The Melting Machine: from PC to Pronit

Reuben Thomas

20th October 1995

Introduction

In the last fifteen years, the personal computer has swept from recondite obscurity, the preserve of hobbyists and hackers, to mundane ubiquity. Many people use one every day; at work, for word processing, at home, for playing games, or at school, for education. Yet within thirty years, personal computers will have disappeared, replaced by a far more pervasive and powerful device, the networked pronit.

Before describing the pronit and its genesis, a few words about prediction are in order. Prediction of the future is at best uncertain, and meteorologists, economists and astrologers are regularly pilloried for the extent of their errors. Prediction of technological advances is no different: most predictions that look more than a few years ahead seem hopelessly naïve in retrospect. In the 1930s, an article was written claiming that in a few years every home would have a large electric motor in the attic to drive machines throughout the house.¹ Instead, the motor was miniaturised, and each machine now has its own. Yet there have been notable successes: Arthur C. Clarke's 1945 paper on earth-orbit satellites not only set out how they could be deployed, but accurately predicted many of the uses to which they would be put.²

In computing, predicting up to ten years ahead is relatively straightforward, because that is roughly the time it takes for a laboratory prototype to become a mass-market product. The main source of error is in guessing which prototypes will be successful and which not; it is easy to predict from current trends how much memory computers will have in ten years' time, and how powerful they will be, but harder to guess what operating system they will run, or which processor they will use. Similarly, it was possible to predict the use of devices such as mice and penbased computers, but not necessarily which would dominate the market. Further than ten years ahead, one has to return to crystal-gazing, and guess what new inventions will occur, and what effect they will have on computing.

The thesis will thus be supported in two ways: first by discussing current products and trends that point to its fulfilment, and secondly by examining the possible fruits of current research.

The pronit

The disappearance of the personal computer will be caused by the marriage of two technologies, the digital computer and the communications network. For decades, networks have spanned the world providing telecommunications services. Shortly

¹Quoted in Fry, S., *Paperweight* (ISBN 0-7493-1397-8), p. 404.

²Clarke, A. C., *Extra Terrestrial Relays* (Wireless World October 1945, pp. 305–308).

after the invention of the digital computer, the advantages of linking computers together were recognised. In the 1970s the Internet was invented; today, it is a worldwide network of tens of millions of computers, currently growing at an exponential rate, with millions of new connexions each year. At the moment, the benefits of networked computers are still largely reaped directly by humans, whether for communication, such as electronic mail and computer facsimiles, or information retrieval, from file repositories, databases and the World Wide Web. Computers also run the networks, deciding the routes along which data flow, compensating for faults, and handling the accounts of commercial networks such as the telephone system. Increasingly though, computers are talking to one another on subjects not directly concerned with the running of the network or the tasks that humans ask them to perform: some computer programs automatically register their users with the manufacturer; complex calculations can spread themselves across networks so as to run on many computers simultaneously and hence complete faster, and financial systems add up and distribute the electronic money exchanged during their users' dealings.³

Traditionally, computers have consisted of several components directly connected together. The processor performs computations according to the instructions in the programs it is given. The memory contains the programs and the data on which they act. The storage devices, such as magnetic disks and tapes, store programs and data that are not currently being used. Input devices, such as keyboards and mice, are used to give commands to computers, while output devices, such as screens and printers, are used to give the computer's response. All the components are connected together by a "bus" along which they exchange information.

As networks have increased in importance, researchers have realised that the internal bus of a computer is itself a network. Programmes such as the Desk Area Network at Cambridge University have built computer systems whose components are connected by a network.⁴ The components of the system are no longer physically connected: this has immediate advantages, as noisy and bulky storage devices no longer need to be on the desk; all that is needed are the input and output devices such as the screen, keyboard and mouse. But once this is done, the computer's identity is lost, and more fundamental changes cry out to be made. If the parts of several computers are all connected to the same network, a variety of "virtual computers" can be dynamically created, according to need: the same set of components could provide ten word processors or two graphics workstations. One might book a set of resources for a particular task; this would then appear as a single virtual computer for the duration of the task, after which the resources would be available for others to use in different configurations.

