THOMAS R. KOCHTANEK

Chair
Department of Information Science
School of Library and Informational Science
University of Missouri-Columbia

Micro Generations: Current and Future Directions

INTRODUCTION

A digital computer is a programmable device which on the broadest level supports the manipulation of symbols aggregated as data. Simply put, the computer is a tool for creating, maintaining, organizing, storing, transmitting and disseminating data of all types. Developments in computing hardware have a certain historical significance and offer a clear portrait of the role of technology in society. Newer technology has its roots in this compact history.

The first section of this paper traces the history of computer hardware in general. The second section focuses on the evolution of microcomputers as a subset of general computing systems. The third section focuses on progress in the area of 32-bit microcomputer architecture. The final section ties those advancements in microcomputing to existing and proposed database applications in libraries and related information agencies.

GENERAL HISTORY OF COMPUTER HARDWARE

A complete depiction of the evolution of computers would include details regarding the simultaneous evolutions of hardware, operating systems and applications software. Hardware development provides the raw resources for computing, operating systems provide real time access to these resources, and applications software guides specific procedures to be carried out by the computing system. The objective of this section

is to focus primarily on hardware generations of computers, but occasional references will be made to developments in the related areas of software design.

The first computers were built exclusively as prototypes and were used primarily to perform highly accurate arithmetic computations for military and related research needs. Most of the earliest computers were invented at universities and were supported by contracts from the War Department, now the Department of Defense. As the benefits of these early computational devices began to be noticed by engineers and scientists alike, additional applications arose which transcended numeric processing. The differences in design and function are due primarily to the manner in which each applications program views the stream of data input to the computer. In some cases the data stream is treated as text, sometimes as formulas, and other times as data from various fields in a record.

The first revolution in the design of computing systems began in 1943 and lasted until 1950. Although there is some dispute as to who can claim to have built the first mainframe or maxicomputer, this event is primarily of historical concern. Of more enduring impact was the introduction of the first mass-produced programmable computer in 1951. The UNIVAC I was so important to the fledgling computing industry that historians refer to it as "the beginning of the first generation of computing." Two short years later, IBM began mass marketing a mainframe computer, the IBM 650, which utilized punched cards in a computer for the first time. IBM quickly dominated the market with its sales of computing systems.

These early computing systems were unique with respect to their applications and users. Initial computers were programmed in machine language only (using base two zeroes and ones) to execute one request at a time, calculating and outputting the results for a single end user. There was no internal or core memory, no keyboards or terminals and no storage devices as we now know them. Central processors were initially composed of large numbers of vacuum tubes, although the transistor had been invented in 1948. By the middle 1950s, programmers were using assemblers in the place of machine language programming to develop applications programs. It was not until 1958 that FORTRAN, ALGOL and a language called APL were introduced and used as high level programming languages.

The second generation of computers (1959-1963) was typified by IBM's 1401 mainframe, developed in 1959 and distributed in 1960. This mainframe benefitted from the introduction of transistors, which replaced the bulky and problem-oriented vacuum tubes of its predecessors. This system also utilized internal memory, supporting

between 1K and 16K RAM. Along with significant improvements to existing high level languages, new languages were developed: COBOL, LISP and SNOBOL, to name a few.

The third generation of computers (1964-1967) was more revolutionary in its rapid development of numerous new architectures, languages and capabilities. An important new language, BASIC, was introduced in 1964. A new architectural design resulted in the introduction of the first minicomputer, the PDP-8, from Digital Equipment Corporation in 1965. An even more significant introduction was the new line of IBM 360s introduced in 1966. Along with the 360 came a promising new programming language from IBM, PL/I. This period witnessed the early developments in time-sharing, whereby certain resources of the computer, particularly main memory and external storage, were optimally allocated among several simultaneous users.

The fourth generation of computers, from 1968-1974, was marked by a steady but somewhat less spectacular growth, especially when compared to the remarkable growth of the previous period. The major introduction was in the form of a new IBM architecture, the IBM model 370. But the post-Vietnam recession had its effects on the computing industry. In 1969 numerous computer firms laid off significant amounts of their work force. Several corporations either folded or sold themselves to other companies. The majority of changes were silent ones, such as the growth in sales of minicomputers and companies such as Digital Equipment Corporation, which sold minicomputers and accompanying services and products. Schools and colleges began to purchase minicomputers for administrative uses and to experiment with their use in classroom instruction. Cathode ray tubes became affordable and supported enhanced access to these systems via screens and keyboards. The use of semiconductors for internal memory became standard. The distribution of smaller, less expensive minicomputers began a silent revolution which was soon to be fueled by the introduction of the affordable microcomputer to the masses.

