

OPENING SESSION: 9:00-10:15

KEYNOTE ADDRESS: TOWARD A HISTORY OF (PERSONAL) WORKSTATIONS

Gordon Bell
Introduction by Allen Newell



Allen Newell (left)
Bill Spencer (right)

The simplest way to introduce Gordon is to say he's one of the world's leading computer designers. However, that doesn't necessarily recommend him to be the keynote speaker at a conference on the history of the workstation. There are a couple of other things to be said about him that make it quite relevant that he give the keynote speech.

First of all, of course, Gordon is a product of that whole MIT-related environment up in the northeast corner of the United States from which so much of the development of workstations has come. He did his master's degree at MIT, working with Ken Stevens on speech recognition. He went from there to Digital Equipment Corporation, where he was a principal in Digital's PDP efforts over many years: the PDP-4, the PDP-5, the PDP-8, the PDP-6, the PDP-10, the PDP-11, the PDP-16, (nobody remembers the '16 except Gordon and myself, and perhaps Wes Clark, but there was something called the PDP-16), the VAX, not the PDP-12 (which was related to the LINC and Wes Clark). All of these machines played important enabling roles with respect to our topic here.

Gordon did come to Carnegie-Mellon University in the mid-1960s, which is when I first met him. He has this love-hate relationship with Digital—he went through a hate phase and showed up at Carnegie for about six or seven years, before going back. He never really left Digital,



Gordon Bell is currently Vice President of Engineering, Ardent Computer Corporation. He was educated at the Massachusetts Institute of Technology, where he received a S.B. (1956) and S.M. (1957). After a year in Australia as a Fulbright scholar (1958), Gordon continued at MIT in the Division of Sponsored Research Speech laboratory until joining Digital Equipment Corporation in 1960.

At Digital, Gordon was the designer of the PDP-1 Message Switch for ITT (1960-1961). He was architect/designer of the PDP-4 (1962-1963), PDP-5 (1964), PDP-6 (1963-1966), and PDP-11 (1969) hardware and software. From 1966 until 1972, Gordon Bell served as Professor at the Carnegie-Mellon University, returning to Digital in 1972 as Vice-President of Engineering. He held this position until 1983 when he became a founder and Vice-Chairman of Encore Computers. In 1986, Gordon expanded his entrepreneurial activities as a founder and Chief Scientist of Dana Computers, manufacturers of single user supercomputers.

Gordon is coauthor of many books and papers on computer structures and design, especially related to multiprocessor design. He is the recipient of the ACM/IEEE Eckert-Mauchley Award and the McDowell Award. In addition, he is a Member of the National Academy of Engineering.

of course, because things like the PDP-11 and the beginnings of the VAX all happened while he was at Carnegie.

There are other things that you need to know about Gordon. One of them is that he was a key member in founding the Computer Museum in Boston, directed by Gwen Bell. As a matter of fact, Gwen is out here and she has a big inverse Santa Claus bag. If anyone has any great workstation artifacts, you are supposed to just dump them in this bag as you leave the room, and she will take them back to the Computer Museum. So, that's another connection with the history of computers.

Back in the early 1970s, Gordon and I wrote a book called *Computer Structures*. To be more precise, because that's the way it always is with Gordon, I helped Gordon to write this book. I cleaned up the prose and did a few things like that on it. That book really reveals Gordon's interest in developing the frameworks within which things happen. As some of you know, our first and preferred title for that book was *Computer Botany*, not *Computer Structures*. We viewed our goal as laying

out the whole structure of the collection of the artifacts we call computers. They were proliferating and needed to be understood, studied, and classified. That same attitude has been with Gordon throughout his career, as I am sure you will see in his talk today. In fact, I could have predicted that Gordon would take the view, which I share, that you can't really write history about things until you have the space laid out in which that history can occur. I think that Gordon's role for us today is to lay out some of the dimensions of the space of workstations so that we can then proceed to deal with the history.

Toward a History of (Personal) Workstations

C. Gordon Bell

Computer and Information Science and Engineering Director
National Science Foundation

I originally accepted this keynote honor for several reasons. First, of course, was to respond to Alan Perlis's request because, ever since he headed the Computer Science Department at Carnegie-Mellon University where he was my boss, I have really had a hard time saying no to Alan. Second, was one to which Allen Newell alluded—to identify the important artifacts that should be preserved in the Computer Museum and, as one of the curators of the museum, to tune up my view of history. The third was to posit a taxonomy of the history of workstations in an evolutionary framework, both in product and in process. How does the technology across all the components drive this evolution? I always insist on knowing the whole picture, going back to the beginning and seeing if I can trace it since 1949 or so. Then, certainly, I feel compelled to write the history into the future, which is no good unless you can get it up to the year 2000. I stop in this history at about 1990. I think it is important when you are doing history to be able to look into the history of the future as well.

It turns out that there were three other reasons to write this paper, all of which were selfish. One was that this might be an interesting way of putting such an evolutionary book together. So I took this as a challenge; this is a breadboard for an outline (I think in terms of breadboards and prototypes and production machines and, ultimately, obsolescence). Then there was some work that I spent six years on, up until 1982, that's now clearly bearing fruit. This is work on VAX and the VAX Strategy, which is the notion of building a hierarchy of compatible machines. I had to put that into an appendix to record this work since I'm not at Digital now. I couldn't help but write that piece of history into the paper in order to clean up that piece of my life. So you may find that this paper is a patchwork quilt, and this is my apology for that patch of the quilt. I believe it is important in the history of workstations since VAX has about 10 percent of the workstation market.

Finally, an idea that has been in my head for many years is that, as an historian, what we really need is an historian's workstation. I hope coming out of this conference will be somebody that will take on the role of building historian's workstations so that, in fact, histories are much easier to do. It's obviously a very important segment of artifi-

cial intelligence, because artificial intelligence has a whole discipline that deals with truth maintenance.

Introduction to the Conference

In June 1976, an international conference was held resulting in a collection of views published in *A History of Computing in the Twentieth Century* (Metropolis, Howlett, and Gian-Carlo, 1980). The participants presented work up to the early 1950s and thus had the benefit of 30 years of hindsight. The late mathematics historian, Kenneth May, outlined the pitfalls of participant-written histories, commenting: "Historical description requires a time-lagged approach and means getting into understanding things as people understood them then, not as we understand them now." He also urged everyone to be open. Most autobiographies rewrite the past from the present perspective. Good diarists and notetakers have reference points. The best we can do is provide grist for the mill of the historian.

The History of Programming Languages Conference held in May 1977 (Wexelblat, 1981) provided a timeline to 1970 as context with detailed coverage of: Algol 60, APL, APT, BASIC, COBOL, FORTRAN, GPSS, JOSS, JOVIAL, LISP, PL/1, SIMULA, and SNOBOL. This covered a period of initial developments from 1954 (FORTRAN) to 1967 (APL) with evolutions up to the present. Ten to twenty years of hindsight gave them the ability to select topics.

If this conference were limited to a 15-year rule, applied by the *Annals of Computing History*, no one would be describing personal workstations. Only a few experiments existed, but people were building systems with multiple graphics terminals and a workstation industry on a profession-by-profession basis started to form. Twenty years ago, most of us were trying to make our newly designed interactive, timesharing system work reliably and economically. In my own case, having also worked on small machines that ultimately became the minicomputer, I saw two independent threads for economical computing, including interactive computing: large shared systems and dedicated small computers.

Personal workstations, like other man-made objects, appear strictly evolutionary, going through the following stages:

- 1950s idea (documented article, proposal, movie) stimulated through early stand-alone use of small computers (e.g., LGP-30, G-15)
- 1960s breadboards to demonstrate the idea and selected use of large computers with graphic displays (e.g., DEC PDP-1, IBM 7090, LINC, TX-2)

- 1970s (early) limited use of interactive shared workstations using graphic display terminals connected to mini-computers; establishment of an industry to supply terminals and professional applications software
- 1970s (late) working prototypes of personal workstations with concept testing through use in a complete environment (Xerox PARC)
- 1981 introduction of personal workstations by Apollo, SUN, Xerox, and first use by early adopters
- 1983 full-scale use for selected professionals; many companies formed; JAWS (just another workstation) term coined
- 1985 healthy industry with evolutionary product cycle and beginning shake out of suppliers
 - ? steady state supply to captive users
 - ? decline through replacement or superposition of functions in some other form of information processing appliance (e.g., conventional personal computers)
 - ? extinction

Hence just as we should concentrate on tracing various workstations through these stages, we may also treat this as a conference about the first phases of development. More hindsight will be needed to write history.

Waiting longer, which gives a future view, has the risk that extinction will come through agglomeration with conventional computing. Alan Perlis urged me to come to this conference because he feared this extinction. Conventional personal computers, that is the evolving IBM PC, will become the de facto workstations for virtually all applications by 2005. The latest chip introduction by Intel strongly permits and supports this to occur, if history is any indicator.

The real workstation phenomena and industry are not well described or recorded at all because most of the products have come from industry, where individuals are not rewarded for writing papers but for producing and selling machines. Examples are, Doug Drane and the ATEX system produced law and typesetting offices, doing essentially what a STAR does, only five years earlier and less expensively; GE's genigraphic system for the graphic artist; Bill Poduska and the Apollo workstations; or Applicon, Calma, Computervision, DEC, Evans and Sutherland, IBM, REDAC, Wang, Datapoint's ARCnet, and so on. No university (other than Utah and Brown) really participated until 1980. Noticeably absent are Dave Evans, whose direct contributions or indirect contributions through products at Evans and Sutherland, and graduate students who went on to create much of the science

behind workstations. Ivan Sutherland's work was recorded. Xerox PARC was a notable exception, which published its work in a delayed fashion and was marginally successful at taking the research into products in a commercially viable fashion.

We now need to focus and agree on some definitions and dimensions, making sure the various threads of personal workstation development are explored so the historical facts can be recorded properly and eventually interpreted by real historians. We need to find the first use of the name. For example, the second edition of the *Encyclopedia of Computer Science and Engineering*, 1983, uses the word three times, but in a somewhat different context. By the summer of 1983, the phrase JAWS (just another workstation company) could be heard in the venture capital community when they realized they had funded too many companies to build personal workstations.


What a (Personal) Workstation Is

Terminals connected to large computers (supercomputers, mainframes, and minicomputers) and utilizing interactive professional-applications programs are the clearest antecedents. These had, as direct descendants, terminals connected to shared, but dedicated, minicomputer systems in the mid-1970s to deliver cost-effective computing. An overview of the classic machines and the classes they formed is given in Fig. 1.