The result of this disintegration of the computer is the "pronit" (short for "processing unit"). I have invented this word to replace the word "computer"; it is supposed to represent the same concept from a different point of view. Like "computer", "pronit" has several connected meanings. First, it means a processor: we currently use one word, "computer", to refer to the processing unit, screen, keyboard, disk, memory and so on; in the future we will think of these as being separate, and distinguish the computing machinery (pronits) from the mass storage (memory and disk) and the terminal (screen and keyboard). Secondly and most commonly, "pronit" will mean "networked processing unit", as all processing units will be networked (a

³It is commonly believed that computers can only do what they are told. That is true, but it is important to distinguish between actions that are initiated by a direct command from a user, for example a word processor printing a document, and those which result from the interaction between the computer and its environment, for example a chess computer making a move. In the latter case, computers can be seen as genuinely taking the initiative.

⁴Hayter, M. and McAuley, D., *The Desk Area Network* (University of Cambridge Computer Laboratory Technical Report 228).

processor is useless unless it has access to memory and input and output devices) and communication will be conceptually joined to computation. Thirdly, the pronit may be used as a rough measure of computing power: a task may be said to require a certain number of pronits; this corresponds to the idea of creating a virtual computer for a particular task. The shift from the first to the third sense is a move from the concrete to the abstract, and mirrors the move from personal computers to the networks of pronits that will emerge in the next few decades.

Past and Present

Until recently, most computers have not even been connected externally to other computers via networks, let alone built around internal networks. Networks have traditionally fallen into two classes: the LAN (Local Area Network), such as Ethernet, which only stretches up to a few hundred yards, and the WAN (Wide Area Network), such as the telephone network, which connects computers over distances of many miles, but is slow and unreliable for data transmission. Both have until recently been expensive, LANs because of high installation costs, and WANs because of their low speed and fidelity. With the increasing availability of ready-installed twisted-pair cabling for use by LANs and the introduction of fibre-optic communications links for WANs, the bandwidth and reliability of the latter have increased enormously, and the costs of both have fallen. The convergence in speed of LANs and WANs and the use of the Internet protocols, which work on both types of network, will make the creation of virtual computers from pronits much easier.

Computers have also been monolithic in construction, as putting all the components on one circuit board resulted in cheaper, faster and simpler designs. In particular, it was much faster for the computer to access its memory directly than across a network, and networks were accessed via memory. However, while network speeds have risen dramatically, memory has not improved as much; fast LANs can now deliver data faster than memory, so distributing a computer's components across a network does not degrade performance. Communication between computers is actually improved, as processors can communicate directly, instead of having to send data via their respective memories. Increasing standardisation of computer components will also mean that it is no longer more expensive to build and sell them separately; on the contrary, greater economies of scale will be possible that way.

In the past programs and data stored on one computer could not easily be accessed from another, and programs that worked on one type of computer would rarely work on another. Computers were also configured to their owners' tastes, and could not conveniently be reconfigured for other users. Recent work at Olivetti's Cambridge Research Laboratory has resulted in "teleporting", which makes a user's data and programs available to him on any computer able to run the teleport software, wherever in the world it is.⁵ Simpler systems such as the World Wide Web allow data to be accessed on almost any computer, and new systems such as Java and Telescript allow programs to be moved in the same way between different types of computer. To allow universal teleporting it only remains to create a universal protocol for configuring computers (setting user preferences such as the colours used in the display, and the language in which the user communicates with the computer) and a large body of application programs that can run on all types of computer.

In the past, personal computers have been obtrusive because of the bulkiness of their main input and output devices, the Qwerty keyboard and cathode-ray tube

⁵Richardson, T. et al, Teleporting in an X Window System Environment (IEEE Personal Communications Magazine Vol. 1 No. 3, pp. 6–12).

screen. Advances in cathode-ray tube technology, as well as the newer LCD (liquid crystal display) and current research into light-emitting polymers, will soon yield affordable screens that can be wall-mounted; light-emitting polymers may even yield disposable screens. At the same time, voice and gesture recognition will render the keyboard obsolete for most tasks. With flat wall-mounted screens, and unobtrusive microphones and video cameras as the main means of communicating with computers, desk space will no longer be needed, as the computing and storage elements can be stored out of sight, connected to terminals which will be mere items of furniture. The personal computer will have melted away.

