The evolution of printers and displays

by A. F. Mayadas
R. C. Durbeck
W. D. Hinsberg
J. M. McCrossin

Printer and display technologies have undergone remarkable changes since the beginning of the computer era. In this paper we trace the evolution of these two types of I/O devices, from the middle of the 1940s to the present, and show how computer system evolution has influenced the designs and technologies of I/O devices.

Much has been written about the remarkable evolution of computers, particularly the dramatic reduction in price/performance over the decades of their existence. To a very large extent, this price/performance reduction has stemmed from a peculiar facet of computer technology—that a good computer becomes an even better computer if the physical space required to store or process a *bit* of information is reduced. It is better not only because it takes up less space but also because it becomes easier to make this smaller computer run faster, use less power, cost less, and run more reliably. The steady, evolutionary miniaturization of logic circuits and memory bit cells has resulted in better processors and memory—the major internal components of computer systems. Magnetic tape and disk files, although not strictly to be considered as internal components of computers, have undergone a similar miniaturization of their magnetic bit cells and electronic elements.

Printers and display devices also have benefited from this same miniaturization, primarily through very large-scale integration (VLSI) of logic and memory. They, however, have an additional very important constraint on size—namely the need to interface with humans, termed the input/output, or I/O, interface. Clearly, because of this limitation, the I/O interface itself cannot simply be reduced to microscopic dimensions as a route to improved price/performance.

It is interesting to note that while I/O devices are classified as *peripheral*, their contribution to the cost or value of a computing system is anything but peripheral. The so-called peripherals, including disks, displays, and printers, account for more than half the cost of a typical IBM Personal Computer (IBM PC) configuration. For example, over 70 percent of the cost of a PC/XT (with 10 megabytes of disk storage, a letter-quality printer, and two displays) is for the peripherals. As to the value: What value would a personal computer have without a display or a printer?

Peripherals have not just recently achieved significance. Herman Goldstine's history[1] of the early days of computing contains a description of what must have been one of the first competitions between commercial machines. An IBM 701 machine and an Engineering Research Associates (ERA) 1103 computer were being considered for use in an application concerned with weather prediction. The two computers were tested with a meteorological calculation, and the evaluators concluded that, although the machines were approximately comparable in computation speed, the I/O equipment for the 701 was significantly faster; thus the IBM 701 was chosen.

Today the I/O function is performed to a large extent (though not exclusively) by printers and displays. In this paper we will trace the evolution of these two types of I/O devices in the broader context of the evolution of computer systems. Our main points will be the following:

- Printer and display developments have been characterized by the steady, incremental evolution of existing technologies, with the occasional appearance of a distinctive new approach. Today we see the result—a number of coexisting and still evolving printer and display technologies.
- The steady, incremental evolution of extant technologies has come primarily from better engineering and the increasing use of VLSI for adding function. (By the term *function*, we mean such attributes as multiple fonts, all-points addressability (APA), color, and gray scale. Usually, in the following, the meaning will be clear from the context.)
- Printer and display evolution has been driven by, and in some cases inhibited by, the pace of the concurrent evolution of computer systems. For example, on one hand, the introduction of the personal computer brought a new need for small and inexpensive displays and printers. On the other hand, although I/O technologies capable of APA have existed since early in the computer era, the widespread utilization of this potential was impeded in the period preceding LSI and VLSI by the high cost of processors and memory elements.
- Advances in printers and displays, unlike those involving CPUs (central processing units) or storage devices, cannot be usefully measured by tracking only one or even a few factors. There are many possible dimensions by which display and printer advancement can be measured. In this paper, we will consider such factors as speed, output quality, reliability, cost, size, acoustic performance, and level of function. A specific display or printer product can be considered to be a unique design point in this multidimensional world. The move from one design point to the next is determined by user requirements, technological advances, and the evolution of the computer system. At one point in time, for example, the goal for a new printer may be better printing quality; at another point, the goal may be higher speed or lower cost.

It is not our intent here to provide a comprehensive description of past and present printers and displays or their underlying technologies. We will illustrate our main points by way of examples, drawn primarily from the IBM experience. We will trace the evolution of printers and displays from the middle of the 1940s to the present and will make some projection of the future on the basis of that short history. Tables 1 and 2 list in chronological fashion some milestones in the history of IBM's printer and display technology.

During the past 40 years, computer system usage has evolved through several distinct eras:

- The early days, including batch processing (1945–1965)
- The development of multiprocessing and time-sharing (1965–1979)
- The arrival of personal computers and workstations (1979–present)

The advent of each era did not signal an immediate change in usage patterns from those of the previous period. In every case, existing usage patterns continued, albeit at a slower growth rate, while the new mode of use saw an explosive growth rate at first. This paper will present display and printer advances from the perspective of these three eras.

The early days, including batch processing (1945–1965)

The period from 1945 to 1965 was characterized by an intense focus on the CPU and on memory optimization. In the development of new computer systems, emphasis was on faster CPUs and memory. Batch processing and later spooling were incorporated into operating systems to allow more efficient use of the CPU. Hardware features, such as the asynchronous channel, were added even later to allow I/O operations to proceed without direct CPU involvement.

Early in this era, little effort was given to the development of printers specifically designed for computers because the existing printer technology, developed for electromechanical accounting machines, was considered acceptable. Later, when printer development did receive significant attention, the overriding goal was to maximize utilization of the expensive CPU resource through improvements in printer speed and reliability. Cost and added function were less important concerns. For practical purposes, displays were not used until late in this period, when a few vector-graphics units and prototype terminals appeared. These displays were generally costly and bulky and had little software support.

The plugboard-programmable accounting machines which were the forerunners of commercial computers were equipped with punched-card input devices and general-purpose electromechanical print-

Table 1 Significant events in the evolution of IBM printers

1873	Key and Ribbon Typewriter	Christopher Sholes, Remington Firearms Co.
1897	Teleprinter	Frederick Creed, Remote Newspaper Publication
1938	Xerography	Chester Carlson, Plain Paper Copying
1950	IBM 407	Accounting Machine, Rotating Typewheels, 150 lpm
1950	IBM 026	Printing Keypunch, Wire Matrix
1955	IBM 720/730	Wire Matrix, 1000 lpm
1960	IBM 1403	Engraved Character Chain Impact, Electronic Timing, 600 lpm
1964	IBM 1403-N1	Engraved Character Train Impact, 1100 lpm
1970	IBM 3211	Train Impact, Front Printing, 2000 lpm
1975	IBM 3800	Laser Electrophotographic, 13360 lpm
1976	IBM 6640	Synchronous Ink Jet
1979	IBM 6670	Laser Electrophotographic, 36 ppm
1984	IBM 4250	Electroerosion, Direct Negative Technology
1984	IBM 4248	Band Impact, Electronic Control, 3600 lpm
1984	IBM 5201	Resistive Ribbon Thermal Transfer, APA, 60 cps
1985	IBM 3812	LED Electrophotographic, 12 ppm, APA

Table 2 Significant events in the evolution of IBM displays

1875	Cathode-Ray Tube	Sir William Crookes
1926	Television Demonstrated	J. L. Baird
1957	IBM 610 Computer	Small CRT on Operator's Console
1964	IBM 2250	Display System, Vector Graphics, 1 MPel, Light Pen
1967	IBM 2260/2848	Alphanumeric Terminal, TV Technology
1971	IBM 3270	1920 Characters, Two-Dimensional Datastream, Local Intelligence
1974	IBM 3600	Banking System, Plasma Gas Panel, 480 Characters
1980	IBM 3279	High-Resolution Presentation Graphics, 7 Colors
1983	IBM 3290/581	AC Plasma Panel, 10K Characters, APA
1984	IBM 5080	Graphics Terminal, Realtime Image Processing

ers. The print mechanism of the IBM 407 Accounting Machine of 1950, for example, used 120 rotating typewheels, one at each print position, to generate printed output at speeds of 150 lines per minute. Even the earliest commercial computer systems, such as the IBM 650 and the IBM 701, could generate output and request input far faster than such I/O devices could handle the input or output. This limitation meant that in many cases the computational throughput capacity of the expensive CPU was substantially reduced by the I/O device. This condition was, and is, commonly referred to as the I/O *bottleneck.* In the extreme, if the printer was out of service, the CPU was effectively shut down. Because of this

I/O bottleneck, an early emphasis was placed on printer speed and reliability. For electronic data processing applications, this emphasis has continued essentially unchanged to the present day.

