TOWARD A HISTORY OF (PERSONAL) WORKSTATIONS (DRAFT)

C Gordon Bell
Page Farm Rd.
Lincoln, Massachusetts 01773

I originally accepted this keynote honor for five reasons: to respond to Alan Perlis' request (he told me I could present anything from a new taxonomy to personal reminiscences); second, to identify the important artifacts that should be preserved in The Computer Museum; third, to posit a framework of the history of workstations that can be written in the next century (we're all too close to create it); fourth, to summarize my own involvement on interactive computing including timesharing and large personal computers (both styles of use led to personal workstations); finally, to record my work on Digital's VAX Homogeneous Computing Environment that was totally announced last year. I led the architecture for this environment over the decade from 1973 to 1982 and was strongly influenced by the personal workstation computing movement including Ethernet. This work hasn't been described outside of product brochures.

INTRODUCTION TO THE CONFERENCE

In June 1976, an international conference was held resulting in a collection of views published in "The History of Computing in the Twentieth Century" (Metropolis, Howlett and Gian-Carlo Rota, 1980). The participants presented work up to the early 50's and thus had the benefit of 30 years of hindsight. The late mathematic's 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

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

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:

- 50's idea (documented article, proposal, movie, etc.) stimulated through early standalone use of small computers (e.g. LGP-30, G-15)
- 60's breadboards to demonstrate the idea and selected use of large computers with graphic displays (e.g. DEC PDP-1, IBM 7090, LINC, TX-2)
- 70's (early) limited use of interactive shared workstations using graphic display terminals connected to minicomputers; establishment of an industry to supply terminals and professional applications software
- 70's (late) working prototypes of personal workstations with concept testing through use in a complete environment (Xerox PARC)
- 81 introduction of personal workstations by Apollo, SUN, Xerox, and first use by early adopters
- 83 full-scale use for selected professionals, many companies formed, JAWS--Just Another Workstation (term) coined
- 85 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 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 workstation 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 is 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 producing and selling machines. For example Doug Drane and the ATEX system that produced law and typesetting offices doing essentially what a STAR does, only five 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, etc. 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 onto 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 Computer Encyclopedia, 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 IS A (PERSONAL) WORKSTATION?

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 70's to deliver cost-effective computing. An overview of the classic machines and the classes they formed is given in Figure 1.

Figure 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

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

line with the salaries of the professional using the station. Four lines, corrected for inflation, depict the LINC price, the cost of a supported professional (e.g. 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

Figure 2. Workstation price versus time for various workstation introductions.

performance or capability. While salaries of clerks are greater than \$10K, the workstation industry believes \$10K to be magic, highly elastic, price barrier which will increase sales. In the later part of the 80's, 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 (>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 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 of 1/4 to 1 megabyte, permitting the use of large programs and construction of high resolution bit-mapped terminals
- disks of >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 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 appearing 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 (eg. 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 and permitting access and sharing of facilities and work. Hence, constant-cost PC's 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% per year. Since it was a constant fraction of system prices, they also declined at the same rate. This 20% decline was based on the learning curves for manufacturing core memory manufacture and distribution. Thus, a computer would drop a factor of 10 in price, per decade.

In 1972, with the introduction of semiconductor memories by Intel, price began to decline 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 development of specific personal computers and personal workstation. The reduction in the price per bit of semiconductor memories amounts to a 36% price decline each year or a factor of 10 every 5 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

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					
1978	16K	8086	16	16K-32K	Apple II, TRS-80, PET
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	?

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. 64Kbytes) 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 Figure 4. This evolutionary model turned out to be quite accurate, predicting both the Apple II (just 3 years away), low-cost shared micros, and the new class of mini-supercomputers that are beginning to emerge this year.

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

Figure 4. Price versus time for each machine class.

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). 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 of 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).

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 (multi-programmed) n-users	65KB - 256KB	Interactive--MUMPS, RSTS, Trans.Proc. Real time--RSX-11M, D	Special purpose
Programmable (multi-programmed) n-users	128KB-1024KB	Interactive--IAS, TOPS 10, RSTS Real time--RSX-11D	Generality

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	knoetop, small personal	1K	5
3*	office	personal, large personal, "workstation"	10K	50
4*	department	shared micro, minicomputer, super-minicomputer, mini-supercomputer	100K	500
5*	center	mainframe	1M	5,000
6*	region	supercomputer	10M	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

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.