One obvious problem with "public" terminals is the identification of the user. In a simple-minded system, the user might have a magnetic card which identifies him to the computer, and allows his data and programs to be accessed. However, this may not be desirable, since if personal data are available to be fetched from any computer, they can be read fraudulently. More sophisticated authentication mechanisms such as voice matching or fingerprinting could be used, but it is likely that the user will prefer to take confidential material with him in a form in which it can be transferred directly to a computer he chooses to use.

Storage devices the size of a matchbox that can store billions of characters of data and programs, using either magnetic disk or memory chips, already exist. These could be carried around and plugged in to public terminals. A more convenient alternative might be the active badge, a credit card sized device that can be worn like a badge.⁶ It contains a microprocessor and an infra-red receiver and transmitter. As it is always switched on it can identify its user to a computer as he approaches it, so that when he arrives his personal computing environment has already been set up, and signal when he leaves, so that his work is saved and his session automatically ended. It can also transmit information on the user's location, allowing automatic routing of telephone calls to the nearest telephone (and video-phone calls to the nearest computer). Combined with a storage device and a higher bandwidth transmitter, the active badge could also act as a personal repository for commercially important or personal information.

Future

The separation of computers into their component parts using a network is only the beginning of a process which will reach its logical conclusion when distributed computing finally fulfils its promise. "Distributed computing" means spreading a program over several processors, each performing part of the program's function. In the past, this has been seen as a way to speed up programs, but when programs can distribute themselves over a network, they can also migrate from processor to processor. This enables "load balancing" to take place: programs migrate away from heavily loaded processors to those with a lesser work-load. Distributed programs can also migrate in search of data, or even follow users around.

Distributed computing turns separate computers each running their own programs into pronits which each run parts of many programs. These pronits will be ubiquitous, embedded in and controlling the network, and part of everything connected to it; input and output devices, and storage devices. Other machines such as printers, cars, microwaves, washing machines and fridges will be connected to the network, allowing them to talk to one another and be controlled remotely. Computing power will no longer be divided into discrete and qualitatively different units, but be a substance rather like electrical power. Computers as objects will be replaced by computing as a facility.

⁶Hopper, A., Communications at the Desktop (Computer Networks and ISDN Systems Vol. 26, pp. 1253–1265), section 2.

The current trend towards integration of different communications media, such as electronic mail, facsimile, telephone and videophone, will continue until the media fuse; printed messages will be able to be read aloud, and spoken messages printed; messages of different types will be accessed and manipulated in the same way, and combined to make multimedia documents. As an example, consider the help function of a word processor. This may give access to help texts, spoken or displayed as the user wishes; if further explanation is required, explanatory video clips might be shown, or even generated on demand; finally, the program could connect the user to a human expert to answer the question. Thus even the boundary between interacting with a human and with a computer will be blurred.

Another example of a multimedia application of a rather different sort is the memory prosthesis, currently under development at Xerox's EuroParc.⁷ It is essentially an active badge with a display, and instead of providing information to the network, it gathers it from the network. It can record where the user has been, and with whom, using active badge information from other users. It can also record associations between people and documents, and events such as telephone calls. The aim is to help answer questions that the user may have, ranging from wondering where he left something (this can also be solved more directly by tagging the object with an active badge!), to with whom he discussed a particular document. The display can list recent events, and focus on particular people, objects, documents, committees, and so forth. In future such devices could also use speech and video, recording everything that the user says and everything said to him, and also recording all his movements.

Another consequence of the integration of different media will be the disappearance of medium-specific appliances: telephones and televisions will cease to exist. Information storage media such as reference books will become far less common, although the forms used for recreation, such as book fiction and newspapers, will survive.

Currently, people store their own data and programs. There is much duplication of programs such as operating systems and word processors, and data such as system documentation and dictionaries. Unless one works with video and sound, one tends to generate relatively little data of one's own. Even with an increasing tendency for people to generate their own video and sound, cheap personal storage devices will be able to hold personal data. On the other hand, resources such as dictionaries and picture libraries may well be accessed remotely, and charged on a pay-per-use basis. This will probably even happen with applications programs such as word processors and spreadsheets, which can already be bought and updated electronically. Much information will also be provided free, by libraries and other academic institutions. Also, as people increase their investment, in information terms, in data storage, they will start to use archive services, which will take regular copies of their data automatically. Such services will have to provide security against unwanted inspection and alteration of the data they store.

