Optical computer technology can speed up the SDI

Computer specialist Kevin L. Zondervan reports on the breakthroughs at hand in high-speed processing.

The advent of the Strategic Defense Initiative (SDI) has dramatically quickened the pace of research and development in the area of high-speed computer processors. Technologies currently under study at Air Force laboratories include VHSIC (very high speed integrated circuits), optical processing, hybrid optical/digital processing, and GaAs (gallium arsenide) semiconductor technologies. Why the acceleration of R&D into these areas? In large part it is because experts have estimated that the processing rates for SDI battle management and sensor systems may approach 1 trillion flops (floating-point operations, such as addition and multiplication, per second). For comparison, the Cray X-MP, one of the world's fastest computers, has been measured at 33 million flops.

Typical functions performed by SDI battle management and sensor systems include sensor signal processing, target designation and discrimination, target trajectory estimation, weapon allocation (deciding which weapons are to shoot at which targets), weapon pointing, target tracking, and weapon firing. To fulfill its mission of shooting down ballistic missiles and nuclear warheads, an SDI system must perform all these functions simultaneously throughout the course of an engagement. However, the processing rates this entails, while allowing an SDI system to perform its primary goal, may not be sufficient also to ensure the defensibility and cost effectiveness of the system.

To achieve these additional goals, and thereby guarantee the success of SDI, the processing speeds may have to be substantially increased. This is true for primarily two reasons. First, the greater the processing speed, the more sophisticated the battle management, strategy, and tactics can be that the system can implement. The ability to implement a greater variety of tactics may be the decisive edge for the SDI if it is pitted against an opposing system of roughly equivalent firepower.

Second, greater processing speeds also can enhance the system's relative firepower and cost exchange. For example, it has been demonstrated that, if a space-based laser boost-phase defense system can achieve laser firing rates of 10 shots per second or better, then the number of space-based lasers required in orbit varies as the square root of the number of boosters. (In other words, if 100 space-based lasers are required to shoot down 3,000 boosters, then only 141 space-based lasers are required to shoot down 6,000 boosters.)

The faster the battle commands can be processed, the more likely the 10 shots-per-second firing rate can be achieved. A square-root scaling effect essentially removes the option of proliferation from the booster force's list of countermeasures, since the space-based laser force can proliferate at a lower rate to maintain its effectiveness. Put another way, the space-based laser inevitably wins the cost exchange under these circumstances. As the laser firing rate falls below 10 shots per second, the scaling becomes a linear instead of a square root relationship, and the space-based laser loses its edge in the numbers game.

Because it is critical to the success of the SDI to explore every promising avenue for increasing the speed of computers, the SDI Office's Innovative Science and Technology Office has formed consortia composed of government, academia, and industry to accelerate R&D in this area. Of particular significance was the formation of a consortium in

April 1985 to investigate high-speed computing using optical signal processing.

According to the SDI Office, the optical signal processing group will conduct research leading to "major advances" in signal and image processing technology. They will focus on "high-risk research aimed at obtaining breakthroughs in the near and mid-term." Two major thrusts have been defined: development of innovative hybrid analog and digital optical/electronic processors and also new techniques for addressing complex nonlinear problems essential for SDI operations.

Participating in this group are the California Institute of Technology, Carnegie-Mellon University, Georgia Institute of Technology, University of Dayton Research Institute, University of Alabama at Huntsville, Stanford University, the U.S. Naval Ocean Systems Center, MIT-Lincoln Laboratories, and Battelle Columbus Laboratories. Dr. Keith Bromley of the Naval Ocean Systems Center is the government program manager, and Dr. John Caulfield of the University of Alabama is the technical director.

The designated consortium associates are Aerodyne Corp., BDM Corp., Harvard University, Houston Associates, MIT, Probe Inc., SAI Corp., the University of California at Irvine, and the University of Southern California. The budget for this work is $9 million over the next three years, with fiscal year 1985 funding being $1 million.

Over the last 20 years, there has been considerable research in the general area of information processing by optical techniques, or to use the more recent terminology, optical computing. Generally speaking, this work involves the idea of using light beams or photons, instead of electrical currents, to perform numerical computations.