Batch processing, spooling, and asynchronous I/O. In the three decades since the introduction of the IBM 701 computer in 1956, CPU speed has increased by a factor of 2000.[2] For printer speed to have kept a comparable pace, a printer capable of producing 300 000 lines per minute would now be required. Such speeds are simply not possible, at realistic cost, with these fundamentally mechanical devices. Fortunately, changes have occurred in the hardware, the software, and the way in which we use computing systems, significantly easing the speed demands on the I/O devices.

Early implementations of spooling used a pairing of computers.

Batch processing was developed in an early effort to increase effective CPU throughput. Prior to this development, computer systems were operated with dedicated user scheduling, where a single user received total control of the machine during a specific time period. In batch processing, many tasks are stacked in an input stream and sequentially processed by the CPU. As each task is completed, execution of the next task in the input stream begins immediately.

This increase in CPU throughput, of course, put an even greater strain on the extant I/O devices, especially printers. *Spooling*[3] was developed as a systems solution to this problem. Early implementations of spooling used a pairing of computers, for example, the IBM 7090 and IBM 1401. Here the 1401 was used to collect an input stream containing *batch jobs* from a card reader. This stream was written onto magnetic tape, which then provided the input data stream for the main CPU (the 7090). After completing its computations, the main CPU wrote its output stream on magnetic tape which was physically returned to the satellite computer (the 1401) for actual printing. Although this operation may seem a bit complicated, it did serve to better optimize main CPU utilization by insulating the CPU from the slower and less reliable printers and card readers.

Later, the development of the *asynchronous I/O channel* and *disk storage* eliminated the need for a separate satellite computer. The asynchronous channel allowed data transfers to proceed in parallel with computing. The asynchronous I/O channel was first implemented in 1957 on the IBM 709 and became a standard feature of IBM computers thereafter. Magnetic disk storage, developed by IBM in 1956, provided a fast, random-access, high-capacity means of holding data until the printer was able to transfer it to paper. Disk storage has been a standard component of IBM's computers since the introduction of the System/360 in 1964.

Such system innovations were only partial solutions to the I/O bottleneck problem. Technological approaches toward greater speed and reliability in the I/O function continued to be pursued. These efforts led to the development of high-speed I/O devices such as the IBM 2540 card reader (800 cards per minute) and the IBM 1403 line printer.

The 1403 and 1403N1 printers. We will discuss here in some detail IBM's 1403 series of impact printers because they are prototypical of the printers of this period.

The IBM 1403, first introduced with the 1401 computer in 1960, was one of the first reliable, high-speed printers and was highly successful. Since output data in this era were restricted almost exclusively to alphanumeric characters, engraved printing with a fixed font was adequate for essentially all computer applications. The 1403 printer set important directions in impact line printer technology and was the first in an evolutionary series of high-speed impact line printers produced by IBM that extends to the present time (see Figure 1). Today's IBM 4248 Line Printer, described in detail later in this paper, is the latest and most advanced member of this family.

The early 1403s used a printhead magnet/armature which, when energized, drove a print hammer (supported by leaf springs) into the paper and ribbon and against the type slug from behind.[4] This method is known as *back printing*. Later versions incorporated a push rod to provide the kinematic link between actuator armature and hammer.[5] Type slugs were linked together into a chain (or assembled into a train in the later 1403-N1 printers) which moved at high speed in an endless loop. (See the inset on chains, trains, and bands.) This chain typically carried multiple 48-character sets. Paper movement was halted during printing to reduce vertical misregistration.[5] Printhead dynamics were optimized for minimum contact time in order to reduce the horizontal slur caused by relative motion between slug and paper during contact. The 1403 printer was capable of printing at speeds of 600 lines per minute, and the 1403-N1, through the use of the improved train technology, could operate at print speeds of 1100 lines per minute.[6]

The 1401 computer used a built-in controller to drive the 1403, and, as described earlier, many 1401s served solely as spooling I/O devices for a larger computer. Later, with the arrival of the System/360 computer series, a separate controller (the Model 2841) was used to drive the 1403 and other peripheral devices. The 2841 could control two printers plus an input device. Its interface to the printer used the then-new IBM I/O System/360 architecture.[7] Much of the 2841 controller architecture used in this I/O attachment remains with us today.

The first use of displays. Today it is hard to imagine working with a computer without a display. Early in the computer era, the potential and promise of computer displays were recognized by John von Neumann, who noted:

"In many cases the output really desired is not digital (presumably printed) but pictorial (graphed). In such situations the machine should graph it directly, especially because graphing can be done electronically and hence more quickly than printing. The natural output in such a case is an oscilloscope, i.e., a picture on its fluorescent screen. In some cases these pictures are wanted for permanent storage (i.e., they should be photographed); in others only visual inspection is desired. Both alternatives should be provided for."[1]

Not only was the potential utility of displays discerned in the early days of computers, but electronic technology suitable for constructing such devices was also known at that time: Television had been demonstrated in the 1920s, high-quality cathode-ray tubes (CRTs), such as those used in oscilloscopes, were commercially available by 1948, and even the electronic circuitry for the display of characters on a CRT had been developed by 1947.[8] Still, displays were a rarity on commercial computers of this period. Ex-

Figure 1 **Evolution of printer speed for IBM high-end printers**

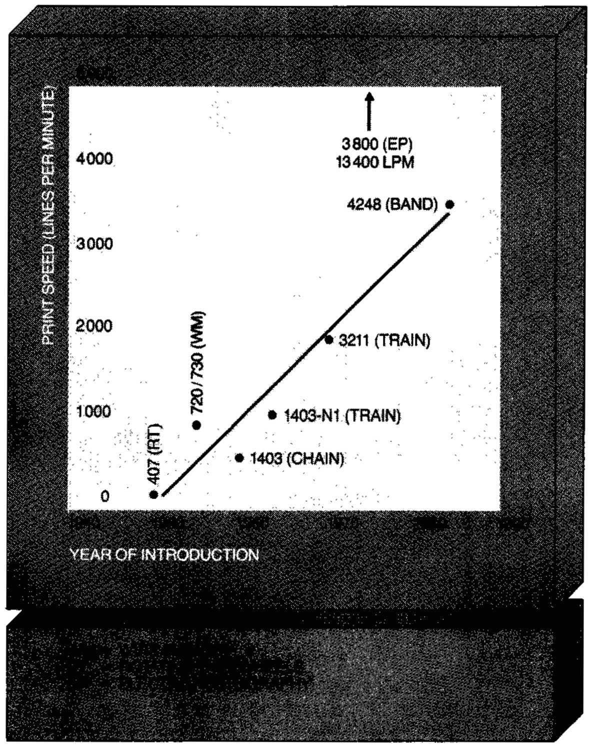

amples of the use of computer displays prior to 1960 are few.

The IBM 610 computer, an early example from the late 1950s, used a small (two-inch) CRT to cryptically[9] display the contents of its 13-digit accumulator—not a very exciting application and very cumbersome by today's standards. The most significant outcome of this display application was the invention by IBM of the cursor, which exists today on virtually every computer display device in use in the world.[10]

One of the first true computer graphics displays was the IBM 2250 display system, introduced in 1964. With over one million screen points, special function keys, and a keyboard plus a light pen and modest supporting software, the 2250 system proved to be a versatile graphics tool for technical applications.[11] The 2250 display had a product lifetime of almost 20 years, until superseded by the IBM 3250 Graphics Display System.

Another early CRT display terminal was the IBM 2260, introduced in 1968. This device incorporated char-

In *train* technology (used in the 1403-N1 and 3211 line printers, for example), the slugs travel in a rigid track and are not mechanically linked to one another. The slugs are moved along the track by a drive gear which engages machined teeth on the back of each carrier (see Figure 3). In the 3211, characters are not engraved directly on the train elements but instead are on pivoting levers suspended below the carriers and between the ribbon and the array of hammers. This arrangement is known as front printing.

