian market premium brands are flooding the market to block the entry of the products of MNCs several premium brands in the last two or three years have been priced up to 10 times higher than the popular brands. This is based on the idea that as incomes of middle class increase there will be demand for products which may have distinction and lifestyle. Thus, the premium brands of Indian producers will pose a challenge to the other products. The premium brand producers are happy as they quote the survey report of National Council of Applied Economic Research (1993) that "there are 6,5 million middle class households earning over Rs. 18,000 per annum out of which 3.7 million earn over Rs. 78,000 per annum. One million households earn more than Rs. 1 lakh a year "In 1993 the purchasing power of average Indians had been equivalent to $ 1,150 (Rs. 35,650) per year. Thus, they have been guessing about the market of premium brands good among these "Maruti Millions" or the status seekers. A host of products like Lacoste's shirt priced for Rs. 750, Park Avenue shoes for Rs. 1600 and so on will meet the demand of the new status seeker middle class consumers. It is true that only a minscule of Indian population will buy. These products will have a core group of buyers and large number of floating buyers. Thus, the premium brand products are posing challenge to other products. Let us call it Intra-country competition. The challenges will be compounded because the consumers may buy the usual brand and also the premium brands for occasional use leading to low frequency of repeat purchase of premium brands, secondly the fast moving consumer goods consumers will float not only between brands in a product category, but also between the categories themselves. Hence even the producers of premium brands will have to face problems in selling their products. Lastly, the introduction of intellectual copyright was introduced in 2000 hence the imitation of product has been beaten down and out. The premium products producers will have to adjust to the socio-economic conditions of the country and produce products for higher price which will give value to consumers for the money paid.

Thus the changing economic scenario opening up the Indian market for the global producers has definitely forced the domestic produces to adopt a new production strategy and a new marketing technique. The change definitely throws some challenges but it provides some opportunities also.

The Liberalisation of Indian economy and going for global and entering into an agreement to be a member of World Trade Organisations, India will now have an Institution-WTO, to negotiate better deal for itself through discussion and creating consensus among the contacting parties. As Mudra's Krishnamurthy puts "the globalisation of the Indian market needs to be seen as an opportunity to become MNC and not as a problem." Building competitive advantage through exposure to international markets should be a conscious strategy. Companies from countries like Japan and South Korea have learnt from the West and then successfully marketed world class brands globally. The Chief Executive of a Japanese car manufacturing company which was established with the help of Technological assistance of General Motars after fifteen years visited General Motors on return he called, his executives meeting and told them that they can go for World Market now. His executives expressed doubts. They were sent to visit General Motors and on return they confessed that they have improved over General Motors and can easily compete in the World Market. Now Japanese cars are in American market and in European Marketing competing successfully and in some cases sharing major portion of automobile market. Japan had imported Technology for Watch manufacturing from Switzerland. It is Japan which added day and dates music and varieties of alarm sounds and finally automatic and quartz and captured major share of world watch market. Indian producers will have to learn from Japanese experience to avail the opportunity of going in foreign market. Clarity of thinking and a single mined focused strategy on improved technology and Superior Products are key to realise the potentials and avail them. In fact, the policy of globalisation has been adopted with an idea to make Indian producers competitive globally. Hence the opportunity for Indian companies to become efficient MNCs has been provided by this situation.

Mr. Balkrishna Zutshi Indian representative at WTO expressed that the changes in economic policy and signing of agreement with WTO provides access to global market. There is tremendous opportunity for India to integrate itself with the Global market. It is now upto the Indian industries and to some extent the government, to adopt the right policy to exploit the countries potential in the world market and to improve its export in different areas.

One of the new areas for achieving some share in world market is the service sector. In the recent agreement some concessions have been provided for the export of services. They form a good basis for further development. This particular area which can interest India, through the movement of natural persons as service providers, has tremendous opportunity. Further, we can increase our exports in textiles, but we have to modernise our textile industry in next ten years. The export of textile will pick up.

India has great potential in agricultural sector. The government can give subsidies upto 10 per cent of the value of production on fertilizer, water, pesticides, seeds and electricity. It can help in improvement of quality and quantity of product. India can export rice and other agricultural products to Japan and Korea fulfilling upto 4 per cent of their demand .

Cashewnuts, tobacco, cereals processed food, sugar and tea can have greater potential for increased exports. The quality, size of packing and value addition is the need for achieving this. If a group of exporters who are quality conscious who can keep pace with international changes and bring in technology make efforts may open opportunity for growth of exports.

There is another area where Indian producers can think of entering into global market, the readymade garments. The Gokal Das Exports have been exporting shirts, to Europe, Australia, Middle East, Singapore, U.S.A., Canada and are regular suppliers to Levi Straus and Wranglar. Even while international brand names like Pierre Cardin, Benetton, Van-Heusen, Lacoste and Arrow are making inroads in domestic market, Indian companies are betting into high technologies to tap emerging global opportunities. Indian companies have no choice but to invest heavily in developing exclusive brands and entering into non-cotton fabrics. However, India can be a potential exporter in this area. There is enough scope of exports of shoes and other leather garments. Liberty shoes and Namaste Exports have shown the way. Korean leather garments have become costly and European markets have started opening up mainly due to the closure of tanneries there. The scope of export of leather products has increased. According to a study conducted by the Federations of Indian Export organisation, "Indian exports are extremely competitive in special fields like leather, garments, handloom, jewellery and engineering goods where it is much easier to adopt to changing trends of a foreigning buyer who suddenly asks for a different auto components, or a different fashion of handloom cloth or leather garments wherein it is much easy for a small producer to change his production process. The only need is to understand the consumer need, improve the quality of the product and add value through packaging and design. The competitive edge can be added by adopting and innovating required technology and a marketing strategy for the changing situation.

Issues

One can now clearly feel that the issues in marketing at the beginning of the new millennium are, to understand the consumer needs and anticipate their expectations whether they are in domestic market or in other countries market. In fact, in an ACME Seminar of top executives the consensus has been that all the successful brands have designed and developed products on the basis of the anticipation of consumers needs and expectations. The second important issue which was emphasised at the seminar by these executives was the need of being "unique", "innovative", "Positioning and therefore becoming competitive, in the way they were "offering" themselves or the way they were presented and communicated to any prospective buyer. Thirdly the executives of the companies of successful brands have confessed that the secret of the success of their company has been the product itself rather than any other marketing technique or inputs such as sales-force, advertising, promotion, quality of overall managements, etc.

If one examines, a few domestic products of wholly Indian companies one may be able to bring home the important issues of marketing at the start of the 21st century such as Nirma (Popular priced washing powder), Titan watches, Onida (TV), Videocon Washing machine and Maruti Cars, Maggi Noodles, and Close Up Toothpaste. Each one of them presents some significant breakthrough in terms of product development and market process.

Nirma, the most often quoted brand in almost any marketing forum in the past ten years entered with almost no backing or background in consumer marketing, but in a period of five years, it emerged as the greatest competitive threat ever faced by Unilever's flagship company in India, Hindustan Lever. Ignored by all those in the detergent market for nearly five years as being not in our segment, not our kind of product, this brand launched at almost third of the unit price of Unilever's Surf grew explosively almost ten fold in a ten years period and overtook everyone else in the process. Titan a company launched by Tata, India's largest private sector entrepreneurial group repositioned the watch and especially the Indian made watch in the eyes of the consumer, and made it into a stylish personal accessory to both men and women. In just four years, it grew to 2.6 million watches and No. 1 in the quartz segment and No. 2 overall, behind the much older and slower, HMT, it also rewrote the 'rule book' of conventional wisdom in watch marketing, that the best watches are imported (legally or illegally).

Some of them created new segments in the market and new usages and users. You may recall Ketchup, for example had long been thought of and used as an essentially western product, which could not easily be blended with Indian cuisine. Maggi by Nestle not only introduced alternatives to the tomato sauce, which was generally too bland for the discerning India palate, but showed it in association with a number of essentially Indian snacks such, as *samosa* and *vada*. An even greater innovation was the 2 minute noodle. Long known both in Europe and the Far East as a meal by itself, the noodle could not be expected to replace rice or the home made forms of wheat such as *pooris* and *chappatis*. However, interestingly presented and positioned as an afternoon, after school snack, it caught the imagination of children and mother to become an almost unique product that virtually owned the territory it developed.

A few of them took competitive standards several niches higher, for example; The Maruti small car (800 cc) gave the Indian motorist a stylish, manoeuvrable and modern alternative with a compact design. Of course, the fact that the technology was several decades ahead of the Fiat and Morris Oxford versions (Premier Padmini and Hindustan Ambassador), ensured that it virtually had no competition. The Maruti 1000 which followed in the late 80's had an almost unique position as "Japanese quality premium personal car". The latest offering is the Maruti Zen, the first Indian car to incorporate the jelly bean shape and rounded edges, to reduce wind resistance and increase fuel economy, along with five speed transmission all with the full impact of Suzuki of Japan now the majority owner of this company. Three out of every four cars sold in India, now come from Maruti. In the home appliances field little known Videocon emerged in a period of 4 years as the major share leader with 27 per cent of the colour TVs market and 50 per cent of the washing machines market. They flooded the market with Japanese style product introduction rate, concentrating on wooing the dealers with substantially higher margins and array of incentives and beating the competition on price. In this last aspect, they had much in common with Nirma in the detergent market and Hero in bicycles and mopeds. For India's largest and most successful consumer product company, Hindustan Lever,, "Close Up", brought success at last to the personal products division in the internal battle against Colgate in the toothpaste market. Relaunched as a gel with variants it attained a 15 per cent market share. Only two of the above brands "Maggi" and "Close Up" are owned and marketed by multi national affiliates in India. All the others are not only from wholly owned Indian companies but in many cases from entrepreneurial first generation business group. Neither the Dhoots of Videocon nor Mr. Patel of Nirma started with any accumulated experience in a related industry. Nor did they have the war-chests from promotional onslaught which a Coke, a Pepsi or a Proctor & Gamble would subsequently bring to the Indian Scene. The above examples of products and their marketing way clearly gives any one to understand that being multinational *per se* cut no ice with Indian consumers that sheer power of resources is not a pre-requisite for success and competitive edge; That what is important for being competitive in the future is to be innovative in product as per need and expectations of the consumer and therefore, there is urgency of improving the product and service quality standard; finally that innovation will have to be applied to every segment of business, retailing, general management, quality structure and even financing over and above this speed through which it is implemented.

The success of the five brands mentioned above proves that there is the need to understand consumer behaviour in the Indian context and to change the approach of strategic marketing. That to succeed in any competition market whether domestic or global the application of the same approach applied prior to 91 will not be effective. Whatever may be the individual ingredients of success strategic style or approach must concentrate on rewriting the 'rule book'. In earlier decades the dictum was that marketing orientation meant not only doing things differently (the USP School) but better and more memorably (as emphasised by David Ogilvy and Ries and Trout in the brand personality and its positioning eras). Today beyond just being different or better thinking about the market and conceiving the future differently from ones predecessors, seems to be at the core of successful strategic marketing. The core and key factor in the above five brands has been innovation to meet the future opportunities and understanding the needs of the consumers also has a high position. Innovation should not be taken as more and more intensive and extensive research in the status quo of consumer; No amount of quantitative analysis of consumer motivation or market segment on quantification can be said to be real alternative of innovation.

