Optics and Supercomputing

P. BRUCE BERRA, FELLOW, IEEE, ARIF GHAFOOR, SENIOR MEMBER, IEEE,
MOHSEN GUIZANI, SLAWOMIR J. MARCINKOWSKI, AND PERICLES A. MITKAS

Invited Paper

We partition our remarks on optics and supercomputing into three major topic areas: storage, interconnection and processing. Under storage we discuss various types of optical disks and page-oriented holographic memories. Optical storage is advancing rapidly and holds the potential of hundreds of megabytes per second data rates from a single storage unit, which can provide many new opportunities for supercomputing. Under interconnection we discuss module-to-module, board-to-board, chip-to-chip and finally gate-to-gate communication. We conclude that optical interconnection is, in many cases, superior to electronic interconnection and holds the key to the development of future electrooptic systems. We present a section on optical computing devices and provide various application areas where optical processing as well as storage and interconnection are expected to play a role in the future. We feel that optical processing, while holding considerable promise, lags behind its electronic counterpart primarily due to the fact that digital optical device development is in its infancy. We believe that near-term systems will be electrooptic, with each technology providing its strength to the problem at hand.

I. Introduction and Background

A. Supercomputing

The term supercomputer has been more or less defined as pertaining to those machines that possess enormous processing capabilities which are required for numerically intensive problems. The general problem that is approached typically involves the simulation of a complex system or phenomenon. These include weather prediction, nuclear research, structural engineering, and, in general, those problems that involve large matrix calculations.

With the advent of the recent artificial intelligence (AI) wave a new variant of the term supercomputing has emerged. This variant places increased emphasis on the nonnumeric aspects of supercomputing. That is, many problems in AI are large and require enormous computer resources but cannot be classified as "number crunching." Instead, for instance, they may be involved with inferencing processes on very large knowledge/data bases for the purposes of generating new knowledge. Other areas of application include full text processing, machine vision, learning, distributed reasoning and, in general, problems that require considerable search capabilities.

With the expanding role of supercomputers in the numeric processing field [1] and the increased emphasis on nonnumeric processing a new and broader definition of supercomputing is emerging. We prefer to describe supercomputing in terms of three necessary hardware components; memory, bandwidth and processing capability as shown in Fig. 1.

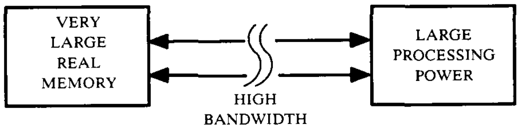

Fig. 1. Supercomputer components.

The memory must be of main memory-type technology with rapid access time and large enough to hold a significant amount of problem data. Since considerable numeric or nonnumeric processing is required, the processors must have enough power to deal with significant problems. Finally, data will be moved back and forth between the memory and processors so a high bandwidth interconnection must be included. We can apply the test of two very different types of supercomputers to this model. The Cray-2 has over a gigabyte of fast memory and gigahertz bandwidth between fast memory and four powerful processors. On the other end of the processor spectrum the Connection Machine has a very different architecture but similar characteristics. From the point of view of memory, each processor has only a few kilobytes, but when aggregated over 64K processors the amount of memory totals one half of a gigabyte. While each processor is limited in its capability, the aggregate capability of 64K processors is considerable. Finally, the total bandwidth between processors and memory is 64K times the bandwidth between one processor and its memory.

B. Optical Storage, Interconnection and Processing

The objective of digital optics is to replace electrons with photons whenever appropriate in a computing environment. The motivation for this is that optics possess some very attractive properties including massive parallelism, high speed, and noninterference of light beams. The pur-

Manuscript received December 17, 1988; revised June 5, 1989.
The authors are with the Department of Electrical and Computer Engineering, Syracuse University, Syracuse, NY 13244-1240, USA.
IEEE Log Number 8932753.

0018-9219/89/1200-1797$01.00 © 1989 IEEE

pose of this paper is to survey many of the aspects of optics as they apply to supercomputing and to indicate what we might expect from optics in the future. We present the ideas in the context of storage, interconnection and processing.