Fig. 2 shows several paths leading to the personal workstation. Large machines (with large memories) provide the major impetus. Large shared systems allow the price per terminal to fall in line with the salaries of the professionals using the station. Four lines, corrected for inflation, depict the LINC price, the cost of a supported professional (e.g., an engineer or scientist), the starting professional's salary, and a clerk's salary. The price of facilities such as a computer are "pegged" to these constants. The justification of equipment is usually based either on productivity or achieving some new level of performance or capability. While salaries of clerks are greater than \$10,000, the workstation industry believes \$10,000 to be a magic, highly elastic price barrier that will increase sales. In the later part of the 1980s, large personal computers will easily have enough memory and capability to take over all functions provided by nearly any of today's personal computers or personal workstations. Constant price, very powerful, large workstations will continue to be built, but the mass market will be served by the rapid price decline inherent in the IBM PC evolution.

A *personal workstation* is a relatively large (greater than 50 pounds) and expensive (\$10,000 to over \$100,000 in 1985) personal computer, with the appropriate transducers, used by a professional to carry out

FIGURE 1
Computer system prices at introduction and classes versus time.




generic (e.g., calculation, mail, and communication) and profession-related activities such as music composition, financial modeling, or computer-aided design of integrated circuits.

Personal workstations are necessarily distributed with the person and interconnected to one another forming a single, *shared* (work and files) but *distributed computing environment*—the *workstation environment*. A workstation's location is either with an individual on a dedicated basis or in an area shared by several members of a group. This choice is dictated by the cost and size of the workstation, relative to the cost and value of the work.

Personal workstations appeared *by name* about 1981 with the concurrent appearance of:

- ▣ microprocessors with at least one-half megabyte of physical memory virtual memory addressing
- ▣ memory chips of at least 64K bits providing primary memories

FIGURE 2
Workstation price versus time for various workstation introductions.



of ¼ to 1 megabyte, permitting the use of large programs and construction of high resolution bit-mapped terminals

- disks of greater than 10 megabytes
- local area networks for interconnecting the stations

Future descendants are likely to evolve from extended, lower-priced personal computers, interconnected via LANs. They would include appropriate transducers and corresponding professional applications programs.

A *profession* is any vocation, occupation or business, associated with work, including: secretarial, commercial, science, engineering, mathematics, and the arts.

A *personal computer* or pc is a self-contained computer with secondary file memory and appropriate transducers to interface with people. A personal computer is used interactively by one person at a time, at a location convenient to the user, and may "belong" either to the person or to a group. A personal computer, for a given use, is self-contained (i.e., requiring no external program or data preparation units) permitting a user to go through various stages without external intervention.

The microprocessor, memory, and mass-storage technology ap-

pearing in 1975 lead directly to the personal computer industry. Early computers utilized the simple, single process, stand-alone operating systems developed for both interactive, timeshared computers and stand-alone minicomputers. Nevertheless the first personal computer, the LINC (Clark and Molnar, 1964) was built in 1962, long before its predicted technological time.

Personal computers can be used either in a

- *personal or private environment* encouraging separation of files, resources, and work, with personal security, whereby the only communication with other computers is via secondary memory (floppy) sharing, or transmission of messages via standard communications lines, or
- *shared, workstation environment* encouraging communication among the pc's, sharing of resources (e.g., files, printers), and working on a single large work assignment or goal.

An *IBM Personal Computer* or PC is a particular personal computer utilizing the Intel 8086 . . . 80386 architectures and the evolving MS-DOS operating systems. A PC may be extended and used in a workstation environment when interconnected via LANs, permitting access and sharing of facilities and work. Hence constant-cost PCs are likely to evolve into the de facto personal workstations.

Bell Model of Memory Price Decline, Forming New Computer Classes

Hardware technology improvements, specifically in silicon and magnetic storage, have been the sole enabling determinants of progress in computing; they are the "technological devils" that drive the formation of our industries. This is the economic basis that forms all computing classes, including all forms of workstations. Since personal workstations evolve from other computer classes, the entire hierarchy of computers must be understood. In effect, we get no product before its (technological) time!

In 1975, I observed that memory price was falling at a rate of 20 percent per year. Since it was a constant fraction of system prices, they also declined at the same rate. This 20 percent decline was based on the learning curves for manufacturing core memory. Thus a computer would drop by a factor of 10 in price, per decade.

In 1972, with the introduction of semiconductor memories by Intel, price began to decline by a factor of four every three years, the cycle for new memory chip introductions. (The Four Phase company produced the first 1K MOS memory chips before 1970.) Table 1 shows the introduction of semiconductor memories and other devices and the

TABLE 1
Memory chip and microprocessor introductions
with resulting personal computer and workstations

Year	Memory	Micro	Width	Memory	Examples
1972	1K	4004	4		
1973		8008	8		
1974		8080/6800	8		
1975	4K	6502	8	4K-8K	Altair, IBM 5100
1976		Z-80	8		
1977					Apple II, TRS-80, PET
1978	16K	8086	16	16K-32K	
1979		68,000	16	16K-64K	Atari
1980		8088	8		
1981	64K			64K-256K	PC, Apollo, SUN 1
1982		186/286	16		Commodore 64
1983					Lisa
1984	256K	68020	32	512K-1M	MAC, AT
1985		386	32		SUN 3
1986					?
1987	1M		32	4M	?

development of specific personal computers and personal workstations. The reduction in the price per bit of semiconductor memories amounts to a 36 percent price decline each year or a factor of ten every five years, twice as fast as the original core memory (and discrete logic) based model.

In 1975, I also observed that memory size is the determinant of computer use and hence computer classes. That is, computer power is proportional to memory size, or memory size squared when a proportionally faster processor is used. Looking at the declining prices for various sized systems, one can see each computer class taking on the attributes and power of its higher order neighbor. Table 2 classifies what I then thought were the capabilities of various computers versus their memory sizes. I felt then, and still feel, that a one-megabyte address space for a process is the minimum size, and OK at least in terms of what we actually use. Anything smaller (e.g., 64K bytes) is sheer hell to use!

My goal in 1975 was to be able to use the past to project the evolution of simple computer classes to 1980 as the basis for the VAX evolution. Figure 3 shows the price of various systems versus their memory size. Note that a new computer class emerges when the price drops an order of magnitude. This evolutionary model was also described for DEC's hardware development (Bell, Mudge, and McNamara, 1978) shown in Fig. 4. This evolutionary model turned out to be quite accurate, predicting the Apple II (just 3 years away), low-cost shared

TABLE 2
System structure, memory size, and resultant use (G. Bell 1975)

Structure	Memory range	Function (use)	
Dedicated (fixed—1 use)	4KB-8KB	Interactive—e.g., POS cash register Real time—e.g., scope, traffic control, automobile	Special purpose, Fixed
Programmable (1 user)	16KB-65KB	Interactive—RT11 (CP/M) Real time—RSX11S, M	Small scale, generality
Dedicated (multiprogrammed n users)	65KB-256KB	Interactive—MUMPS, RSTS, Trans. Proc. Real time—RSX-11M, D	Special purpose
Programmable (multiprogrammed n users)	128KB-1024KB	Interactive—LAS, TOPS 10, RSTS Real time—RSX-11D	Generality

FIGURE 3
System price versus time for various memory sizes and system types (Bell 1975).

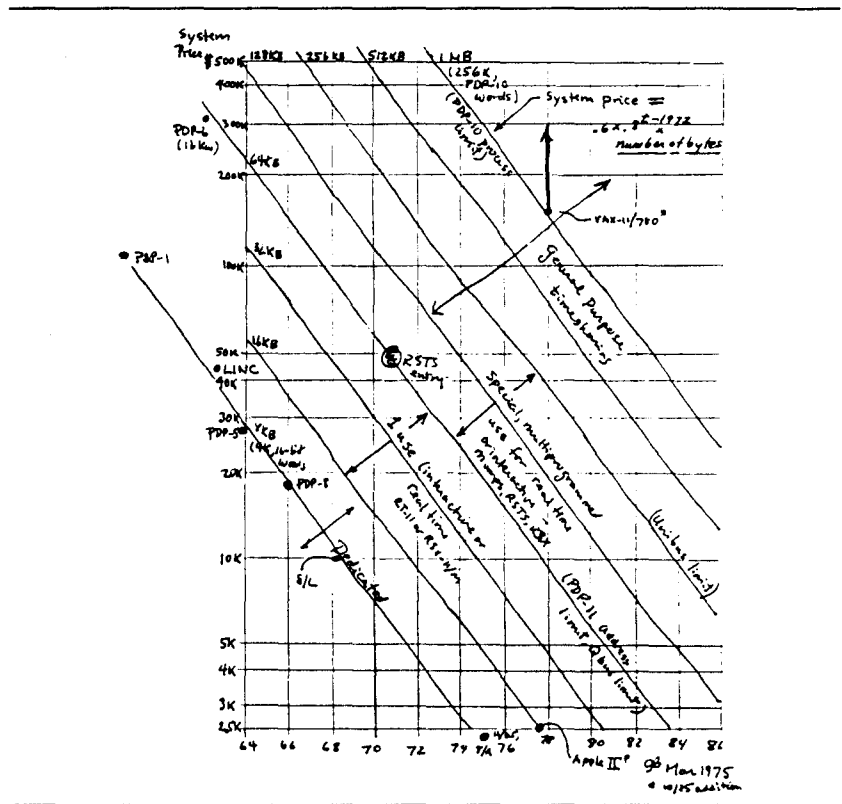



FIGURE 4
Price versus time for
each machine class.



micros, and the new class of mini-supercomputers that are beginning to emerge this year.

The chart labeled Fig. 3 has just the one variable, memory price. The first thing about prediction is to be able to predict the past. This chart allowed me to predict the past; I was able to predict things back to the PDP-1, the LINC, the PDP-5, and the PDP-8. I observed that the only thing that really matters is memory price and that memory size is a single variable. Allen Newell and I played this game in the *Computer Structure* book, that is, if you only had one number to talk about a computer, what would you say about it? And I think it's like any animal: How big is its brain size? This is simply the brain size of the computer, from which you determine all the computer classes. If you want to explain computers in one graph, this is the basic graph. It stopped working as well as I wanted it to in 1980.