Innovation; therefore, seems to come more from internal thinking daring and serendipity than analysing the industry in two by two matrix in mind numbering detail.

Very recently the executives have started talking about the fifth 'P'. The pace with which you adopt and improve a product launch and relaunch a product, use a technology, how speedily your plans implement any strategy is more important. Now they feel the fifth P only can form the base of new marketing strategy.

I wish to point out at this stage that student and teachers of marketing should develop a technique of marketing forecast to foresee or have understanding and perception of what marketing in India in the new millenium and after, would demand of them. We find that the need to understand consumer behaviour in the India's context leads the rest of things behind. Assessing the relative effectiveness of different forms of marketing inputs on a cost benefit equation. Conventionally marketers have always tended to think of most marketing inputs other than the soft area of advertising as fairly measurable in terms of sales response. Every company has its own rules of thumb developed over the year from their own or their international affiliates based on experiences on how much is desirable to spend on promotion distributions, product launched, or percentage of sales required to sustain market shares of stable on going brands, and so on. In the changed situation the question obviously is whether the old rule will be valid in changing context and how far such rules will apply in the changed situation. Here then is a rich area for constructive and continuous collaboration in research between the thinking manager and the user-friendly academics.

Strategic Marketing is a dynamic concept and it is still under evolution. In India situations have been changing, liberalisation policy, opening up of Indian market for multinationals, and the Government vacating the board rooms, the customer being ushered in effect being given a place of pride, has brought changes in economic environment. It has posed three major challenges. Increasingly competitive market with new entrants providing superior product and service, a subsequent change in most industries from a seller market to buyers market with quality and price conscious consumers; and necessity to succeed in globalisation economy by exploiting potential sales market.

Marketing even in a global situation is a process of perceiving, understanding, stimulating and satisfying the needs of specially selected target market by channelling an organisations resources to meet those needs. Marketing is thus matching dynamic interrelationship between a company's products and services, the consumers want and needs and the activities of competitors.

Peter F Druker has also mentioned that old system of marketing cannot be useful today. He said "But the right marketing knowledge won't be of much help in these turbulent, competitive fast changing decades. It requires the right marketing actions." Japanese define market by the use of their product. The customer rather than the makers determine the market and Japan became a forceful competitor in the global market of automobiles, T.V. etc.

Hence even in India most of the Chief Executives have now been focusing on developing consumer oriented marketing strategies and brands. One is reminded of Peter F Drukers view "Marketing is business seen from the view of the final result, that is, customer point of view". The advocates of this view have been now talking of customerisation. There is another approach also for competitive strategy. As in a competitive market there is fight between the rivals in every retail shop for each customer but from domestic and international producers some executives feel watching the rivals move only can help in developing a competitive strategy. Wayne-Calloway has pointed out "Nothing focuses your mind better than the sight of a rival who wants to wipe you off the face of the earth". Can Druker and Calloway both be correct? Can the classical marketing which involves studying the consumer alone be a tool for successful marketing in the liberalised market?

There is a battle for the market, every marketer is trying to retain or gain market. Hence in their own way they are adopting various strategies. It is now time for Indian producers to think new about their strategy, the way they conduct their business. The Indian marketing has come of age. In 90s the marketers often seen as intent on getting each other as they are on getting consumers. Some ten years ago there was hardly any competition.

Each market was dominated by one producer. In some fast moving consumer goods there was some kind of an armed truce. Rival marketers kept themselves out of one anothers way. By the late 80s the stream of new brand launch and flood of MNC products has gathered force. And the flood gate opened after 1991 upset the whole status quo. Each year an established brand had to deal with many new factors that could affect its fortune.

Though in a number of seminars executives have emphasised that firstly the basis of success is the product and its quality rather than any other marketing input, secondly the executives have emphasised the improved technology as the basis of the successful strategy of marketing and quote the success of a few Indian products e.g. Bajaj and Maruti. The continuous improvement and updating of technology by Bajaj has proved, "Nothing can beat Bajaj" and the continuous change in technology has made Maruti as one of the most desired car and now it occupies a place of pride. Out of every four cars sold today in Indian market three are the Maruti cars. The product and technology has placed products in a unique situation and made them competitive whether they were offering themselves or the way they were communicated to any prospective buyer. Thirdly executives have been quoting the strategy of distribution adopted by Videocon washing machine as the key of success of the product in occupying 50 per cent share of the washing machine market. The most successful products; Nirma, Titan, Maruti, Maggi, Videocon washing machine and 'Close Up' have been quoted to prove that product, technology, price, the process of offering the products and understanding the need of consumer have been the cause of their success. Fourthly the executives have expressed that the product innovation and the technology innovation have been to satisfy the consumer. One cannot deny these facts.

Now it is no longer limited to watch the consumers only. It is also essential to watch your rivals. Today marketing has become a battle of perception and not just a product which explains why competitive marketing is gaining ground in the overall strategic process. The product innovation or technology adoption may be one aspect but not as important as 'Watching your rival' and 'the speed of action' the right package at the right price, the right upgradation at the right time. Hence in a competitive marketing 'watching the rivals' and the 'pace' with which you react are two strategies seem to be important and have been deliberated here in this chapter.

Watching the Rivals: Now it is not enough to watch the consumer only it is essential to watch your rivals. Product differentiation is lowering its sharpness hence watching the action of rivals about distribution and such other things have become crucial, and the 'pace' of marketing is quick really quick, like combatants in power tennis, marketers have to anticipate where the opponent will place the ball. Trying to get it after the ball has been struck is too late.

Intense competition has forced the companies to be more specific in watching their rivals. Now companies try to have resource audit, inter-firm comparison, checking rivals financial cost-price analysis, profit margins, man-power cost, in fact everything under the microscope. The route to consumer is always through competition. Since rivals are targeting the same consumers, marketers must anticipate possible retaliation in the execution of their own strategy.

By watching the rivals offerings you will ensure that you have not missed out something that may affect sales. For example; if Bajaj Auto would have reacted earlier to ailing LML attempts to fragment the scooter market by LN4L, might have been unable to 'turn it round' and share some market of scooter from bajaj. There is another more vivid example of not watching the rival and facing a disastrous consequence by 'Farex' of Glaxo when Nestles baby food Cerelac came. Cerelac being milk based was preferred by doctors and mothers. Cerelac offered a range of variants. Cerelac was preferred because milk shortage can be managed. Doctors recommended Ceralac because if

Farex upsets the baby's stomach it was difficult to identify which cereal has caused it, and Ceralac offered a range of variants. There are other examples of not watching the rivals such as; 30 year ago Milton appeared, by continuous innovation, segmentation and sub segmentation of market. Milton has acquired the position of No. 1, in Rs. 250 crore, thermo ware market. In 1986-Cello joined, and is now no. 2. Eagle lost the market of thermoware, except flask to Milton and Cello.

We have example of BPL which anticipated the arrival of global brand and quickly expanded its consumers electronic range to cover a widespread size and price segments and block as many entry points of MNCs as possible. Another example is of the Proctor and Gambles which in August '95 quickly flooded the Baby market with its 'Paper Brand' of diapers to pre-empt the launch of Kimberly Clark Huggies marketed in India by Kimberly Clark Lever Ltd. by watching its rival and by acting quickly, really quickly Proctor and Gambles threw its rival marketer off its strides. If you are not watching rivals you are doing only half the marketing.

Unfortunately in India it is still informal, not database. 'Indian Companies are still like babies-tend to fiddle with their navels, instead of studying the external environment; rivals move rather than theirs. In fact any strategy cannot be made in isolation of environment, existing market situation. Broadly speaking a marketer has to monitor the three 'Cs' — Consumer, Competition, Cost. In responding the competition the company's cost can be affected and in turn influential company's try to damage the competition example is of *Times of India* changing its advertising rates every time. In fact extensive attention by a rival-specially if it is a larger one can unnerve the marketers example being of Standard Batteries Ltd. coming with 'Signature' and two years guarantee unnerved the Exide. The Exide people were in so much hurry that they did not even change the warranty card and directed the wholesale and retail outlets to issue two year warranty. When Standard introduced Calci charge, a maintenance free battery, Exide tracked its performance for four months and it brought down its new product, 'Exide Freedom' for Rs. 200 below Calci charge. Exide is compelled to scrutinise standard's activity. When Jhonson and Jhonson was planning to relaunch SAVLON an antiseptic liquid brand bought from ICI Dettol felt its monopoly threatened. It's marketer RCI created a media blitz and launched Dettol medicated plaster which coincided with Savalon's launch. The strategic objective was to heighten brand salience for Dettol while forcing Jhonson & Jhonson to splutter for defending its leader Band Aid.

However rival watching is a complex phenomena. Number of parameters on which they have to be watched have grown. Rivals have to be watched and their move is to be anticipated well in advance. They have to be watched cautiously and properly otherwise many a time marketers end up in watching the wrong kind of competition. In any case, quality product and improved technology alone can not help the marketer in a competitive market. Watching the rivals is very essential. Further watching the rivals may be important; but the product attribute the technology difference and rival watch wont be as important as the 'Speed of Action' — the right package at right price, the right upgradation at the right time, and the right action to put the right blockade for the rivals at a right moment i.e. the pace strategy.

Pace Strategy: The fundamental four 'Ps' have been found insufficient for developing any marketing plan and marketing strategies in the changing competitive market. The smart and intelligent executives have discovered that 'PACE' is another strategum by which they can have an advantageous position in the competitive, crowded market created by the policy of liberalisation.

In their opinion 'Pace' or Speed is the many edged weapon to be used on all fronts; whether product innovation, technology upgradation, promotion or pushing the product in the market. "Every marketing man can match the four Ps of marketing. But the ultimate differentiater is the speed to market".

This was never realised by corporate sector in India before liberalisation and globalisation of Indian market they were contended with sporadic launches of products. A company like Hindustan Lever had never relaunched its product until four years after its introduction knowing that it would milk its money making potential. With the open economy and open market, traditional tools have lost their significance. All marketers are trying to adopt 'Pace strategy' to protect their market share Speed strategy is being recognised as the important element in a competitive market of the new millennium specially because it is becoming easier for corporate sector to catch upon other fronts. Peter F Druker has also mentioned that the winner in competitive world economy is going to be the firm that most effectively shortens the product life of its own product, that is, one which organises most systematic abandonment of its product. It must be done quickly without loosing any opportunity. Indian producers in the present situation of global competition should identify PACE as one of the important element in marketing strategy.

What PACE Means?

PACE means increasing the speed of launch of new products and re-launch of old ones. It indicates that be innovative in product quality, but be fast.

Secondly, it means reduction of gap between one launch and another launch and making several brand launches in a short period of time.

Thirdly, it means quick response to changes in consumer needs and preference by creating brand variations, time consuming test marketing and product fine tuning may be anti PACE but consumer need had to be judged. The use of qualitative research can offer quick insight into buyers perception for quick response and help in reducing marketing time.