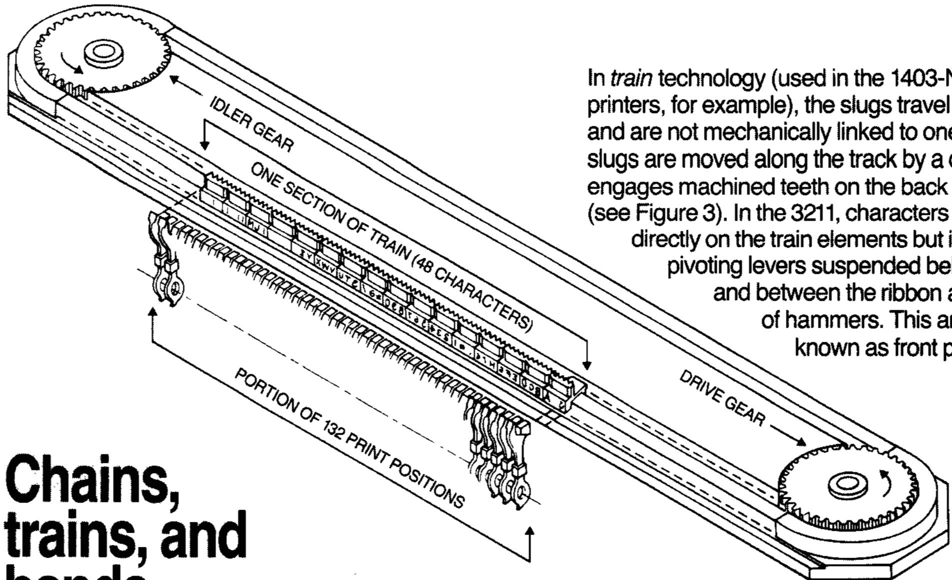

Figure 3 Train printing mechanism[4]

Chains, trains, and bands

In the evolution of IBM's high-end impact line printers, three different character carriers have predominated. *Chain* technology was used in the 1403 printer (Figure 2). Here a series of metal slugs, each engraved with alphanumeric characters, is interconnected by attachment to a continuous flexible metal loop. This "chain" is driven horizontally by a sprocket past the paper and ribbon. In the 1403, characters are printed by striking the paper from behind to contact the chain with an array of hammers (back printing). Chain technology throughput is limited by mechanical failure of the chain at high speed.

Figure 2 Chain printing mechanism[4]

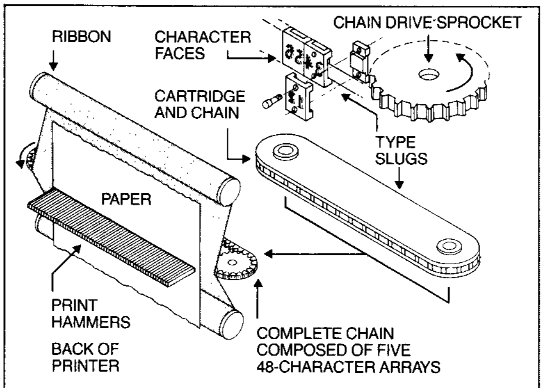

In *band* technology (such as is used in the 3262 and 4248 printers) raised characters are formed directly on a continuous flexible metal band (see Figure 4). This approach is less complex (a single part rather than hundreds of parts), less costly, and more reliable than train or chain carriers. The band is driven over pulleys rather than gears, and synchronization of hammer movement with character movement is established with etched timing marks. In the 4248, hammers strike the paper from behind, compressing the paper and ribbon against the engraved band.

Figure 4 Band printing mechanism

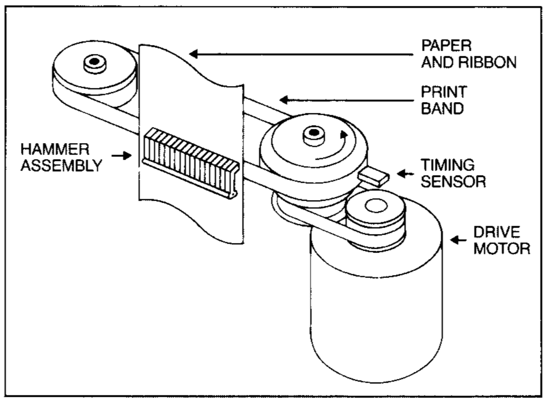

acter-generating hardware capable of displaying 960 characters on a full screen. The 2260 generally was used as an inquiry terminal for commercial applications and as a status-monitoring device for System/360 computers (replacing the ubiquitous typewriter terminals).

By the mid-1960s, the experimental use of CRT terminals for interactive computing was under investigation.

Such pioneering displays were derived from television technology, and the evolution of computer display devices has continued to be closely tied to technological developments in the consumer entertainment industry. Today the CRT remains the dominant display technology. The new functions appearing in today's displays, such as high resolution, APA, and color, are in part due to evolutionary improvements in CRT materials and design but are also the result of improvements in integrated circuit technology.

By the mid-1960s, the experimental use of CRT terminals for interactive computing already was under investigation, e.g., by IBM and by several universities.[12] The commercial implementation of on-line processing and time-sharing in the late 1960s would result in even greater focus on the I/O interface. The computer display would increase in importance as an I/O device, and the growth of distributed printing would make printer characteristics such as size, cost, and acoustic performance increasingly important.

Time-sharing and on-line processing (1965–1979)

In the era from 1965 to 1979, speed and reliability continued to be the primary characteristics demanded of printers used for data processing applications; in general, however, additional function remained a secondary concern. As computers and dedicated word processors began to be used for document preparation, improved print quality grew in importance, with Selectric® typewriter-quality print the goal. The nature of the time-sharing environment brought a need for printers positioned outside the central computing complex. As a result, such printer attributes as size, cost, and noise became increasingly important for such applications.

In the previous section we showed that displays originally were perceived to be graphics devices. In the period we are about to discuss, the use of displays became widespread, and they were used largely as alphanumeric devices *without* graphical capabilities.

The beginnings of multiprogramming. The typical computer system before 1965 had one main task at a time: *the job*. There was usually little need for more than one card reader to provide input and one alphanumeric printer to record the results. A typical large machine, at that time, might perform half a million instructions per second and have 512 000 bytes of memory for program and data. By 1978, both speed and memory size had increased by a factor of five. Such enhanced performance allowed computers to execute more than one job at a time, a process that became known as *multitasking* or *multiprogramming*. Multitasking was used in batch systems as a way of simply running two, three, or more jobs at once. More card readers and line printers were added to carry the extra load, and some of these additional card readers and printers soon were located outside the computing center, near the users' work locations. The IBM 2780 card reader/printer workstation, announced in 1967, provided card input and hardcopy output at moderate distances from the computing center via communication lines which operated with data transmission rates of 4800 bits per second. Although this remote output meant printing at speeds of only forty lines per minute instead of thousands, the value to the user of close access to the computer was significant. However, keypunches and card readers still were the dominant means of data input, and line printers were the usual means of output.

Progress in impact printing technology continued to provide improvements in print speed and quality in support of this mode of use. The latest improvements were evident in the IBM 3211 printer, introduced in 1970. Much effort and scientific study were put into understanding the physics of print-hammer dynamics to make these improvements. Electronic control was used extensively to monitor the high-speed movements of paper, train, and hammer. The 3211 featured printing speeds of up to 2000 lines per minute and improved reliability, in addition to providing better print quality.

In this era, minicomputers appeared in large numbers. However, they did not trigger a significant increase in I/O technology activities or products as did microcomputers in the late 1970s and in the 1980s. Minicomputers of this period typically used lower-performance devices based on I/O technologies developed for large computers.

The 3800 illustrates well the occasional appearance of a new I/O technology.

The IBM 3800 and 6670 printers. There is little question as to the motivation behind the IBM 3800 Printing Subsystem, introduced in 1976: It was simply the need for significantly greater speed than impact line printer technology was capable of delivering. The 3800 illustrates well the occasional appearance of a new I/O technology; i.e., it was quite clearly neither an evolutionary nor an incremental technology development, but rather an entirely new approach to printing.

The 3800 introduced laser electrophotography as a new printing technology for very high-speed applications. In the 3800 printer, the beam of a helium-neon laser scans a charged photoconductor drum by means of a rapidly rotating polygonal mirror, and by modulation of the light beam a latent image is constructed. This electrostatic image is then developed with thermoplastic toner particles, which are transferred to the paper and fused with heat and pressure. The printer is capable of producing 167 sheets (11-inch length) per minute. The high speed and throughput of the 3800 demanded a level of reliability much greater than that required of office copiers, and the 3800 has set new standards in this regard.

The 3800 produced output faster than any impact line printer could. More significantly, inherent in electrophotographic print technology is its greater capability for advanced function, which has moved it into application areas that are difficult for impact printing to address. The combination of APA and high speed, for example, makes electrophotography a very important printing technology, one that will assume even greater significance in the future.

The IBM 6670 Information Distributor, introduced in 1979, is a smaller, less costly, and lower-speed version of this laser electrophotographic technology. This printer was intended to be used as a distributed, departmental printer handling mainly word processing and/or document preparation applications. The 6670 offered very high-quality print, high speeds (up to 36 pages per minute), and intermixing of fonts. Later the capabilities of this printer were enhanced in the limited-edition APA6670, which allowed the printing of graphical and image data.

