Computers Then and Now—Part 2

Maurice V. Wilkes
Olivetti Research Laboratory
Cambridge England

Abstract

In 1967, when Maurice Wilkes delivered his Turing lecture under the title 'Computers Then and Now', the computer field had emerged from its early struggles and was beginning to take on its modern aspect. Programming languages had ceased to be attached to a particular make of computer and were becoming increasingly mobile. However, the merits of time-sharing were still an issue and the ARPANET had yet to be announced. Wilkes commented on the achievements and problems of the day, and their implications for the future. In this second lecture, 29 years later, he will continue in a similar vein.

By 1967, when I delivered my Turing lecture 'Computers Then and Now', the computer field was well on its way to assuming its modern form. Progress had been made with high level languages (including the scientific study of their syntax), operating systems (including time-sharing systems), and CRT displays (including their application to Computer Aided Design). I referred to all these topics in my lecture.

One feature was missing. There was as yet no computer networking, although it was in the air. Three years were to elapse before the ARPANET came into operation.

My interest in computer communications was sparked off by a conference organized by the Benelux branch of the Institute of Radio Engineers (one of the precursors of the present IEEE) and held at Delft in the Netherlands in 1960. The conference was largely concerned with technical problems as seen from the transmission engineer's point of view. Much emphasis was put on the supposed need for data lines to be virtually noise free. It was not then appreciated, at any rate by those at the conference, that a very low rate of error can be achieved with a relatively noisy line if appropriate error control techniques are used.

A False Prediction

In 1964, in an article published in the New Scientist, a London based journal, I referred to computer networking as a coming development. The article was written in response to a request from the editor, made to me along with other engineers and scientists, to give my views on where technology would be heading by 1984. That year was chosen because '1984' was the title of a satirical novel by George Orwell, which had attracted much attention.

In my article, I wrote: 'Once a network of computers exists—and by 1984 it will be an international network—it will be the obvious thing to use computers for the transmission of messages'. I went on to say that many messages contained much unnecessary redundancy and that, in order to save bandwidth, the computer at the sending end would compress a message before sending it. At the far end the message would be restored to its original form before delivery. Speech would also be transmitted as data and reconstituted at the far end. I went on (remember that I was writing in a British journal):

> When you talk to a business associate in the United States and would just as soon have your English accent suppressed (along with any emotion that might creep into your voice and give away your business secrets) then you will save money by asking for the maximum compression possible, and will allow the computer at the far end to speak with its own accent. If, on the other hand, you are ringing up your loved one in a distant country you will order less compression and pay the bill gladly.

At the time the above was written, it was only eight years since the first transatlantic telephone cable had come into operation. Until then we had been entirely dependent on radio and only a handful of circuits across the Atlantic, for example, were available. Telstar, the first experimental communication satellite, had been launched two years earlier and was still being evaluated. Stripped of journalistic embellishments, what I was saying in my article was that a severe shortage of bandwidth for long distance telephony would persist for at least another 20 years, and that there would be an increasing emphasis on compression as a way of making best use of what bandwidth there was. This was an utterly false prediction and by 1984, nothing was heard of speech compression for telephone conversations.

On the other hand, the compression of television signals is an area of active interest. Should one predict that this too is a passing phase and that in 20 years time there will be no difficulty in providing ample bandwidth for all the television channels that the world needs without compression? Perhaps this may be so for television of present quality. But what will be the effect of the increasing demand for higher definition? I have learned my lesson on the folly of trying to predict the future and I will say no more.

The Economic Consequences of Technological Advance ...

When a commodity, that has long been in short supply or expensive, suddenly becomes easily available at very low cost, severe problems are created for the industry concerned. This has happened recently in the computer industry where the improvement in the performance of CMOS chips went with such speed that it caught everyone by surprise and led to the destabilization of the older section of the industry. Some hitherto successful manufacturers of minicomputers were forced to abandon the manufacture of hardware. Others, and these included the most famous companies, had to embark on a process of down-sizing and re-orientation that has been both difficult and painful.

Fiber optic cables are not only greatly reducing the cost of wide bandwidth transmission, but are making available circuits of a very much wider bandwidth than could possibly have been contemplated in the past. This change in the cost and availability of wide bandwidth circuits has called for changes of attitude in the telecommunications industry. For example, it is no longer appropriate to express the bandwidth of a circuit in terms of the number of telephone channels that it can provide and to estimate its potential revenue earning capacity accordingly. It is tempting to compare the impact of optical fibers in the telecommunications industry with that of silicon chips on the computer industry.

Email is rapidly establishing itself as an alternative to the use of the telephone. However, the public demand for telecommunications increases every year. In spite of email, I should be surprised if conventional telephone traffic does not continue to grow for the foreseeable future.

It has long been recognized that only a small part of the cost of carrying a long distance call is attributable to the long distance transmission, most of the cost being attributable to the local networks at the two ends. Long distance callers have in effect been subsidizing subscribers making local calls. This has now become more true than ever and if tariffs were to reflect the true facts, telephone users would pay very nearly as much for a local call as for a long distance call. The restructuring of the telephone industry now taking place may lead to this becoming a reality.