Fourthly, it means adopting a fast system design to keep fast track and PACE with the emerging global trends and control the quality of manufacture. It will help to squeeze the time in product development for effective speed marketing. A new product must be ready almost as soon as it is launched globally. Lastly, it means increase in the pace of promotion to maintain high frequency. The high frequency of promotion keeps the consumers interest high in the product.

The PACE means the rapidity of quality management, product development, technology adoption and creating fast brand wave and increased frequency of promotion. Hindustan Lever — Denim after Shave lotion, Lux White, Surf Easy wash. Domesto floor clean, Confident tooth brush, Walls-Ice cream, relaunches of the entire range of Kissan, Life Boy Liquid, are the few examples. Even Britannia launched six brands of biscuits between January 94 to July 94. e.g. 50-50, Little Heart, Bakers Choice, Joshesps Thin, etc. PACE means the right product of right quality, the right Package at right price, the right promotion at right time. Executives in India have identified PACE as a priority strategy to meet the global challenges.

In Indian market the PACE setters are Hindustan Lever which has floated a variety of new products and relaunched its old products and has gone into hyper-drive in a very little time. Lever has been followed by Proctor & Gambles, I.T.C., Reckett and Colman, Britannia Industries, Coco-cola India, and Videocon International. (Except Videocon all the above companies are multinationals having global market).

These firms are not only pushing the pace, they are also increasing their risk. Under researched brands launched merely flop. Obsession for speed however should not affect adversely the contents or substance and quality of product. However, occasional the failures are and price that the new generation of speedsters is willing to pay, it is also known that in flooding the brands not all products will succeed, and the winners will carry the burden of failures.

What Pace Can Do?

PACE knocks down your competitor out even before he can get going. Let us recall the Hindustan Lever's rushing of its liquid 'Lifebuoy' anti bacterial soap in market and at the same time the launch of Gel verson of Liril and Lux International. Let us also examine why did Proctor and Gamble race to launch its Whisper sanitary napkins in August '94 choosing to dispense with its elaborate test marketing?

Both were proactive because launches from competitors were imminent. Hence they used PACE to blunt their competitors edge for instance the Lifebuoy Liquid soap earned an edge over the RCI 'Dettol' handwash liquid which came later. Similarly Proctor and Gamble's rapid fire Launches we are pre-empting the paper products that the Lever-Kimberly Clark joint venture planned to launch soon. There are examples of Meslos and Reebock Shoe companies Meslos used space to block Reebock Sports Shoe market.

As soon as opportunities arise the producer must go out into the market place to exploit them, for the faster you react to the consumer the better it is for the bottom line. A series of quick launches hurts your rival in many ways (a) It raises his cost of entry since he must catch product you have launched, (b) forces the rival to dedicate additional resources, which is bound to affect his profitability.

PACE Helps to Create a New Segment and Dominate in that Segment

In consumer products where there is fierce competition quality and tactical move at marketplace have been dominating. So far, the B.C.G. Strategy of segmentation and product market matrix helped producers to find segment or adopt segmentation. Now the producers are adopting 'Appeal Segmentation' for example, Maggi the two minutes noodle created an appeal segment. The school going children prefer it as an after school meal. It has become very popular Indian dish.

The Hindustan Lever has been designing a series of product attributes taking each one to a brand extension for creating an appeal segment. Occupying a new segment first is the best way to ward off the competition. Similarly the first innovation of 2 minutes 'Maggi' created a segment of its own and ruled over.

PACE Wins Over the Promiscuos Consumer

The example of Britannia Co. illustrates it very clearly. In 93 it did not launch any brand. In 1994 Britannia launched one brand every three months e.g. Little Hearts, 50-50, Bakers Choice, Bake Sandwitch bread, and Josheps Thin are example of it.

As increasingly fickle consumers show low brand loyalty particularly for premium or Low involvement product. PACE is providing a panacea for the harried marketers. Instead of fighting the customers propensity to experiment with new products that offer novel appeal, savvy companies are launching string of products at regular intervals different from one another to ensure that consumers stay with them. The example of 'Classic' Cigarette of ITC is before us. It is launching flavour based premium brands at rapid pace, four brands of classic brands have already been launched. In fact, using brand extensions at a quick pace is an effort to push down a customer every time when he gets restless and fiddles with the idea of switching to other brands. Speedy launching the variants of a strong brand, instead of new brands satisfies the customers penchant for loyality as well his urge for new experiment. Thus, it helps a producer in creating an appeal segment and also in retaining the segment even in the situation of global competition.

PACE Revives Dropping Brand

It is but natural that even the best brands may go on the downward life-cycle stage in the market. In an open competitive market with global competition brand life-cycle is being more compressed today. It may be further compressed in the next millennium. A brand must have multiple lives today. Attempts for speedy action of rejuvenating your brand is essential. Velocity alone will ensure that the revitalisation of the brands will hammer away at the consumers consciousness without allowing his interest to slacken in the brand due to long intervals between relaunches. This can be observed by the relaunches of Cadbury's four chocolates, Cadbury Dairy Milk, 5 Stars, Eclaire and Gems. Cadbury has used speed of its relaunch to bounce into the consumer awareness just in seven months. In 1994 Cadbury redefined 'Dairy milk' as an adult food and 5 Star as a snack. The pay-off from this was rewarded by renewed interest of the consumers.

PACE Can Compensate for Being the Last One

An organisation which comes into operation at a later stage can definitely become a formidable competitor by setting dizzying pace of product launches. For late entrants PACE is the only antidote. If competitors have deployed pre-emptive strategies of their own. It is trade-mill. Once you have stepped on to it you cannot get off. Tour the front line of Cola wars for a first hand taste of Pace. Anticipating the launch of Coke, Pepsi Foods unleashed a blizzard of promotional activity preceding its entry in every market possible. When Limca 300 ml was launched in August '94 Thums up 300 ml was also launched alongside the Limca in the market the very next week. President of Coca Cola remarked that Coca Cola India has changed its strategy to PACE. He remarked "Our actions and reactions time has been reduced to nil." Earlier the marketeer worked to the slogan of 'ready aim, and fire.' Today it has changed to ready, fire, aim. Thus, PACE not only helps an organisation to earn a segment but also to be one of major shareholder of the market even the organisation is a late starter.

PACE Lends Products a Life Long Technological Edge

Present comprises the age of technology. A number of products are based on technology and are flooding the market e.g.. Timex watches- MNV's televisions, pump dispensers for liquid soap. It is essential to catch the hi-tech wave before rivals adopt it. Titan watches did it. Titan made watches as personal accessories and it became leader of the Quartz segment of watch market, pushing HMT to second position. It also rewrote the Rule Book of Conventional Wisdom in watch market that the best watches are imported one, legally or illegally. Videocon International has adopted the PACE strategy in this area. It has rushed into market the first product in every segment created by technology, from picture-in-picture, to frost free refrigerators. Videocon international believes in being first with a new concept, using a tactical and strategic marketing in compressed time frame. Videocon by adopting speed in technological adaptation has become the technology leader. Thus, speed marketing has carved out an advantage to Videocon and has provided a life long technology to its product.

Conclusion

The success stories of Nirma (Popular priced washing powder), Titan Watches, ONIDA (T.V.) Videocon Washing machine, Maruti Cars, Maggi Noodles and Close-up tooth paste, each one of them presents some significant breakthrough in terms of product development. Their success also proves that there is need to understand consumer behaviour in the context of global competition and change the approach to strategic marketing. That to succeed in any competitive market whether domestic or global, the application of the traditional approach will not be effective. The strategic approach must concentrate on looking at the market differently than what was being done before liberalisation. The companies must be innovative from the point of future opportunities and understanding of the needs of consumers also. It is something more than USP which sells. Innovation should not only be quantitative. It should be quick, daring, dynamic than analysing the industry in two by two matrix. It should be qualitative and must be based on watching the performance of rivals.

Now it is clear that the product innovation, the quality management and technological updating can matter only when you watch your rivals action and they are adopted with speed. Pace only can increase market share, increase turnover, enhance dealers motivation, improve product viability, improve corporate image, lead to new technique, a new position in the market.

TWENTY THREE

The New Generation

The current supercomputers are only at the threshold of what computer designers think can be achieved; the next generation of advanced supercomputers will make today's machines look like handheld calculators. "We have problems that would take 500 to 1,000 hours to solve [on today's supercomputers]," says David Nowak, division leader for computational physics at Lawrence Livermore National Laboratory, where a cluster of seven supercomputers-known as "Octopus" are used for nuclear weapons research.

Before the end of the twentieth century, computer scientists had developed machines that not only crunched numbers at high speed but also exhibited artificial intelligence— computers that think and reason somewhat like human beings and that could understand information conveyed by sight, speech, and motion. The question was, which nation's scientists would get there first? The Japanese had announced a two-pronged plan to build advanced computer technologies.

One project was the $100 million, eight-year National Superspeed Computer project, which aimed at producing machines 1000 times faster than the earlier Cray-1 supercomputer built by Cray Research of Minneapolis. The other, the $500 million, 10-year Fifth Generation Computer project, was focussing on artificial intelligence. Both were being countered by American efforts, including a Pentagon request for up to $1 billion over the next five years for superspeed and artificial intelligence (AI) technologies. Although behind, Great Britain and France have also launched national supercomputer projects. The great danger for the losers in the race-and the opportunity for the winners was that whoever built the next generation of computers would have a huge technological and commercial advantage: these computers would be used for computer and microelectronics design-to build even smarter, and more powerful machines. They won't be self-replicating machines, but they would be close. "It takes you a long time to catch up," says computer scientist Rai Reddy of Carnegie-Mellon University, one of the top US computer-research centers. "In some of these areas, that is the difference between a first-rate power and a second-rate power-from an economic point of view and from a security point of view. " The losers in the race will fall farther and farther behind. The leading edge of computer science is still a black art; there are no fixed laws, and the field is highly experimental. That is what worries US scientists about Japan's approach—some success is inevitable. "Because the field is experimental, [the Japanese] will come out with something," says Dertouzos. "It may not be what they wanted, but they'll come up with new architectures, new insights, and new design techniques.

To build the computers that dominated the 1990s, both Japan and the US are depending on the onrushing technological advances in microelectronics. Japan's Fifth Generation project will use faster, denser circuitry to create a new class of superintelligent computers.

The 24 projects in the Fifth Generation concentrate on artificial intelligence, a goal of American computer scientists for more than a quarter century. "We are trying to catch up to you, and not the other way around," says Tokyo University Professor Tohru Moto-oka, who organized the project for the Japanese government. Although the Japanese are highly regarded as superb engineers, Japanese computer scientists have often been faulted for failing to develop innovative computer software.

In a break with tradition, Kazuhiro Fuchi, the Fifth Generation project director, has deliberately assembled a young team: "The question was who would adapt most easily to this research," says Fuchi. "Young people have fewer fixed ideas." The project, headquartered in a downtown Tokyo skyscraper, will focus on computer architectures, software, and the symbolic logic necessary to build thinking computers. In the US, meanwhile, three huge new programs in electronics and computing are getting under way.