The beginnings of time-sharing. Christopher Strachey is credited[13] with first proposing time-sharing. The Compatible Time Sharing system was demonstrated at the Massachusetts Institute of Technology on an IBM 709 computer in late 1961. In 1963 this concept was made into a usable system on an IBM 7094 computer.

Time-sharing systems added a new dimension: direct access to computation from the user's normal work location. Time-sharing also reinforced the need for printers or terminals near the user. Typewriter terminals such as the IBM 2740/2741 were available for direct connection to the main computer. Systems now typically kept track of 30 to 50 jobs at once, with each job controlled interactively from a different terminal. The advantages of this approach over batch processing were immediately apparent. No more waiting in line for small calculations; the results were available immediately and right there in the office! The keyboard/printer combination could now be used for direct input of data to the computer.

At the beginning of this era, typewriters, not displays, remained the standard time-sharing terminal. In 1966, R. M. Fano and F. J. Corbato made an important observation:

"We must mention a bottleneck for which a practicable solution is not yet in sight. The output devices still leave a great deal to be desired. The teletypewriter is a frustratingly slow means of communication—and it cannot draw a picture. The graphical display devices that are currently available are ex-

Figure 5 Data stream of IBM 3270 display

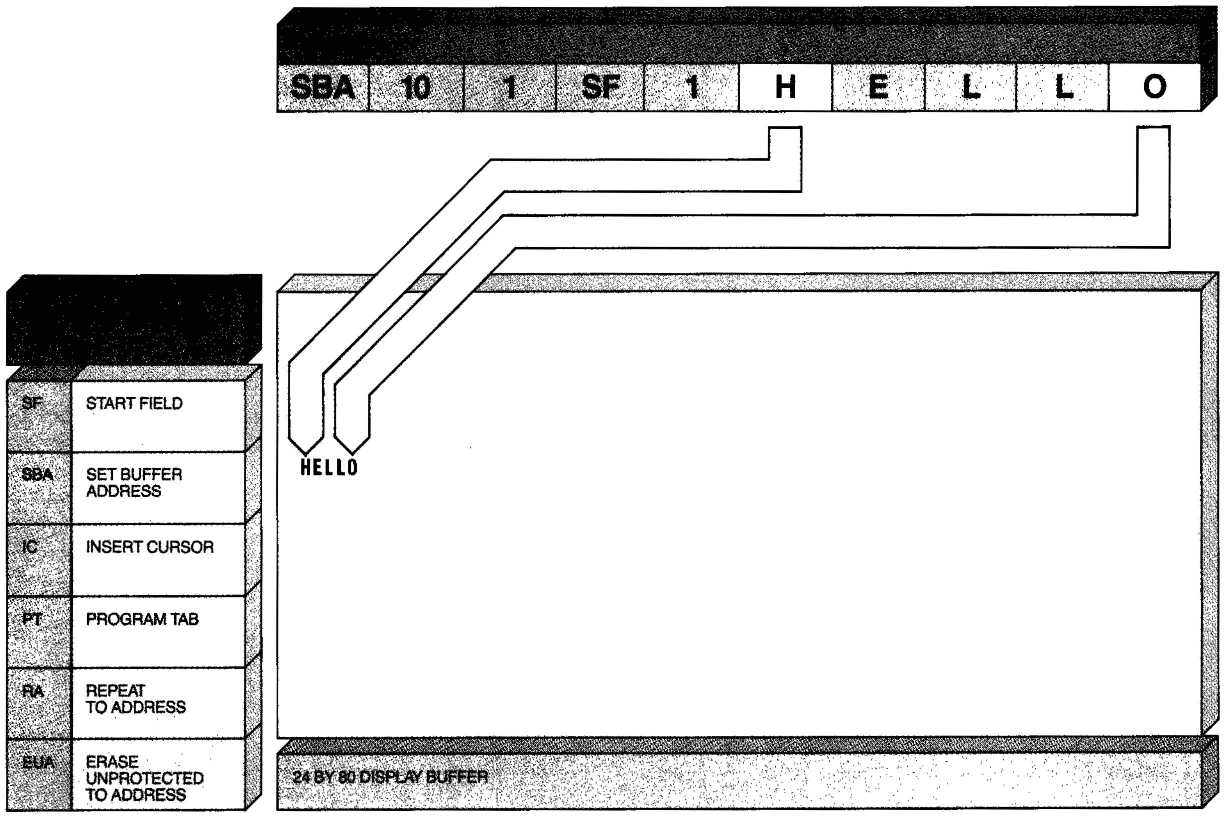

pensive and require elaborate communication facilities. Inasmuch as, from the standpoint of convenience and of economics, efficient communication between the time-sharing system and its users will become at least as important as the operation of the system itself, this problem presents a crucial challenge to designers."[13]

3270 Information Display System. The IBM 3270 Information Display System was announced in 1972. It offered the time-sharing user a new way to view text output—one that exhibited attributes of speed and silence, and which utilized two dimensions of information rather than the single-dimension, line-at-a-time method characteristic of card input, typewriter terminals, and line printers.

Typewriter terminals were attached via slow-speed communication lines, but the control unit of the 3270 attached directly to the channel of the computer with a high-speed coaxial cable connecting the control unit to the actual display head. An entire screen (24 lines of 80 characters) could be displayed in less than one second, whereas the typewriter terminal might take a couple of minutes to display the same information. This speed enhancement permitted a true conversational mode of computer usage. Not only did the typewriter take a long time to produce a page, it made a lot of noise in the process. The quietness of the 3270 display was an important added benefit.

Although the 3270 display could be used in the *line-of-input-line-of-output* mode characteristic of card readers and typewriter terminals, the structure of its data stream encouraged its use as a *two-dimensional* I/O device (see Figure 5). Each character could be individually addressed on the screen and updated. The display could be changed by the computer, but it could also be modified by the user through the keyboard, with the computer then reading the result. This capability made the 3270 screen an efficient two-way communication vehicle for both user and computer. The basic architecture of the data stream of the 3270 is still in use with today's displays (IBM's 3179 Color Display Station terminal, for example).

In Figure 5, Buffer Control orders allow the computer to address individually any position on the screen and to control certain characteristics of the display. For example, the Start Field order is used to set attributes of a particular portion, or field, of the screen buffer (e.g., normal intensity, highlight, or protect against erasure). The Set Buffer Address order is used to define the location of that field on the display. This figure schematically shows a data stream that writes an alphanumeric output message at a specific position on the screen.

The two-dimensional structure of the 3270 made possible new ways of interfacing with the user.

The two-dimensional structure of the 3270 made possible entirely new ways of interfacing with the user. For example, menus, which are two-dimensional structures, were used in IBM's Structured Program Facility, announced in 1977. The Time Sharing Option (TSO) Session manager, announced the same year, also made use of the two-dimensional capability of the 3270 to realize the concept of subwindows on the display screen. Full-screen editors used the two-dimensional aspect of the 3270 display, and in addition made good use of its *local intelligence*, which included buffering, keystroke-capture, and editing. The contents of the screen could be changed without involving the central computer. This use of local intelligence represented a first step toward actually distributing computing power physically in the office.

In summary, the widespread use of displays in this era is regarded as a major advance in the I/O area. It is one of the factors that made possible real computer time-sharing, and it set the stage for the display interface of the microcomputers and workstations that followed.

The arrival of the personal computer (1979–present)

The period from 1979 to the present has seen profound changes in printer and display design. The ever-present needs of speed and reliability are still with us, and they continue to drive the evolution of impact line printers in the machine room. However, for the first time there also is a demand for printers attached to large systems to incorporate advanced function. Further, the appearance of microcomputer systems in the business office in the late 1970s and early 1980s has had an important impact on I/O technology. The emergence of personal workstations brought a need for printers suitable for operation in a business office setting and in the home. To quantify the impact of the microcomputer, consider that in 1976, 16 percent of the U.S. printer market was for office and communications applications. By 1984 the market had grown by more than a factor of three, and more than half of this market was for office systems applications. Here size, noise, and cost are factors of primary importance, with the speed of printing assuming a secondary role.

The microcomputer also has had a major impact on display technology. For example, because display cost is an important consideration for workstation applications, substantial effort has been directed toward reducing the cost of CRT monitors. The compactness of the personal computer suggests portability, and this idea has spurred significant activity in the development of display technologies with form factors and ruggedness suitable for transportable and portable computers.