The lines in the chart are the 20 percent cost reduction lines that I derived from fundamentals based on DEC markups and what it cost to string core memories. So we were able to show how one went through dedicated machines up through very general purpose machines. What that phenomenon does, what the parallel lines show is that, as you go

down in price for a given function, what tends to happen is that the same computer gets reinvented by different people. This then forms a new market for a new computer class, such as a workstation or a personal computer or a pocket calculator or a pocket computer or what have you. The other thing that happens is the people who built that computer then conspire with all their market friends, as engineers, because it's a different problem to solve—namely, how do you build a lower priced computer? By this time you've got a market established and you don't want to build a low-priced computer because you don't know the customers; you know the old customers because they are constant budget customers. The large computer center people always have money. Originally they used to have 3 to 4 million dollars to spend, now they've got about 15 million dollars to spend on a computer and they buy one every five years. You don't know what the computer does, but you know how much you're going to pay for it.

I have plotted what I think are the seminal machines here; these are the production machines or the machine of first use, where a number of people can use them for experimental work. You have to go back a number of years to identify various people that were involved in that experimental use. There is a line of computers that people used as personal computers, starting with the G15 and then the PDP-1; that line of \$100,000 seems to be always a reasonable amount, independent of the fact that \$100,000 then bought quite a lot more than it does now. It is a kind of a magic number of how much one would spend now to let it be used by an individual before you felt compelled to share it and put it in a batch mode and have to have an organization.

Since computer price determines the performance and therefore the economics and style of use, I began to observe at about the same time that there were roughly three levels of computer use: central (mainframe), departmental (minicomputer), and personal. Dave Nelson, formerly at DEC, Prime, and now at Apollo, extended this view to several other computer classes and introduced a model for price and weight (Table 3).

You buy computers by the pound. You say how much it weighs, you know what its price is, you know the technology at the time, and you know roughly what it's going to do or could do. If you look at the evolution of personal workstations, the ideas were really done on small computers in the 1950s. We saw the Whirlwind being used that way. In fact, a recent book on graphics said that EDSAC was the first personal computer. It was probably the first useful machine, and it was operating in 1949 at Cambridge University. It was a one-of-a-kind the university had built, but the claim was its CRT memory could display bits. It had bit-mapped graphics, and you could photograph the memory bits, and people did and got out functions; it was often used as a personal computer. A lot of the experiments were done there in the 1960s.

TABLE 3
Nelson⁺/Bell computer classes

Tier	Location	Names	Price (\$)	Weight (#s)
0	wallet	calculator	10	.05
1	pocket	calculator, special function unit	100	.5
2	briefcase	kneetop, small personal	1,000	5
3*	office	personal, large personal, "workstation"	10,000	50
4*	department	shared micro, minicomputer, super-minicomputer, mini-supercomputer	100,000	500
5*	center	mainframe	1,000,000	5,000
6*	region	supercomputer	10,000,000	50,000**

* Founder and Vice-President, Research, Apollo Computer Corporation

* Based on Bell's Distributed Processing Computing Model (circa 1975)

**Perhaps, if heat exchanger is included


People used large computers as though they were personal computers. Out of the work on the TX-2, came a real personal computer, the LINC. When people ask me what was the first personal computer, I have to say, it's LINC, even though individuals used LGP30s and G15s.

These emerging smaller classes put significant pressure on interconnecting all the computers and forced the need for local area networks (LANs). Figure 5 shows the style of use and environment for various computer classes. Note that personal workstations provide roughly the same capability as larger shared computers and are large personal computers interconnected via LANs. IBM appears to have recently adopted this hierarchical view of computing and assigning meaningful work to all the levels, so it must be right (at least for a few years).

Networks and the Shared and Personal Workstations Evolution

Although technology is a major variable that creates computer classes, the development of networks and changing patterns of use also influenced the evolution of personal workstations. The impact of the network and local area network on workstations is based on my perspective as chief architect of the ensemble of DEC products until 1983.

FIGURE 5
Computer use style
versus user
environment for
various computer
classes.



DECNET AND THE INFLUENCE OF ARPANET

The various DECNET architectures were initially created in 1974 by Stu Wecker while a member of DEC's research group. They were based on the notion of hierarchical computing according to use, concurrent with Digital's work on distributed processing. The imperative was to interconnect the hierarchy of DEC's minis and mainframes as tightly as possible, forming a single computing environment.

The need for DECNET occurred because Digital was proliferating timeshared and real-time minicomputers that had to have a way to communicate with one another and with their own and other central mainframes. DECNET was introduced in the following phases:

- I 1975 task-task, file transfer, file access; point-to-point; 8-bit net address; operating systems: RSX-11 D, M, IAS
- II 1977 improved system generation; used for first DEC engineering network (now several thousand machines) at tens of sites; operating systems: rt-11, rst
- III 1980 routing, virtual terminal, full network management; operating systems: DECsystem 20, vms

- IV 1983 large routing address, full remote rms files; c-terms; ethernet support
- 1985 portable DECNET (by Wecker's company)

The basic model for DECNET was ARPANET (Roberts, 1970), with the important difference that a network need not have specific IMPs or TIPs to do the packet switching for the worker computers. That is, the worker computers could do the packet switching for the network, and when the packet switching load became too high, one inserted packet-switching computers.

XEROX PARC ETHERNET/ALTO ENVIRONMENT, CARNEGIE-MELLON SPICE PROPOSAL

In 1979, Carnegie-Mellon published a proposal soliciting vendors for a personal workstation environment. Their statement, which had the effect of stimulating the design of various workstations was:

The era of time-sharing is ended. Time-sharing evolved as a way to provide users with the power of a large interactive computer system at a time when such systems were too expensive to dedicate to a single individual . . . Recent advances in hardware open up new possibilities . . . high resolution color graphics, 1 mip, 16 K word micro-programmed memory, 1 megabyte primary memory, 100 megabyte secondary memory, special transducers, . . . We would expect that by the mid-1980's such systems could be priced around \$10,000.

Although the concept was correct, the details of the machine were wrong because evolving technology was not considered. This led to selecting the wrong vendor, the late Three Rivers or PERQ Systems. Three Rivers's founder, Brian Rosen, a CMU grad who had worked at Xerox Palo Alto Research Center (PARC), designed PERQ as his version of the Xerox ALTO and D-series machines. Carnegie-Mellon may have taken on a leadership role in personal workstation environments by building its SPICE (Scientific Personal Integrated Computing Environment) on the PERQ and then contracting to build a similar environment for IBM, provided the SPICE derivative product is marketed.

The prototype for SPICE, and virtually all other distributed personal workstations environments (e.g., MIT's Chaosnet and the personal LISP workstations that were the basis of LISP Machines Incorporated and Symbolics), was PARC's first 3-megabit-per-second Ethernet interconnecting Altos, operational as an environment in 1975 (Perry and Wallich, 1985; Pake, 1985). Datapoint's ARCnet with personal computers using the Intel 8080 was clearly the first commercial workstation environment. The Alto processor was a near derivative to the Data General Nova with an integrated bit-mapped display. It is

important to note that the PARC environment also included a time-shared computer based on the PDP-10 architecture, MAXC, which could be used for shared files and significant computation. In April 1981, Xerox introduced a product version of Alto, the STAR workstation, and various file and print servers utilizing a 10 megabit per second Ethernet, but without a real central computer. The PARC Ethernet was the forerunner of today's 10 megabit per second IEEE 802.3 standard (as originally drafted by Digital, Intel, and Xerox).

THE DIGITAL VAX, HOMOGENEOUS COMPUTING ENVIRONMENT

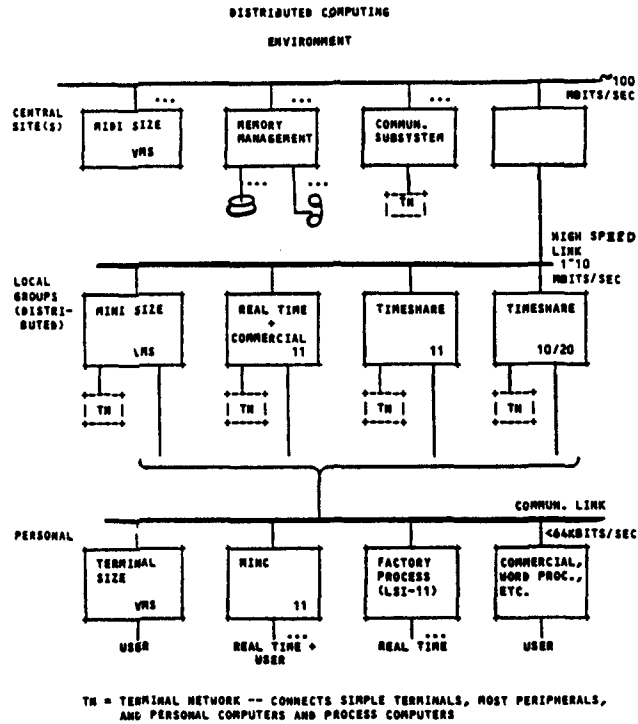
The idea of a complete, distributed computing environment, based primarily on the VAX architecture (but including the PDP-11 and DECsystem 10/20 computers) was presented to DEC's Board of Directors in December 1978 and approved as the strategy for guiding future product development.

The goals and constraints of this strategy is outlined in Appendix 1. The impetus was based on the early success of VAX, including plans to make a wide range of products to fill each of the three classes from MOS micros for personal computers to clusters of high-performance ECL VAXen. Furthermore, it was desirable to limit future evolution of the 10/20 and 11 lines, since both were providing essentially the same capability. The need for interconnection and the availability of LANs were critical as described above. Finally, computing was evolving from a centralized computing style to a more distributed computing environment as demonstrated by the Alto environment (but including high-performance and departmental level computers).

The environment, shown in Fig. 6, provided a wide range and hierarchy of compatible computers and ways of interconnecting them to provide users with generic (e.g., word processing, mail) program development, and profession-specific computing using timeshared minicomputers or personal workstations. The goal was to provide the widest range of choices by having complete compatibility for where and how computing was to be performed without having to make a priori commitments either statically (purchase or installation) or even run-time to a particular computer system class (i.e., mainframe, minicomputer, team, or personal workstation). By 1981, the original structure evolved to require the local area network (Ethernet) for connecting all computers, initially called NI for Network Interconnect, as part of the hierarchy of interconnects, as shown in Fig. 7 (CI = Computer Interconnect, NI = Network Interconnect, BI = Backplane Interconnect, and II = Interchip Interconnect; II did not materialize).

The design of the environment is substantially more than the design of a single range of compatible computers because different styles of user are required depending on the machine class, and all the com-

FIGURE 6
Digital distributed
computing
environment (Bell,
12-78).




puters must be interconnected and work together in a multi-level hierarchy.