The reasons for using light instead of electrical currents are many. Because light is at the upper end of the electromagnetic spectrum, a light signal can be encoded with much more information than lower-frequency forms of radiation. This property can be exploited to increase the density of independent channels in a computing system. Optical signals can also propagate through each other in separate channels with essentially no interaction, and can propagate in parallel channels without interference. Optical signals have also been demonstrated to interact with certain materials on a subpicosecond (a picosecond is a trillionth of a second) time scale, offering the potential for high throughput.

The research into the application of these properties of light for computers has followed two general paths: analog processing and digital processing. Analog computers work by measuring physical properties—numbers might be expressed as voltage levels or degrees of brightness in a laser beam, for example. Digital devices, on the other hand, work by measuring different "states," for example *on* versus *off* or *high* versus *low*. These states are then used to perform binary arithmetic, that is, arithmetic based on ones and zeros.

Since analog computers represent quantities of interest with physical properties that can be altered at speeds approaching the speed of light, the result to a complex computation can be achieved almost instantaneously. In a digital computer, these same quantities are represented as binary numbers. A complex computation typically requires manipulating these binary numbers billions of times. Consequently, analog computers can be substantially faster than digital computers.

Analog devices

An analog system is one in which the inputs and outputs can take on a continuum of values. Typically, in optical systems, these inputs and outputs are intensities or brightness. In order to utilize the parallel processing capabilities of optics, intensity profiles in two or three dimensions can be utilized, as is done in optical image processing.

Operations of the analog variety have been widely developed over the last 20 years. There is a fairly well-defined repertoire of operations that can be performed, including such important operations as addition, subtraction, and multiplication of images; Fourier transformation (representing a signal with a finite number of sinusoids), correlation (comparing two quantities), convolution (multiplying two functions) and other operations that can be derived from these. The last three are especially significant as they are useful for solving differential equations.

Historically, there has always been a problem maintaining the accuracy of analog processors because of component stability, approximation error, and noise (unwanted signal) propagation.

The first drawback stems from the fact that analog devices represent numbers by physical conditions such as voltage levels or brightness. If the components generating the physical condition change in some fashion over time, the calibration of the analog device can be altered. This, in turn, can effect the computational accuracy of the device.

The second problem, approximation error, is due to the fact that physical phenomena are being used to model mathematical operations, which is the converse of the usual situation. Thus, because mathematics can only approximate reality, the mathematical calculations performed by analog devices are necessarily approximations. In addition, these approximations usually hold only over a limited operating range.

Finally, the operating principles underlying analog devices generally result in unwanted signals being propagated and amplified throughout the entire sequence of their mathematical operations. Depending on the amplitude of these signals, the computational accuracy of the analog device can be substantially degraded.

Optical analog devices

In large part because digital computers did not suffer from the above operational problems, computer users preferred them over electronic analog devices for their computational needs. However, current progress suggests that optical analog devices will not suffer the same fate.

Exemplary of the progress made in analog devices is the 1981 design of a vector-matrix processor, a device that performs highly repetitive mathematical operations on large amounts of similar data, by Dr. John Caulfield of the University of Alabama. Caulfield used Bragg cells—a transparent crystal that combines light waves and sound waves in such a way that the path of the light beam is dependent on the amplitude and frequency of the sound wave—to construct his device. A key feature of the implementation was the use of feedback (the routing of the output back to the input) to improve the accuracy of the device.

Harper Whitehouse of the Naval Ocean Systems Center in San Diego subsequently showed how the use of a convolution technique could improve the accuracy of Caulfield's device even further. Following this, Peter Guilfoyle, head of GuilTech Research in California's Silicon Valley, combined Caulfield and Whitehouse's ideas into an optical processor that apparently possesses both analog speed and digital accuracy. Exactly how the device works, Guilfoyle will not reveal.

Both Caulfield and Guilfoyle see the initial use of their processors as add-on devices for existing digital computers. The resulting hybrid computer is expected to perform faster than the digital computers alone and still maintain digital computer accuracies.

Optical circuits increase throughput of computers

The introduction of light, as a medium of communication in computers, will increase the number of mathematical operations that may be performed per second. The most obvious reason for this, is that an optical circuit has zero resistance to flow, whereas an electronic circuit, such as the one shown here (left), always presents some inertia that must be overcome for current to flow, and calculations to be carried out. As a result, the speed of data transmission in an electronic computer is, inherently, considerably slower than the speed of light, and consequently the speed of mathematical transformations that it carries out is necessarily slower than an equivalent optical circuit.