These examples of foreseeable developments, though interesting, are unimpressive when set against the general changes that will take place in society's use of computers. The use of computers will no longer be a distinct activity. Computers, being out of sight, will be out of mind; people will think in terms of information and communication rather than in terms of computation. It is not only buildings that will become active, controlled by and permeated by communications equipment; clothes will contain active badges and pronits, and screens will be built into glasses, like a head-up display on the world. We will never be out of touch, using satellite and radio links to communicate to the ends of the earth.

⁷Lamming, M. G. et al, The design of a human memory prosthesis (Computer Journal Vol. 37, pp. 153–163).

Not only the new active objects, but also the old familiar ones will be more accessible. Computers will be able to read books, making translations and looking up commentaries on the spot, and to identify paintings and birds that the user sees. At any given moment, his computing environment will offer him information from the totality of that available, according to what seems relevant to his current needs. This will increase our mental capacities and alter the way we think, as the digital computer on a much smaller scale freed mathematicians and engineers from the dull necessity of arithmetic, and allowed them to concentrate on more abstract problems.

Obstacles

Pronits will bring problems as well as benefits, and there are two main difficulties which will have to be mastered for the benefits to be fully realised: the retrieval of information, and its abuse.

Of the various problems faced by computer science today, one of the greatest and most pressing is that of information retrieval. We now have access to vast amounts of useful (and useless) information, and that will continue to increase exponentially for the foreseeable future. What is not really known is how to organise and search the virtually unlimited amount of information available to extract that which is relevant to a particular user with a particular query. Even in a specialised domain such as that of providing a library catalogue, there are problems: how can the user be allowed to specify his query in a manner that is both natural and comprehensible by the computer, and so that the information returned by a search will be relevant yet complete? To solve this problem in general will require advances in artificial intelligence, on which I shall comment later.

The second problem is that of information abuse. The initial reaction of many people to the active badge, when its uses are described, is that it represents the incarnation of Orwell's Big Brother, the ultimately intrusive State. This is certainly a possibility if the information generated by such devices is misused; the obstacles currently posed to intelligence-gathering services by the sheer volume of information available will be removed by the invention of the searching mechanisms outlined above.

The problem is not insoluble, however, provided that the principle of reciprocity is observed, that is that government and citizens should be subject to the same rules. With that proviso, a balance can be struck between the government and its citizens anywhere from total secrecy to complete freedom of information.⁸ Secrecy is hard to achieve; the challenge is to develop mechanisms which are effective, yet simple enough for everyone to use: by far the largest cause of security breaches is human error, either in the design or operation of the system. Security systems must also be flexible: it might be desirable to grant different levels of access to one's data to family, friends, colleagues, government agencies, the general public and even companies with which one deals.

As well as access to information, the question of who owns and maintains the networks, especially the worldwide Internet and its successors, is important. At the moment it is owned by a multitude of private and public bodies and profit-making companies with different and often conflicting interests. In principle, nearly everyone will be able to afford access to the Internet, just as currently most people have a

⁸I favour as much freedom as possible. Freedom of information not only promotes a fairer society, but encourages citizens to take an interest in government. More fancifully, complete freedom of information would allow everyone to know almost anything about anyone. This would not destroy privacy, but redefine it as something people accorded one another rather than being granted automatically. Society would become more open, although there would be definite disadvantages such as intrusion by advertisers and insurers.

telephone, but this must be ensured. Freedom of information is worthless if a large section of society cannot afford to access it. Similarly, if access to a network can be restricted or even subtly sabotaged at will, rights of access to information made available through that network are of little use.

Omissions

In my predictions I have omitted several emerging areas of technology, notably nanotechnology, self-replicating machines and artificial intelligence. There are two connected reasons for this: first, these technologies are in their infancy, and will probably not greatly affect society in the next thirty years, largely because of the time required to bring products to market and gain a significant market share; secondly, as these technologies are so new, it is even more difficult to predict what fruits they will bear. Nevertheless, they merit comment.