A MORE DETAILED HISTORY OF MICROCOMPUTERS

In 1971 a company named Intel began shipping the first microprocessor, the 4-bit 4004, a complete central processor on a single silicon chip. In 1973 an improved version of that first chip, the 8080, was shipped by Intel. In 1975 another company, MITS, manufactured and offered the first popular mail order microcomputer kit for $395. The Altair was based on the 8080 central processing unit, had 256 bytes of RAM, no ROM, no CRT, no keyboard, no printer, and no external

disk storage. The Altair incorporated twenty-five switches for input and thirty-six blinking lights as output in support of the 8080 machine language. This offering compared to those mainframes of thirty years prior, in that the computing system was limited in its communication capabilities and only supported a single user possessing a high level of computing expertise. The main difference was the fact that this microcomputer was affordable and was designed with a more "open" architecture, allowing the addition of specific peripherals to the base system. Eventually, other companies built add-on boards, disk controllers, keyboard interfaces and the like. Bill Gates invented a way to make this personal computer kit handle BASIC. Gary Kildall wrote a single user disk-based operating system for the 8080 called CP/M. Although the Altair never made it as a full production offering (MITS went bankrupt) the impact of this new microcomputer system was resounding. For a few thousand dollars, one could have true personal computing at one's fingertips!

By 1977 there were four companies offering microcomputers with built-in keyboards: Radio Shack (TRS-80 @ $499), Commodore (Pet @ $795), Apple (Apple II @ $970) and Processor Technology (Sol 20 @ $1850). Commodore's and Radio Shack's offerings were considered the bargains since they included a monitor with the unit. In 1979 Texas Instruments and Atari entered the fray. Then in 1980 and 1981, the Timex Sinclair and the portable Osborne were introduced.

On August 12, 1981 the IBM PC was announced. To many, this introduction signalled legitimacy for the fledgling microcomputer industry. Business began an unprecedented mass purchase of millions of personal microcomputers offered by IBM and its competitors. End users of computing systems proliferated as these industry standard microcomputers took their place on the desk tops of corporate America. A new development in computing, mass access to personal workstations, invited literacy, efficiency and productivity to individuals across many walks of life. While large centralized systems operated primarily by data processing departments still dominated corporate operations, millions of end users experienced for the first time personalized computing, an experience which has made its mark in computing systems as we now know them.

Another significant development in industry standard microcomputers began with the 1984 introduction of the Apple Macintosh, a derivative of an earlier Apple product called Lisa. The first Macintosh used a different microprocessor, the Motorola 68000, and incorporated a graphics-based operating system capable of supporting easy-to-use applications software. Most of the ease was for the end user who navigated pull down windows with a mouse, selecting among various icons to interact with the machine. Developers of applications software

have found the Mac to be somewhat slow and cumbersome in developing business and related software. One major development, known as HyperCard, offers promise for future developments of this offering from Apple.

IBM introduced a new line of personal computer systems on April 2, 1987 known collectively as the PS/2 line. While the "low end" of the line (models 25 and 30) offered minor improvements such as swift processing speeds and smaller desktop "footprints," the upper end of the line (models 50 through 80) offered extended memories, even faster processing speeds and a new bus architecture called "micro channel" or MCA for short. This architecture holds the promise for more sophisticated multi-tasking applications development, an important consideration in the design of many current database management applications running on microcomputer systems.

MS-DOS Based Personal Computers: A Review

A review of events limited to the IBM line of PC-DOS products may prove helpful to chronicle one segment of the personal microcomputer arena. As mentioned above, the original IBM PC was announced to the public on August 2, 1981. The next developments included increased internal (RAM) memories and fixed disk storage capabilities, released together as the IBM PC/XT. In 1984 IBM announced its IBM PC/AT (for Advanced Technology) based on the Intel 80286 microprocessor. Three years later the AT outsold the number of systems produced as the original PC and PC/XT. Eighteen months after the PS/2 line was introduced, the first one million model 30s were sold. Other models of the PS/2 line are now being purchased in greater quantities, some as high performance personal workstations, others as file servers within local area networks. If a decision-maker were to limit one's selection solely to MS-DOS based personal computers from 1981 to present, Table 1 might reflect personal purchases made from year to year.