Optical fibers also have significance for local area computer communications. In particular they make possible a revolutionary approach to the disposition of computer equipment within a building. I shall return to this subject later.

. . . and the Political Consequences

Until recently, the telephone industry in all countries was a monopoly, either owned and operated by the state, or closely regulated by the state. With minor exceptions, the only service provided to the public was a telephone service and the only device ever connected to a telephone line was a telephone. As soon as devices other than telephones began to be connected to the telephone network, it became apparent that major changes in the regulatory position would eventually come about. This is the restructuring in the telecommunication industry that I referred to above.

The situation varies from country to country, but in most western countries there has been substantial liberalization of the environment in which the telephone industry operates. In the United States and Britain for example, long distance services are provided by competing companies. In some countries, notably in Britain, liberalization has also gone a long way at the local level. Here the situation is complicated by the existence of cable television companies which are fully able to take telephone and data services into subscribers' premises—in particular into private homes—if they are permitted to do so. Here there are major variations between individual countries. The economic prospects of taking fibers into the home are greatly improved if telephone, data, and entertainment services can all be carried through the same cable. In some countries, there is a danger that, at the local area level, the pattern of the industry will be determined by the whims of regulatory authorities, rather than by technical and economic forces.

Programming Languages . . .

I said, in my Turing lecture, that in the future people would look back on the period around 1967 as one in which the principles underlying the design of programming languages were just beginning to be understood, but that the time had not yet arrived to standardize for ever on one or two languages, as some people were suggesting. No-one would now dispute the truth of that remark. Even now we are far from approaching finality in the design of programming languages.

In 1967 it seemed that all attempts to design a programming language that would be ideal both for numerical calculation and symbol manipulation had proved disappointing. I ascribed this to the fact that, with the programming languages of the period, if you chose a particular language you had to go along with the data structures that the designer of the language had provided. For example, if you used FORTRAN or ALGOL you would naturally use arrays, whereas, if you used LISP then you would use lists. Nowadays, programming languages give the programmer the opportunity of constructing a wide variety of data structures to suit his own needs. This has been a big advance and the major modern programming languages are efficient for both symbol manipulation and numerical calculation.

The early programming languages had all been designed with a particular machine in mind and they ran on that machine only. By 1967, a number of these languages had been reimplemented to run on other machines. ALGOL 60 had from the outset been conceived as being machine independent.

A few people were experimenting with the problems encountered when compilers themselves were written in high level languages and then compiled to run on an arbitrary machine. I referred to this in my Turing lecture as a development that was not receiving the attention that it deserved. The problems were subtle rather than difficult and it turned out that insights obtained by practical experience were not easily communicated to others. In consequence, it was some years before the appropriate techniques became widely understood and applied.

One tiresome problem arose from the fact that there was no standard binary code for the reprsentation of letters and figures. Care had to be taken to ensure that literals, embedded in the source code of the compiler, were compiled in the correct binary code for the target machine. This problem went away with the general adoption of the ASCII representation.

. . . and Operating Systems

In my Turing lecture I made the remark that “there is reason to hope that the new found mobility [of programming languages] will extend itself to operating systems, or at least to parts of them”. For reasons which I will try to assess, this did not happen in the case of the major proprietary operating systems used on minicomputers and mainframes. However, a number of small scale operating systems were developed by research workers in computer science and these ran on a variety of machines. UNIX was originally of this kind. The story of how it developed, from small beginnings into the first major machine independent operating system, is a fascinating one which I hope will one day be told in detail.

UNIX came into its own when the first RISC workstations were designed in the 1980s. These had non-standard instruction sets and needed an operating system. UNIX was ready to hand. Not only was it written in C and capable of being compiled for any computer, but it had a comprehensive set of libraries and utility programs.

There were business reasons why vendors generally did not see it as being in their interests for their proprietary operating systems to be capable of being run on machines marketed by their competitors. Nevertheless, machine independent forms of those operating systems would have been useful to the vendors themselves, and would probably have come into existence if severe technical difficulties had not stood in the way. The principal difficulty arose because no clear interface, between the operating system and software designed to run under it, had been imposed on the writers of the software. This made it difficult to meet the requirement that existing software should run without modification.

Protocols: the Key to Progress

An implementer sees an interface in terms of a set of protocols that he must observe. Progress often occurs when a set of protocols emerge that the computer world as a whole is prepared to adopt.

The evolution of computer networking is a case in point. We had reached a situation in which the old ARPANET existed side by side with a variety of other networks, each having its own internal protocols. Bridges were established between the networks, but these did not always operate efficiently, particularly as regards the passing back of error messages. All the networks have now been merged into one comprehensive system known as the internet. This was made possible by the various authorities concerned seeing that their best way forward was to sink their differences and adopt the internet protocols. Once this had happened, nothing stood in the way of an explosive growth of computer networking.