Nanotechnology is the building of extremely small machines, whose finest structures can be measured in nanometres, a thousand times smaller than the micrometre machines we have at present.⁹ Such machines, used in teams of hundreds or even millions, could perform surgery from inside a patient without incision or anaesthetic, maintain large buildings, and report and correct faults in networks, to give a few possible applications. The machines would contain tiny nanopronits, which need not be networked, running small programs, and could obtain their power from blood sugar, the sun, or batteries as appropriate. Nanomachines could also perform everyday tasks, such as dusting and washing clothes, using small-scale rather than the current large-scale techniques, and could form part of an "active building" which automatically cleans itself.

By allowing the construction of smaller digital circuits, nanotechnology could also give us faster and more powerful computing devices, such as processors and memories. Current digital memories require several tens of thousands of electrons to store one binary digit or "bit" of information (either a zero or a one); current work in the Cambridge University Engineering Department and at Hitachi's Cambridge Laboratory aims to produce memories that use a single electron to represent a bit.¹⁰

Self-replication will be extremely important in the near future, but applied to programs rather than machines: distributed computing relies on programs being able to spread copies of themselves across networks. When self-replicating machines are constructed, they will be especially useful in places that are hard to reach; one of the main thrusts of current research is aimed at making such machines for extra-terrestrial exploration. Self-replication will be most useful when applied to nanotechnology; because many possible applications for nanomachines are in places inaccessible to humans, machines that could repair themselves or manufacture replacements on site would be far easier to manage than those which required human monitoring and intervention.

Earlier, it was mentioned that advances in artificial intelligence (AI) would be necessary to enable intelligent filtering and presentation of information by computers to their users. In fact, techniques developed in AI research underlie some of the other advances mentioned, including voice and gesture recognition. Of even greater potential importance is the ultimate goal of AI, the creation of a self-conscious intelligent machine, but I do not think I or anyone else can predict what will happen when it is finally reached.¹¹ Whatever does happen will have more far-reaching con-

⁹Dewdney, A. K., The Magic Machine (ISBN 0-7167-2144-9), pp. 85-93.

¹⁰Nakazato, K. and Blaikie, R. J., Single-electron memory (J. Appl. Phys. Vol. 75 No. 10, pp. 5123–5134).

¹¹It is almost certainly possible; doubters such as Roger Penrose (*The Emperor's New Mind* (ISBN 0-09-977170-5) and *Shadows of the Mind* (ISBN 0-09-958211-2)) and J. R. Lucas (*Minds*,

sequences than anything I have discussed, and will be the true achievement of the information age.

Conclusion

The next thirty years will see the focus of attention on computing shift from the technical to the social; while there will undoubtedly continue to be fascinating technical advances (my predictions rely on them), the social changes effected by technology will become of greater interest and importance. Computing will no longer be a fancy new plaything, but become much more, insinuating its way into our lives until we rely on it almost as a symbiote. Though the digital computer is our newest tool, it will soon become our most fundamental.

Postface

I was gratified to note, on 27th January 2001, the existence of *The Invisible Computer* by Donald A. Norman, published in 1999, which, it seems, argues in the same direction as this essay.

On 7th July 2006 I noticed *The Next Few Decades of Computing* by Linus Vepstas (http://linas.org/theory/eternity.html), an occasionally revised essay started a few years after mine with a similar theme.

Acknowledgements

Eugenia Cheng, Donald McFarlane, James Martin, Lina Christopoulou, Martin Richards and Paul Phillips criticised the essay; Andy Hopper and Bill Newman provided much of the inspiration for it in their lectures as part of the Computer Science Tripos, and also provided references, as did Jeremy Douglas.

The essay was written for the 1995 Master's Prizes competition at St John's College, Cambridge, in which it won a prize.

Machines, and Gödel, Philosophy Vol. 36 No. 112) have generally been misguided about the nature of computers and determinism (Hofstadter, D., Gödel, Escher, Bach: An Eternal Golden Braid, ISBN 0-14-017997-6, pp. 471–479).