We can look at this difference by comparing the nature of the waveguides required in electronic and optical circuits. The low frequency of the electromagnetic radiation known as electric current, requires a conducting substance for its ordered transmission, as indicated even in the microcircuit shown here. (That's the edge of an American coin pictured in the background.) Furthermore, electromagnetic interference phenomena require that the waveguides be sufficiently distant from each other, or insulated. On the other hand, coherent light in the optical range does not require a waveguide for linear transmission; and distinct signals may intersect each other with no practical interference. Optical waveguides, for example, the optical fiber shown on the right, are required only for turning the direction of a signal, or for preventing its dissipation over long distances.

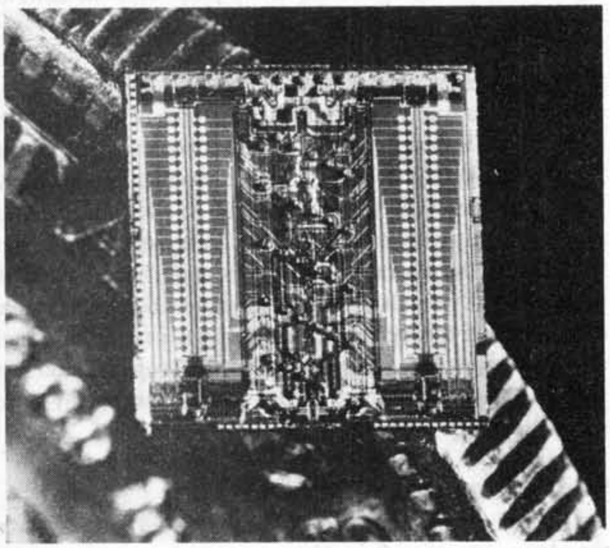

Electronic circuit technology—the IBM memory chip.

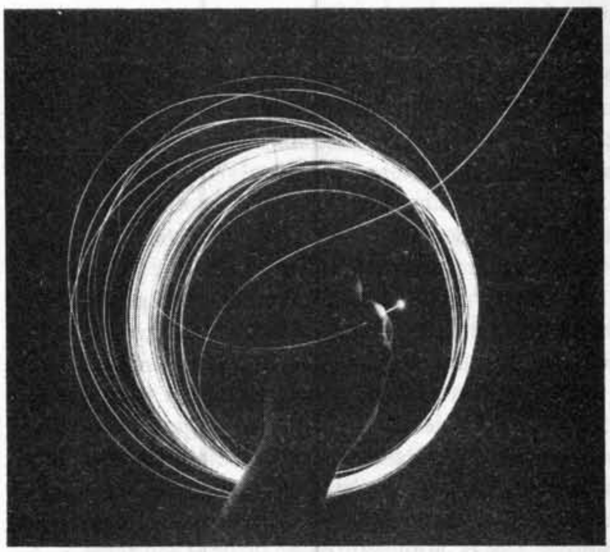

Optical fiber technology will allow even more rapid data processing.

Digital devices

Optical techniques not only have potential for transforming the design and construction of analog devices, but also hold tremendous promise for revolutionizing digital devices.

When computers were first developed, the hardware for gates (typically a logic circuit) and for gate interconnections was expensive. Thus a computer design evolved that minimized the hardware requirements. These requirements have led to the classical von Neumann computing system architecture shown in **Figure 1**. In a von Neumann machine, all of the processing logic is contained in the central processing unit and the memory is located almost entirely in a separate unit.

FIGURE 1
Classical von Neumann computing system architecture

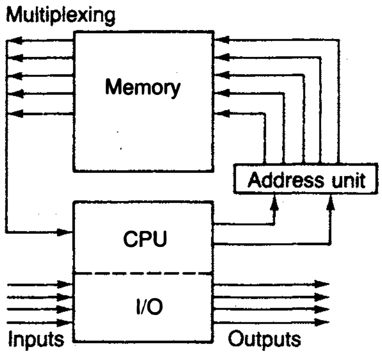

Input-output is generally performed by the central processing unit or by a system closely coupled to the central processing unit. The central processing unit accesses the memory through a binary addressing unit, and memory contents are returned to the central processing unit through a single or small number of lines.