The Microelectronics and Computer Technology Corp (MCC)

Last year William Norris, the founder and chairman of Control Data Corp., convened a meeting of top computer and semiconductor industry executives at the Grenelefe Golf and Tennis resort in Oriando, Florida, to discuss setting up a hugh research cooperative.

The companies agreed to form a nonprofit joint venture so that they could pool their resources and share the cost of doing long-range research. Twelve major US corporations joined the new organization, including Honeywell, Motorola, RCA, and Control Data. MCC is a bold departure from the way research is usually done in American universities and corporations, and it will probe the limits of the nation's antitrust laws. The venture partly follows the Japanese model: the companies will donate scientists and researchers to MCC, loaning them for up to four years.

The corporate co-owners will also put up the money to fund MCC' research, in return for the rights to use the results. Whether the scheme will work is an open question. The 12 companies in MCC are competitors in fast-moving, high-technology markets, and ordinarily they jealously guard any technological edge they gain. In fact, many top US "firms—Cray Research, Texas Instruments, Intel, and others—chose to stay out of MCC. "That's not our style," says John A. Rollwagen, chairman of Cray Research, an 11-year old company proud of its entrepreneurial creed. "We don't want to participate." The biggest market force of all-IBM-reportedly stayed out of MCC because it feared antitrust action against it if it joined. So far, however, the creation of MCC has not provoked any such suits.

San Francisco antitrust lawyer Joseph M. Alioto did write to the chief executives of the companies that were about to form MCC: "In my opinion, your contemplated conduct is an unequivocal combination in violation of the antitrust laws of the United States." But the threat did not deter MCC's co-owners and, for the time being at least, the Justice Department has allowed the MCC plan to stand. To run the new corporation, MCC's directors chose retired Admiral Bobby Ray Inman, former director of the National Security Agency and former deputy director of the CIA. Inman is widely respected for his managerial abilities and is an adept politician besides. "The day they picked Bob Inman to head MCC," says George W. Keywork II, Ronald Reagan's top science adviser, "any concern about its success diminished in my mind."

Over the past five months, Inman orchestrated a competition among 57 cities for the MCC headquarters; the winner was Austin, Texas, after private donors, the state, and universities put together a generous package of incentives.

The consortium will have a budget of about $75 million a year and a staff of 250. Its first projects include programs in semiconductor packaging and interconnect technology, advanced software engineering and computer-aided design and manufacturing (CAD/CAM) for the electronics and computer industries. Most ambitious is a 10-year program aimed at breakthroughs in computer architecture, software, and artificial intelligence.

MCC will own the licenses and patents to the technologies; the manufacturing and marketing will be left to the companies that sponsor the projects. MCC win give them a competitive edge on the market—they will have exclusive rights for three years before the research is published and other firms are allowed to buy licenses.

The Semiconductor Research Corp.

Over the past few years Japan has captured a vital segment of the world semiconductor industry, the market for so-called RAM (random access memory) chips, a technology invented in America.

Japan now supplies 70 per cent of all 64K RAMS sold, and it appears that, as the next generation of memory chips, the 256K RAMS, is being readied for market, the Japanese semiconductor companies are threatening to take a big share of sales. It makes for a grim reminder: a decade ago, before an all-out government project to build up the industry, Japanese semiconductor firms lagged far behind American and European chip companies. Last year 13 US chip manufacturers and computer companies banded together to form a nonprofit research consortium, the Semiconductor Research Corp. (SRC), to share the spiraling costs of advanced research and development. SRC's founders include Control Data Corp., Digital Equipment Corp., Hewlett-Packard, MM, Intel and Motorola. Unlike MCC, however, SRC, which is headquartered at Research Triangle Park in North Carolina, does not carry out its own research. Instead, it sponsors research at universities. SRC is spending $12 million this year and has allocated $30 million next year. The goal: "To assure long-term survival in the market," says Larry Sumney, SRC's executive director.

DARPA

The Pentagon's Defense Advanced Research Projects Agency (DARPA) is, more than any other single agency in the world, responsible for the shape of advanced computer science today-and for many technologies now in widespread commercial use.

Over the past 20 years, DARPA has poured half a billion dollars into computer research, in the process virtually creating the science of artificial intelligence.

The first supercomputer, built in 1964, was a DARPA project. Computer time sharing, a fundamental advance, came out of work sponsored by DARPA. So did packet-switched networks, the workhorses of today's telecommunications data networks. And computer graphics—now used on desktop computers and video-arcade cockpits—is a DARPA-sponsored invention. DARPA's next priority is a push for advanced supercomputing and artificial intelligence technologies that may cost as much as $1 billion.

DARPA plans to do everything the Japanese have set out to accomplish-and more. Earlier this year DARPA proposed a "Strategic Computing and Survivability" project, which it hopes will lead to a variety of new machines. "We want some architectures that are good for building semantic memories, memories that can hold knowledge," says DARPA's computer director, Robert Kahn. "Other kinds of systems are good for logic processing. We want architectures that can do very rapid signal processing [and] structures that can handle very, very large amounts of data in communications."

Once they are in place, these technologies will make possible an astonishing new breed of weapons and military hardware. Smart robot weapons—drone aircraft, unmanned submarines, and land vehicles—that combine artificial intelligence and high-powered computing can be sent off to do jobs that now involve human risk. "This is a very sexy area to the military, because you can imagine all kinds of neat, interesting things you could send off on their own little missions around the world or even in local combat," says Kahn.

The Pentagon will also use the technologies to create artificial intelligence machines that can be used as battlefield advisers and superintelligent computers to coordinate complex weapons systems.

An intelligent missile guidance system would have to bring together different technologies-real-time signal processing, numerical calculations, and symbolic processing, all at unimaginably high speeds—in order to make decisions and give advice to human commanders. While the national security needs are driving supercomputing technologies, there is a growing market for commercial spinoffs.

The same technology that can be used to simulate an anti-tank missile smashing into a heavily armored tank can also be put to work on less martial arts. In Los Angeles, Digital Productions, Inc., is using a Cray supercomputer to produce television commercials for Mattel, rock video segments for Turner Broadcasting, and special effects for Lorimar's "Star Fighter" space epic, which will be released next summer. Instead of shooting a commercial the conventional way—a costly photo session—Digital can create the "pictures" it needs in detail so precise that it's impossible to distinguish between the supercomputer's graphic image and the real photo. For inanimate objects, that is: the Cray-1/S can simulate a car right down to the glint of sunlight on the windshield, but not a human being. As the commercial market grows, the Japanese believe that their strategy of building supercomputers for the general-purpose market will pay off.

During the past year Fujitsu, Hitachi, and NEC have each announced supercomputers faster than the most powerful American machines now on the market, the Cray X-MP and Control Data's Cyber 205. Moreover, all three are designed to use the standard Fortran language, lifted from ordinary mainframe computers, but at much faster speeds. "Although supercomputers right now are mainly being used by specialists," says Takamitsu Tsuchimoto, Fujitsu Corp.'s development manager for supercomputers, "we believe that in the next 5 to 10 years they will be used by a lot of ordinary people, so we wanted to design a machine [they could use] without making great efforts." Designing a supercomputer is no easy task. America's premier supercomputer designer is Seymour Cray, 57, who designed the Cray-I and the basic architecture of Control Data's Cyber 205.

The Cray x-mp contains a dense pack of 240,000 silicon chips arranged to shorten the distances the electrical signals must travel, thereby decreasing the time it takes to perform an operation; the new Cray x-mPs will run at 400 million operations per second. And Japan's superspeed goal is 10 billion operations per second. (An Apple IIe computer contains 31 chips and can execute 500,000 operations per second.) Because of the cost of supercomputers—Cray's x-mp will sell for $11 million—the market will remain limited. Japan's Fifth Generation project, however, plans to build "super personal computers."

These huge increases in raw computer power are just the first step. The most profound changes brought by the new technologies will be the development of reasoning computers that will use superspeed symbol manipulation to simulate human thought. "We are about to see the next explosion, which is the application of computers to reasoning," says Stanford University computer scientist Edward A. Feigenbaum, who is a founder of two artificial intelligence companies, Teknowledge Inc. and IntelliGenetics, and co-author of a new book on Japan's challenge, *The Fifth Generation*.

Limited forms of artificial intelligence are already enjoying commercial success. Digital Equipment Corp. the nation's second largest computer firm, uses an AI program called X-CON to make custom designs of its, computer systems.

Using a set of more than 2500 rules programmed into the system, X-CON examines a customer's specifications, determines whether all the necessary components are included, then draws a set of diagrams showing the proper spatial relationships among the components. Scientists studying artificial intelligence have been aided by advances in computer hardware, too. In 1980 a group left MIT's Artificial Intelligence Laboratory to found a company, Symbolics Inc., of Cambridge, Mass., to build computers specially designed to run LISP, a language used to develop artificial intelligence programs; and Xerox has begun to sell similar computers. Symbolics has been selling its machines, the Symbolics 3600, to a broad range of university and industrial research labs.

The company has another customer as well: the managers running Japan's Fifth Generation project have bought 10 Symbolics machines—and they have 15 more on order. In exploring the brave new world of artificial intelligence, computer scientists are concentrating on several important problems. "Knowledge engineers" are building the so-called expert systems that can mimic human expertise in a narrowly defined area. The firm of Teknowledge, in Palo Alto, California, for example, has built an expert system for the French Elf Acquitaine Oil Co., a system that will give advice on one of the industry's most costly technical problems—what to do when a drill bit gets stuck thousands of feet below the earth's surface.

To build the system, Teknowledge engineers interviewed Elf Acquitaine's top troubleshooter, Jacques-Marie Courte, and then programmed his answers into a computer. The program is, in effect, a computer replica of Courtes expertise; the computer will ask the drilling-rig foremen questions, just as Courte would.

Once it gathers the information it needs, the computer will make recommendations by drawing images on the screen and giving suggestions on how to retrieve the bit. Because daily drilling costs are high, Elf Acquitaine may well recover the program's development costs the first time it is used successfully. General Electric is building a software program that will provide expert advice on repairing locomotives.

The Pentagon would like to build artificial intelligence programs that could serve as a pilot's assistant in the cockpit. Stockbrokers and insurance agents may also soon get help: "Some people are beginning to see a gold mine in [building artificial intelligence programs for] financial services," says Patrick Henry Winston, chairman Of MIT's Artificial Intelligence Laboratory. And savvy software designers for the personal computer industry are beginning to look at artificial intelligence as the next big wave that will sweep the market. Despite the successes, there are problems ahead as researchers attempt to move beyond the building of narrowly defined "experts." "There's a lot of hard stuff out there we just don't have the answers to," says Roger Schank, director of Yale University's Artificial Intelligence Laboratory. Computers have always been maddeningly literal machines, subject to the absolute tyranny of the binary codes they use to do their calculations. (The switches are either on or off, simulating ones or zeros, nothing else.)

That literal-mindedness can transform a problem that would seem trivial to humans into a nightmare. Consider the simple statement "Mary had a little lamb."