The configuration shown in Fig. 1 has enormous capabilities, provided that enough data can be supplied to it. But, at some level data must be supplied to this system from disks. However, disk access times are on the order of a million times slower than main memory access time. Thus, one must idle the processors when problem data are being loaded into fast memory or some shared arrangement is utilized. While considerable work is being done to generate larger data rates from magnetic disks, the characteristics of optical media are such that massive data rates are potentially possible.

We discuss two types of optical storage; optical disks and page-oriented holographic memory (POHM). Optical disks of various types are in wide use. They are attractive because of their large storage densities even though their access times are slower than magnetic disks. However, with suitable modification to read multiple tracks simultaneously, data rates on the order of hundreds of MBytes/s are possible [2], [3]. This is also the case with POHM. For comparison, the vast majority of current magnetic disks have data rates on the order of 3 MBytes/s. Notable exceptions include multichannel disks and parallel reading from multiple disks where data rates on the order of 50 MBytes/s could be obtained.

Given massive data rates from optical disks or POHM (i.e., 500 MBytes/s), electronic computers would be hard pressed to accept the data without hardware modification since they have been designed for magnetic disk rates. This situation is depicted in Fig. 2(a). That is, massive data rates are available from optical memory, the data are converted from photons to electrons (D-detection) and then fed to suitably modified supercomputers. A slightly different approach, as shown in Fig. 2(b) is to distribute the data to many locations thus making use of the high bandwidth of optics. Another approach, shown in Fig. 2(c) is to provide optical processing of the data stream prior to conversion from photons to electrons. In this case the data rate will be reduced because of the processing operations and the input to the electronic computer will be much richer in content. While we direct our remarks on optical computing to supercomputers, we must point out that the technology applies equally well to less powerful computers.

We begin the discussion with optical disks and page-oriented holographic memories. This is followed by a discussion of optical interconnections at the module, board, chip, and gate levels. We devote considerable attention to storage and interconnections since these are the most well developed. We then discuss the least well developed aspect of optical computing, processing, and provide several application areas where progress is likely to be made in the future.

II. OPTICAL DISK STORAGE

Supercomputers are capable of processing data at rates reaching or exceeding hundreds of GFLOPS. These rates can be achieved when the data involved reside in the main memory of the system, which is often the case with numerical applications (matrix operations, etc). For input/output intensive applications, however, the performance of the system is usually determined by its ability to efficiently transfer data to and from secondary storage. Therefore, storage devices allowing access to terabytes of data in a relatively short time will form an essential part of any future supercomputer.

Secondary storage devices offering large capacity include magnetic disks with fixed or moveable heads, magnetic tape, optical disks, and, in the future, holographic mem-

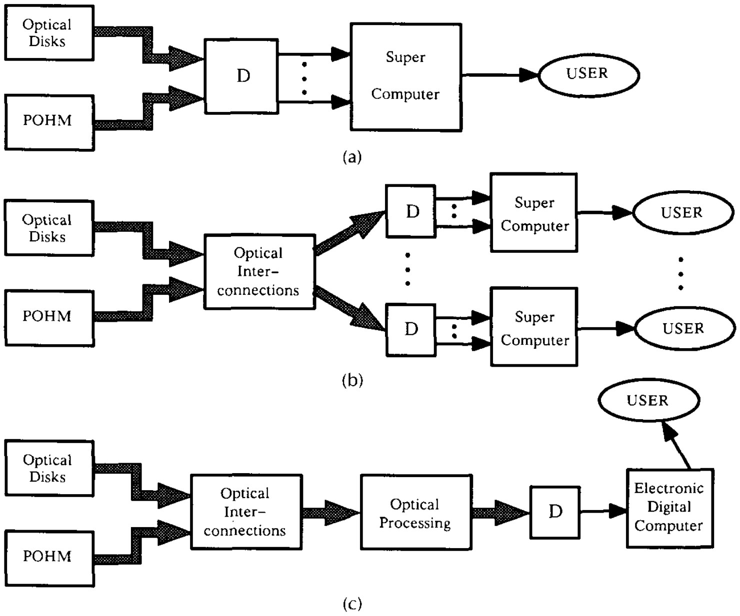

Fig. 2. Optical storage, interconnection, and processing. (a) Massive data rates from storage. (b) Storage and interconnections. (c) Processing of optical data.