Although the serial addressing and multiplexing scheme of Figure 1 reduces interconnection and communication requirements and minimizes the number of lines, it has the disadvantage of limiting the time it takes to transfer information between memory and the central processing unit. The eventual limitation of this tradeoff on computing speed is known as the von Neumann bottleneck.

Other limitations on processing speed

This is only one of several fundamental problems constraining the processing speed of today's and future electronic computer systems. The most important of these other problems are interconnection bandwidth, clock skew, electromagnetic interference, and crosstalk.

Bandwidth: As system cycle time and pulsewidths shrink for the purpose of increasing processing speeds, the bandwidth (the range of frequencies that a circuit can carry above a specified percentage of its maximum power level) needed to preserve the square-wave shape of these signals increases. (A square wave pulse is the sum of many different frequency sinusoids.) This forces the need for bulky and expensive terminated coaxial interconnections. Thus, the amount and speed of information transfer is restricted. Furthermore, the terminating resistors on the ends of the cables reduce system energy efficency.

Clock skew: This problem occurs when signals from different parts of a circuit arrive at a gate at different times. If this skewing is severe, the gate can generate an erroneous output. As computational cycle times become shorter, the amount of clock skew that can be tolerated diminishes. To keep clock skew within acceptable limits, all interconnection lengths between gates must be kept constant. If this is not possible, shorter lengths must be padded with gate delays to make the propagation time equivalent to that of the maximum interconnection length.

The problem of clock skew also makes it difficult to exploit the performance advantage of extremely fast logic gates in traditional electronic circuits. Inputs to a logic gate must be allowed to settle before the output of the gate can be considered reliable. The input settling time is dependent on the time it takes to fully charge a connection. In most circuits this settling time is longer than typical transistor switching times. Clearly, gates with switching times less than the interconnection settling times are of little benefit under these circumstances. Very large scale integration, VLSI, provides no solution to this problem because the scaling down of circuits does not change the settling time properties of interconnections.

Electromagnetic interference and crosstalk: The problems of EMI and crosstalk stem from the fact that electrical signals affect each other over a distance. Consequently, adequate shielding and separation between signal paths to maintain the integrity of each signal necessarily limits the number of channels that can be supported per given volume. This, in turn, imposes a minimum separation distance, and therefore a minimum signal propagation time, between gates, ultimately restricting computing speed.

All the above problems are interrelated. While increasing cycle times (that is, diminishing both separation between pulses and pulse widths) circumvents the von Neumann bottleneck, it aggravates the bandwidth, clock skew, and electromagnetic interference problems. Since the root cause of these problems stems from the inherent properties of the electromagnetic radiation used to perform the task, solutions ultimately lie in developing techniques to perform the task using a "higher quality" form of radiation. Thus, we are naturally led to consider the binary computing problem from the standpoint of light and optics.

The advantage of optics

What do light and optics have to offer? First, optics is capable of communicating many high-bandwidth channels in parallel. Lenses, prisms, and mirrors can convey images consisting of millions of resolvable spots. Each spot is capable of supporting a very large bandwidth channel. In addition, optical beams do not interfer with each other and can cross each other without interaction.

Although these attributes have been exploited in analog optical devices, they have not yet been fully applied to digital systems. The main reason for this has been the lack of suitable optical logic and memory devices, but this situation has recently changed.

Transphasors: The fundamental components of any digital logic or memory device are switches capable of two different states of transmission. A dominant factor in the speed of the component is the time required for a switch to change states. Currently, computers use transistors to perform the switching function. The fastest these devices can be made to change states is about a nanosecond, or a billionth of a second.

Researchers at several laboratories (for example, Bell Labs in the United States and Heriot-Watt University in Scotland), however, have developed an experimental optical device analogous to the transistor. This is called a transphasor, and it can achieve switching times on the order of a few picoseconds (a picosecond is a thousandth of a nanosecond).