ETHERNET

While DEC had prototype LANs and proposals for LANs internally, I felt it was essential to have a standard. We hired Bob Metcalfe as a consultant, and his missionary role was to make Ethernet the standard, using Intel's chip design capacity and Xerox's patent and LAN experience. The IEEE got involved once they heard the idea of the LAN. They went on to facilitate the design of many other "standards," which, of course, had the effect of diminishing the notion of a standard. Fortunately the development of Ethernet crept into existence slightly more rapidly than the IBM Token Ring LAN, which was introduced in October 1985, otherwise I doubt if there would be any significant use of LANs today. I understand about 30,000 Ethernets exist today.

In January 1981, when Ethernet was being introduced by various vendors and being completed as an international standard, several of

FIGURE 7
Digital homogeneous
computing
environment 1984.



us from the developer companies (myself from Digital, Noyce from Intel, and Liddle from Xerox) made a broad appeal to the U.S. and European press and manufacturers on the importance of the standard. My theme was: "Ethernet is the Unibus of the Fifth Generation." We all argued that the standard was essential

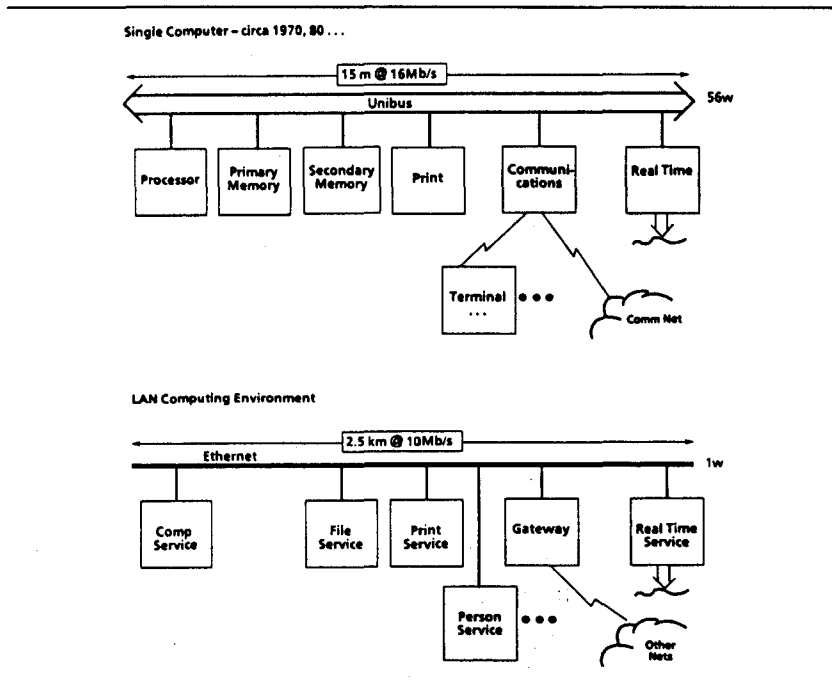
- to interconnect department-level minicomputers and mainframe computers to one another, permitting file-to-file transfers, remote-terminal emulation, and distributed processing among the shared computers;
- as a switch for interconnecting terminals to the higher level computers within the environment;
- to centralize the gateway function in one place, versus requiring every computer to have protocol conversions to communicate outside the Ethernet environment;
- as a bus for building computer systems; and
- for building fully distributed computing environments. The

PARC model of a computing system was to decompose computing into a series of functional servers such as printing, filing, and with individual workstations. The analogy with Unibus (circa 1970) as a computer versus Ethernet (circa 1981) as a computing environment is clear when comparing the two structures (see Fig. 8).

THE ENCORE CONTINUUM

With the formation of Encore Computer Corporation in 1983, our goal was to provide a complete environment, the "Continuum" (Fig. 9), like the VAX environment, but based on a single microprocessor (Bell et al., 1985). The Continuum provided both the distributed personal workstation style (described in the evolution, Appendix 1) and shared central computing using a multiple, microprocessor architecture that I have named the Multi (Bell, 1985). Ethernet provides the switch for interconnecting terminals to the Multimax using concentrators that preprocess terminal requests for communicating with other environments and for distributed-workstation environments. There are no terminal connections to the shared computers (something I have been trying to eliminate for years).

FIGURE 8
Unibus (c. 1970) and Ethernet (c. 1982) computer structures comparison.






FIGURE 9
The Encore computing continuum.

APOLLO DOMAIN

Dave Nelson, founder of Apollo and Domain architect (May 1980), was in the research group at DEC during the formation of DECNET and VAX. In the late 1970s, he went to Prime to work on the Prime token local area network and then left to develop the Apollo Domain architecture. Domain, introduced in March 1981, most typifies today's personal workstation because it provides a single addressing scheme for accessing all the memory (files) of a system. Dave credits me with stimulating this unifying architecture, which is like the Unibus in terms of addressing.

Workstation Events, Dimensions, and Timeline Evolution

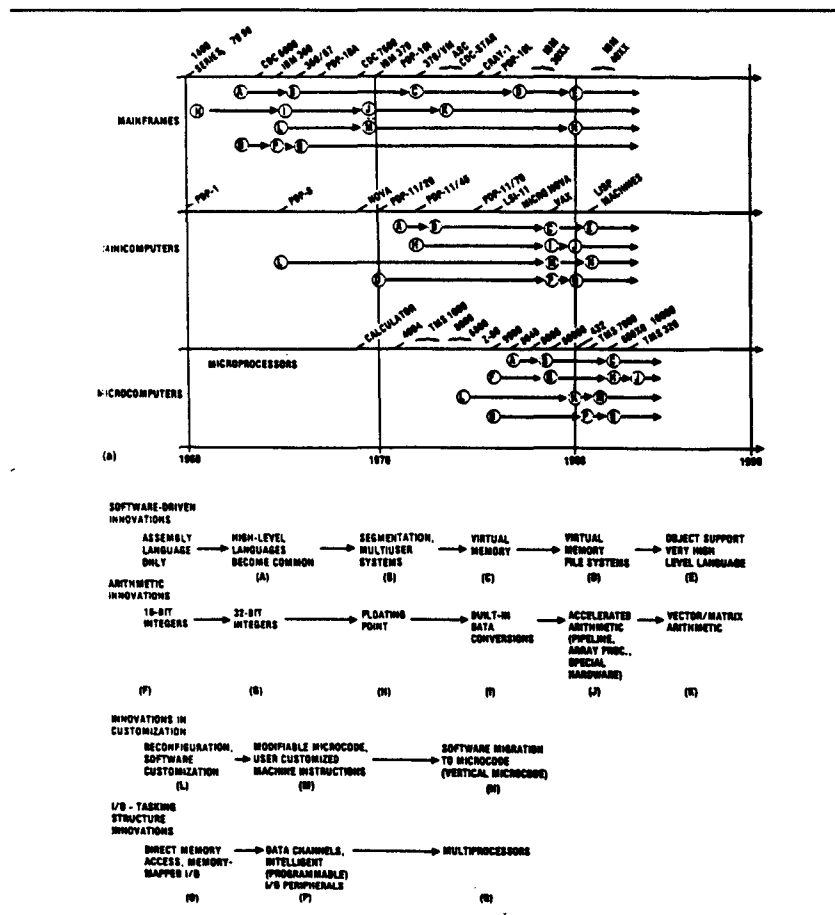
This section serves as a model for future historical research because it outlines, albeit incompletely, some critical events in workstation evolution.

The evolution of the physical workstation hardware is almost indistinguishable from the computer, and hence the evolution of personal workstations parallels, and is directly coupled to, the development of mainframes, minicomputers, and microprocessors. The principal difference is that workstations include a significant set of transducers to interface with humans and communication links and additional user interface (software) advances.

PROCESSOR AND MEMORY COMPUTING ENVIRONMENT

The reinvention of the wheel and features within each of the computer classes has been nicely depicted by Burger, Caving, Holton, and Sumney (1984) as a timeline of the critical processor features (see Fig. 10). The single most critical dimension of processing power is virtual mem-

FIGURE 10
Computer architectural innovations for mainframes, minicomputers, and microcomputers from 1960 to 1990 (a) and the descriptive legend (b). (Taken with permission from Burger, Cavin, Holton, and Sumney, 1984.)



ory size. The reinvention of the wheel of virtual memory size is shown in Fig. 11 from Siewiorek, Bell, and Newell (1982) as each architecture reinvents and evolves the concepts of virtual memory.

VIDEO ENVIRONMENT

The evolution of displays has occurred along a set of dimensions that are relatively closely correlated with evolutionary time as follows:

- performance as measured in ability to display objects, including the ability to transform the objects;
- hardware generation of the picture process (Myers and Sutherland's Wheel of Reincarnation shows the evolution of how pictures are controlled [Fig. 12 reproduced from Bell, Mudge, and McNamara, 1978.] Note the evolution: direct control of the dis-


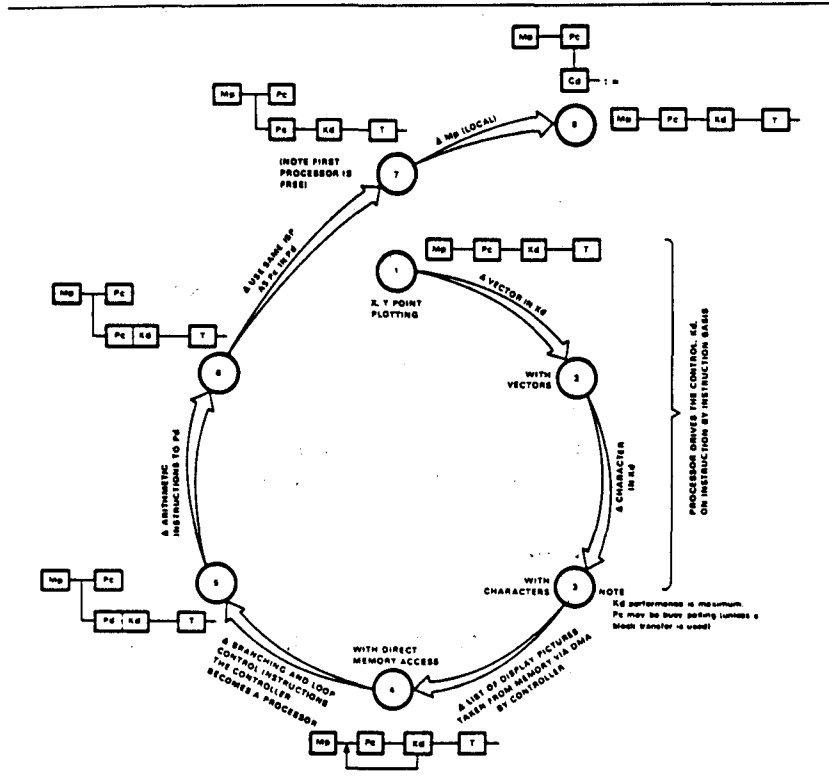


FIGURE 11

Virtual address space evolution for various computers and computer classes. (From Siewiorek, Bell, and Newell, 1982.)