ories. The fact that the read/write head of a magnetic disk has to be extremely close (about 0.5 μm) to the disk surface, poses some physical limitations to the density of magnetic technology. Optical memories (disks and holograms) are being developed as an alternative to magnetic devices in order to meet the constantly increasing requirements for massive low cost, high density secondary memory.

Data are recorded on the optical disk surface as a series of pits on multiple concentric or a single spiral track. In order to read from the disk a light beam generated by a laser diode is focused on the reflective layer of the disk and one or more of the optical properties (intensity, polarization, phase, etc.) of the reflected beam is modulated according to the pitch length of the pits. The reflected light is then detected by photodetectors. Appropriate servo-mechanisms are employed for precise focusing and tracking.

The relatively large Head/Medium gap (order of millimeters or about 10 000 times that of magnetic disks) enables the optical head to focus into very small spots, eliminates surface wear, and makes head crashes virtually impossible. The diameter of a pit is set by the limit of diffraction at the selected wavelength and the numerical aperture of the optics. The information density lies in the order of hundreds of Mbits/cm^2 while the most advanced Winchester disks can carry up to 6.7 Mbits/cm^2.

The major disadvantage of optical disks is the low input/output transfer rate (sustained) due to high access times and the relatively undeveloped nature of read/write disks. As the field develops, however, we anticipate a significant improvement in the access mechanisms. Some possible solutions are discussed in following paragraphs.

A. *Classification of Optical Disks*

Optical disks can be classified into three categories: a) Read-only (ROM), b) Write-Once-Read-Many (WORM) and c) Erasable/Rewritable. The best-known representative of the first type is the CD-ROM, a compact disk with a standardized format. All the information is prerecorded by the manufacturer and cannot be altered by the user. Optical disks of the second category are not standardized and come in different sizes. As their name suggests, data can be written only once by the user and cannot be erased afterwards. They are suitable for storage of less volatile data or applications that require a complete history of updates. For a discussion of CD-ROM and WORM see [4], [5], and [6].

Research efforts in Erasable/Rewritable optical disks have focused on three different approaches as discussed below.

1) *Phase change of the material (amorphous to crystalline):* Certain thin film metal alloys based on tellurium or selenium can exist at room temperature in two distinct states, amorphous or crystalline, and can be switched from one state to the other with a laser. Since the optical properties (reflectivity, etc) of the film are different for the two states, binary encoding is possible and usually the amorphous state corresponds to a logical 1 while the crystalline to 0. The recording speed depends largely on the crystallization time of the material, and ranges from about 200 ns to above 1 μs [7]. Many phase-change media suffer from a degradation in performance during the reading process, which places a limit on the number of rereads possible before the material must be refreshed [6].

2) *Plastic deformation [8], [9]:* This approach employs two thin layers of organic dye-polymer materials, permanently bonded to each other, which possess different thermal, optical, and mechanical properties. When a spot on the disk is temporarily heated it can change its form and keep it until another heating cycle is imposed. Therefore, bumps can be generated or destroyed on demand.

3) *Magnetooptic combination [10], [11]:* This is the most advanced method and appears to be the most viable one. It is based on the Kerr effect in which a polarized laser beam experiences a rotation of its polarization upon reflection from a magnetized surface. The direction of the rotation is determined by the orientation of the magnetic field on the reflective surface. The recording medium is a thin layer of vertically oriented magnetic material which is sandwiched between a transparent polycarbonate protective coating and a reflective substrate. Originally the orientation of data bits in a blank disk is the same (i.e. north-pole-down) and corresponds to the logical 0. At room temperature the magnetic field required to reverse the orientation of a bit to a logical 1 is extremely high but drops to almost zero when the material is heated to its Curie point. To record a bit, the temperature of the corresponding point is raised momentarily to about 150°C by illumination with a high-intensity laser beam and the magnetic vector of the medium is aligned with a relatively weak bias magnetic field generated by a coil underneath the disk. Erasure is accomplished with the same high-intensity laser beam, but with the imposed magnetic field reversed, which corresponds to writing all zeros. Data are read by illuminating the recording surface with a low-intensity polarized beam. Current products use separate rotations for write, read, and erase.