A Changed Situation?

In 1967 there were already a few minicomputers in the world, but nearly all computing was done in computer centers equipped with large mainframes. These generally operated a batch computing service, but here and there time-sharing was coming in. Now we have both time-sharing systems and workstations—under which latter term I include the larger PCs.

A time-sharing computer is located centrally, whereas workstations are located on their owners' desks. However, sometimes these workstation are networked in such a way that other users can log-in and use them in a time-shared fashion. Nevertheless, the user on whose desk a workstation rests will have privileged access via his display, and he will usually regard himself as its owner.

It is generally taken for granted, and with good reason, that a display should be located physically near to its processor. This is because of the very wide bandwidth link that is required between the two. Other components, especially the high speed memory and the hard disk, must also be near the processor. These considerations lead to the configuration we are familiar with, namely a workstation containing processor memory and hard disk connected by a short cable to the display and keyboard. The workstation is situated on or near the user's desk and is connected to other workstations and servers via a local area network.

It appears that we are within sight of being able to design a fiber optic link of sufficient bandwidth to handle the whole traffic that passes from a processor to the display in a high performance workstation. If this is the case, and in what follows I will assume that it is, the processor and display need no longer be physically near to each other. The workstation can be moved from the user's office to some other convenient location, for example, to a central "equipment room", leaving the display behind. There is no technical reason why the workstation and display should be close together, or even in the same building. The keyboard must also remain behind, but in this case a low bandwidth connection is all that is required[1].

The user will see no differences resulting from the change, except that his office will be quieter and cooler, and he will not be troubled by technicians coming in to service his workstation. Similarly, the management will have fewer problems with air conditioning in offices and will enjoy the obvious advantage of having all the computer equipment, except the keyboards and displays, in one place.

On this model, the central equipment room will contain a set of regular workstations.

[1]There is at least one precedent for this separation. In the early 1980s at Xerox PARC it was found convenient to site a number of hot and noisy Dorado personal computers in a room of their own and to connect them to bit-mapped displays and keyboards in nearby offices.

It will also contain a variety of servers, for example, file servers and communication servers, placed there for convenience. These workstations and servers can be connected to a local area network in the normal way, with some advantage arising from the fact that all the nodes are in the same room. Cables providing connection to the outside world would come into the room and would be connected to the appropriate servers.

One could stop at this point. However, it would be natural to go on and reconsider the packaging of the items in the equipment room and their interconnection. Economy might be secured and space saved by dispensing with separate boxes for the workstation and bringing their contents together in a rack. Again, it is possible that a local area network might not be considered the most suitable way of interconnecting the workstations and servers now that they are all together. Possibly, some form of bus-like interconnect would be both faster and lower in cost.

The above configuration is merely an alternative way of implementing essentially the same architecture as we have at present and I commend it to your consideration as something with possible merit. To the user, it would be just the same as if he had a workstation in his office. Its possible advantages would be from the point of view of planning, installation, and maintenance and, I might add, of physical security. The ultimate cost of the fiber link will, of course, be an important consideration.

With all the equipment now in one room, it becomes possible to consider fundamental changes to the architecture, although my interest is not primarily in this possibility. One might add additional memory shared between the processors, converting what were workstations into a parallel computer with shared memory. Some readers might like to follow this approach and consider the merits of installing, in the equipment room, a version of their favorite highly parallel supercomputer which had remote displays connected to some or all of the processors.

An architectural enhancement of a different kind would be to provide facilities for switching the incoming fibers from one workstation or processor to another. The fast opto-electronics switches for fibers that are now beginning to be used in the telecommunications industry would be very suitable for this purpose.

At the Olivetti Research Laboratory in Cambridge, England, we have designed a fiber optic link that can go between a workstation and its display. In its first form, this link, or 'workstation adapter', will be plug compatible with a regular workstation connection. One will simply unplug the display from the workstation and reconnect it through the adapter. The workstation and the display can then be sited as far apart as the fiber allows.

We hope that the workstation adapter will enable some experience to be obtained with the new system. In the first experiments it is likely that, as a matter of convenience, a conventional local area network will be used for the connection of the keyboard.

Later, we plan to implement a fiber link more efficiently interfaced to the workstation and providing also a low bandwidth channel for the keyboard.

Acknowledgement

The work described in the final section of the paper is being done in the Olivetti Research Laboratory, Cambridge, England, in active collaboration with the Director, Dr A. Hopper, and I am grateful to him for allowing me to refer to it. I also wish to thank Dr R.M. Needham for reading the manuscript and commenting on it.

References

Wilkes, M. V. *Computers Then and Now.* ACM Turing Lecture, 1967. Journal ACM, vol 15, p. 1 (1968). Reprinted in: ACM Turing Award Lectures: The First Twenty Years, p. 197. ACM Press, New York (1987).

Wilkes, M. V. *A World Dominated by Computers?* New Scientist, London, 1964. Reprinted in: The World in 1984, ed Nigel Calder. Penguin Books, London (1965).