For a computer to translate the text into another language—a function scientists are now trying to develop-it would have to sort through what is, by one count, 28 possible interpretations (Mary owned the lamb, Mary ate the lamb, Mary had sexual relations with the lamb, Mary gave birth to the lamb, and on and on).

The kind of understanding humans experience as a "flash of recognition" is also difficult to instill in a computer. The statement "Ronald Reagan is president" carries a number of immediate meanings to a flesh-and-blood American, but a computer would have to rummage through its silicon memory chips in search of dozens of facts—what the word "president" means and biographical facts and details that tell who Ronald Reagan is. And the biggest challenge of all—teaching computers to learn to acquire knowledge on their own—is nowhere near being solved.

Mankind has long been enchanted—and frightened—by the prospect of creating machines that think. "I don't see any limitations to artificial intelligence," says Nobel laureate Herbert Simon, professor of computer science and psychology at Carnegie-Mellon. "All the mechanisms for human intelligence are present. Already machines can think just like people—in a limited sense. Man isn't unique in that respect." Moreover, many machines may soon possess sight, touch, hearing (in the form of voice recognition), and speech, thus imitating humans' sensory capacities along with their intellectual ones. But whole areas of the human thought process—volition, emotion, the creative uses of error—still lie well outside a computer's experience. And scientists doing research on artificial intelligence are far from their ultimate goal—a computer-based analog of the human brain.

Still, the race to build superintelligent computers—Japan's challenge to the US—will almost certainly push the technology to new levels. "While at the beginning I was angry at Japan because I felt that they were swiping our best ideas, really, on second thoughts, I couldn't blame them," says Dertouzos of MIT. "And I don't blame them today at all. I think they're doing exactly what they should." And the Japanese challenge has at last spurred the US into action. "If we really wake up, I'm very optimistic," says Dertouzos. "We could beat the daylights out of them." No one knows where the competition will ultimately lead.

TWENTY FOUR

The Next Generation

Japan's Next Generation project will use faster, denser circuitry to create a new class of super intelligent computers. The 24 projects in the Fifth Generation concentrate on artificial intelligence, a goal of American computer scientists for more than a quarter century. Although the Japanese are highly regarded as superb engineers, Japanese computer scientists have often been faulted' for failing to develop innovative computer software. But in a break with tradition, Kazuhiro Fuchi, the Fifth Generation project director, has deliberately assembled a young team: "The question was who would adapt most easily to this research," says Fuchi. "Young people have fewer fixed ideas." The project, headquartered in a downtown Tokyo skyscraper, will focus on computer architectures, software, and the symbolic logic necessary to build thinking computers. The Microelectronics and Computer Technology Corp (MCC). Last year William Norris, the founder and chairman of Control Data Corp., convened a meeting of top computer and semiconductor industry executives at the Grenelefe Golf and Tennis resort in Orlando, Florida, to discuss setting up a huge research cooperative.

The companies agreed to form a non-profit joint venture so that they could pool their resources and share the cost of doing long-range research. Twelve major US corporations joined the new organisation, including Honeywell, Motorola, RCA, and Control Data. MCC is a bold departure from the way research is usually done in American universities and corporations, and it will probe the limits of the nation's antitrust laws. The venture partly follows the Japanese model: the companies will donate scientists and researchers to MCC, loaning them for up to four years. The corporate co-owners will also put up the money to fund MCC' research, in return for the rights to use the results. Whether the scheme will work is an open question. The 12 companies in MCC are competitors in fast-moving, high-technology markets, and ordinarily they jealously guard any technological edge they gain. In fact, many top US firms—Cray Research, Texas Instruments, Intel, and others—chose to stay out of MCC.

MCC's directors chose retired Admiral Bobby Ray Inman, former director of the National Security Agency and former deputy director of the CIA. Inman is widely respected for his managerial abilities and is an adept politician besides. Over the past five months, Inman orchestrated a competition among 57 cities for the MCC headquarters; the winner was Austin, Texas, after private donors, the state, and universities put together a generous package of incentives.

The consortium will have a budget of about $ 75 million a year and a staff of 250. Its first projects include programs in semiconductor packaging and interconnect technology, advanced software engineering and computer-aided design and manufacturing (CAD/CAM) for the electronics and computer industries. Most ambitious is a 10-year program aimed at breakthroughs in computer architecture, software, and artificial intelligence.

MCC will own the licenses and patents to the technologies; the manufacturing and marketing will be left to the companies that sponsor the projects. MCC will give them a competitive edge on the market—they will have exclusive rights for three years before the research is published and other firms are allowed to buy licenses.

Over the past three years Japan has captured a vital segment of the world semiconductor industry, the market for so-called RAM (random access memory) chips, a technology invented in America. Japan now supplies 70 per cent of all 64K RAMS sold, and it appears that, as the next generation of memory chips, the 256K RAMS, is being readied for market, the Japanese semiconductor companies are threatening to take a big share of sales.

Over the past 20 years, DARPA has poured half a billion dollars into computer research, in the process virtually creating the science of artificial intelligence. The first supercomputer, built in 1964, was a DARPA project. Computer time sharing, a fundamental advance, came out of work sponsored by DARPA. So did packet switched networks, the workhorses of today's telecommunications data networks. And computer graphics—now used on desktop computers and video arcade cockpits—is a DARPA—sponsored invention.

DARPA plans to do everything the Japanese have set out to accomplish and more. Earlier this year DARPA proposed a "Strategic Computing and Survivability" project, which it hopes will lead to a variety of new machines. "Other kinds of systems are good for logic processing. We want architectures that can do very rapid signal processing [and] structures that can handle very, very large amounts of data in communications." Smart robot weapons-drone aircraft, unmanned submarines, and land vehicles-that combine artificial intelligence and high-powered computing can be sent off to do jobs that now involve human risk. "This is a very sexy area to the military, because you can imagine all kinds of neat, interesting things you could send off on their own little missions around the world or even in local comba. The Pentagon will also use the technologies to create artificial intelligence machines that can be used as battlefield advisers and super intelligent computers to coordinate complex weapons systems. An intelligent missile guidance system would have to bring together different technologies-real—time signal processing, numerical calculations, and symbolic processing, all at unimaginably high speeds—in order to make decisions and give advice to human commanders.

The same technology that can be used to simulate an antitank missile smashing into a heavily armoured tank can also be put to work on less martial arts. Instead of shooting a commercial the conventional way, a costly photo session, Digital can create the "pictures" it needs in detail so precise that it's impossible to distinguish between the supercomputer's graphic image and the real photo. As the commercial market grows, the Japanese believe that their strategy of building super-computers for the general-purpose market will pay off. The Cray X-MP contains a dense pack of 240,000 silicon chips arranged to shorten the distances the electrical signals must travel, thereby decreasing the time it takes to perform an operation; the new Cray X-MPs will run at 400 million operations per second. And Japan's super speed goal is 10 billion operations per second. Because of the cost of super-computers, Cray's X-MP will sell for $ 11 million, the market will remain limited. Japan's Fifth Generation project, however, plans to build "super personal computers." These huge increases in raw computer power are just the first step.

The most profound changes brought by the new technologies will be the development of reasoning computers that will use super speed symbol manipulation to simulate human thought. The Fifth Generation field is not new: scientists in the US have been studying artificial intelligence for more than 25 years, struggling to understand the nature of knowledge and how to represent it in forms adaptable to computer usage.

Digital Equipment Corp the nation's second largest computer firm, uses an Al program called X-CON to make custom designs of its, computer systems. Using a set of more than 2500 rules programmed into the system, X-CON examines a customer's specifications, determines whether all the necessary components are included, then draws a set of diagrams showing the proper spatial relationships among the components. Scientists studying artificial intelligence have been aided by advances in computer hardware too. In 1980 a group left MIT's Artificial Intelligence Laboratory to found a company, Symbolics Inc., of Cambridge, Mass., to build computers specially designed to run LISP, a language used to develop artificial intelligence programs; and Xerox has begun to sell similar computers. Symbolics has been selling its machines, the Symbolics 3600, to a broad range of university and industrial research labs. The company has another customer as well: the managers running Japan's Fifth Generation project have bought 10 Symbolics machines and they have 15 more on order in exploring the brave new world of artificial intelligence, computer scientists are concentrating on several important problems. "Knowledge engineers" are building so-called expert systems that can mimic human expertise in a narrowly defined area.

The program is, in effect, a computer replica of Courtes expertise; the computer will ask the drilling-rig foremen questions, just as Courte would. Once it gathers the information it needs, the computer will make recommendations by drawing images on the screen and giving suggestions on how to retrieve the bit. Because daily drilling costs are high, Elf Acquitaine may well recover the program's development costs the first time it is used successfully. General Electric is building a software program that will provide expert advice on repairing locomotives.

The Pentagon would like to build artificial intelligence programs that could serve as a pilot's assistant in the cockpit. Stockbrokers and insurance agents may also soon get help: "Some people are beginning to see a gold mine in financial services," says Patrick Henry Winston, chairman of MIT's Artificial Intelligence Laboratory. And savvy software designers for the personal computer industry are beginning to look at artificial intelligence as the next big wave that will sweep the market. Despite the successes, there are problems ahead as researchers attempt to move beyond the building of narrowly defined "experts."

Computers have always been maddeningly literal machines, subject to the absolute tyranny of the binary codes they use to do their calculations. (The switches are either on or off, simulating ones or zeros, nothing else.) That literal-mindedness can transform a problem that would seem trivial to humans, into a nightmare. Consider the simple statement "Mary had a little lamb." For a computer to translate the text into another language-a function scientists are now trying to develop-it would have to sort through what is, by one count, 28 possible interpretations.

The statement "Ronald Reagan is president" carries a number of immediate meanings to a flesh-and-blood American, but a computer would have to rummage through its silicon memory chips in search of dozens of facts-what the word "president" means and biographical facts and details that tell who Ronald Reagan is. And the biggest challenge of all-teaching computers to learn to acquire knowledge on their own-is nowhere near being solved.

Mankind has long been enchanted and frightened by the prospect of creating machines that think. All the mechanisms for human intelligence are present. Already machines can think just like people, in a limited sense. Man isn't unique in that respect. Moreover, many machines may soon possess sight, touch, hearing (in the form of voice recognition), and speech, thus imitating humans' sensory capacities along with their intellectual ones. But whole areas of the human thought process volition, emotion, the creative uses of error-still lie well outside a computer's experience.

Japanese planners view the computer industry as vital to their nation's economic future and have audaciously made it a national goal to become number one in this industry by the latter half of the 1990s. They aim not only to dominate the traditional forms of the computer industry, but to establish a "knowledge industry" in which knowledge itself will be a saleable commodity like food and oil.

The Japanese plan is bold and dramatically forward looking. It is unlikely to be completely successful in the ten-year period. But to view it therefore as "a lot of smoke," as some American industry leaders have done, is a serious mistake. Even partially realised concepts that are superbly engineered can have great economic value, pre-empt the market, and give the Japanese the dominant position they seek.