These selections are based on those assumptions of financial restrictions which might apply to personal situations. Certainly, more powerful systems can be configured if cost is no object. The attempt is to portray mass selections, not optimal designs. The appearance of non-IBM equipment in the list is indicative of a trend in the industry to offer lower cost, more powerful clone systems as competition to one industry giant, the IBM standard.

The major trend is clearly in favor of the 80386- and 80486-based microprocessor as a machine which poises the end user for developments in both operating systems and applications programs. Worldwide

estimates of sales of 386-based systems were projected to be 4.4 million systems in 1989 alone. This most significant trend deserves further explanation.

TABLE 1
MS-DOS BASED PERSONAL COMPUTERS, 1981-PRESENT

Date	Product Description
Summer 1981	IBM PC w/64K RAM 1 floppy
Summer 1982	IBM PC w/256K RAM 2 floppies
Summer 1983	IBM PC XT w/640K RAM and 10 Mbyte hard drive
Summer 1984	IBM PC AT w/640K RAM w/slow 30 Mbyte hard drive, 1.2 Mbyte floppy
Summer 1985	IBM PC AT w/speedup crystal and large, quick, reliable hard drive
Summer 1986	IBM PC clone 80286—much progress was made by the competition
Summer 1987	IBM PS/2 model 30 or 50
Summer 1988	IBM PS/2 model 50 zero wait state or any competitive 80286 clone
Summer 1989	Clone 80386 with large, quick drive (Dell 386 and Everex Step 386 were big sellers)
Summer 1990	Faster 80386 based systems addressing increased RAM and sophisticated peripherals in support of multi-tasking or "windowed" applications
Summer 1991	80386 or possibly an 80486 system addressing even larger, faster storage devices; possibly an external CD-ROM drive.

CURRENT STATE OF MICROCOMPUTING

A Review of 32-Bit Technologies

The development of the 8-bit microprocessor and accompanying peripherals was but an initial seed in the harvest of microcomputer products. Soon, demanding users moved up to 16-bit technologies based on two microprocessors: Intel's 8086 and Motorola's 68000. These and their powerful descendants (the 80286 and 68010) have made their mark while dominating the microcomputer industry over the past several years. In the quest for even more processing capabilities, the 32-bit processor platform is emerging as a new force in the industry. Although these processors have been available for some time, high demand and mass production have continued to lower costs well inside of the $10,000 mark for a configured system. That amount is often considered the high water mark of personal computing cost. Business applications involve different cost considerations.

With these improved technologies comes a whole flood of terminology new to the microcomputer user: memory caching, virtual memory, pipelining, RISC, CRISP, MIPs, and so on. These not-so-new concepts (many were originated in the mainframe world) can be examined in light of additional facts regarding microcomputer-based products.

What Is It about Word Length that Is So Important?

Word length is a term used to designate the number of binary digits that can be processed at one time inside the computer. Just as an automobile with eight cylinders can deliver more useable horsepower than one with four cylinders, so a central processing unit with 16 or 32 bits can provide more raw computing power than one with 8 bits. The wider the bit path, the more work a CPU can do at one time. Some liken word length to a highway of signals, where a highway of four lanes can provide higher transportation rates than one with just two lanes. In addition to these increased processing capabilities, CPUs with larger word lengths also possess such features as higher clock rates, larger internal registers and increased addressable memory.

What Applications Need the Kinds of Power Reserved for 32-bit Chips?

Faster, more powerful processors are needed not so much for applications such as recalculating a spreadsheet, but for those types of applications that require high resolution graphical interfaces and large amounts of high speed memory. Databases are a prime example of these kinds of needs, since they occupy large portions of fixed and volatile memories. Many applications can benefit from improvements in user interfaces supporting graphical interfaces. As operating systems continue to develop, they appear to be following the trends set by the Macintosh, incorporating pull-down windows, overlayed graphic windows and higher resolution screens. These features all require extremely fast and powerful processors such as those 32-bit CPUs currently under refinement.

How Does One Measure Internal Clock Speeds and Cycle Times? What Are Wait States?

As soon as a CPU receives a set of data or instructions, questions of timing arise. How long does the CPU store that data? When does the CPU refresh dynamic RAM? When does the CPU move it? How are signals synchronized? These issues are so critical that logic with memory is called sequential, as opposed to combinatorial logic of memory-less computers. Sequential logic is kept synchronized with an internal clock.