Being nondestructive, the Record/Read/Erase cycle can be repeated millions of times. The same laser diode performs all three operations and erasure takes place either sector by sector, or in bulk with a large magnetic field. Also, magnetooptic recording will benefit significantly by the advanced technology of magnetic disks.

B. *Designs for Larger Capacity and Higher Transfer Rates*

Magnetooptical disks have the potential to dominate the secondary memory market. Predictions for tenfold increase in storage capacity and enormous transfer rates over the next few years are no longer considered optimistic. Future improvements include faster drives and smaller optical heads along with more efficient encoding and error-correcting techniques. Higher recording densities will result in higher transfer rates. Second generation erasable disks are expected to further increase throughput by recording data in a single pass using two spot Direct-Read-After-Write (DRAW) or Direct-Read-During-Write (DRDW) mechanisms. In this section we present some approaches that indicate possible future developments that would have a significant impact on supercomputers.

1) *Optical Jukebox:* A device that provides huge capacity and considerably high I/O bandwidth is the optical "jukebox" [12], [13], [14]. In such a system a large number (60–150) of 14-inch diameter WORM optical disks are arranged in an on-line library configuration and accessed via an automated handling mechanism. Access is provided to any data in the store of 10^{13} bits within seconds. The disks are mounted onto a turntable and spun up to speed so that data can be transferred at user rates up to 25 Mbits/s per channel

using focused laser beams. The number of channels per jukebox varies from 1 to 8. Other projects [13] include the Optical Disk Buffer with on-line data capacity greater than 1 Terabit, I/O transfer rate greater than 1.6 Gbits/s and access time less than 100 ms and the High-Data-Rate Optical Tape Recorder with a projected capacity of 20 Terabits.

2) Multiple-Beam Read: Solutions to the problem of relatively low access times and transfer rates can be provided by using more than one laser beam simultaneously. The noninterfering nature of light and the relatively large distance between the optical head and the disk surface make parallel multibeam data access feasible. In fact, commercially available CD-ROM drives can read many tracks without head movement. A single laser beam can be expanded to multiple parallel coherent beams and each of them focused on a different track on the disk. The readout beams can be further separated and directed to an array of photodetectors. This technique (Fig. 3(a)) can provide data rates on the order of hundreds of MBytes/s but it is not suitable for parallel recordings, since independent modulation of a particular beam is not possible.

Another approach [3], depicted in Figure 3(b), involves a 1-dimensional semiconductor array of sources and detectors placed above the disk so that each source-detector pair corresponds to a single track. This configuration allows individual beam modulation and is suitable for recording as well. Since it would be extremely difficult to fabricate detectors only 1 μm apart (the typical separation between tracks), the detector array can be divided into multiple segments staggered in a fixed pattern with more widely spaced elements. Up to 10 000 bits (the number of tracks in a typical optical disk) could be read in parallel or the entire disk could be scanned in about 20 msec which is the rotation time of the disk. This technique, however, has considerable power requirements because the rate at which information can be retrieved increases at the expense of higher optical power. Another problem, that of controlling the relative position of the recorded bits across the different tracks in order to achieve proper data alignment, has been overcome with the development of an optical drive capable of recording a bit in any specified location on the 2-D surface of the disk [3]. A logical extension of these approaches is shown in Fig. 4 in which multiple heads are uniformly spaced across the disk surface.

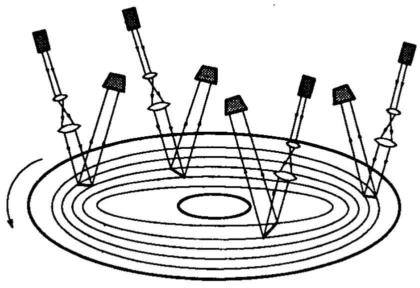

Fig. 4. Multiple head reading from optical disks.