Like any switch, the transphasor is based on a physical phenomenon whose input-output relationship takes the general form of the characteristic curve shown in **Figure 2**. The curve's nonlinearity is the key attribute that produces a binary (that is, high versus low) output based on a binary input and permits the implementation of logical functions.

FIGURE 2
Characteristic nonlinear curve of binary switch

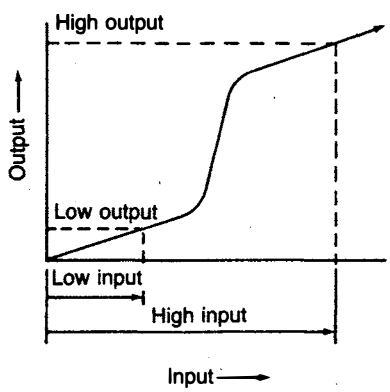

For example, if we desire a device that outputs a high signal only if its two inputs are simultaneously high (the classic *AND* function), we can scale the inputs such that only when they are both high will the total input signal exceed the kink in the curve of Figure 2 and produce a high output.

Similarly, if we desire a high output when either one of the two inputs is high (the *OR* function), we can scale the inputs such that a single high input exceeds the kink in the characteristic curve and produces a high output. Given the *AND* and *OR* functions, all of the logical functions required of binary digital computers can be constructed.

Most transphasors are constructed using a Fabry-Perot interferometer and a material with a nonlinear refractive index in the interferometer cavity. In its simplest form, the Fabry-Perot interferometer consists of two plane, partially reflecting mirrors placed parallel to each other and separated by a cavity or space. If a coherent beam of light is input through one of the mirrors, the intensity of the beam output from the other mirror is dependent on the interference pattern set up in the cavity by the incident and reflected beams.

At full destructive interference, the intensity in the cavity is almost zero and transmission through the output mirror is negligible. However, at full constructive interference, the intensity of the cavity can be as much as 10 times the intensity of the incident beam. Because of the losses through the output mirror, the intensity of the output beam is roughly equal to the incident beam intensity for this situation. By placing a material in the cavity whose refractive index (ratio of speed of light in a vacuum to its speed in the material) is dependent on the intensity of the light passing through it, the conditions under which constructive or destructive interference take place can be made dependent on the intensity of the incident beam.

For example, suppose, at low incident beam intensities, destructive interference prevails in the cavity. The intensity of the output beam is therefore negligible. As the intensity of the incident beam is increased, it causes the material in the cavity to achieve a refractive index that results in constructive interference, and the intensity of the output beam increases

Analog computers will transform computation

Digital computers are incapable of truly representing any process in nature, from the propagation of sound to the growth of an economy. The reason for this is that they internally represent all data or information with discrete, linear objects: integers (that is, the numbers, 1, 2, 3, . . .) in the form of binary numbers (ones and zeros). The concept of an **analog computer** is based on the idea of representing the "data" of nature, in forms close to those in which they occur in nature. Digital computers are unable to directly represent rotational action, exponential functions, and even square roots. It is not unrelated that the principal use of computers today is in the limited domain of accounting.

Pythagoras showed over 2,000 years ago that it was impossible to represent the square root of two with either a single ratio of integers (that is, a so-called rational number), or a finite sequence of them. Archimedes showed some hundreds of years later that it was impossible to dramatically. The resulting input-output relationship looks generally like that shown in Figure 2.

By altering the length of the interferometer, the wavelength of the incident beam, or the material in the cavity, a hysteresis loop can be created, as shown in **Figure 3**. In simplified form, the occurrence of a hysteresis can be explained as follows: As the intensity of the incident beam is reduced, enough light remains in the cavity to keep the refractive index near the value corresponding to constructive interference. Thus the "decreasing incident intensity" curve is to the left of the "increasing incident intensity" curve.

FIGURE 3

Characteristic curve with hysteresis, or bi-stability

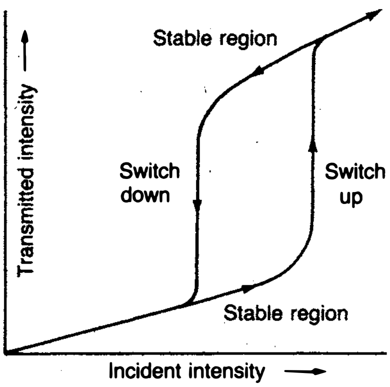

An optical device that exhibits this behavior is said to be optically bistable because it has two stable regions where the transmitted intensity changes very little with variations in the incident intensity. Either the high state or low state can be maintained indefinitely with an incident beam of intermediate intensity. The obvious application of such a device is as a binary memory element.