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 its own and other central mainframes. DECNET was introduced in the following phases:

I	75	task-task, file transfer, file access; point-to-point; 8-bit net address; operating systems: RSX-11 D, M, IAS
II	77	improved system generation; used for first DEC engineering network (now several thousand machines at tens of sites; operating systems: rt-11, rsts
iii	80	routing, virtual terminal, full network management; operating systems: decsystem 20, vms
iv	83	large routing address, full remote rms files; c-terms; ethernet support
	85	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 Kword microprogrammed 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' founder, Brian Rosen, a CMU grad who had worked at Xerox Palo Alto Research Center (PARC), designed PERQ as his version of the 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 (eg. 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

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

processor was a near derivative to the Data General Nova with an integrated bit-mapped display. It is important to note the PARC environment also included a timeshared 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 3 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 Figure 6 provided a wide range and hierarchy of compatible computers and ways of interconnecting them to provide users with generic (eg. word processing, mail) program

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

Figure 7. Digital homogeneous computing environment 1984.

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 (CI=Computer Interconnect, NI=Network Interconnect, BI=Backplane Interconnect, and II=Interchip Interconnect; II did not materialize) as shown in Figure 7.

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 computers 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.

Single Computer - circa 1970, 80...

LAN Computing Environment

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

In January, 1981 when Ethernet was being introduced by various vendors and being completed as an international standard, several of 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,
- 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 Figure 8).
- as a bus for building computer systems.

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, 1985a). The "Continuum" provided both the distributed personal workstation style (described in the evolution, Appendix 1) and shared central computing using a multiple, microprocessor architecture, which I have named the Multi (Bell, 1985b). Ethernet provides the switch for interconnecting terminals to the Multimax using concentrators which pre-process 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).

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 70's 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 principle 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, Cavin, Holton and Sumney (1984) as a timeline of the critical processor features (see Figure 10). The single most critical dimension of processing power is virtual memory size. The reinvention of the wheel of virtual memory size is shown in the Figure 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 which 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. (Figure 12 reproduced from Bell, Mudge and McNamara, 1978.) Note the evolution: direct control of the display object by the central processor; display list issued to a DMA controller or external controller; separate processor with loops, subroutines, etc.; 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.

Figure 9. The Encore Computing Continuum.

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.)

- . resolution (in pixels or analog vector resolution)
- . number of bits per pixel including use for monochrome/color
- . display data-types
 - . points, characters, lines, curves,
 - . 3D line drawing
 - . wireframe meshes, 2D surfaces, 3D surfaces,
 - . shading, sculptured surfaces, constructive geometry
- . translation of data-types
 - . fixed
 - . translation, scaling, rotation,
 - . 2D clipping, 3D perspective, 3D clipping
- . standard interface evolution: direct control, Siggraph CORE Library) VDI, CGI, NAPLPS, etc.
- . windowing of output images
- . 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, mouse, etc.

PAPER OUTPUT

This section must include the evolution of: conventional printing and plotting along with early film recorders to capture images and 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.

Figure 11. Virtual address space evolution for various computers and computer classes.

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.

Figure 12. Myer and Sutherland wheel of reincarnation for display processor evolution.

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, scientific experiments).

OPERATING SYSTEM, DATA-BASE, 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 which 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 word processing
 - database access
 - transaction processing
- **SHARED WORKSTATION AND PERSONAL WORKSTATION ORIENTED (FAST LINKS)**
 - multiprocess control including control, windows
 - spread sheets
 - drawing and painting

PROFESSION SPECIFIC WORKSTATIONS

In addition to generic applications programs (eg. 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 (eg. 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 (eg. radiologist, cardiologist, bone surgeon)
 - . geologist, petroleum engineer
 - . physics experiment control and analysis
 - . statistician and mathematician
 - . signal analysis, encryption/decryption
- . ENGINEERING WORKSTATIONS
 - . electrical or electronic schematic (2D line drawings)
 - . VLSI physical design 2-1/2 D
 - . mechanical drafting
 - . design of sheet metals or piping (2-1/2 D)
 - . architectural and structural engineering
 - . general mechanical design of arbitrary surfaces for analysis and fabrication
- . SOFTWARE AND KNOWLEDGE ENGINEERING WORKSTATIONS
- . ARTS (TYPESETTING, GRAPHICS ARTS, MUSIC, ETC.) WORKSTATIONS
 - . publisher of: newspaper, magazine, book
 - . scholars including: historians, linguists, literary scholars
 - . graphic artist, illustrator, cartoonist, 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, aircraft, etc
 - . 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-0, 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 which 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 SAGE (Semi Automatic Ground Environment) 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 60's. 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 (1/59 - 6/60) 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, which took 50 microseconds to display a point on the 10" x 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 which was controlled by loading a 36 bit, 5 x 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 was a 5" high precision display with 4K point resolution. Only two? were built, the first being 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.