Clearly, the generation of massive data rates from optical disks will have a profound effect on supercomputing. But the situation becomes even better if there is no need for photon-to-electron conversion. Instead, the reflected or transmitted laser beams could be guided through optical waveguides directly to optical processors, thus eliminating any contention or saturation problems. These ideas will be elaborated upon in subsequent sections.

III. HOLOGRAPHIC MEMORIES

In principal, holographic storage offers a potential improvement in both density and access time. In comparison to magnetic and optical disks a higher density can be achieved, because there is no need for mechanically repositioning the read/write mechanism. Access to the data is made by steering a laser beam through an acoustooptic reflector, which is an inertialess operation, unlike that of moving a disk arm. Also, due to the nature of holography small imperfections on the recording media are relatively unimportant, because the recorded data is spread out over the recording medium.

Holographic memories also offer the possibility of providing an inexpensive optical content addressable memory (OCAM). OCAMs have the desirable search processing properties of electronic CAMs [15], yet offer a considerably larger storage capacity (10^8–10^{10} bytes) and parallel output. Their main drawback is that they are read only. However, with the advent of new recording media, holograms can be erased and rewritten a limited number of times. Generally, selective changing of information is not possible, unless

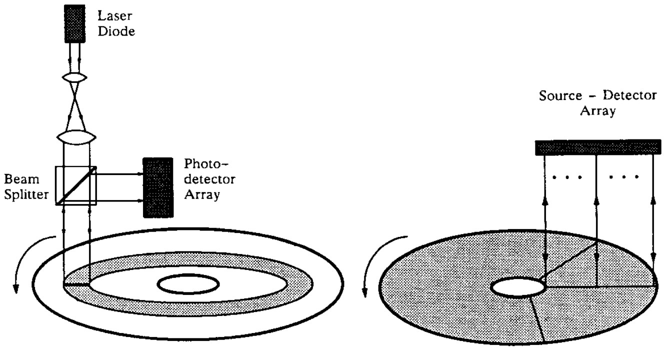

(a) (b)

Fig. 3. Multiple-beam read from optical disks. (a) Expanded beam. (b) Source-detector array [3].

the hologram is recorded in a page format. In that case the entire page must be erased prior to rewriting.

As shown in Fig. 5, a *page hologram* is made up of many nonoverlapping subholograms or pages. Recording data in a page format is analogous to storing data in pages in an electronic computer. A page hologram can contain hundreds of pages and arrays of page-oriented holograms can hold large amounts of data in small areas. For example, a page one millimeter on a side can have a capacity of about 10^6 bits. Assuming that the hologram array can hold pages on two-millimeter centers, a 3 by 5 inch card can hold 2400 pages. With a capacity of 10^6 bits per page, the card can hold 300 MBytes of data.

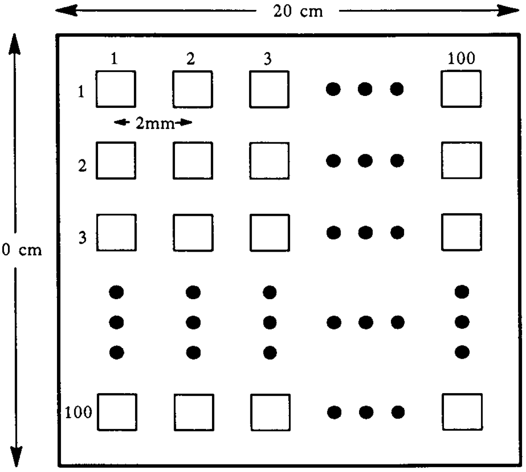

Fig. 5. Page-Oriented Holographic Memory (POHM).