Technology Targets

The Japanese Fifth Generation project aims to design and produce computer hardware and software for knowledge engineering in a wide range of applications-including expert systems, natural language understanding by machines, and robotics. To accomplish these goals, the Japanese must improve present computing capabilities dramatically, but they must also make major innovations in existing technology that will enable Fifth Generation computers to support very large knowledge bases, allow very fast associative retrievals, perform logical inference operations as fast as current computers perform arithmetic operations, utilise parallelism in program structure and hardware to achieve high speed, and develop a machine-user interface that allows significant use of natural speech and images.

All experts systems built by knowledge engineers to date consist of three main parts. First, is the subsystem that "manages" the knowledge base needed for problem solving and understanding. Secondly, is the problem-solving and inference subsystem, which discovers what knowledge is useful and relevant to the problem at hand, and with it constructs, step by step, a line of reasoning leading to the problem solution, the plausible interpretation, or the best hypothesis. Thirdly, are the methods of interaction between human and machine, in modes and languages that are "natural" and comfortable for the user. Ordinary human natural language is often preferred, but the stylised notations of some fields like chemistry are also desirable for specific groups of users. Knowledge—based management, problem solving and inference, and human interaction-these have all been approached in our present expert systems via software innovations, innovations that have pressed traditional von Neumann hardware architectures to their limits.

The Fifth Generation plan organises its work around these three subsystems, but with a critical added dimension: for each component subsystem there is a hardware level and a software level. And between the levels the Japanese designers must define a "language" with which the software and hardware interact. The knowledge in the knowledge base must first be represented in symbolic form, and in memory structures, that can be used efficiently by the problem-solving and inference subsystem. This representation can take many forms.

A typical kind of associative network is the taxonomy, known as "The... is-a (hierarchy)." For example, "The sparrow is-a kind of bird." In this case, both sparrow and bird are objects within the knowledge base. If the knowledge base is informed that "The bird is-a kind of animal that can fly," the knowledge base management system must automatically propagate the deduction that sparrows can fly.

A rule consists of a collection of statements called the "if" part, and a conclusion or action to be taken called the "then" part. For example, "If the fog ceiling is below 700 feet and the official weather forecast calls for no clearing within the hour, THEN landing is dangerous, will violate air traffic regulations, and diversion to a neighbouring airfield is recommended." To find out if a rule is relevant to the reasoning task at hand, the problem-solving program must scan over the store of "ifs" in the knowledge base. That search can be immense in the size of knowledge base the Japanese plan to make possible. Here again, the knowledge-base management subsystem will be designed to organise the memory in ways that will reduce the amount of processing to be done.

Parallel processing capabilities in both the software and hardware levels of the system will also speed associative retrievals. In the Fifth Generation plan, knowledge will be stored electronically in a large file known as a relational data base. The job of automatically updating the knowledge in the file, and of organising appropriate searches for relevant knowledge, will be performed by the knowledge-base management software. The Fifth Generation prototype knowledge-base subsystem will handle a modest knowledge base thousands of rules and thousands of objects-about the size needed for current expert system applications. Each object will be allotted 1000 characters of file storage space. Within the ten-year trajectory of their plan, the Japanese goal is to develop knowledge-base capacity in their systems that will be able to handle tens of thousands of interference rules and one hundred million objects! What could so much knowledge encompass?

Since the Fifth Generation is so farreaching, it demands dramatic improvements in other technologies that support the main-line knowledge information processing systems (KIPS) goals. Essential to the future of the enterprise, for example, are extremely high-speed processors, capable of processing by orders of magnitude faster than anything now available. Artificial intelligence made its debut on first generation machines and has been implemented subsequently on second- and third generation- machines, but not yet on fourth generation super-computers.

The Japanese aim for chips with 10 million transistors. Chips in current production carry a few hundred thousand transistors at most. Such processors are being developed in the course of another MITI effort, the SuperSpeed Computing Project, and will be adapted into the Fifth Generation machines. In addition, the Fifth Generation depends on access to knowledge bases in many locations, so its technology will ultimately be fused with the most advanced communications technologies.

The whole area of intelligent interfaces—the ability the machines will have to listen, see, understand, and reply to human users—will require extensive R&D in natural language processing in speech understanding, and in graphics and image understanding. All these have been concerns of artificial intelligence research from virtually its beginning some 25 years ago, and basic research in each of these fields has made reasonable progress. Because non-experts will be the largest group of users, natural language processing is one of the most important research goals of the Fifth Generation.

Research here will cover speech wave analysis, phonetic and syntactic analysis, semantic analysis, and pragmatic analysis, which derives understanding by extracting themes or foci in a given sentence, detecting focus shifts, and so on. For speech output, sentence generation will also be studied. Text analysis is also considered a part of natural language processing by the Japanese, although they are quite aware that the techniques used for large-scale text analysis are different from the techniques needed to smooth the way for an individual user to talk to his machine.

The Japanese see knowledge systems as a potential solution to the steadily increasing amount of text and documents that must be handled by computer. "Our research on intelligent man-machine interface will help to solve this problem," reported a group of Japanese scientists at a Fifth Generation project meeting. Present artificial intelligence research suggests this can be done in a prototype system, intelligent automatic analysis has been successfully applied to a wire news service in the US, for instance-but the sheer scale of the automatic analysis planned by the Japanese dwarfs any existing systems.

All this research in natural language processing will proceed in three stages, beginning with an experimental system, followed by a pilot model implementation stage that is connected with the inference and knowledge base machines, and concluding with prototype implementations. At that point the machines will be expected to understand continuous human speech with a vocabulary of 50,000 words from a few hundred or more speakers with 95 per cent accuracy. The speech understanding system is also expected to be capable of running a voice-activated typewriter and of conducting a dialogue with users by means of synthesised speech in Japanese or English.

The machine's capacity to respond intelligently to users, known as its question-answering system, will first be designed to handle queries in the computing field, but it is expected to be a prototype for such systems in many professional fields: in addition to the query system's 5000 or more words of vocabulary, it will have 10,000 or more inference rules.

Images

Picture and image processing are considered almost as important as language processing, especially as they contribute to computer-aided design and manufacture (CAD/CAM) and to the effective analysis of aerial and satellite images, medical images, and the like. Here again the research will take place in three phases, beginning with an experimental phase to tackle ' such topics as the hardware architecture of "feature extractors "-for example, to distinguish the boundaries of objects—display generators, and the image data base.

The second phase will produce a pilot model, and in the third and final phase a prototype will be built, integrated into the Fifth Generation machine, and applications will be studied. One obvious use is in constructing robots that can see, understand, and act under novel circumstances. The bulk of robotics R&D, however, will be done in a Robotics National Project. Eventually, the image understanding system is expected to store about 100,000 images.

There are both positive and negative sides the Japanese project. First, its flaws. The science upon which the audacious Fifth Generation plans are laid lies at the outermost edge and beyond what computer science presently knows. The plan is risky; it contains several "scheduled breakthroughs." There are major scientific and engineering challenges in every aspect, from artificial intelligence through parallel architectures and distributed functions to VLSI design and fabrication. The project demands early successes to maintain its momentum and funding, and that could be a problem.

Conversely, meeting or exceeding the goals of the first three-year periods might well propel the Japanese ahead of their timetable, bringing increased support from the participating companies. Central to the success of the project are the Japanese managers, both governmental and industrial. Generally conservative and risk-averse, they are now to be charged with managing a very ambitious, high-risk project based on technology they hardly understand. Most of the breakthroughs the Fifth Generation project must achieve are basically innovations in software concepts.

The key ideas in the approach to knowledge information processing systems came out of the software world—ideas about the creation, maintenance, and modification of large and complex symbolic data structures in computer memories and the discovery of symbolic lines of reasoning. A quick fix to the problem is to work on the intermediate territory of the so-called firmware—intricate and detailed "programming" of the hardware-switching functions that sit at the bottom of the computing process. This is not a desirable final solution, however, since interpreting and executing the "firmware program" consumes time and slows down the machine.

Japanese computer specialists and managers are not, and never have been, comfortable with software—it's intangible to them, and its production is notoriously difficult to manage "on schedule and on budget." The Japanese lack the experience base in knowledge engineering and expert systems from which to draw as they begin to work out the details of what to build. In addition, the Japanese lack a large corps of university

Finally, from the Al viewpoint, two elements of the plan are questionable. First, the priority given to the highspeed logic processor. Are all those millions Of Lips needed? Second, is PROLOG the best choice for the machine language of the logic processor? In the American engineering experience, few applications have been limited by the number of inference steps per second that could be performed. Rather, limitations in performance arise from limitations in the quantity and quality of knowledge available to the machine (too little and not well defined); the facility with which it can be managed and updated; and the speed with which it can be searched and accessed.

A second plus for PROLOG is that it solves problems by proving theorems in first-order predicate calculus using computationally fast methods. The user never has to be concerned with the details of the problem-solving process. But PROLOG detractors see this as a serious flaw. The major successes of Al have come from mastering the methods by which knowledge can be used to control the search for solutions in complex problems.

The last thing a knowledge engineer wants to do is abdicate control to an "automatic" theorem-proving process that conducts massive searches without step-by step control exerted by the knowledge base. Such uncontrolled searches can be extremely time-consuming. The parallelism that can be brought to bear is a mere palliative, because the searches become exponentially more time-consuming as the problem complexity increases. One can't keep up with exponential growth simply by lining up a hundred or a thousand more parallel processors.

Creating the knowledge industry, with hardware, software, and knowledge system applications, is a great bet. Indeed, it is one of the few great bets sitting out there now in the information processing industry, ready for a major push toward exploitation. Of course the traditional modes of numerical calculation and data processing will continue to develop and prosper. But these will see steady incremental growth, not explosive growth.

The exponential growth will be seen in symbolic computation and knowledge-based, reasoning by computer. MITI's key economic insight is correct. For an island trading nation, exports create wealth and in the knowledge industry the value of exports is enhanced by indigenous resources-the intelligence, education, and skill of people. Further, knowledge information processing systems will significantly enhance the productivity of many other industries, thereby indirectly contributing to the value added.

The creating of ICOT, the pooling of talent in a cooperative endeavor, plus the well-coordinated transfer of technology between ICOT and the parallel labs of the firms, seem inspired. The ten-year planning horizon is excellent. Ten years is a long time in the information processing industry. Ten years ago pocket calculators cost hundreds of dollars, video games were primitive laboratory toys, and the Japanese had yet to produce their first viable microelectronic chip. As we live through it, we tend to underestimate the speed of technological change. MITI's concern for nurturing the innovative talents of Japanese computer scientists appears well placed. The US and the advanced European nations have become wary of providing the leading technologies upon which Japanese technical achievements have hitherto relied.

Trade wars are under way, and blockades are inevitable. Though solutions to the technological problems posed by the Fifth Generation plan may be hard to achieve, paths to possible solutions abound. The Japanese are rich with excellent engineering talent and have an adequate supply of cutting-edge computer scientists. The mix of talents enables (but does not guarantee) a good chance of success.

What's Real?