FIGURE 12
Myer and Sutherland
wheel of
reincarnation for
display processor
evolution. (From
Bell, Mudge, and
McNamara, 1978.)



play object by the central processor; display list issued to a DMA controller or external controller; separate processor with loops, subroutines, and the like; arithmetic instructions added to aid dynamic displaying; use of virtual memory [beginning to occur to manage large pictures]; and removing the display processor to form a separate computer.);

- resolution (in pixels or analog vector resolution);
- number of bits per pixel including use for monochrome/color;
- display datatypes
 - points, characters, lines, curves
 - 3D line drawing
 - wireframe meshes, 2D surfaces, 3D surfaces
 - shading, sculptured surfaces, constructive geometry
- translation of datatypes
 - fixed
 - translation, scaling, rotation
 - 2D clipping, 3D perspective, 3D clipping

Keynote Address: Toward a History of (Personal) Workstations

- standard interface evolution: direct control, SIGGRAPH CORE Library, VDI, CGI, NAPLPS, and the like;
- windowing of output images; and
- character, line, picture, and 3D image input.

TACTILE AND MANUAL POSITION OR LINE INPUT

This category includes keyboard, 2D and 3D light pens, direct touch screen, tablet, digitizing tablet, knob, joystick, track ball, and mouse.

PAPER OUTPUT

This section must include the evolution of conventional printing and plotting along with early film recorders to capture images, extending to modern dot matrix printers of all kinds for monochrome and color.

VOICE I/O

While today's personal workstations do not permit substantial voice communication, voice input such as the Kurzweil voicewriter promises to add a significant dimension for human communication. The critical voice "firsts" for use with personal workstations are relatively sparse at this time.

LINKS, SWITCHES, AND NETWORKS

Links, switches, and communications networks have influenced the formation of computer networks and computing environments. The previous section described their influence on the formation of the DEC computing environment. Others should be traced.

SPECIAL PURPOSE I/O AND ROBOTS

The ability to handle arbitrary signals from other environments is crucial to the formation of various profession-specific environments (e.g., music and scientific experiments).

OPERATING SYSTEM, DATABASE, LANGUAGE, AND GENERIC ENVIRONMENTS

Traditional computer science history has laid claim on operating systems, databases, and languages, including their application to personal workstations. For example, the History of Programming Languages conference constructed an excellent timeline of the various languages,

including those for specialized environments such as interactive computing and specific application domains such as statistics, science, and engineering.

Generic environments that users of traditional, timesharing computers have required include:

- Timesharing oriented with text and graphics terminals (slow links)
 - text editing, mail, and network mail
 - drawing and graph plotting
 - shared work processing
 - database access
 - transaction processing
- Shared workstation and personal workstation oriented (fast links)
 - multiprocess control including control and windows
 - spreadsheets
 - drawing and painting

PROFESSION SPECIFIC WORKSTATIONS

In addition to generic applications programs (e.g., word processing and mail), each profession requires an environment to support and "understand" the work of the professional. This may take the form of special hardware (e.g., high-resolution mega-color scopes for the graphic artist, high precision analog output for the musician, or fast simulation processors for the VLSI designer) in addition to specific programs that understand and have knowledge or expertise of the work being carried out. The following taxonomy provides an outline of the environments that have developed quite independently and that historians should research.

- Scientific, medical, and mathematics workstations
 - life science with transducers appropriate to the measurement task
 - medical (e.g., radiologist, cardiologist, bone surgeon)
 - geologist and petroleum engineer
 - physics experiment control and analysis
 - statistician and mathematician
 - signal analysis and encryption/decryption
- Engineering workstations
 - electrical or electronic schematic (2D line drawings)
 - VLSI physical design (2-½ D)
 - mechanical drafting
 - design of sheet metals or piping (2-½ D)

- architectural and structural engineering
- general mechanical design of arbitrary surfaces for analysis and fabrication
- Software and knowledge-engineering workstations
- Arts (typesetting, graphics arts, and music) workstations
 - publisher of newspaper, magazine, or book
 - scholars including historians, linguists, and literary scholars
 - graphic artist, illustrator, cartoonist, or choreographer
 - composer
 - artistic film production
- Commercial workstations
 - general clerical and transaction processing, office word processing, executive
 - financial personal workstations for banking, insurance broker
- Special control-oriented workstations
 - air traffic control
 - process and plant control
- Training workstations (and are these workstations?)
 - simulation of industrial plants and aircraft
 - computer aided instruction such as Plato

Some Displays, Personal Computers, and Workstations I Have Known

This section describes various displays connected to computers, which were used as either personal computers or workstations. The MIT and Lincoln Laboratory displays embedded in Whirlwind, TX-O, TX-2, and LINC were among the earliest personal computers. Since these computers (except LINC) were one of a kind, their impact was as prototypes for subsequent developments. One direct impact was in the formation of the DEC computers, many of which were used as personal computers.

The first DEC display, Type 30 for the PDP-1 introduced in 1961, had instructions that permitted direct plotting of points and, optionally, characters and vectors. As DEC evolved and built timesharing computers with operating systems that protected the user from accessing input/output equipment directly, the ability to have highly interactive terminals and workstations decreased.

Nearly all of the DEC computers were used as components for building interactive, "workstation-style" systems. DEC and Original Equipment Suppliers (OEMs) built and sold workstation-style systems for use in the sciences, printed circuit layout, mechanical and architectural drawing, typesetting, and office automation. These pioneering

applications (circa 1970), operating on dedicated timesharing systems, served as the basis for much of today's evolving workstations.

**MIT WHIRLWIND, LINCOLN LABORATORY TX-0, TX-2,
AND IBM AN/FSQ7 (SAGE)**

During Whirlwind's first year of operation, 1950, Bob Everett published a paper that showed a picture of its display (with an attached camera) and listed some actual problems carried out on the machine.

1. An industrial production problem for the Harvard Economic's School
2. Magnetic flux study for Whirlwind's magnetic storage work
3. Oil reservoir depletion studies
4. Ultra-high frequency television channel allocation for Dumont
5. Optical constants of thin metal films
6. Computation of autocorrelation coefficients
7. Tape generation for a digitally controlled milling machine

Whirlwind was initially established as a large specialized simulator to the Airplane Stability and Control Analyzer (ASCA). The actual use turned out to be a prototype for the semi automatic ground environment (SAGE) system for air defense and SAGE, in turn, was the basis for air traffic control. The first successful use of Whirlwind for air defense occurred on April 20, 1951.

MIT's Lincoln Laboratory, established in 1951, moved into its quarters in Bedford, Massachusetts, in the summer of 1953. Lincoln built the TX-0 computer to test the use of transistor circuitry and a large core memory and then the TX-2 for large-scale computing experiments. Both had a point-plotting display with 10" x 10" area and 1K point resolution, light pen, camera, switch input, and abilities to interconnect arbitrary I/O devices. Hardware innovations of the TX-2 include addressable magnetic tape for a filing system; the Lincoln Writer, a typewriter for engineering/scientific use; and the TX-2 multiple sequence operation for rapid context switching. While the TX-2 was initially a personal computer, it operated under control of a timesharing operating system by the mid-1960s. The TX computer circuitry was virtually identical to the logic and laboratory modules that DEC sold in its first four years, prior to the introduction of the PDP-1 in 1961.

These machines, used as personal computers, pioneered numerous applications, the most famous being Sutherland's Sketchpad.

I spent a brief but wonderful period (January 1959 to June 1960) using TX-0 for speech research and writing the program to do

“Analysis-by-Synthesis” (Bell, et al. 1961), which took as input speech spectra directly from a filter bank connected to TX-0. I also spent six months designing an IBM-compatible tape TX-0 and exploring hybrid computation (an analog computer used with TX-0). Peter Deutsch, then age 12, acted as a relatively patient user consultant and helped me debug my macros. Gwen Bell entered data on land use in the Boston area, and I wrote a program to display and explore this data. The program was demonstrated to the city-planning faculties at MIT and Harvard. Ten years later, Harvard was able to turn the clock back over 30 years when they established their computer graphics laboratory to do a simpler version of computer mapping, using punched cards and line printers.

DEC PDP-1 DISPLAYS: TYPE 30, 31, AND THE FIRST COLOR DISPLAY

The PDP-1 continued in the tradition of the “MIT personal computers.” The PDP-1 had three displays. The Type 30 point plotting display took 50 microseconds to display a point on the 10" × 10" tube with 1K resolution. This was one of the first displays connected to a commercial computer, converting it into a personal computer for various scientific data analysis applications. Type 30 was extended to include a character generator controlled by loading a 36-bit, 5 × 7 bit raster specifying the character. The circuitry, designed by Ben Gurley, formed the basis of the 3XX displays. The DEC Special Systems Group designed an elaborate display, PEPR, for bubble chamber photograph analysis for MIT using the basic display circuits.

Type 31 designed by Gurley with a 5" high precision display with 4K point resolution. Only one was delivered to Lawrence Livermore Laboratory (LLL, now LLNL—the “N” for National) for use on their PDP-1. The Livermore PDP-1 had every option that DEC had proposed to build including Remington Rand tapes, and was used as a front end and exploratory PC for the large computers including Stretch and LARC. Ben Gurley left DEC and with Ed Fredkin established Information International Inc. (III). The first III product included a high-resolution display for various image input/output applications. The basic high resolution display is still in existence.

One PDP-1 Color Display was delivered to the Air Force Cambridge Research Laboratory (the military counterpart to Lincoln Laboratory). The color display was built by modifying a standard RCA color television for point plotting.

At the Fall Joint Computer Conference in 1961, Ed Jacks, from General Motors, spent almost the entire conference watching people use the PDP-1/Type 30 to draw and doodle using a program I had written. The display influenced Ed to tackle Computer Aided Design, later using the IBM 7090 with a connected display. The classic program

SPACEWAR! was written on the PDP-1 in 1962, by Steve Russell. Today SPACEWAR! continues to be the center to settle litigations about the design of computer games.