Each page of the hologram is recorded separately. To record a page, data is first loaded into a page composer (an electronic input-optical output device), then a small area of the recording medium is exposed by an aperture, and the light coming from the page composer is directed by an x-y deflector to the exposed area, thereby recording the hologram. To record another page the aperture is moved, new data is loaded into the page composer, and the beam deflected to the newly exposed area. This process is repeated to record each page of the hologram. Depending on the recording medium used, information can be erased from a page, and the page can be rewritten with new data a limited number of times.

Reading data from the array depends on the time it takes to direct the read beam to the page of interest and the response time of the output photodetector array. If we assume that the maximum value of these times is 10^{-4} s then we can attain a readout rate of 1 250 MBytes/s assuming 10^6 bits/page.

Holography offers the potential of providing a fast memory system for supercomputing, since accessing the data does not involve mechanical movement. Once the data are accessed they can be loaded into main memory for subsequent use. Furthermore, since the data are read from memory in parallel, holograms can be utilized as content addressable memory. We believe that holographic memories can provide the major component of a low-cost high-bandwidth mass storage system for supercomputers.

IV. OPTICAL ELEMENTS

In order to compete with electronic technology, optical elements must be able to duplicate semiconductor functions. Bistable optical switching devices similar to diodes and transistors, and spatial light modulators will be necessary for any digital optical implementation. These devices must offer a sharp threshold in order to distinguish between a logic zero and a logic one, operate at the same wavelength and at standardized levels for ease of cascading, and have a high gain for fan-out purposes [16], [17]. They must have low switching delay times, the capability of restoration after each switching action, low power consumption, relative immunity to phase variations, large temperature operating range, and low cost.

A. Optical Bistable Devices

If an element has two stable output states for the same input over some range of input values, it is bistable. For such an element to be bistable, positive feedback is required in addition to nonlinearity. The output of the element depends upon how the present input level was achieved; by lowering from a higher level or raising from a lower level. This phenomenon of optical hysteresis is known as optical bistability and may be used in constructing optically operated switches, memories, logic gates, and amplifiers.

The Fabry-Perot resonator, is a device made by sandwiching a nonlinear electrooptic material, such as lithium niobate ($LiNbO_3$), between two partially reflective surfaces (mirrors) as shown in Fig. 6. Light entering the cavity bounces back and forth between the mirrors. If the spacing between the mirrors is related to the wavelength of the optical signal in precisely the right ratio, then the combination of mirrors will act as a transparent window. Changing the spacing results in switching the device from the reflective to the transparent mode. The "apparent" length of the gap is dependent upon the index of refraction of the nonlinear material and can be changed by applying a control optical beam which causes the device to switch.

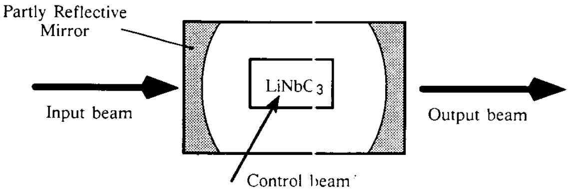

Fig. 6. Fabry-Perot cavity.

Two similar bistable devices: *etalons* [18], [19] and *optical interference filters* [20] are based on the Fabry-Perot resonator. With low intensity input light, the cavity will not resonate but when the input intensity exceeds a certain threshold, the device switches and most of the light is transmitted instead of reflected.

The quantum well is a thin layer of a material of smaller band gap, such as gallium arsenide (GaAs), sandwiched between layers of material with a larger band gap, such as aluminum gallium arsenide (AlGaAs), with electrons being confined in the smaller band-gap material. Multiple quantum-well (MQW) materials often have physical properties not found in bulk quantities of either of the constituent materials. The *Self-Electro-Optical Device*, (SEED) [21], [22], [23] requires both optical and electrical energy, and consists of a resistor connected in series to a p–i–n diode. The intrinsic region of the diode is made of an MQW material typically formed of hundreds of thin (10 nm) alternating layers of GaAs and AlGaAs. The SEED has a switching response time of 30 ns, low switching energy of 4 fJ/μm^2 optical and 16 fJ/μm^2 electrical. Parallel 2 × 2 arrays of integrated SEEDs have been demonstrated [24] and it is anticipated that within 10 years arrays of 10 000–100 000 individual SEEDs/cm^2 will be possible [25].