All computers have internal clocks. The clock's pulse is the computer's heartbeat. One clock pulse is the burst of current when clock output = 1. One cycle, also known as one Hertz, is the interval from the beginning of a pulse to the beginning of the next. Depending on the computer, the clock frequency may be hundreds or thousands or millions of cycles per second. Megahertz, or Mhz, is a measure indicating the number of millions of cycles of a CPU per second. Mhz is one measure of a CPU's capability to perform.

The idea of using a clock is that the computer's logical state should change only on the clock pulse. Ideally, when the clock strikes one, all signals move, then stop on clock = 0. A condition known as zero wait states means execution occurs at the conclusion of a single cycle. One or two wait states implies pauses in the transfer of binary data to and from registers within the CPU.

Certain operations within the CPU take more than a single instruction to perform. Some mathematical operations normally take many different instructions to execute on a 16-bit CPU. A computer with a 32-bit word length may process an operation in one or two instructions, thereby increasing throughput efficiency. Megahertz alone is not a perfect measure of the raw computing speed of a central processing unit. MIPS is often used to designate millions of instructions per second, as opposed to millions of cycles per second (Mhz). This measure is dependent on the type of instruction under consideration. For example, certain often-used instructions execute in a single cycle of the CPU while others require hundreds of cycles to execute. MIPS is calculated by determining the average number of clock cycles a chip's machine level instructions take to execute and dividing the CPU clock speed, measured in Mhz, by that number. If a CPU can perform each of its binary instructions in one cycle of the clock and that chip has a clock speed of 10 Mhz, then it would process 10 million instructions per second, or 10 MIPS.

What follows in Table 2 are some of the processing speeds of Intel's 8, 16 and 32 bit central processors. Each is measured in millions of cycles per second or Megahertz (Mhz) and in millions of instructions per second (MIPS).

In comparison, the Motorola 68020 CPU is rated at 4.0 MIPS and the 68030 at 6.8 MIPS. A change in clock speeds, instruction sets or wait states can render these comparisons useless for any specific case. For example, the recently announced NeXT computer uses the Motorola DSP56001 CMOS chip operating at 20 Mhz. However, this chip's instructions execute (on the average) every two clock cycles to give the CPU a 10 MIPS rating.

Comparing these ratings to other, more powerful computers is interesting as well. For example, an early DEC VAX 11/780 minicom-

puter processes at approximately 1 MIPS. Powerful workstations such as Sun, Apollo and the DEC MicroVAX house CPUs ranging from 5 to 15 MIPS. An IBM 3090 mainframe operates at 100 MIPS. This indicates that the current generation of microprocessors are capable of providing up to five times as much computing power as the previous generations of minicomputers. Parallel processing supports the linking of multiple

TABLE 2
PROCESSING SPEEDS OF INTEL CENTRAL PROCESSORS

CPU	:	*Mhz*	:	*MIPS*
8088	:	4.77 Mhz	:	0.5 MIPS
8086	:	8-10 Mhz	:	0.5-1.0 MIPS
80286	:	10-20 Mhz	:	0.5-3.0 MIPS
80386	:	20-35 Mhz	:	4.0-8.0 MIPS
80486	:	50 Mhz	:	speculative

CPUs within a single system unit. Imagine linking ten 10 MIPS CPUs together in a single microcomputer workstation, one that is capable of providing as much computing power as a current generation IBM mainframe!

How Can A CPU Be Sped Up? What Is RISC and Clock Speed?

There are two ways to increase any particular chip's processing speed. One is to step up the clock speed. New improved versions of existing CPUs often offer increases in clock speed. For example, the 80286 as it was first introduced in 1984 had a clock speed of 8 Mhz. In subsequent updates, that same processor was bumped to 10 and then 12 Mhz. Another method is to decrease the number of cycles needed to execute instructions. Two distinct approaches have been used to accomplish this decrease:

1. Some of the more predominant CPUs, such as the 80386 from Intel and the Motorola 68030, support an approach called Complex Instruction Set Computing (CISC). This design uses an on-chip microcode program (software) to pre-process certain instructions before actual CPU execution occurs. This type of design reduces each instruction to 2 to 6 CPU cycles, resulting in a full set of instructions executing at higher speeds than previous conventional designs.