LABORATORY PERSONAL COMPUTERS: LINCOLN LABORATORY LINC, DEC PDP-12, AND MINC

Wes Clark and Charles Molnar continued the MIT tradition of building personal computers, and LINC was designed for wider scale use for the life sciences, by being lower cost (about \$40K in 1962). The LINC was designed using mostly DEC modules. The original LINC had two 5" displays, a keyboard, 1 or 2 Kw, 12-bit memory, and two LINC-tapes. Each unit had approximately 256 Kbytes of addressable tape memory. I believe LINC was the first production personal computer for scientific use (about 50 were produced, 21 by DEC). Its significance was its completeness at the low cost, yet being open-ended.

DEC went on to build 140 LINC-8s beginning in 1967 (designed by Wes Clark and Dick Clayton) and around 1000 PDP-12s beginning in 1970 (Bell, Mudge, and McNamara, 1978).

DEC 338, 339, AND 340 DISPLAYS

The 338/PDP-8 display/personal computer, introduced in 1967, was used as a front-end computing terminal. The 338 and 339 (for the 18-bit computers) were, I believe, the first display processors (Bell, Mudge, and McNamara, 1978). They took a display program in the PDP-8's primary memory, which specified points, character strings, vectors, subroutines and other program-control instructions and permitted completely autonomous operation of the display, independent of the central processor. The 339 was used with the 18-bit computers, especially PDP-9. Both systems were used for stand-alone data analysis and for display terminal front ends to large computers.

A special version of the 340 was designed on a Saturday morning with Ivan Sutherland for his use within the government just prior to joining ARPA. This PDP-7/340 had every conceivable keyboard and input device we could think of, including switches, trackballs, joysticks, and lighted buttons.

DEC PDP-15/VT15, AND GT4X "WORKSTATIONS"

The PDP-15/VT15, introduced in 1970, costing about \$90K was used as a standalone for one or two users (to reduce the cost to \$50,000 per terminal using a \$15,000 graphics terminal). About 30 percent of the systems had two terminals. It took about 100 microseconds to draw a vector using the DDA technique. The PDP-15 had from 8 Kw to 32 Kw

of memory. Computer Aided Design system for Printed Circuit layout written and marketed by REDAC (U.K.) was a key application, although it was used extensively for scientific-research applications.

In January 1973, the GT40 was introduced at a price of \$25,000 with 16 Kw of memory based on the lower cost PDP-11/05. The well-known program, Lunar Lander, was written by Jack Burness for the product introduction. The relatively slow speed of the display, coupled with a limited memory and configuration of the 11/05 limited the application to a preprocessor for time-shared systems doing CAD.

The GT44 was introduced in June 1973 at a price of \$38,000, using the same display processor as the GT40, but with an 11/34 computer and RK05, 5 Megabyte disk. The GT44 was used as a complete, stand-alone computer.

The GT48, introduced in October 1975, selling for \$55,000, was a high performance personal computer, using the 11/34, with memory up to 128 Kw. Vectors were generated using analog circuitry at a rate of 100,000 vectors per second.

THE PDP-8 WORD PROCESSING PERSONAL COMPUTERS

In addition to the early use of the PDP-8, with the 338 and various storage tube displays, the PDP-8 began to be used extensively 10 years after its birth for the DECmate series of word processing terminals. The PDP-8 was based on the PDP-5 (1963) 12-bit architecture, and influenced by the CDC 160 and LINC, addressed up to 32 Kw (later expanded to 128 Kw). The following word processing personal computers used the PDP-8:

<i>Year</i>	<i>Model</i>	<i>Price</i>	<i>Comments</i>
1975	DS310	\$16,000	8/A, desk, 8" floppy and VT52 terminal
1976	WS200	\$40,000	8/A, cabinet, 1-8 users, 5 Mbyte disks
1977	VT78	\$15,000	Intersil 6100 PDP-8 chip, embedded in VT52, 8" floppies
1981	DECmate I	\$12,000	6120 chip, embedded in VT100, 8" floppies
1983	DECmate II	\$3,500	integrated unit, 5" floppies, bit mapped
1984	DECmate III	\$2,400	Z80/8086 and 5" winchester disk options

By 1984, the twenty-first birthday of the basic architecture, over 100,000 PDP-8 connected terminals or personal computers were used

for word processing. More PDP-8s were built in 1984 than in any other year. In a very large fraction of the use, the personal computers connected with larger departmental-level VAX computers for electronic mail, filing, and auxiliary word processing for a true, shared workstation environment.

LSI-11 BASED PERSONAL COMPUTERS: TERAK, PDT 11/150, AND VT103

The LSI-11 (11/03), introduced in 1976, with a smaller printed circuit board form factor (approximately 10" square) and using the Qbus proved to be an important building block for Digital's subsequent personal computers.

Terak was founded in 1975 to exploit the notion of a bit-mapped graphics computer and used the LSI-11 (11/03) and subsequently the 11/23. The 8510, first delivered in 1976 to Ken Bowles at the University of California, San Diego, for the development of UCSD Pascal, was developed independent of the Xerox Alto. The system operated with up to 28 Kw of memory, used 256 Kbyte floppies, and the bit-map display's resolution was 320 × 240. About 3000 were built and sold at a price of \$5000 to the education market.

DEC introduced the PDT 150 as a small computer with two 8" floppies and 32 Kw at a price of less than \$10,000. While all the DEC marketing groups debated the price and their role in distribution, the personal computer industry formed. DEC was able to discourage Dan Bricklin from using PDT to build Visicalc. (He went on to use the newly introduced Apple II.) The VT103, introduced in 1981, was based on the VT 100, which had been introduced to accept the Qbus, LSI-11/23 modules.

DEC TYPESETTING TERMINALS: VT20 AND VT71

The PDP-11, introduced in 1974 and 1977, was used as a base for the VT20 and VT71 typesetting terminals. These had a programmable character set of 256 characters, used 15" portrait-oriented tubes, and communicated with the host via 9600 baud asynchronous communication. The VT20 was based on the 11/05 and could handle two CRTs, while the VT71 used the 11/03 and sold for \$8000. These terminals connected to host PDP-11s that carried out the main typesetting editing functions.

DEC VAXSTATIONS: 100, I, II, AND 500 SERIES

The introduction of high resolution, bit-mapped terminals connected to VAX occurred quite late in VAX's life. A number of attempts were

made, including SUVAX (for single user VAX) providing both monochrome and color, operating in 1980.

The VAXstation 100, introduced in May 1983, was a high resolution 19" terminal with bit-mapped graphics that connected to any of the VAX computers by a high speed fiber optic link connected to the VAX's I/O bus (the Unibus).

VAXstation I and II, based on the MicroVAX I and II, were introduced in October 1984 and June 1985, respectively. Both units sold for approximately \$20,000, with 19" bit-mapped CRT, using the Qbus for interconnecting I/O equipment.

The VAXstation 500 family was introduced in October 1985 with 1280 × 1024 pixel resolution color permitting 2D and 3D wire frame and 3D shaded solid images.

Conclusions

If you've ever used paper tape, you know you had to do off-line preparation. Have you ever used cards? (I hope some of you are old enough to know people who used cards.) I used cards one year when I was a Fulbright scholar in 1958. I went to Carnegie and found key punches, and that's why I wrote a book—I didn't want to compute. The 360/67 had just arrived, but the people were using those machines as breadboards for timesharing. And then, in the 1960s, a lot of these ideas created a whole set of industries using minicomputers and putting one or two high-performance graphics terminals on them, such as the Adage and the Raster Graphics. There are about a hundred companies doing that now, some of which formed in the early 1960s. The goal was to get the cost per workstation seat to on the order of \$50,000. You could do that with a \$50,000 computer and a couple of \$25,000 graphics terminals. The other thing that happened in the early 1970s was that we had the prototypes of the distributed workstations forming at PARC. Datapoint built a distributed workstations environment using an 8080, or an 8008, on a local area network, which may have preceded most everything else. In 1981, with the devil providing technology, namely the right memory sizes and powerful microprocessor, workstations were formed, including those from Apollo, Sun, and Xerox. And then, by 1983, everybody realized there were too many of these companies and there were one hundred that built workstations. So the word "JAWS" was coined for "just another workstation company."

By 1985, large companies like IBM and DEC began to embrace the concept. In a few years, a bunch of these companies will begin to shake out and you'll get a steady state supply. Alan Perlis said "we've got to hold this conference because workstations are going to disappear," and I think, to a large extent, they are going to, a large number of them, at least the evolutionary form of what we think of as worksta-

tions will disappear because by then the personal computer, as I defined it, will be the workstation.

Figure 1 shows the memory size of a workstation of given capability. These were the lines I drew in 1975 that I said would predict the past, but I found out they were starting to fail. These lines are all 20 percent per year decline lines. That's every ten years you have a decade drop in price. So, if you look at those computer class models, they identify what was a computer in one class ten years ago. Now, can we build an order of magnitude cheaper and form a new machine class? In 1980 we saw a factor of ten decline in five years, and this is 36 percent per year. And that's, of course, because the memory density increases and you build a factor of four bigger chips every three years.

The other phenomenon is how much do you pay for a workstation? That's a function of who uses them. For a few people, a Cray supercomputer is a reasonable personal workstation, if you get that kind of use. In fact, some Crays are used virtually in a personal workstation form. It's a little bit expensive, and off these scales. The thing that I've observed over time is never to judge exactly how much people are willing to spend as single users for a computer; it's what benefit you get that is the important thing. So, starting with the LINC, one could argue, that it was a little bit expensive at below \$50,000. But, in fact, if you take the constant price line, LINC today would be about three times the price of that, or \$150,000 today. And we do find a personal workstation in that range, the Pixar for example. Other lines show the cost of various professionals per year, and you ask: Would you provide the same amount of computing for that person as you pay for the individual?

So this tries to show what Alan was worried about, that, in fact, you have a machine here today and suddenly that machine is going to drop by a factor of 10 in five years. So workstations will disappear, and the PCs will cause their demise unless software exploits larger machines. To avoid annihilation, the trick is to increase the power, at high price levels, and make it possible for the individual to get more done. Then you run into the problem that since this is the synergy between a system of a human and a computer, if you can't get enough or the right software there, you may just be person-limited rather than computer-limited.

I believe we need a definition of what a personal workstation is now, but I think the definition will change over time. First, there is a professional personal computer. What's a profession? Any vocation, occupation, or business that is associated with work. So playing doesn't count here. Maybe searching for the workstation doesn't count either. That rules out the PARC work, for example. Let me not do that. I just want to raise questions. Second, what is a personal computer? It is a self-contained computer with its own file memory and appropriate

transducers, used by one person at a time, at a convenient location. It's self-contained. And then what people say is the personal computer (that's big PC not little pc) is an IBM compatible PC, so this was the definition. Today we think of it as first a personal computer, but it's in a distributed interconnected environment, because work had the nature of being interconnected versus being in a private environment. It's large, about 50 pounds; it's more expensive than a PC; and it has a bunch of generic applications. Also it understands the profession; it has all of the tools that the professional uses and needs.