B. Spatial Light Modulators

Spatial light modulators (SLM) are active optical devices having the ability to store an electrical or optical pattern on a one- or two-dimensional array [26], [27], [28]. They can spatially modify or amplify some of the optical characteristics (phase, amplitude, intensity, or polarization) of a readout light distribution as a function of space and time. SLMs operate in either transmissive or reflective mode. They can be classified as optically or electrically addressed (O-SLM and E-SLM) according to the nature of the control or write signal, and can process optical signals in 1-D, 2-D or 3-D formats.

Optically addressed SLMs can be liquid crystal devices (Hughes Liquid Crystal Light Valve, (LCLV) [29]), nonholographic photorefractive signal-multiplying devices (Pockels Readout Optical Modulator, (PROM) [30]), microchannel plate devices (Microchannel Spatial Light Modulator, (MSLM) [31] and Photoemitter Membrane Light Modulator, (PEMLM) [32]), deformable surface devices (Ruticon [33]), thermoplastic modulators [34], volume holographic devices or semiconductor O-SLMs (Silicon/PLZT modulator [35]). The Photo-DKDP light valve [36] is based on a single-crystal ferroelectric-photoconductor made of KD_2PO_4 and must be cooled down to −51°C to operate properly. One-dimensional E-SLMs can be magnetooptic (LIGHTMOD [37]), micromechanical [38], electrooptic, or acoustooptic devices.

Acoustooptic (AO) Bragg cell deflectors are well developed and widely used SLMs (Fig. 7). An electrical input signal drives a piezoelectric transducer, creating an acoustic wave that propagates through an optically transparent elastooptic material. The induced spatially periodic variation of the material's index of refraction modulates or deflects an incident beam.

These devices are able to modulate up to 1000 individual beams according to the modulation of the spatial frequency grating. In an AO beam deflector the sine of the angle of diffraction is proportional to the frequency of the acoustic wave in the cell. Thus, the direction of the diffracted beam can be modified by choosing the appropriate driving signal frequency, which makes AO deflectors suitable for the implementation of arbitrary interconnection patterns. Typical electrical signal bandwidth of these devices is 1 GHz, with diffraction angles up to 10°.

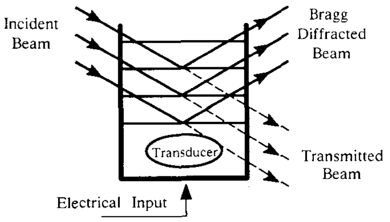

Fig. 7. An acoustooptic Bragg cell operating as a beam deflector.

A number of physical characteristics that are critical to an SLM's performance are indicated in Table 1. *Spatial Resolution*, in line pairs per mm, must be high since it determines the number of individually addressed pixels in the array. The time needed for a complete write/read/erase cycle, the *Framing Speed*, must be high also. However, speed and resolution cannot be maximized simultaneously since a trade-off relationship exists between them. A figure of merit for an SLM is given by its *Time-Space Bandwidth*, in pixel operations per second. Time-Space bandwidths of 10^{10} operations/s, which compare favorably with electronic processors, have been experimentally demonstrated [39]. *Storage Time*, varies considerably with different designs. *Exposure Sensitivity* must be high to minimize power requirements and heat dissipation problems. Finally, *Contrast Ratio* can be critical in some image-processing applications but is not that important in digital information processing.

While SLMs are either too small or too slow for current applications, considerable research is being performed in order to increase the time-space bandwidth of these devices [40]. Later in the paper we discuss optical processing, and the important role of SLMs will become clear.