2. Certain chips use a modified and highly optimized internal instruction set called Reduced Instruction Set Computing (RISC). Several 32-bit chips under development use RISC technology to approach the theoretical limits of processing: single cycle instruction processing efficiency, where each instruction is processed within a single cycle of the CPU. In RISC architectures, the CPU's machine language instruction set is pared down to a subset of fundamental

and frequently used instructions. The instructions themselves are optimized to execute directly in hardware, without the need of microcode. The results are extraordinary: RISC processors generally execute each basic instruction in 1.25 to 2 cycles. A RISC chip can operate at up to three times faster than its non-RISC counterpart. These chips are being developed primarily by high end workstation companies such as Sun and Hewlett Packard. IBM has offered a RISC workstation for the past few years. Apple computer is rumored to have been developing a RISC chip for its high end Macintosh line. Motorola has introduced an entire line of RISC-based CPUs in the 88000 family. Intel is also developing a RISC chip called the 80960. Experts envision the development of CRISP, for Complexity Reduced Instruction Set, a combination of optimized CISC chips using RISC technology. Recently released chips such as Intel's 80486 and Motorola's 68040 are likely to lean in this direction.

What Are Some Roadblocks to Performance?

The performance enhancements of microprocessors are truly remarkable feats when considered independent of other potential bottlenecks in a computing system. But when any CPU is sped up beyond 20 Mhz, the main impediment to performance becomes RAM memory, specifically, dynamic RAM. While DRAM is an excellent bargain in terms of price per megabyte, the fastest DRAM chips available cannot keep up with the relentless increases in CPU clock speeds. One solution to the problem would be to replace all DRAM chips with their speedy counterparts, static RAM, or SRAM, but the cost for several megabytes of SRAM would be prohibitive.

Instead, chip makers are focusing on developments such as Single In-line Memory Modules (or SIMM) used to store up to 4 megabytes of inexpensive DRAM on a system board. This is counterbalanced with a small (between 4K and 256 Kbytes) amount of very high speed static RAM installed as a buffer or "cache" used to feed the DRAM and to help that inexpensive mass memory keep pace.

Cache memory is a small but high speed holding area for data that a CPU is using or about to use. Consider the situation where one is attempting to prepare a meal in one's own home. Perhaps a key ingredient is missing from the cupboard and a trip to the grocery is required. At the grocery, one has the ability to purchase not only the specific items necessary for the preparation of that particular meal, but also additional items which may be needed for other preparations. So it is with cache memory. High-speed static RAM is utilized to store anticipated data which the CPU is likely to require in ensuing processing.

Many of the current generation 386 PCs use this technique to improve database access and performance. Since database queries are disk-intensive activities, cache memory can be used to temporarily store frequently accessed data likely to be requested by the CPU for processing. Intel makes a 82385 controller which uses 32K to 256K of 35 nanosecond static RAM for use as cache memory. The 80386 CPU in combination with this controller can locate data in the cache with 95 percent "hit rate."

The IBM PS/2 model 70-121 does not support cache memory. The model 70-121 took sixteen seconds to query a 1000 record Paradox 386 database. The Everex Step 386/20 uses cache memory and the 82385 controller mentioned above. The elapsed search time for the Paradox search was eight seconds, or half the time of the uncached system. A Dell System 310 using cache memory recorded the same eight-second response time. All hard drives for the three systems had the same twenty-five millisecond seek rating.

Motorola uses a slightly different approach to cache memory. Instead of using separate SRAM chips and a controller, the 68020 and 68030 CPUs use a 256-byte instruction cache built into the CPU. By caching both the instruction and the data and eliminating the external controller and cache chips, these CPUs display even higher clock speeds.

Another technique is known as *pipelining*. It is a known fact that a CPU is idle during certain processes it must perform. Basically, a CPU is very routine in its procedures. It first reads an instruction from memory and decodes that instruction. Then the CPU reads data from memory and processes that data in accord with the instructions, writing the results to memory. This cycle continues until all instructions are processed or the process aborts. During this prescribed cycle, the bus between memory and the CPU sits idle, waiting for the CPU to access additional instructions or data. During an idle moment, the CPU can be instructed to peek at the next instruction or chunk of data, parking its location or contents in a special address within its register. This "look ahead" technique, or pipelining, increases the throughput performance of the chip. Both Motorola and Intel have introduced pipelining in their current 68030 and 80386 CPUs. Certain RISC chips can be doing up to five tasks at once, thereby increasing the efficiency of the chip.