For both workstations and terminals, the use of CRT displays with advanced function (e.g., color) has grown rapidly in this period. The introduction of advanced microcomputers and workstations has made the display of high-quality images and graphics more available; no longer do such graphics-oriented applications as computer-aided design (CAD) require a mainframe or minicomputer.

Printers. Today's printer world can be characterized by the following observations: (1) in the mainframe and minicomputer areas, the 1403 line printer and its descendants have continued to evolve for data processing applications, and the technology exists in modern versions today; (2) electrophotographic printing, which started with the Xerox 1200 and the IBM 3800, is achieving more importance in mainframe and office applications; (3) microcomputer needs in the office and home environments have produced a number of competing technologies for low-end printers; and (4) the rapid evolution in silicon technology is integrally related to progress in printers.

The IBM 4248 Line Printer is a band printer, introduced in 1984, and is the most recent entry in the evolution of high-end, high-speed impact line printers that began with the 1403. In the 4248, characters are engraved on a single continuous band of steel, rather than on a linked chain or train as in the 1403 series, and a back printing method is used. (See inset on chains, trains, and bands.) The 4248 incorporates new refinements that allow improved reliability and availability while printing at higher speeds of up to 3600 lines per minute (corresponding to a printer band velocity of 44 miles per hour). To a large extent, these refinements involve the use of electronics to control mechanical printer functions. For example, semiautomatic hammer flight calibration has been implemented in the 4248, allowing corrections for drifts in flight time. In earlier line printers such an adjustment required a service call by a customer engineer and was a two-and-a-half-hour procedure. The calibration routine of the 4248 takes only five minutes and can be initiated by the machine operator. In another feature, the 4248 periodically tests its paper carriage dynamics and automatically retunes the carriage for optimum performance. Significant effort has also gone into further optimization of the print mechanism and driver circuitry. Through improved packaging and soundproofing, the 4248 is also one of the quietest impact line printers: A 4248 printing at 3600 lines per minute is quieter than a 3211 printing at 2000 lines per minute. The characteristics of the 4248 are compared to those of the earlier 1403 and 3211 line printers in Table 3.

The rapid progress in miniaturization of semiconductors has had an important impact on printers across the entire speed/performance spectrum. The availability of small and inexpensive memory and logic components has brought about improved reliability and greater function while reducing size and weight. All-points-addressable graphics, enhanced character sets, and additional fonts are commonplace today even in low-end printers, all made possible by the availability of low-cost memory devices and microprocessors. As we have pointed out in the case of the 4248 above, operations such as timing and tabbing, previously carried out by mechanical means, are now executed electronically.

Table 3 Comparison of IBM high-speed impact line printers

Model	1403	3211	4248
Year of introduction	1960	1970	1984
Character carrier	Chain	Train	Band
Carrier velocity	100 ips	225 ips	775 ips
Hammer head mass	2.2 gm	0.50 gm	0.25 gm
Impact velocity	105 ips	200 ips	250 ips
Maximum print speed	600 lpm	2000 lpm	3600 lpm
Acoustic performance	63.5 dB(A)	63.0 dB(A)	62.0 dB(A)
Price/performance*	2.8	1.8	1.0

* Price/performance is defined here as the rental price divided by the print speed.

In the Model 3 version of the IBM 3800 laser electrophotographic printer, the incorporation of semiconductor logic has brought about a fundamental change in the relationship between the host processor and printer. The 3800-3 has the ability to use a large number of fonts in a single document, to print raster images, and to merge output with stored electronic forms. The printer receives a series of commands from the host, composes the output in a "page build" process, and stores the completed page in a buffer until its print function is available. Further, the 3800 incorporates microprocessor control of system characteristics such as print contrast and error recovery which are the key to maintaining consistently high output quality and machine reliability while minimizing operator intervention.

The IBM 3812 Pageprinter represents an evolution from the high-end 3800 and APA6670 laser printers to low-end printers. In the Pageprinter a low-cost light-emitting-diode (LED) printhead replaces the laser imaging system. With the exception of speed, the 3812 incorporates most of the key functions of its more expensive predecessors: 240 × 240-pel-per-inch resolution, all-points-addressable graphics, merging of text and graphics, multiple fonts, and more.

As the applicability of microsystems to word processing requirements in the office environment became established, the demand for small, quiet, and inexpensive printers producing high-quality print on a variety of papers and having advanced features gained strength. The ensuing activities in printer

R2T2 printers

In conventional thermal transfer print technology, the printhead is an array of thin-film resistive heating elements. This array is placed in contact with the back side of a ribbon coated with a low-melting ink. To print a character, the ribbon is held in contact with the paper and the correct patterns of electrodes are energized as the printhead moves across the ribbon. The result is transfer of the melted ink to paper. Since the heated elements must cool down before moving to the next print position, the long thermal cycle times of the print electrodes effectively limit the speed of such thermal transfer printers. Relatively low-melting-point inks are used in efforts to reduce the cycle times, but such materials are subject to unintentional transfer by pressure or contact. Further, with such inks the print quality is dependent on paper roughness, and special, highly calendered paper is required for high-quality print.

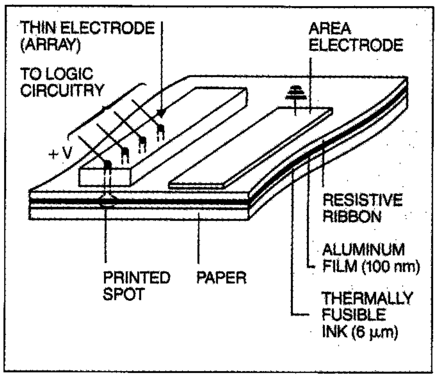

Figure 6 Schematic of R2T2 technology[14]

Resistive ribbon thermal transfer (R2T2) print technology has evolved from the above thermal transfer approach. R2T2 eliminates most drawbacks associated with thermal transfer printing. In R2T2, the ribbon is a multilayer structure that itself is a resistive heating element (see Figure 6). Localized heating by an array of small electrodes causes melting of a polymeric ink film on the opposite side of the ribbon which then transfers to the paper. Since the ribbon and not the printhead is heated, printhead thermal cycle time is a much less important factor in determining the maximum print speed. The heated zone of the resistive ribbon is moved away from the print area by advancing the printhead across the ribbon. Higher-melting-point inks with improved mechanical properties can be used in R2T2, allowing high-quality printing on a wide range of papers. The basic components of the R2T2 technology as implemented in IBM's Quietwriter printer are shown in Figure 7.

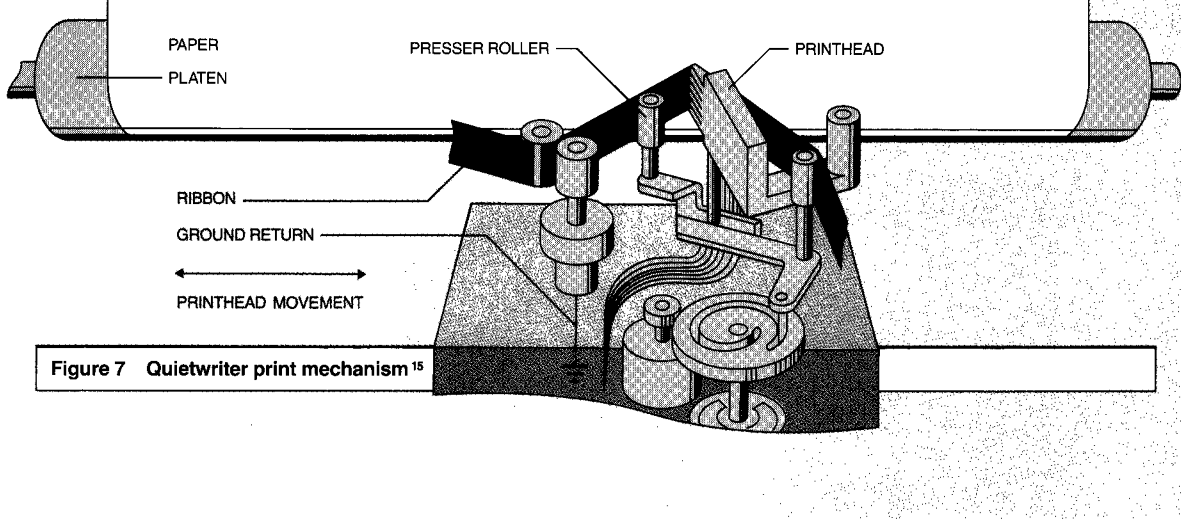

Figure 7 Quietwriter print mechanism[15]

development have brought significant and rapid progress to impact printing (e.g., in the development of daisywheel technology and in the refinement of wire-matrix impact printing) and to such nonimpact technologies as electrophotography, ink jet, and thermal transfer. Inherent in all these marking technologies are trade-offs among speed, print quality, and cost.