The reason I take this holistic view is that I look at the interactive systems we built from 1965 to 1985, either as single-user or shared systems, as similar. The important thing is what systems do, not how they are built. For some of the graphics, there's no way you can do some of those functions unless you put them in a single personal computer.

ACKNOWLEDGMENTS

Dave Fanger, Encore Computer, did the figures on a Xerox Star. Designers of DEC products supplied product information: John Clarke and John Kirk (PDP-8), Herve Lavoie (VT15 and GT4x), Jim Milton (typesetting), and Stu Weckar (DECnet). Dave Nelson, Apollo, contributed a table and much to my understanding of workstations. Gwen Bell helped clarify my thoughts well enough to get a first draft, which she then edited. Mary Jane Forbes prepared the final manuscript.

APPENDIX 1. The Digital VAX, Homogeneous Computing Environment—Technological and Market Background

[Author's Note: This Appendix is solely the view of the author and is not necessarily the view of Digital Equipment Corporation or its engineers.]

Batch Mainframes for Central Services

In the first two computer generations, 1950-1970, computers were used in batch processing under the name of mainframe computing. During the 1970s the mainframe began to be used almost interactively from remote job entry terminals as "glass key punches." The general direction is to have larger mainframes and larger terminal networks that interconnect to a single computer by an array of front-end computers. When more power is required, more switching computers are connected to several mainframes, each of which performs a particular function. Attached dual processors are used to provide increased power for what is fundamentally a single system. Over time the evolution will be to small-scale multiprocessors for incremental performance and higher availability.

Minicomputers and Timesharing for a Group

In the mid-1960s, both minicomputers and timesharing were developed at Digital around the PDP-8 and PDP-10 computers, respectively, to lower the cost of computing. Minicomputers were initially used as components of real-time systems and for personal computing. The LINC minicomputer, developed at MIT's Lincoln Laboratory, was the first personal computer, providing a personal filing system and the ability to write and run programs completely on line.

Timesharing started out as a centralized mainframe facility for a large group with the early demonstration at MIT with CTSS, followed by the introduction of products by DEC (PDP-6), GE (using Dartmouth's BASIC), and Scientific Data Systems 940 (based on the Berkeley Timesharing System). Access was via individual teletypes which were eventually replaced by cathode ray tube terminals, or "glass teletypes." By the mid-1970s, low cost PDP-11 timeshared computers began to be used by separate groups and departments to provide "personal computing." By the mid-1970s, a number of minis used high performance graphics terminals for CAD/CAM, typesetting, and other

applications requiring graphical output or fast response. These minis were the forerunner of today's distributed, personal workstation environment.

In the early 1980s, low cost disks and large memories permitted two evolved computer structures: the 32-bit supermini, and the microprocessor based "team computer." The supermini had all the power of its mainframe ancestors, especially the critical 32 bits to access memory. The "team computer," based on modern, powerful microprocessors, is simply much lower priced (e.g., \$15,000) providing "personal computing" at a price below personal computers.

Workstations

The fourth generation appeared in 1972 with the microprocessor. With the second 8-bit generation microprocessor, floppy disks and 16 Kilobit semiconductor memories (circa 1976), personal computers were practical and began to be manufactured by Apple, Commodore, Radio Shack, and others. With 16-bit microprocessors (measured by datapath) and 64 Kbit rams, the second generation of PCs that appeared in the early 1980s were suitable for building personal workstation environments. Table 4 summarizes the characteristics of these styles of computing.

TABLE 4
Computing style characteristics (circa 1978, G. Bell)

Style*	Central	Departmental Group/Team	Personal
machine	mainframe	mini	micro
price range	\$500,000-\$5,000,000	\$10,000-\$500,000	\$1,000-\$10,000 \$10,000-\$50,000
communications	coupled to terminal network	to mainframes, peers and PCs	terminal emulation
database	organization's archives as service	organization, project function	personal files
terminal access	"glass keypunch"	"glass teletype"	direct CRT
typical uses	corporate accounting, electronic mail	project CAD, order entry, small business "mainframe"	word processing, financial analysis, CAD**
caretaker	a staff providing service	distributed with user group	the user

Comments (1985):

* Not including regional-styles supercomputers.

** Workstations used on a one at a time basis.

In 1979, Carnegie-Mellon University wrote a proposal for personal computer research to implement a workstation environment, and declared: "The era of timesharing was ending." Today's powerful personal workstations, such as the Apollo or SUN workstation, provide a large virtual address, connected with shared facilities on a local area network and characterize this type of machine and computing environment.

PERSONAL COMPUTERS CLUSTERS (I.E., THE WORKSTATION ENVIRONMENT) AS AN ALTERNATIVE TO SHARED COMPUTERS

In the mid-1970s, Xerox PARC researchers developed and provided themselves with a personal computing environment consisting of powerful personal computers all linked together via the first Ethernet cable (3 Mbits), and created the notion of the local area network. Their network had various specialized function servers, including a shared central computer that was compatible with the DECsystem 10, for archival memory and large scale computation.

Figure 13(a) shows the hardware and software of a multipro-




FIGURE 13(a)
 Departmental timesharing (1965-1985) and personal computers with terminal emulation (1978-1985).

grammed computer used for timesharing, and the corresponding structure of a personal computer cluster consisting of functional services and interconnected by a common interconnect that provides basically the same capability. The timeshared system has a central memory containing various jobs connected to terminals and an operating system that attends to the users and handles the particular functions (e.g., real time, files, printing, and communication). Personal computers are connected to timesharing systems as terminals. By comparing the shared system with the systems formed from functionally independent modules, one would expect two design approaches:

1. decomposing systems to provide shared LAN services, and
2. aggregating personal computer to form PC networks and clusters.

Decomposing Systems to Provide Shared LAN Services

As shared computers become more complex and more centralized, it's desirable to decompose the functions for execution on smaller computers that can be distributed to be nearer the use. Thus, the decomposition of a shared system into various boxes, each of which perform a unique function permits the evolution of the parts independent of the whole, the physical distribution of a function and the ability of several computers to share a function (see Fig. 13b). While we have described the evolution of local area networks (LANs) as a decomposition of a single system, LANs are generally an aggregate of heterogeneous systems that access a shared service of some kind as described below.

LANs differ from wide area networks (WANs) in that they assume a low latency, high bandwidth interconnect. This permits file access as well as file transfer applications. With file access, it is possible to remotely locate part or all of a system's mass storage to a file-serving computer. File access requires bandwidth and latency that are roughly equal to that of a disk (i.e., 10 Mhz rates); file transfer can be done at substantially slower rates (56 Khz to 1 Mhz).

Using the reasoning that allowed the formation of the file server, we continue the decomposition of a large central system into servers or stations and then combine these servers into a LAN. The major servers include:

1. *Person server* (the personal computer used as a workstation)—local computation and human interface, possibly private storage of files
2. *File server*—mass storage
3. *Compute server*—batch computation or existence of particular programs
4. *Print server*—printing and plotting of graphics images

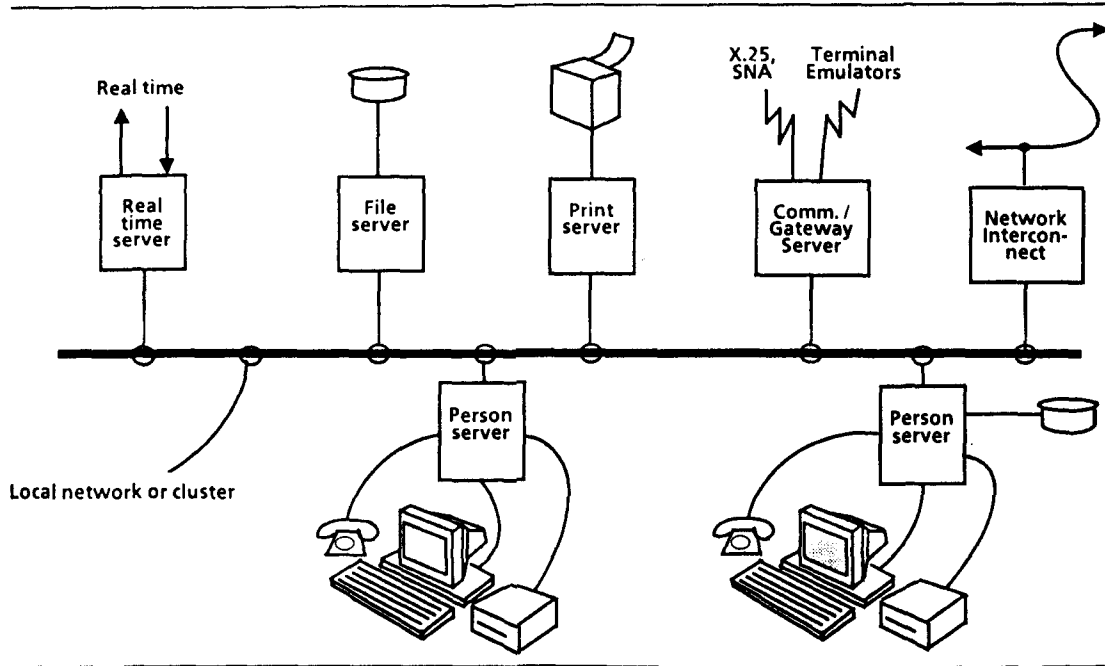


FIGURE 13(b)
Personal workstation computing environment using LAN clusters (1980s).

5. *Communication server*—terminal, telephone and PABX, wide area network access including interface to international standards, other companies (e.g., SNA)
6. *Name/authentication/directory server*—naming the network's resources and controlling access to them

A LAN formed as a complete decomposition of a single system and containing no other incompatible servers would be defined as a homogeneous cluster of personal computers or workstations.

Aggregating Personal Computers to Form PC Networks and Clusters

As personal computers require more facilities (e.g., printing, communication, and files), and the number and type of PCs grow, the need to directly communicate for sending messages and sharing files grows. Furthermore, as a collection of computers in one place forms, economy is gained by sharing common facilities such as printers, phone lines, and disks. Appletnet and Corvus Omninet are relatively short and low data-rate local area networks used to permit the construction of what might best be called a network of personal computers because of the

heterogeneity of type. The 3Com system for interconnecting IBM PCs is more characteristic of the homogeneous network, or cluster.