Addressable Internal Memory: What Good is All That RAM?

Earlier CPUs had severe RAM memory limitations. Eight-bit CPU architectures were limited to 64K RAM and 16-bit CPUs had 1M memory limitations. Addressable or user memory of any particular system offering varied depending on operating system characteristics and address paths of the internal bus architecture. For example, the original

IBM PC, with its twenty address lines on its systems bus, could theoretically address 2 to the 20th power or one megabyte of main memory. Because of the storage requirements of the operating system (and other features), the actual processing capacity of the original IBM PC is reduced to 640K of useable RAM for the end user.

Why would a user need any more than 640K RAM? There are several reasons. First is the fact that complex applications programs can exceed such a limit. As software is developed and enhanced, its code can readily exceed that barrier. Another reason was referred to earlier: cache memory can take up additional addressable memory of the CPU. A third reason for extending internal memories is due to the rise in popularity of multi-tasking in the microcomputer environment. Operating systems such as UNIX and now OS/2 have the capability to run multiple tasks or applications in memory simultaneously. This load requires a much larger internal capacity to store and run these multiple applications. An example of multi-tasking would be the simultaneous loading of a word processor, a spreadsheet and a database manager into memory. A single user could look up information in the database, calculate something from that data using the spreadsheet program, and transfer the result for inclusion in the word processor.

There are two primary means for providing large amounts of RAM memory in addition to the "base" memory. One is to "extend" RAM by using a second segment of RAM chips and "bank switching" between base and extended memory. This is how the early Apple IIs could address 128K RAM when their architecture permitted only 64K of directly addressable RAM. Another method is to incorporate "virtual" memory features similar to the manner in which mainframes and minicomputers have done. Virtual memory is a technique that allows a CPU with a small amount of "real" memory to act as if it has even more than that amount of real memory. A special chip is used which responds to a request for more RAM than is physically present by generating an "interrupt." The operating system is then asked to swap certain contents of currently unused (but currently storing data or instructions) RAM to physical disk, thereby freeing up extra RAM for the requested instructions. By dedicating a segment of a hard disk to virtual memory, large RAM-intensive applications programs can be run on computing systems with relatively small memories.

Both Intel and Motorola incorporate virtual memory options into their 16-bit CPUs: the 80286 and 68020. For example, Motorola's 68020 and 68030 chips can access a full 4 megabytes of virtual RAM, even though only 1 or 2 megabytes are present. Intel's 80386 theoretically can access 64,000 gigabytes (64 terabytes) of virtual memory. The main usage of virtual memory will be in support of multi-tasking processes for single and multiple users of such systems. Currently operating systems

such MS-DOS, PC-DOS and Apple/Finder do not support virtual memory. Operating systems such as UNIX and refinements such as the new versions of OS/2 and Apple's System 8.0 are in support of virtual memory on microcomputers.

What Are Some Emerging Hardware Developments that Offer Promise in the Evolution of Microcomputer Technology?

Great progress has been made in the development and refinement of existing components that make up the microcomputing system. Specific areas include new chip designs, including CPU refinement and RAM developments; marked progress in storage media development; the continued refinement of "supermicros"; and new architectures for future hardware platforms.

Developments in CPU refinement and RAM improvements have been discussed earlier. One of the more explosive growth areas across all levels of computing systems involves mass storage devices. The range of devices for microcomputers begins with floppy disks and culminates in such mass storage devices as CD-ROM. The cost per unit of stored byte has been reduced drastically since mass produced microcomputers were first introduced. The optical storage technologies associated with CD-ROM are purely microcomputer-based and CD-ROM access is not associated with computing systems beyond the microcomputer. While larger scale computing systems can access laser disk storage devices, only micros have been used to control access to the more popular CD-ROM products and devices.

A new line of powerful top end microcomputers, referred to as LAN (Local Area Network) servers, has been introduced into the marketplace. These systems support very large addressable internal and external memories, process data at very high rates, and are capable of hosting an interconnection of micros across a network via cabling and data exchange protocols. The LAN design holds great potential for the refinement of small scale microcomputer-based distributed applications.

Another major development is in the area of multi-processor systems, specifically in the areas of parallel processing. While personal computing will most likely continue to utilize a single 16- or 32-bit processor, perhaps working in tandem with a co-processor for mathematical computations, higher demands for computing power will likely be met using processors linked together in a parallel configuration. Each CPU is dedicated to a specific task, such as video display, general input/output or printer output. The CPUs share a centralized internal memory. Excellent processing benchmarks are associated with such designs.