The transphasor described above is known as an intrinsic system because its performance is based on the intrinsic refractive properties of the material in the interferometer cavity.

Hybrid systems: Work is also being done on so-called hybrid systems, which typically employ a crystal whose refractive index depends on an applied voltage rather than incident light. The voltage is generated by a light detector that detects a portion of the output beam of the interferometer via a beamsplitter. The resulting feedback loop causes the electro-optic crystal to behave in the same manner as a crystal with an intrinsic nonlinear refractive index. Depending on the tuning of the feedback loop, a characteristic curve of the form shown in Figures 2 or 3 can be produced.

Transphasors have also been constructed without the use of Fabry-Perot interferometers. Most of these hybrid devices are based on liquid crystal light valves. This device is a small wafer sandwich essentially composed of two different mate-

represent a circle (and therefore, rotational action), with such a finite number sequence, in the form of a multiple-sided polygon. Later geometers showed that *e*, a characteristic of the logarithmic spiral (as π is a characteristic of the circle), is also unmeasurable using sequences of rational numbers.

The point of changing over to the use of analog computers, is not to find a different, "more efficient" way to produce digital representations, but to dispense with digital representation altogether. Analog (i.e., nonlinear) representations are primary. When the need arises to convert an analog representation into a digital form, this is trivially accomplished.

The first attempts to construct analog devices used the waveforms of the low-frequency electromagnetic radiation known as electric current, to directly represent the nonlinear "arbitrary functions" or data, characteristic of all natural processes. Our increasing mastery of coherent light and optics, enables a more direct path to the analog computer, unhampered by the limitations of electrical circuitry.

The figure shows the approximate form of coherent laser light, under consideration as the building block of optical computers. The superiority of this form of data representation over the digital, is obvious. First, nonlinearities, such as rotational action, are the form of the representation itself. Progressive frequency upshifting enables the direct representation of the logarithmic spiral. Second, light is primarily a hydrodynamic phenomenon. As such, it will provide a way to represent hydrodynamic differentiation (e.g., solitons or shock waves).

This direct representation of nonlinearities in nature will result in analog computers which can solve several nonlinear problems in a timeframe in which a digital machine would have just gotten started. In short, the introduction of optical analog computers based on these principles, opens up a whole new domain of possibilities.

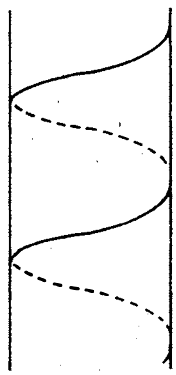

The simplest form of circular action: the cylindrical helix.

rials positioned back-to-back—a photoconductor crystal and liquid crystal material much like that found in watches and calculators.

An electrical current applied to the device is modulated by the intensity of an incident beam directed on the photoconductor crystal. The modulated current, in turn, affects the optical polarization properties of the liquid crystal material. When a secondary beam, or a reflected portion of the input beam, is reflected off the liquid crystal, it becomes polarized. By passing the output beam through a polarized lens, its intensity can be made dependent on the current applied to the liquid crystal, and ultimately dependent on the intensity of the incident beam.

The input-output relationship of the device has the general form of the curve shown in Figure 2. Use of an optical feedback loop can create the curve shown in Figure 3. The utility of liquid crystal light valves is being investigated primarily at the University of Southern California, University of California at San Diego, and Ohio State University.

Optical arrays and computer architectures

In principle, two-dimensional arrays of transphasors can be made from either intrinsic or hybrid devices. The reason for doing so is to permit parallel operation of a large number of transphasors. Even if the switching speed of each element is relatively slow, the large number of parallel operations can produce very high total data rates and processing speeds.