New approaches derived from earlier print technologies can bring together quality print and high speed at lower cost. In thermal transfer printing, for example, the printhead is an array of small resistive heaters held to the back of a ribbon coated with a waxy ink; when the ribbon is locally heated, the ink melts and is transferred to paper. The speed of printing is inherently limited by the time required for the heater to cool down between print cycles. To overcome this speed limitation, IBM's Quietwriter printer uses resistive ribbon thermal transfer (R2T2) technology,[14,15] in which the conductive ribbon medium itself is resistively heated to effect ink transfer. (See the inset on R2T2.) The printhead itself does not become appreciably hot, so print speed is not limited by its thermal cycle time. The Quietwriter can print at speeds of up to 60 characters per second.

Displays. Displays have undergone a parallel evolution during this time. Several areas of significant progress can be identified:

- Higher information content displays
- Enhanced graphics capabilities
- Improved form factors and weight reduction
- Lower costs

Advances in display technology generally appear first in special-application products such as computer-aided engineering/computer-aided design (CAE/CAD) workstations and military systems, to be incorporated later at some level in more general-purpose display products.

As in printers, progress in LSI and VLSI can be identified as a key ingredient in this evolution (see Figure 8). Through the use of dedicated microprocessors and controllers specifically designed for graphics applications, manipulation of graphical data can be carried out largely at the terminal, without heavy interaction with the host computer. Enhanced function and greater speed are the benefits of this local intelligence. These improvements in controller electronics have been accompanied by a parallel evolution in CRT materials and design which has led to enhanced resolution and color capability. An advanced display such as the IBM 5080 Graphics System incorporates multicolor capability and large amounts of semiconductor memory as screen buffers (in this case a 1024 × 1024 × 8-bit plane buffer). Its control unit provides real-time manipulation of displayed graphical data, performing functions such as image rotation, scaling, and area-fill. Even at the low end, the monochrome alphanumeric terminals of the 1970s have advanced to more sophisticated APA displays with full color capability. (See inset on display architecture and memory use for further discussion.)

Figure 8 Evolution of video memory in IBM displays (the size of the video buffer in the IBM 3277, 3279, and 3179 display products is shown)

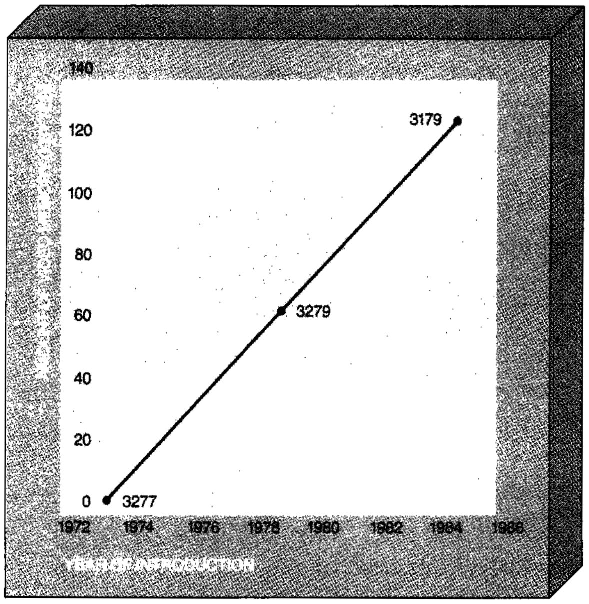

A unique driving force for the refinement of display technology has been its close connection with the consumer market. Both technical enhancements and cost reductions realized in the development of a consumer product can frequently be applied in computer displays. For example, the vector-scan and storage-tube display technologies, widely used for computer graphics in the past, are now giving way in many applications to full-color raster-scan CRT technology derived from the television industry. In a similar vein, we can now see the technical expertise gained through (a) the development of liquid-crystal displays for watches, calculators, and automotive applications, (b) the use of gas-panel, electrolumi-

Display architecture and memory use

With the exception of those monitors using storage tube technology, most CRT displays incorporate a memory buffer that contains, in some form, the current contents of the screen. This buffer provides for the periodic redrawing or *refreshing* of the screen image. It is of interest to examine how the data stored in this memory are related to the screen image, and how this relation has changed over time. As we have pointed out, an important factor in the evolution of displays has been the increasing availability of inexpensive semiconductor memory—both Read/Write (R/W) memory for the screen-refresh buffer, and Read-Only Memory (ROM) for fonts and for display logic.

Vector displays

Early CRT displays often used a pattern of small vertical, horizontal, and diagonal strokes, or *vectors*, of the electron beam upon the screen to create characters and symbols. Here the beam acts as an "electronic pencil," leaving a visible trace of its path on the screen phosphor (see Figure 9A). To display a character, a data element representing that character in some coded format is placed in the screen-refresh buffer. At the appropriate time, a character generator within the display circuitry reads this buffer and transmits to the CRT driver a predetermined set of instructions (stored in ROM) defining the beam deflections which draw the character on the screen. Line drawings, graphs, and other symbols can be drawn in an analogous manner by directly generating a list of beam movement instructions. For the refresh buffer, such displays used magnetic core memory or delay lines, which were both expensive and limited in performance when compared to today's semiconductor memory.

Alphanumeric raster displays

For alphanumeric applications, raster-scan CRT displays have become the dominant technology. Here the electron beam repeatedly traces a path of closely spaced parallel lines across the screen, creating figures by modulating the electron beam intensity at specified positions during the trace. To display characters on the screen, data elements representing the characters (again in a coded format) are placed at the correct offset in the refresh buffer. Then, in synchronization with the screen refresh, a character generator reads the buffer, translates these data elements into pel patterns using information stored in ROM, and sends the patterns (which define the on or off state of the beam as it sweeps across the screen) to the CRT driver. The screen representation of an alphanumeric character is composed of a two-dimensional array of spots or pels (see Figure 9B). In a typical implementation, this array might be 9 by 14 pels, requiring 16 bytes of ROM to define its screen image. The data element representing a single character in the refresh buffer is often one byte, so a typical 80-column by 25-line display might incorporate two kilobytes of ROM for character gen-