Many of the emerging architectures for microcomputer systems units utilize a bus standard which supports such multi-processor designs. These architectures include IBM's Microchannel (MCA) as well as the microcomputer manufacturer's EISA standard. Some envision that future computers of all sizes may ultimately be composed of 32-bit CPUs operating in parallel. To increase system performance in one area, a single CPU is added with specific processing domains. If this becomes a reality, then the lines distinguishing micros from minis from mainframes will become increasingly blurred. For example, Intel has reportedly developed a prototype system which supports the parallel connection of thirty-two 80386 CPUs, yielding the kinds of performance associated with a top-of-the-line Cray mainframe. Parallel computing systems offer the potential for the development of very specialized, intelligent, shared database applications.

The field of library automation has at least one vendor currently offering a product based upon parallel computing. One CLSI turnkey system utilizes a Sequent parallel processor which supports the installation of multiple CPUs to accommodate growth as it relates to demand for increased processing capabilities. The specifications and benchmarks for this system indicate a marked increase in performance over more conventional single processor systems. Yet another vendor, The Library Corporation, states the following in its brochure for a linked circulation control module. "Dual 386/20 computers operating in parallel under a DOS or UNIX environment are supported by a network of distributed processors."

CONCLUSION

Link to Database Management Systems Applications

Database applications programs were first introduced when applications developers decided to treat individual blocks of data as self-contained units and further divide those units (called records) into named and addressable fields. In this manner, many similar records pertaining to a certain application (e.g., online catalogs) could be easily created, stored, edited and retrieved for various display or print purposes. Database applications which support those functions conducted by information professionals tend to be extremely demanding on computer resources, both internal and external.

Database management systems allow the user to create, edit, store and manipulate data of various forms in electronic files, much as one would create and maintain manual files on any given subject of interest. The major difference between manual and electronic data files lies in

the fact that electronic files are much more readily manipulated and searched than their manual counterparts. Combined with the fact that electronic files require much less physical storage space than manual files, one can begin to see numerous situational advantages of these automated database applications.

Many of the hardware developments discussed previously have had a major impact on the design of resulting microcomputer-based database applications. Chip technology is setting a rapid development pace, improving CPU performance and at the same time providing massive amounts of high speed internal addressable memory. Database applications tend to require very fast processors and have the need to address large internal memories, especially those applications which involve high transaction situations. Library applications are replete with high traffic opportunities such as online public access catalogs, circulation control systems, and automated reference services. In addition to requiring speedy resolution of events and procedures, these applications also demand very large storage capacities.

The development of mass density devices, such as high speed magnetic drives and high density optical drives, is directly in support of these requirements. Libraries and related information agencies are using microcomputers with high density magnetic hard drives capable of storing up to 314 megabytes. Many also access optical storage media capable of storing over 600 megabytes per unit. Multiple configurations of these units can currently provide gigabytes of external storage capacity.

Limitations

It seems as though these "chip and disk" implementations are well ahead of developments in systems and applications software. This is not an unusual phenomenon. But what are some limitations of microcomputers as hardware platforms for library-related applications? While the mass production and purchase of small computing systems certainly bring the cost per unit down, these personal computing systems do suffer from reliability and durability constraints. Initial systems were developed for use by a single person running a single application in RAM. When put through the paces of multi-user and multi-tasking applications which dominate the library marketplace, such systems perform differently than in the personal workplace. Library automation systems are required to run night and day in faultless fashion, without skipping a beat. They may be called upon to perform literally thousands of transactions in a very short period of time, say one day. The mean time between faults of such systems must be very long. While microcomputers are fairly simple to repair, their ability to perform day in and day out is suspect. Most personal-based systems were not designed

to take that sort of computing punishment. Despite all the efforts and progress made in developing such systems, the fact remains that microcomputers may not be the best performance purchases on the market today, at least not for larger library automation projects. Many library applications are multi-tasking, multi-user situations that require tremendous processing prowess, more than most current personal workstations have to offer.

Nevertheless, continued refinement of 32-bit technologies and maturity in terms of connectivity issues will provide increased alternatives for information professionals in years to come. As microcomputers improve their track record, they may evolve as a stable hardware platform for various database storage and retrieval applications such as those associated with the automation of library processes.