Two? PDP-1 Color Displays were built, one of which 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 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 PC'S: 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 2Kw, 12-bit memory and two LINCtapes. 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-8's beginning in 1967 (designed by Wes Clark and Dick Clayton) and around 1,000 PDP-12's 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) was, I believe, the first display processor (Bell, Mudge, and McNamara, 1978). It 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 standalone 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, lighted buttons, etc.

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 \$50K per terminal using a \$15K graphics terminal). About 30% 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 \$25K 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 pre-processor for time-shared systems doing CAD, etc.

The GT44, was introduced in June 1973 at a price of \$38K, 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, standalone computer.

The GT48, introduced in October 1975, selling for \$55K, 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:

Year	Model	Price	Comments
1975	DS310	\$16K	8/A, desk, 8" floppy and VT52 terminal,
1976	WS200	\$40K	8/A, cabinet, 1-8 users, 5 Mbyte disks,
1977	VT78	\$15K	Intersil 6100 PDP-8 chip, embedded in VT52, 8" floppies
1981	DECmate I	\$12K	6120 chip, embedded in VT100, 8" floppies
1983	DECmate II	\$3.5	integrated unit, 5" floppies, bit mapped,
1984	DECmate III	\$2.4	Z80/8086 and 5" winchester disk options

By 1984, the 21st 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 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: TERA, 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 x 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 \$10K. 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 VT100, 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 \$8 K. These terminals connected to host PDP-11's which 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, which 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 84 and June 1985, respectively. Both units sold for approximately \$20K, with 19" bit-mapped CRT, using the Qbus for interconnecting I/O equipment.

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

CONCLUSIONS

Personal workstations have evolved from many directions and along many dimensions. This conference barely scratches the surfaces of a few dimensions. We should endeavor to be as complete and accurate as possible in what we provide and in going from these papers to a more widely available publication.

The conference has served to convince me that the critical dimensions of personal workstation evolution should be explored and documented by historians.

ACKNOWLEDGEMENTS

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

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 70's 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 perform 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-60's 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 M.I.T.'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-70's low cost PDP-11 timeshared computers began to be used by separate groups and departments to provide "personal computing". By the mid 70's 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 80's, 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 (eg. \$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, etc. With 16-bit microprocessor (measured by data-path) and 64 Kbit rams, the second generation of PCs appeared in the early 80's which were suitable for building personal workstation environments.

Table 4 summarizes the characteristics of these styles of computing.

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 providing a large virtual address, connected with shared facilities on a local area network 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 70's, Xerox Research Park researchers developed and provided itself 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 13a shows the hardware and software of a multiprogrammed computer used for timesharing, and the corresponding structure of a Personal Computer Cluster consisting of functional services and interconnected by a common interconnect which provides basically the same capability. The timeshared system has a central memory containing various jobs connected to terminals and operating system which attends to the users and handles the particular functions (eg. real time, files, printing, communication). Personal Computers are connected to timesharing systems as terminals. By comparing the shared system with the systems formed

Table 4 - Computing Style Characteristics (Circa 1978, G.Bell)

STYLE*	CENTRAL	DEPARTMENTAL GROUP/TEAM	PERSONAL
machine	mainframe	mini	micro
price range	500K-5M	10K-500K	1K-10K 10K-50K**
communications	coupled to terminal network	to mainframes, peers and PC's	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.

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 Figure 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 which 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 which are roughly equal to that of a disk (i.e. 10 Mhz rates); file transfer can be done at substantially slower rates (56Khz to 1 Mhz).

Using the reasoning which 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
5. Communication Server - terminal, telephone and PABX, Wide Area Network access including interface to international standards, other companies (eg. 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. Furthermore, as a collection of computers in one place forms, economy is gained by sharing common facilities such as printers, phone lines, and disks. Applenet 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 PC Cluster, one would expect to have a single

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

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

File Server which can supply records at random to any of its constituency. Table 5 summarizes what timesharing, PC's 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 to mainframes. Thus, dispersed computing became distributed computing. Store and forward wide area networks evolved from the ARPA-net, which was used to interconnect timeshared mainframe computers (mostly PDP-10's).

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	low-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 heterogeneous

*Personal workstation environment

Campus Area Networks

When a collection of Local Area Networks are connected together in a single area which 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 (eg. Omninet, Ethernet (802.3), IBM Rings (802.4), Appletan, 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, it's 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,

- 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 '78 statement of the Distributed Computing Environment (Figure 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) computer, [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 which 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;
- 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, 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 \$1,000, 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 homogenous computing environment.

By introducing "VAX Clusters" (Figure 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 8600's 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 not requiring 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).

Note: this Appendix is solely the view of the author and is not necessarily the view of Digital Equipment Corporation or its engineers.

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.

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, ppl725-1736.

Bell, C.G., J.C.Mudge, and J.E.McNamara, Computer Engineering, Digital Press, Bedford, Massachusetts, 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.