For a personal computer cluster, one would expect to have a single File Server that can supply records at random to any of its constituency. Table 5 summarizes what timesharing, PCs, and PC clusters provide.

DISTRIBUTED PROCESSING USING CAMPUS AND WIDE AREA NETWORKS

The proliferation of timeshared computers required the development of networking in order for various systems to communicate with one another and with mainframes. Thus dispersed computing became dis-

TABLE 5
Timeshared, PC, PC cluster, and PC network characteristics
(circa 1981, G. Bell)

What	Timeshared system	Personal computer	PC cluster* networks
processing	highest peak	lo-med, guaranteed	= PC
programs size	very high peak	small to medium	= PC
filing	large	small, guaranteed (+ off line)	= PC and TSS
communication	network	term, emulation	= PC and TSS
CRT	slow response "glass teletype"	fast response, screen oriented	= PC = PC
cost	fixed, can go to lowest \$/terminal	lowest entry	f(no. of PCs)
secure	shared, public access	totally private	contained/TSS
pros	explicit costs, shared programs, large jobs	low entry cost, "owned" by indiv., security, SW publishing = low cost	ability to expand, shared facilities, better match to org. structure
cons	shared, poor response for terminals, higher entry cost, security	limited capability, but increasing most rapidly	limited proc./prog., shared facilities
heterogeneity	one homogeneous system	micro computers	one system/can be heterogenous

* Personal workstation environment

tributed computing. Store and forward wide area networks evolved from the ARPAnet, which was used to interconnect timeshared mainframe computers (mostly PDP-10s).

Campus Area Networks

When a collection of local area networks are connected together in a single area that extends beyond a typical LAN, we call this a campus. Universities clearly typify the campus as does a collection of buildings. Gateways are used to interconnect LANs of different types (e.g., Omninet, Ethernet [802.3], IBM Rings [802.4], Applenet, Arcnet, Broadband token bus [802.2] PCnet), whereas bridges or repeaters are used to interconnect networks of the same type to form one larger network. Bridge technology is necessary in connecting multiple LANs together on a wide-area basis using high-speed links (e.g., satellites).

Wide Area Networks

WANs are characterized by low bandwidth (up to 56 Kbits), high latency, errorful transmission, and autonomous operation of the nodes. The applications typically include: mail, file transfer, database query, and low interaction remote terminal access. Wide area networks can be constructed in several ways: direct dial up using conventional circuit switching with voice-grade circuits, an intermediate store and forward network such as Telenet, or a hybrid approach where various worker computers do store and forward switching.

The VAX, Homogeneous Environment or "E"

Although the specific design of the VAX Environment began in December 1978 with the approval of the corporation, its origins include:

- the original VAX-11 goals for a 1000:1 range of computers;
- evolution of distributed processing minicomputer networks, in wide areas, "campuses," and local areas;
- the appearance of powerful personal computers and local area networks, permitting the aggregation of tightly coupled "PC networks and clusters" that provide some of the benefits of timeshared minicomputers and mainframes; and
- the ability to aggregate minicomputers and mainframes into multiprocessors (the 78.2) and multicomputer clusters (VAX clusters) that appear to be a "single" system in order to provide higher reliability, higher performance, and incremental performance.

The December 1978 statement about the distributed computing environment (Fig. 6) and subsequent evolution [shown in brackets] was:

Provide a set of homogeneous, distributed-computing-system products based on VAX-11 so a user can interface, store information and compute without re-programming or extra work from the following computer's system sizes and styles:

- via [a cluster of] large, central (mainframe) computers or network;
- at a local, shared departmental/group/team (mini) computers [and evolving to a minicomputer with shared network servers];
- as a single user personal (micro) computer within a terminal [and evolving to PC clusters];
- with interfacing to other manufacturer and industry standard information processing systems; and
- all interconnected via the local area Network Interconnect, NI [i.e., Ethernet] in a single area, and the ability of interconnecting the Local Area Networks (LANs) to form Campus Area and Wide Area Networks.

Figure 7 shows the origin of the "E" shape that characterizes the VAX Homogeneous Computing Environment. The three horizontal segments of the "E" provide the different computing classes that roughly correspond to different-priced computers; the functions are described in Table 4. In order to implement the environment, many requirements were initially posited, and several developments evolved from necessity:

- a range of VAX-11 and 11-compatible computers to meet the requirements of the various computing styles based on different classes of computers; and
- interconnection schemes and the corresponding protocols for building multiprocessors, tightly coupled centralized VAX clusters, LAN-based PC clusters, LANs, campus area networks and wide area networks.

GOALS OF THE ENVIRONMENT

Product Range

The important goals and constraints of the environment are contained in the original statement about what the environment should do, which is simply "to provide a very wide range of interconnectable VAX-11 computers." The original goal of VAX was to be able to implement the range (for what appears to be a single system) of a factor of 1000 price range . . . with no time limit given. Since a given implementation tends to provide, at a maximum, a range of 2-4 in price and 10 in performance if performance is measured as the product of processor speed times memory size, then many models and ways of interconnection were required.

At the time the 780 was introduced, the total range of products for

both the VAX-11 and 11 family was almost 500, beginning with \$1000, LSI-11 boards and going to a \$500,000 VAX-11/780. The VAX 8600 increased the performance range by another factor of four. If the LSI-11 is included as a personal computer, the price range is reduced to only a factor of 50. While the two ends of the system were "compatible" and could be interconnected via DECNET, they lacked the coherency necessary for a fully homogeneous computing environment.

By introducing "VAX Clusters" (Fig. 7), the range can be extended by a factor of up to the number in the cluster, or about a factor of 10 more. For VAX, Digital now provides a price range of from about \$20,000 for a MicroVAX II to about \$7.5 million for a cluster of twelve 8600s and a corresponding performance range of several hundred. The VAX cluster, shown at the highest level of the hierarchy appears to the user as a single, high availability system.

The VAX cluster work was initiated concurrent with the VAX strategy to address the fault-tolerant, and incremental expansion needs for shared computing including a common disk and file system for the cluster. The interconnection of the clusters, using the computer interconnect (CI), was the beginning of a standard's activity for interconnecting computers on a standard basis.

Static and Dynamic Assignment of Programs to Nodes

Ideally, a user can decide on how to compute on a completely variable basis at the following times:

- At system purchase or rent time ranging from outside facilities reached via gateways, to a central facility, to a shared department or team computer, to a user's own personal computer.
- At system-use time, ranging from access via a terminal, or personal computer interconnected to the system LAN, or a particular shared computer. Here work is bound statically to a particular set of system resources. Most likely, particular nodes would execute special programs on data located at the node.
- At task time on the basis of reliability. VAX clusters provide for dynamic allocation of work among the computers in the cluster to effect load balancing since files are centralized.
- At task-use time on a completely dynamic basis, ranging from computing on his own local system to being able to collect any resources and move work dynamically while programs are in execution. With this ability, as a program goes through its various stages of development, it might be moved from small system to large system to take advantage of increased computational power at higher level nodes.
- At task time on a dynamic basis with the ability to acquire arbitrary resources to engage in parallel computation.

CONCLUSIONS

The VAX Homogeneous Computing Environment is an important computer structure for providing a computing environment. This environment encompasses all styles, sizes, and computing classes from the traditional mainframe and the high availability cluster (utilizing a fault tolerant architecture) to distributed minicomputers, and now the distributed personal workstations. It was stimulated by a number of factors including the distributed-workstation-environment idea. In terms of architectural scope, it is much broader than simply a range of computers as typified by the evolving IBM 360 -> 370 -> 43xx -> 30xx series because it provides full compatibility through homogeneity for a complete range of computing styles and does not require separate architectures for shared-minicomputer systems (System 36, System 38, and Series 1), personal computing (IBM, PC, word processor), or terminal switching and gateways (e.g., IBM 8100).

REFERENCES

- Bell, C. G., "The Multi: A New Class of Multiprocessor Computers," *Science* Vol. 228, No. 4698, 26 April 1985, 462-467.
- Bell, C. G., H. Burkhardt, S. Emmerich, A. Anzelmo, R. Moore, D. Schanin, I. Nassi, and C. Rupp, "The Encore Continuum: A Complete Distributed Workstation-Multiprocessor Computing Environment," *AFIPS Conference Proceeding*, 1985, Vol. 54, 147-155.
- Bell, C. G., G. Fujisaki, J. M. Heinz, A. S. Stevens, and A. S. House, "Reduction of Speech Spectra by Analysis-by-Synthesis Techniques," *The Journal of the Acoustical Society of America*, Vol. 33, No. 12, 1961, 1725-1736.
- Bell, C. G., J. C. Mudge, and J. E. McNamara, *Computer Engineering*, Digital Press, Bedford, Mass., 1978.
- Bell, C. G., and A. Newell, *Computer Structures: Readings and Examples*, McGraw-Hill, New York, 1971.
- Burger, R. M., R. K. Caving, W. C. Holton, and L. W. Sumney, "The Impact of ICs on Computer Technology," *Computer*, Vol. 10, No. 10, October 1984, 88-96.
- Clark, W. A., and C. E. Molnar, "The LINC: A Description of the Laboratory Instrument Computer," *Annals of the New York Academy of Sciences*, Vol. 115, July 1964, 653-668.
- Metropolis, N., J. Howlett, and R. Gian-Carlo (Editors), *A History of Computing in the Twentieth Century*, Academic Press, New York, 1980.
- Pake, G. E., "Research at Xerox PARC: A Founder's Assessment," *IEEE Spectrum*, Vol. 22, No. 10, October 1985, 54-61.
- Perry, T. S., and P. Wallich, "Inside the PARC: The Information Architects," *IEEE Spectrum*, Vol. 22, No. 10, October 1985, 62-76.
- Ralston, A. (Editor) and E. D. Reilly (Assoc. Editor), *Encyclopedia of Computer Science and Engineering*, Second Edition, Van Nostrand Reinhold Company, 1983, New York.

Redmond, K. C., and T. M. Smith, *Project Whirlwind: A History of a Pioneer Computer*, Digital Press, Bedford, Mass., 1980.

Roberts, L. G., "Computer Network Development to Achieve Resource Sharing," *AFIPS Conference Proceedings, SJCC*, Vol. 36, 1970, 543-549.

Siewiorek, D. S., C. G. Bell, and A. Newell, *Computer Structures*, McGraw-Hill, New York, 1982.

Wexelblat, R. L. (Editor), *History of Programming Languages*, Academic Press, New York, 1981.