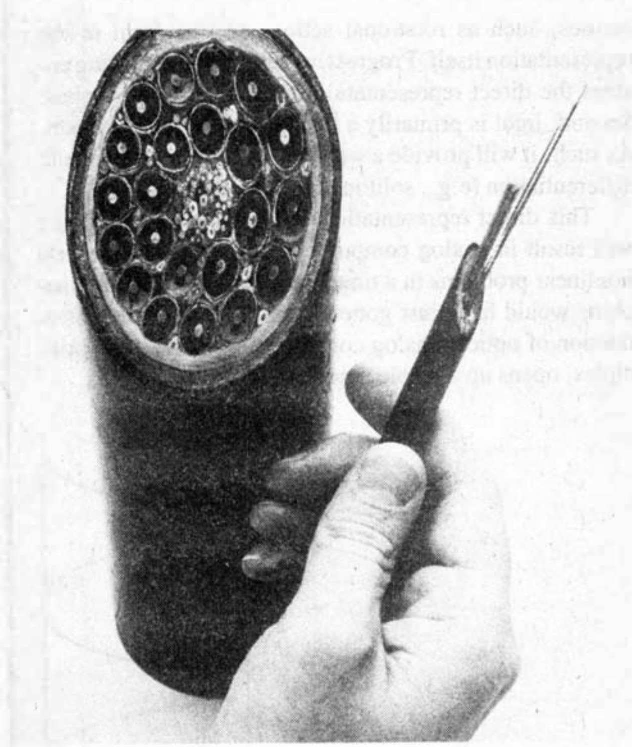

The fiber optic cable used in the Bell System's Northeast Corridor project is 100 times lighter than the copper coaxial cable on the left.

The availability of optical logic and memory arrays, plus the freedom to make arbitrary interconnections, makes possible totally new computer architectures. For example, the ability to update all memory elements in parallel eliminates the need for a memory addressing unit (see Figure 1) and therefore obviates the problem of von Neumann bottleneck. An illustration of such an architecture is given in **Figure 4**.

Theoretically, it is possible to implement such an architecture using a single integrated optical gate array or "chip." This array would combine many transphasor arrays to create an input/output unit, central processing unit, and a memory unit on a single chip. The communication among transphasor arrays comprising the chip and the three functional chip units would fully exploit the properties of light.

FIGURE 4

Non von Neumann computing system architecture

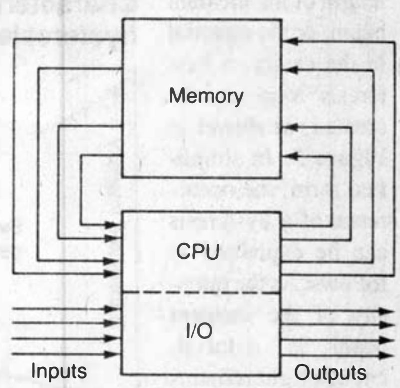

FIGURE 5

Non von Neumann optical computer schematic

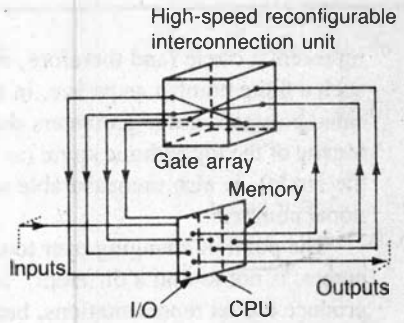

A schematic of this implementation provided by researchers at the University of Southern California is given in **Figure 5**. A key feature of this implementation is the ability to redirect communication paths during computer operation. Reconfigurable holograms under computer control is one method available for doing this.

Multilevel logic

Perhaps even more promising for increasing the utility and speed of computers than the list of innovations available for binary logic computers is the development of computers with multilevel, or multidiscrete, logic. This becomes possible if switching and memory devices with more than two states can be developed.