Managers of the Fifth Generation project have said that it would not disturb them if only 10 per cent of the project goals were achieved; others have remarked that the ten year planning horizon should not be taken too seriously; that the project goals are so important that an extension over another half or full decade would not be unreasonable. Partially realised concepts that are superbly engineered can have great utility and be of great economic benefit. At the very least, a partial success can pre-empt the area and make it not worthwhile for others to enter playing catch up.

The first 20 per cent of the technical achievement may skim off 80 per cent of the potential economic gain. If true, firms in the American industry might never find it in their economic interest to enter the arena. Being late might put them out of the contest. Consider this: though video-taping was invented in the US, the lengthy and expensive R&D process for the consumer-oriented video cassette recorder led to an all-or-none market share result, with American industry getting the "none." No matter how partial its success may be, the Fifth Generation project will provide a decade-long learning experience for a new generation of Japanese computer scientists.

They will be called upon to confront and perhaps solve the most, challenging problems facing the future of information processing, rather than re-engineering traditional systems. They will be learning advanced software concepts in a way that has never been done before in Japan and has never been widely done in the US or Europe. The Fifth Generation project, in its short life, has emplaced the technology transfer mechanisms necessary for Japanese industry to move effectively to bring its developments to market. Right now, the US has a substantial lead in virtually every area of Fifth Generation work.

Of course, the answer depends on what one substitutes for "XXX." Readers who filled in, say, "national socialism" could have supposed that these words were spoken, as a warning to German intellectuals who had not yet appreciated the glory of the Nazi revolution, by Josef Goebbels on the occasion of the book burning in Berlin on May 10, 1933. Readers could substitute "the ideas of the great leader and teacher" for "XXX" and leave open what particular revolution is being talked about. Leaders who come to mind, and whose names would render the quoted paragraph plausible, are, to name just a few: Karl Marx, General Pinochet, Stalin.

The Germans, by the way, had a word for what intellectuals are here being warned to do: Gleickschalffing, which is translated as "bringing into line" or "coordination." But, implausible as it may seem at first glance, "XXX" in the quoted passage stands for "this new instrument," meaning the computer. The authors of The Fifth Generation maintain that intellectuality, the creative use of the mind engaged in study and in reflection, will soon become inevitably and necessarily dependent on the computer. They are astounded that American intellectuals aren't rushing to enlist in their revolution.

A professor of computer science at Stanford University and a co-founder of two commercial companies that market artificial intelligence software systems, and Pamela McCorduck, a science writer, give the reader an idea of what's happening in the world of computers. They make the following claims. First, certain American computer scientists have discovered that, if computers are expected to intervene in some activity in the real world, then it would help, to say the least, if they had some knowledge of the domain of the activity in question. For example, computer systems designed to help make medical diagnoses had better know about diseases and their signs and symptoms.

Second, other American computer scientists have described designs of computers, "computer architectures," that depart radically from the industry's traditional design principles originally laid down by the pioneer computer scientist. The new architectures allow computational chains to be decomposed into steps which can be executed as soon as the data for executing them are ready. They then don't have to wait their turn, so to speak. Indeed, many steps can be executed simultaneously. Computation time is in a sense "folded" in such machines, which are consequently very much faster than their orthodox predecessors.

Third, still other computer scientists, mainly French and British, have created a computer language which they believe to be well suited for representing knowledge in computers in a form that lends itself to powerful logical manipulation. The Japanese hope that the conjunction of this computer language with the new architecture will allow ultra rapid computation of "inferences" from masses of stored knowledge. Fourth, these developments have taken place at a time of continuing dramatic progress in making computers physically smaller, functionally faster, and with increasing storage capacities.

Finally, the Japanese, who already dominate the world market in consumer electronics, have seized on the resulting opportunity and decided to create entirely new and enormously powerful computer systems, the "Fifth Generation," based on the developments described above. These systems, as the book jacket puts it, will be "artificially intelligent machines that can reason, draw conclusions, make judgments, and even understand the written and spoken word." This appears to be a very ambitious claim. But not to seasoned observers of the computer scene, who have long since learned to penetrate the foggy language of the computer enthusiasts.

In other words, the most recent ambitions of the Japanese were already close to being realised according to leaders of the American artificial intelligence community a quarter of a century ago! All that remained to be done-and it would be done within the "visible future"-was to extend the range of the problems such machines would solve to the whole range of the problems to which the human mind has been applied. That ambition remains as absurd today as it was twenty-five years ago. In the meanwhile, however, much progress has been made in getting computers to "understand" the written word and even some words spoken in very highly controlled contexts. Is the Japanese project then really not very ambitious?

A geriatric robot that frees old people from the murderous instincts of their children and is programmed to lie to them systematically, telling them that it understands their petty stories and enjoys "listening" to them. Mechanical doctors we can be utterly candid with and which won't disapprove of us, as human doctors often do. Technical devices in our own homes that gather information about us and determine to what group we ought to belong and which ones we should hate.

Technical systems that permit us to exchange knowledge "engagingly" with our fellow creatures while avoiding the horror of having to look at them or be looked at let alone touched. This is the world Feigenbaum and McCorduck are recommending. In fact, they haven't told us the whole story. Professor Tohru Moto-oka of Tokyo University and titular head of the Japanese Fifth Generation project promises even more:

> ... first, Fifth Generation computers will take the place of man in the area of physical labor, and, through the intellectualisation of these advanced computers, totally new applied fields will be developed, social productivity will be increased, and distortions in values will be eliminated [emphasis added].

We are well on our way to the kind of world sketched here. Already, we are told, and it is undoubtedly true, many people prefer to "interact" with computers. Schoolchildren prefer them to teachers, and many patients to doctors. No one seems to ask what it may be about today's doctors and teachers, or with the situations in which they work, that causes them to come off second-best in competition with computers.

The computer has long been a solution looking for problems-the ultimate technological fix which insulates us from having to look at problems. Our schools, for example, tend to produce students with mediocre abilities to read, write, and reason; the main thing we are doing about that is to sit kids down at computer consoles in the classrooms. Perhaps they'll manage to become "computer-literate"-whatever that means even if, in their mother tongue, they remain functionally illiterate. We now have factories so highly computerised that they can operate virtually unmanned.

Devices that shield us from having to come in contact with fellow human beings are rapidly taking over much of our daily lives. Voices synthesised by computers tell us what to do next when we place calls on the telephone. The same voices thank us when we have done what they ask. But what does "the same voices" mean in this context? Individuality, identity, everything that has to do with the uniqueness of persons—or of anything else!-simply disappears. No wonder the architects of, and apologists for, worlds in which work becomes better and more engaging to the extent that it can be carried out without face-to-face interaction, and in which people prefer machines that listen to them to people, come to the conclusion that intellectuals are irrelevant figures. There is no place in their scheme of things for creative minds, for independent study and reflection, for independent anything.

Aside from whatever advantages are to be gained by living our lives in an electronic isolation ward, and aside also from the loss to "them" of a market "we" now dominate, what good reasons are there for mounting our own Fifth generation project on a scale, as Feigenbaum and McCorduck recommend, of the program that landed our man on the moon? -if you can think of a good defense application.

One chapter of their book is devoted to "AI and the National Defense." However, the chapter is only six pages long, and is mainly a song of praise-perhaps gratitude is a more apt word-for the Pentagon's 11 enlightened scientific leadership." Laid on thick, the praise tends to betray its own absurdity:

> Since the Pentagon is often perceived as the national villain, especially by intellectuals, it's a pleasure to report that in one enlightened corner of it, human beings were betting taxpayers' money on projects that would have major benefits for the whole human race.

The corner mentioned is the Defense Department's Advance Research Projects Agency, DARPA, also often called just ARPA. Well, there is reason for workers in Al to be grateful to this agency. It has spent in the order of $ 500 million on computer research, and that certainly benefited the AT community (the artificial intelligentsia) enormously. But "major benefits for the whole human race".

The military, however, has good reason to continue, just as generously as ever, to provide funds for work on the fifth generation, as Feigenbaum and McCorduck make clear: The so-called smart weapons of 1992, for all their sophisticated modern electronics, are really just extremely complex wind-tip toys compared to the weapon systems that will be possible in a decade if intelligent information processing systems are applied to the defense problems of the 1990s.

The authors make five points. First, they believe we should look "with awe at the peculiar nature of modern electronic warfare," particularly at the fact that the Israelis recently shot down 79 Syrian airplanes with no losses to themselves. "This amazing result was achieved," they write, "largely by intelligent human electronic battle management. In the future, it can and will be done by computer.

Second, we cannot afford to allow the "technology of the intelligent computer systems of the future . . . to slip away to the Japanese or to anyone else.... Japan, as a nation, has a long-standing casual attitude toward secrecy when it comes to technological matters." Third, with the ever-increasing cost of military hardware, "the economic impact of an intelligent armaments system that can strike targets with extreme precision should be apparent . . . fewer weapons used selectively for maximum strike capability."

Fourth, "it is essential that the newest technological developments be made available to the Defense Department." Finally, "the Defense Department needs the ability to shape technology to conform to its needs in military systems. Here, perhaps more than in any other argument of the book, we are close to what it's "all about." A much shorter, and better, account of the Japanese Fifth Generation project appeared recently as a cover story of *Newsweek* *magazine*.. People generally should know the end use of their labor.

Students coming to study at the artificial intelligence laboratories of MIT, my university, or Stanford, Edward Feigenbaum's, or the other such laboratories in the US should decide what they want to do with their talents without being befuddled by euphemisms. They should be clear that, upon graduation, most of the companies they will work for, and especially those that will recruit them more energetically, are the most deeply engaged in feverish activity to find still faster, more reliable ways to kill ever more people-Feigenbaum and McCorduck speak of the objective of creating smart weapons systems with "zero probability of error" (their emphasis). Whatever euphemisms are used to describe students' Al laboratory projects, the probability is overwhelming that the end of their research will serve this or similar military objectives. To underline the importance of the computer as the "main artifact of the age of information," Feigenbaum and McCorduck instruct us just wherein the importance of the computer lies: [The computer's] purpose is certainly to process information-to transform, amplify, distribute, and otherwise modify it. But more important, the computer produces information. The essence of the computer revolution is that the burden of producing the future knowledge of the world will be transferred from human heads to machine artifacts [emphasis in the original].

How are we to understand the assertion that the computer "produces information"? In the same way, presumably, as the statement that a coal-fired electric power station produces energy. But that would be a simple and naive falsehood. Coal-fired power stations, transform energy, they do not produce it. Computers similarly transform information, generally using information-losing operations. For example, when a computer executes an instruction to add 2 and 5, it computes.

Social Impact of Machines

Since then, it has had some notable successes, enabling computers to perform-albeit in a very limited way some of the tasks normally done by our minds. Some Al workers see Al as a way of helping us understand human psychology; they try to write programs that tackle their tasks in the sort of way in which we do. Others see it as an approach to a theory of intelligence in general, human (and animal) intelligence being a special case. Still others simply want to write programs to do something (to understand language, to describe visible objects, or to solve problems, of various kinds), irrespective of how we do it.