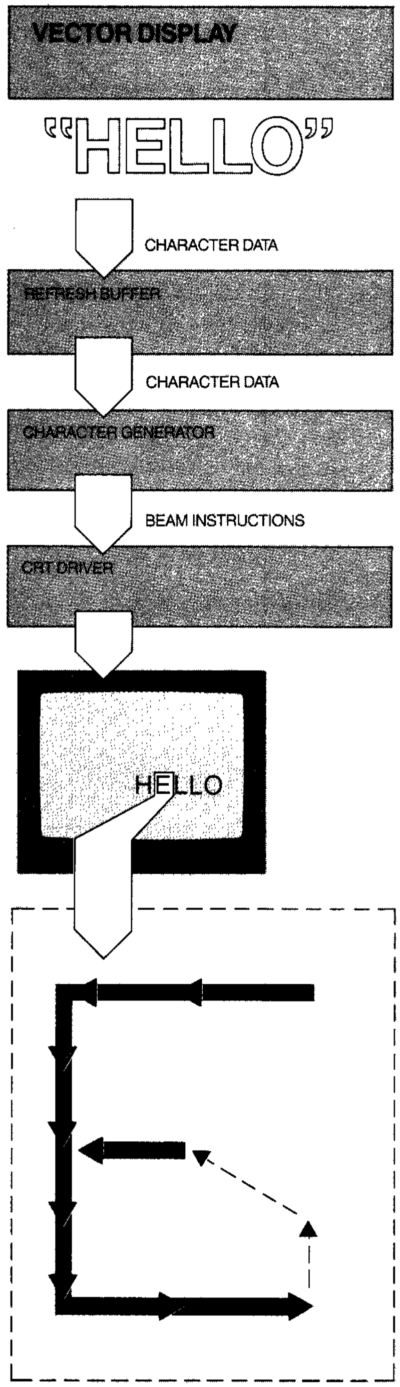

Figure 9 Types of displays

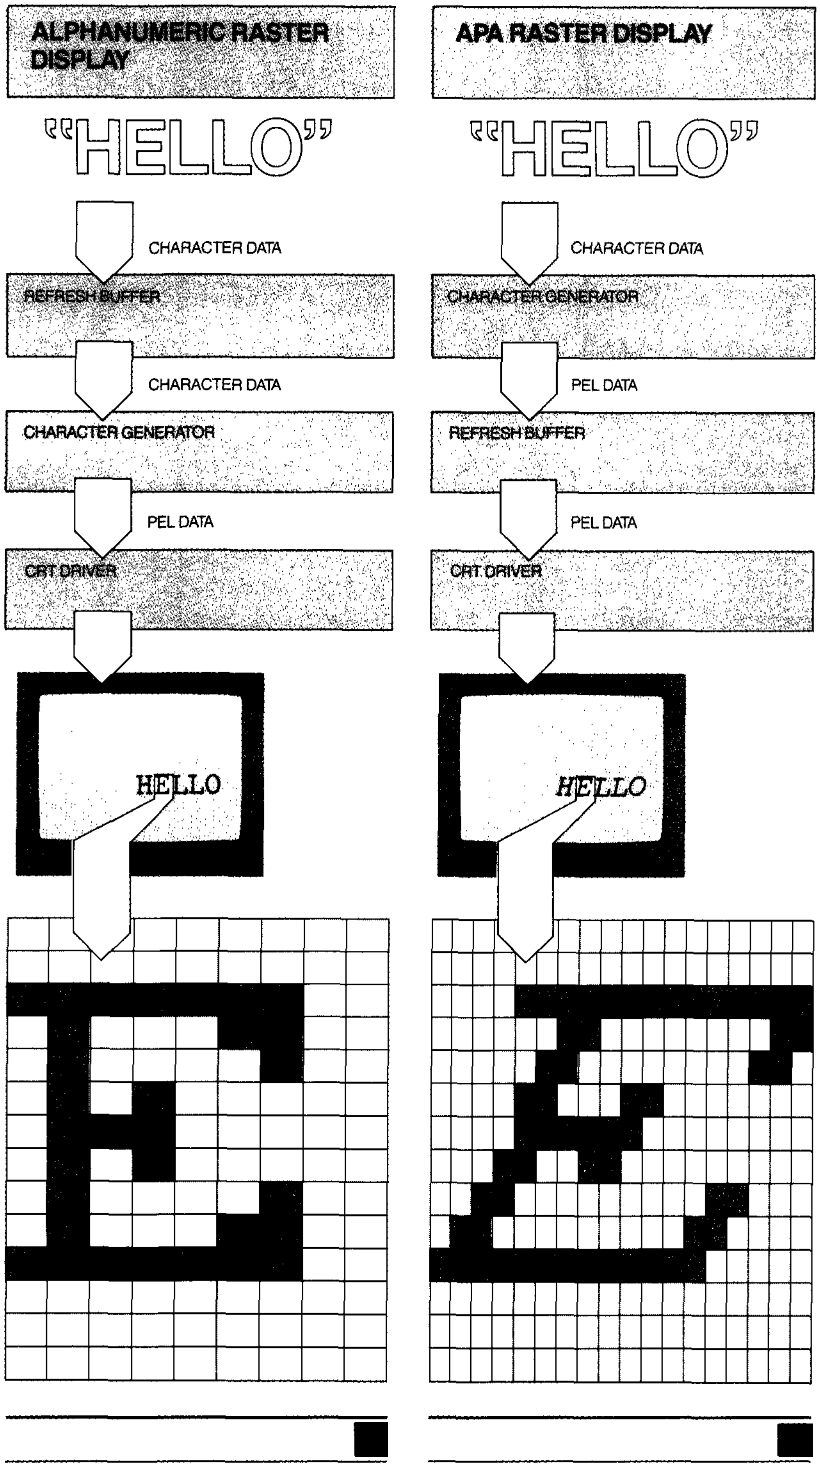

eration and a two-kilobyte refresh buffer of R/W memory. The extension of this approach to a limited form of graphics is achieved by the use of a *programmable character set,* which replaces the pel patterns stored in ROM used for character generation with a set of user-defined pel patterns in R/W memory.

APA displays

In many of today's raster-scan CRT displays, a value for each point on the screen is stored as one or more bits in the screen-refresh buffer. Termed a *bit-mapped* display, in this case the refresh buffer stores a true image of the screen rather than a character representation that must be translated into bit patterns to produce the screen image. The translation of character data into a bit array in the buffer is carried out by a character generator external to the *screen-refresh circuitry* (see Figure 9C). Because each pel can be set on or off independently, any possible screen pattern can be generated, allowing full APA graphics and characters to be freely intermixed. In general, an APA display of this sort requires significantly more R/W memory than an alphanumeric display–typically from 32 kilobytes for a small-screen monochrome implementation up to two megabytes for a large-screen color APA display. Another 8K of ROM might be used for character and graphics generation logic. The increased availability of APA displays is directly related to the reduced cost and increased performance of semiconductors.

nescent, and vacuum-fluorescent display technology in military and other special applications, and (c) progress in "flat" television technology beginning to emerge in future computer flat-panel display technologies.

The desire for true portability in microcomputer systems, and for improved form factors of office workstations, has heightened activity in the development of flat information displays.[16] A noteworthy advance in this field has been the IBM 3290 Information Panel, a high-information-content display using AC plasma technology. In this display, luminescence is produced by an electrical discharge in a neon gas mixture. Horizontal and vertical sets of parallel electrodes within the device define a two-dimensional array of discharge points which can be individually energized to generate the screen image. With a capacity of 10K alphanumeric characters, the 3290 allows concurrent viewing of up to four standard full-screen applications. Features such as a total of 737K pels and all-points addressability make the 3290 display suitable for technical applications requiring simultaneous text and graphics.

Trends

We have traced the evolution of printers and displays over four decades and have shown how computer system evolution has influenced the design and technology of I/O devices. We have seen that incremental improvement in established technologies is an important part of this evolution. We have noted that important new technologies occasionally appear, and that their significance is enhanced if they offer a new capability or an increased function. We have also pointed out the important impact of VLSI in I/O technology. Finally, we have noted that, although the presence of minicomputers had little influence on I/O technology, the introduction of microcomputers into the office environment has been responsible for the growth in importance of a number of new technological approaches in the printer area.

In general, we can expect in the near future that all categories of I/O technology will continue migrating toward increasing function. To a large extent this migration will be linked to advances in other areas of computer technology. For example, current trends in silicon technology will lead to further cost reductions and increasing computational power, and the impact of this on future I/O technology will be severalfold. First, both large and small computer systems will continue to evolve into faster, more powerful machines incorporating significantly larger memory and storage. The capabilities of these systems to deal with larger amounts of information with greater speed will require that future I/O devices have the ability to receive and transmit larger and more complex data streams. Second, more powerful data-handling capabilities will be available within the I/O device itself. Printers and displays will, and for some advanced applications already do, incorporate their own computing capability. Already at the low end there are cases where the printer has more computing power than the system to which it is attached. Third, we will see a further shift toward electronic control of printer mechanical functions to improve performance and reliability.

Electrophotographic print technology will grow in importance.

Printer technologies. In the future, electrophotographic print technology will grow in importance. For large computer systems, the significant speed advantage of electrophotography compared to impact printing is an important characteristic. Equally important is the ability of electrophotography to accommodate advanced functions such as APA and gray scale or halftone more readily than impact printing. Nonetheless, impact line printers will continue to be important components of the high-end data processing printer market, a consequence of the great reliability of this technology and its ability to print multipart forms.

A continuation and expansion in the use of distributed and departmental printers is also expected for the future. Electrophotographic technology is likely to be dominant here, but ink-jet technology may also find an application, particularly for color printing.

In the workstation printer segment of the market, which is expected to show the fastest growth, there currently is tremendous conflict between competing marking technologies, and at this time it is not clear which technologies will dominate. Each print technology has unique strengths and weaknesses, the relative importance of which varies with application. Nonimpact technologies have the advantages of being intrinsically better adapted for higher resolution and APA and of being quiet, but impact printing continues to improve in these factors also. Again, for applications requiring multipart forms, impact printing will have continued importance.

Display technologies. The CRT display will continue to be an important technology in the future. Incremental improvements in electron beam optics, phosphors, and resolution, and new approaches such as multiple-beam devices, will improve CRT usefulness for advanced applications such as computer-aided design and engineering and in-house publishing. In such applications, display resolution and information content comparable to those of a high-quality APA printer are desired. This is a specific requirement for true WYSIWYG ("what you see is what you get") in complex document composition applications. Several flat-panel technologies, in particular those based on liquid crystals, appear to be becoming even more commercially attractive. Flat panels will continue to find application primarily as compact, rugged displays for use with truly portable microcomputers, and for use on the executive's desk. The wide-scale use of flat-panel display technology in office workstations and general-use terminals will probably require flat-panel costs to be close to those of the CRT.