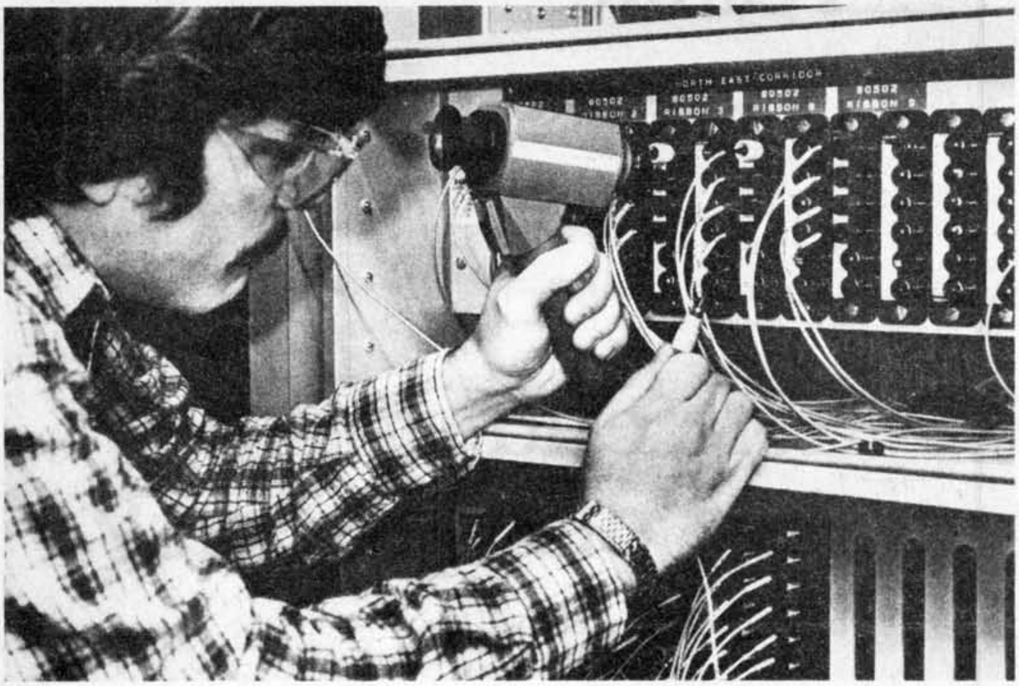

An AT&T technician uses an infrared viewer to test light transmission on the hair-thin glass fiber cables in New York City's electronic switching system.

Intrinsic transphasors such as the Fabry-Perot have been designed to produce characteristic curves like that in **Figure 6**. Note that only a single region of bistability exists, limiting the use of the device to a multilevel switch. Several multilevel regions of b before a multilevel memory is possible. Of course, multilevel switches and memories can always be developed by combining several binary devices of differing operational input levels. When this is done using today's electronic devices, the interconnection complexities that result make implementation of electronic multilevel devices impractical. However, as reported above, this problem can be circumvented with light and optical devices.

FIGURE 6
Multilevel characteristic curve with single hysteresis

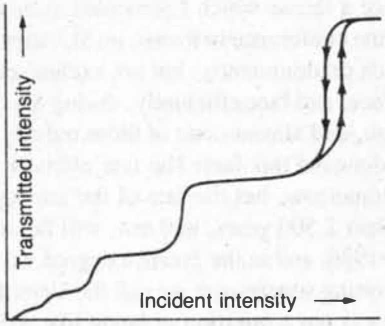

Another obstacle that now affects multilevel computing is the lack of a well-developed theory for its implementation. Most of the work in this area concerns the theory of residue arithmetic, which still has several difficulties: Division of two numbers, sign testing, and relative magnitude comparisons are not easily and conveniently handled. Suffice it to say that much work remains to be done. However, if this work is brought to a successful conclusion, the payoff could be significant. Several experimental residue processors have been built and tested and have performed many times faster than binary devices performing the same computations.

Implications for the future

Current R&D in the area of information processing and computing is examining virtually every premise underlying the design of today's computers. This return to fundamentals has come on the heels of major advances in the field of lasers and optics. Results to date make it clear that only time and money stand in the way of major breakthroughs in all areas of information processing.

istability must e

Without question, the new generation of optical computers will far surpass the capabilities of the computers available today. These new computers not only will fulfill the needs of strategic defense systems, but also will provide many other benefits.

It is no coincidence that the electromagnetic phenomena that can be used to satisfy the computational needs of a successful SDI system also provide its required kill mechanism. Advances in tools such as computers and lasers (or anything else that does useful work) are ultimately dependent upon advances in our understanding of ever more highly organized forms of electromagnetic action. Tools do not create action; they simply redirect it.

The progress being made in optical computing is but one example of the true promise of the SDI as not merely a defense against ballistic missiles, but the spawning of a new industrial revolution.