And most of these hope that what their programs do will be not only interesting, but useful. The technological aspects of Al have suddenly become more visible. Public interest in Al, and media coverage of it, have increased enormously over the last two years. More and more people view it as an incipient technology of great potential power and social significance. The public interest dates from the announcement in 1981 of Japan's ten-year national plan for developing "Fifth Generation" computers. These are defined as incorporating Al techniques (as well as large-scale parallel processing).

Since then, large sums of money for Al-research have been made available also by governments and industry in the Western industrialised nations.

Developments

Several core research areas are likely to make solid progress within the next decade. Each of these is already being worked on in various countries, and progress does not depend upon the success of Japan's ambitious "Fifth Generation" project. One is low-level vision, based on techniques using parallel hardware and cooperative processing.

Current "connectionist" research in this area differs in its approach from work on two-dimensional pattern recognition by "property lists," and from-the-top-down /Iscene analysis" of three-dimensional scenes. Based on detailed studies of image formation, it is able to extract from the ambient light information about three-dimensional features (such as shape, depth, texture, and surface orientation) which in previous approaches could have been computed only, if at all, by way of high-level knowledge of the expected scene. Some of this work is being done in the context of human psychology and neurophysiology, some in a more technological context.

Dedicated (massively parallel) machines are being designed for this research, and major advances depend upon such hardware. A second area in which we can expect significant progress is robotics. This includes problems of movement control, trajectory planning, and visuomotor coordination (and will take advantage of advances in low-level vision). As in the case of vision, some projects will rely on "artificial" means to ensure success (such as light stripes for automatic welding machines, capable of recognising different sorts of weld joint and guiding the welder accordingly), while others will relate more closely to psychophysiological theories of motor control and visuornotor coordination in living organisms. Knowledge-based "expert" systems will multiply enormously in the next decade, not least because there is considerable commercial interest in them. Different domains of human expertise may require different approaches to knowledge engineering. In domains less fully covered by an explicit scientific theory. It may be easier to extract knowledge from human experts who are competent but who have not yet achieved the "intuitive" mastery of the domain that top-flight experts enjoy.

The latter give the right answer more often, but cannot easily introspect their reasoning processes, which happen very fast and are not consciously accessible. The former take time to come to a decision, after consciously weighing distinct considerations against each other and verbally identifying areas of unclarity. Domains (such as medical radiology) that depend on the comparison and interpretation of complex visual images are especially difficult to automate, since low-level visual processes are not open to voluntary inspection or control.

Indeed, experts often give highly misleading advice about how they may be carrying out the relevant comparisons. (Eye movement studies show, for instance, that expert radiologists do not scan x-ray photographs in the way they say that they do.) In tandem with the increasing experience of Al-trained knowledge engineers, further psychological studies of the organisation of knowledge in different domains should be useful. Research on expert systems will also focus on the computational architecture required to deal with large, complex knowledge bases.

Current systems are relatively simple and inflexible, and are restricted to very narrow domains. They can be incrementally improved, but only up to a point. Eventually, the interactions between the increasing number of independently added rules become too difficult to control, and the system's reliability and intelligibility are jeopardised. Current systems have no access to higher-level representations of the knowledge domain and their own problem-solving activity. Special problems arise if a system has to work in real-time, where unexpected events can require quick switching from the current activity to some other.

The next ten years will see some general work on powerful (IKBs) architectures (as well as the production of more examples of specific commercially useful systems), including parallel-processing devices. Progress can be expected also in natural language processing, both of individual sentences and of texts. Key issues include syntactic parsing, the integration of syntax with semantics, and the understanding of connected text.

Machine translation could in principle benefit from advances both in single-sentence parsing and in text analysis. Current work on parsing is motivated both by theoretical (linguistic) interests and by the hope of improving the man-machine interface so as to make it possible for nonspecialist users to communicate with programs in (some reasonable subset of) natural language. Where a program is used for some specific purpose, semantic factors can be more readily used to help in the parsing and disambiguation of queries and instructions input by the user.

Verbal interchanges about lunar geology, or about airline reservations, are already reasonably "natural" because of the exploitation of semantic constraints, and further domain-specific semantics will be developed over the next decade. More generally applicable (theoretical) research will continue into the best point at which to use semantics in parsing: from the beginning of the sentence, or spreading out from the middle, or only after an initial parse of the entire sentence? Text-analysis programs can already give a precis of most short news stories about specific topics (such as earthquakes, hijackings, and road accidents). But they rely on rigid, preprogrammed schemata, which provide the semantic skeleton of the types of stories concerned.

Some recent research is aimed at enabling a text analysis program to learn new schemata for itself, to integrate one schema with another so as to understand a story combining both, and to use a given schema to reason analogically in an unfamiliar context. A high degree of success cannot be expected within the next ten years, but our understanding of the relevant problems should be advanced. A variety of educational applications is already receiving attention.

Some are focused on particular curricular subjects, and require both a model of the theory of that subject and a model of the student's knowledge of it (which varies in level and in organisation, from person to person and from time to time). Others are less specific, and aim to use AI-based techniques to improve the pupil's attitude to intelligence in general. There is some evidence that both normal and handicapped students can attain greater self-confidence and intellectual achievement by experience with these specially designed programming environments. Controlled research into the classroom effects of Al-based systems has recently been initiated, and this can be expected to bear fruit within the next decade. An extremely important area, which is increasingly being studied because of recent hardware developments, concerns the computational properties of large parallel systems. At present, we understand very little of the potential and limitations of such systems. Some of the connectionist work mentioned above suggests that cooperative processing may have some highly surprising properties.

For example, the number of individual processors required to make the "human" range of visual shape discrimination appears to be markedly less than one would naturally assume. Again, making a connectionist system stochastic rather than deterministic improves its chance of finding an optimal solution. The computational properties of parallel machines will not be well understood for a long time, but experience with these new systems in the near future will doubtless lead to some advance.

Sciences and Technologies

The impacts of AI on other technologies will include many different examples of applications to individual problems. For example, an olfactory chip is being designed using AI techniques of pattern recognition. Given advances in very large-scale integration (VLSI), instruments and products of many different kinds will come to include chips whose design makes use of Al methods. Any commercial-industrial task that could benefit from even a limited degree of intelligence could in principle be performed better with the help of Al, so the technological applications of Al will be extremely diverse. AI will influence other sciences in their general philosophical approach as well as their specific theoretical content. Indeed, psychology and (to a lesser degree) biology have already been affected by computational ideas. And, contrary to what most people assume, AI has had a humanising effect in psychology. Al's influence will be especially strong in the psychology of vision and language, and, as noted above, it is likely that robotics will engage with the psychophysiology of movement. Psychological research will feed back into Al; for example, insofar as psychologists arrive at a better understanding of the organisation of knowledge, their work may be useful in designing computerised expert systems.

Social impacts will be of various types. First, there will be effects on individuals and institutions brought about by specific applications of Al, such as expert systems for medical diagnosis, legal and financial advice, or educational help. These programs will not merely provide a service (whose adequacy should be very carefully monitored), but will very likely change the social relations of the profession or institution concerned.

For example, if general practitioners, or nurses, can use an AI program to aid in various aspects of patient care, the social image of the specialist physician may be profoundly affected. (And legal responsibilities for medical decisions may be assigned in a way very different from today.) Likewise, legal programs may undermine the status of lawyers, and alter the nature of their work.

In both cases, while the mystique, of the human experts may be lessened, their opportunity for exercising their specifically human powers may be increased. The general public might come to be less dependent on human experts than they are today. Reducing the power of professionals such as doctors, lawyers, and teachers would certainly have advantages. But replacing human professional advice by computer programs is dangerous to the extent that Al systems in public use are inadequate-and/or ill-understood. Systems that have taken several man-years to develop (and whose original programmers may be retired, or dead) are often very difficult to evaluate or alter, because even computer scientists do not fully understand how they work.

A second type of social impact concerns general social trends brought about by applications of Al and information technology (IT). These include changes in the proportion of the workforce in service and leisure industries, changes in the division of labour and sexual roles, and changes in general lifestyles and patterns of interaction.

But other potential consequences of Al point in the opposite direction. The widespread use of home terminals, for instance, threatens to have an isolating influence even more powerful than that of television. If people are encouraged to work, and to shop, from their sitting-rooms, there may be unfortunate psychological effects in terms of personal stress and loneliness. Community computer centres could offset these effects to some extent, providing a social meeting place outside the confines of the home and nuclear family.

Some writers even predict that commercially available (and highly profitable) Al systems will be heavily used not only in task-oriented ways, but as surrogates for human contact. On this view, the strong tendency to anthropomorphism that most of us share will result in patterns of interaction being skewed away from human beings, and towards quasi-human computer systems (with naturalistic, "voices," and sometimes even "bodies").

Long-range Al research

There will be "more of the same," in that the areas mentioned above will provide perplexing problems for many years to come. Especially hard problems include learning, high-level vision, naive physics, and abstract work in computational logic.

For instance, expert systems are at present unable to "plain their reasoning except by "backwards-chaining": giving a resume of the chain of inferences (rules) that led up to their conclusion. They cannot relate their conclusion to the domain in general, or rely on an overview of the problem to assess the relative theoretical reliability of different hypotheses.

Nor can they monitor and adjust the structure of their own problem-solving, for they have no high-level representation of it. They are unable, too, to integrate different knowledge domains, and to use concepts and patterns of inference taken from one domain to reason (analogically) in another. Nor can current systems explain their conclusions differently to different users, taking account of the specific user's knowledge.

The user can ask for a "deeper" explanation (a more detailed inference resume) but the program has no user model in terms of which to adjust its explanations to the human's particular range and level of knowledge. For this reason also, the pattern of interaction between user and system is at present very limited. The user cannot offer his own conclusions for comment and criticism, for example, as students can do with human teachers. A special case of human knowledge is "naive physics," one's everyday knowledge of the properties and behaviour of different sorts of physical substances, and the nature of the causal relations between them. This knowledge enters into vision and motor control, and also into natural language.

For example, a language-using program would have to understand the differences in meaning between verbs such as pour, flow, spill, drop, and the like, if it were to give instructions or understand texts about activities dealing with liquids. Similarly, a robot capable of seeing that a container was just about to spill its contents onto the object below, and of adjusting its movements accordingly, would need some representation of the behaviour of fluids. Very little work has been done on these issues so far, and they are likely to provide a challenge for many years.

"Computer-aided design" is typically thought of as involving the graphical display of precise three dimensional specifications of various products (from machine tools through cars to buildings), taking into account a wide range of values of many parameters. But a recent form of computer-aided design involves suggestion rather than specification, in the sense that the design program originates novel ideas—ideas that are not merely quantitatively different from previous specifications.

For example, heuristic programs are already being used to suggest novel experiments (described at the intramolecular level) in genetic engineering, or to help design new sorts of three-dimensional silicon chips. The potential of systems like these should be further explored. The computer modelling of creative thinking will require long-term research, especially with respect to domains whose crucial concepts cannot be so readily defined as the concepts of molecular biology, chip circuitry, or set theory. Machine learning is a pressing problem for the future. If a program cannot learn for itself, its development is limited by the time and ability of the programmer to provide it with new information and ideas.