I/O systems. The movement toward true WYSIWYG and the trend toward interconnection of both large and small systems are creating a need for standardized, compatible data streams across these system types and between I/O devices. In addition, data compression/decompression algorithms and hardware will be important in the future for transmitting these more complex data streams between devices. Another important goal for the future will be the development of software that will improve the ease of use of I/O devices.

Summary

Until recently, the quality of printed or displayed material generated by computer has generally been less than that found in books, photographs, and magazines. We accepted that poor quality because the limitations of the computers, printers, and displays simply did not allow the creation of high-quality output. Today, computer and I/O technologies are close to the point where they will become associated with quality text and graphics and with the realistic representation of images.

In general, the printer and display technology areas continue to be dynamic and evolving. We expect to see improvements in cost, quality, and function continue at a rate at least as great as we have experienced over the last decade. This improvement, in turn, will lead to greater ease of use and greater productivity, and will mean greater acceptance of these devices in the everyday world.

Acknowledgments

We thank the many people within IBM who shared their knowledge and experience with us in the preparation of this overview, and those who critically read drafts of the manuscript. These include R. Williams, J. King, and D. Rosenheim of the IBM Almaden Research Center, J. Zable, M. Clark, and D. Jensen of IBM Endicott, R. Aylsworth of IBM San Jose, and John Lentz, formerly at IBM's Thomas J. Watson Research Center.

Cited references and notes

1. H. H. Goldstine, *The Computer from Pascal to von Neumann*, Princeton University Press, San Francisco (1972).
2. Actually, this estimate is conservative. The IBM 701 was rated at 14000 instructions per second, and the IBM 3090 is rated at 28 to 30 million instructions per second, an increase by a factor of 2000. In addition, the 3090 instruction set is more powerful than that of the 701, so the 701 could not do the 14000 3090 instructions per second.
3. In spooling, an image of the data to be printed is received from the CPU by a high-speed storage device such as a tape drive. At a later time, this image is directed to the printer, which is operating independently of the main computer. Since the CPU directly interacts only with a high-speed device, the time spent waiting for output data to be read from its buffer is reduced, and utilization of the CPU is improved. A similar scheme can be implemented for the input function.
4. T. Y. Nickel and F. J. Kania, "Printer technology in IBM," *IBM Journal of Research and Development* **25,** 755–765 (1981). Figures 2 and 3 reprinted with permission from this source.
5. *IBM 1403 Field Engineering Manual: Theory of Operation* (Second Edition), IBM Corporation (1973).
6. Earlier wire matrix printers, such as the IBM 720 and 730 printers, could produce up to 1000 lines per minute. These machines, first used with the 700 series computers in 1955, had a 5×7 array of print wires at each print position.[4] Though this technology had speed and the potential for all-points addressability, it was set aside in the late 1950s for reasons of reliability and print quality.
7. The System/360 architecture used eight channel control words (CCWs): Sense Status, Control, Conditional Transfer/Branch, Write, Read, Read Backwards, Invalid, and No-op. Each eight-byte Write CCW consisted of a command byte, a three-byte

data address field, a flag byte, an unused byte, and a two-byte count. The data were almost entirely simple strings of characters, sent one line at a time. I/O operations began with a START I/O and a pointer to the first CCW. CCWs could be chained together to improve operating efficiency by using a control bit in the CCW. An interrupt signal was directed to the computer when the channel was completed, and also for device end. At channel completion, the Channel Status Word (CSW) was stored in a fixed memory location in the CPU where it was available for interrogation.

8. H. Lukoff, *From Dits to Bits: A Personal History of the Electronic Computer*, Robotics Press, Portland, OR (1979).
9. Numeric values were not displayed as Arabic numerals; rather, the horizontal positions of dots on the screen were used to indicate the value in the register.
10. It is interesting to note that one early computer, the IBM 701, depended on a CRT not as a display but as its internal memory. Binary *1*s and *0*s were stored electrostatically by manipulating spots on the face of the CRT screen. The screen was mounted behind glass such that one could actually watch the memory locations change in real time.
11. I. E. Sutherland, "Computer input and output," in *Computers and Computation*, W. H. Freeman and Company, San Francisco (1971), pp. 42–69.
12. P. Suppes, "The uses of computers in education," in *Computers and Computation*, W. H. Freeman and Company, San Francisco (1971), pp. 249–259.
13. R. M. Fano and F. J. Corbato, "Time-sharing on computers," in *Computers and Computation*, W. H. Freeman and Company, San Francisco (1971), p. 79.
14. K. S. Pennington and W. Crooks, "Resistive ribbon thermal transfer printing: A historical review and introduction to a new printing technology," *IBM Journal of Research and Development* **29**, 449–458 (1985). Figure 6 reprinted with permission from this source.
15. A. Bohnhoff, D. Croley, S. Dyer, T. Green, R. Maddox, and L. Struttmann, "System controls for a resistive ribbon printer," *IBM Journal of Research and Development* **29**, 494–508 (1985). Figure 7 reprinted with permission from this source.
16. L. E. Tannas, *Flat-Panel Displays and CRTs*, Van Nostrand Reinhold, New York (1985), pp. 1–30.

A. Frank Mayadas *IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120.* After receiving a Ph.D. in engineering physics from Cornell University in 1965, Dr. Mayadas joined IBM at the Thomas J. Watson Research Center, where he held assignments in basic and applied research areas. In 1979 he transferred to IBM's San Jose Research Laboratory to start up and manage a new function called Storage Systems and Technology. In 1981 he returned to the Watson Research Center as Director of Technical Planning and Controls for the division. In July 1983 he was named Research Division Vice President, Storage and I/O, and Director of the San Jose Research Laboratory (now known as the Almaden Research Center). Dr. Mayadas has published in the areas of thin film physics and magnetic and superconducting devices. He is a member of the IEEE Computer Society and the American Physical Society. He is also a member of the Advisory Board for the Georgia Tech Research Institute and of the Advisory Committee for the Department of Engineering and Applied Physics at Cornell, and an alternate member of the Industrial Research Institute.

Robert C. Durbeck *IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120.* Dr. Durbeck joined IBM in 1958 to develop advanced tape library systems at Poughkeepsie, New York. In 1961, he was awarded an IBM Fellowship to pursue his doctoral studies. After he joined the San Jose Research Laboratory in 1964, his initial technical assignments included research in magnetic recording, large-scale integrated power systems, and automatic control of large industrial processes. He received an Outstanding Contribution Award for his work in computer control of power system networks. He has managed projects in computer control, magnetic recording media and devices; ablative, phase-change, and magneto-optic optical storage media; liquid-crystal, electroluminescent, and TFT-LCD displays; ink-jet and electrophotographic laser transfer printing; and micromechanical silicon devices, plus other selected fields such as semiconductor module cooling and image processing for printing. He is now the area manager of I/O Science and Technology. Dr. Durbeck received his Ph.D. in electrical engineering at the Case Institute of Technology in 1965. He is a member of the IEEE Computer Society, the Society for Information Display, the Optical Society of America, and the American Society of Mechanical Engineers. He is presently Chairman of the SID International Conference Advisory Committee and a member of the Hard Copy Committee.

William D. Hinsberg *IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120.* Dr. Hinsberg is a Research Staff Member in Polymer Science and Technology at the Almaden Research Center. He received his B.S. in chemistry from the University of Michigan in 1975 and his Ph.D. in organic chemistry from the California Institute of Technology in 1979. A postdoctoral position in the Chemistry Department of Stanford University followed. He joined IBM in 1982 as a staff engineer in the Materials Laboratory of the General Products Division. He has published in the areas of mechanistic organic chemistry, bioanalytical chemistry, and microlithography. Dr. Hinsberg has received an IBM Outstanding Technical Achievement Award for his work in photoresist technology. He is a member of the American Chemical Society and the Society of Photo-Optical Instrumentation Engineers, and is also a member of the Advisory Board of the Institute of Photochemical Sciences of Bowling Green State University.

James M. McCrossin *IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120.* Mr. McCrossin is a manager in the Computer Science Department with experience in systems programming and management. Management responsibilities have specifically been in the areas of computing centers, product development, and research. He has been extensively involved in software for alphameric and APA displays. Currently, he is the manager of a group active in the area of I/O servers, which is examining the interfacing of I/O devices, such as printers and scanners, in the local-area network environment. Mr. McCrossin has held assignments in all three of IBM's research laboratories in San Jose, California, Yorktown Heights, New York, and Zurich, Switzerland.

Reprint Order No. G321-5283.