V. OPTICAL INTERCONNECTIONS

Interconnection networks serve as an important component of a multiprocessor system since they provide the mechanism for efficient transfer of information among processors and memories. The crossbar switch provides general switching capability without contention for *N* processors [41]. However, as *N* grows large, the crossbar becomes

Table 1 Some Characteristics of O-SLM Devices

	LCLV	PROM	DKDP	MSLM	PEMLM	Si-PLZT
Exposure Sensitivity (μJ/cm^2)	1–5	10	10	0.01	0.01	1
Storage Time	msecs	hours	long	days	days	10s
Contrast Ratio	<100:1	1000:1	70:1	1000:1	1000:1	10:1
Framing Speed (Hz)	40	30	30	200	1000	10^4
Spatial Resolution (cycles/mm)	20	5–10	75	5–15	20	10
Time-Space BW (pixel op/sec)	10^7	$9 \cdot 10^6$	10^7	$2 \cdot 10^9$	10^{10}	10^{10}

expensive since it requires N^2 switches. Also, due to VLSI pin limitations, the size of the largest crossbar that can be integrated into a single chip is limited. Although large crossbars can be partitioned into many single-chip smaller ones, several important issues must be addressed before multichip crossbar networks can be fabricated. These include fast reconfiguration time, reduced control complexity, and minimization of propagation delay through the partitioned networks. Furthermore, the pin limitation constrains the number of data channels that can be switched in parallel and thus restricts broadcasting ability in the network.

Multistage interconnection networks (MINs) or shared bus architectures [42], [43] are used to reduce the number of switching elements in many leading multiprocessor systems. However, the interconnection network can cause a major bottleneck resulting in poor performance for various applications [44], [45], [46]. Many fine grained systems such as the Connection Machine, shared bus architectures such as the Encore Multimax, and systems based on MINs such as BBN's Butterfly generally suffer from this drawback [47], [48].

There is a growing demand to increase the throughput of future supercomputer systems, especially in real-time applications. The current emerging technologies for building high performance microprocessors are targeted to meet this challenge to build massively parallel systems [49], [50], [51]. To meet this demand and to overcome the limitations inherent in electronic technology, a great interest has developed in exploring optical technology. In optical interconnection networks, extensive research has been pursued and a number of experimental systems have been developed. The interest in optical switching arises from many promising features offered by the emerging semiconductor optical technology as compared to the electronic technology [52], [53], [54]. The most important features are listed below.

High Speed: The most prominent advantage of optical interconnection is the speed of propagation of the optical signal, which is essentially independent of the number of components receiving that signal. On the other hand, the speed of the electrical signal in a metallic interconnection network can be quite slow due to slow charging time of the loading capacitance of the gates at the far end of the network.

Large Bandwidth: An optical network with considerable bandwidth is comparatively cost-effective to an electronic network of the same size. This is especially true if the network uses multimode fibers which can provide a data rate of several Gbits/s. Fibers offer a low cost interconnection technology and are easy to couple. For example, an optical fiber based 4 × 4 crossbar switch has been demonstrated which can provide a data rate in excess of 1 Gbits/s with small channel loss [55].

Less Interference: Optical networks with low crosstalk and mutual interference can be easily fabricated due to the absence of stray capacitance. However, care must be taken that scattering light does not introduce a similar effect. Less interference results in reduced power requirement for optical networks.

Flexibility in Routing Signals: During their propagation in waveguides or free space, optical signals do not encounter significant cross-coupling or interaction, thus their routing within a chip or across a board is not constrained. This is a major advantage over electronic circuit integration where planar or even quasi-planar layouts encounter severe constraints. Consequently, optical layouts with extremely high density and a large degree of parallelism are feasible.

These features have resulted in an increasing acceptance of semiconductor optical technology for designing high speed switching systems. A number of high-speed optical switching devices have been developed within the last decade which can be controlled using electric voltage, or an acoustic signal or even by a beam of light. These include the optical directional coupler, the SEED, the OEIC and the OLE, as shown in Fig. 8 [25]. With the advent of these devices, the achievable data rates and the number of parallel channels available are well beyond the capability of any electronic system. Based upon these devices, a new generation of optical switching networks has emerged within the last five years. These networks are based on architectures ranging from fiber optics to reconfigurable holographic interconnects. Most of these networks use the space-division (SD) configurations [56] in which the signals are physically isolated and do not share common paths while being routed.

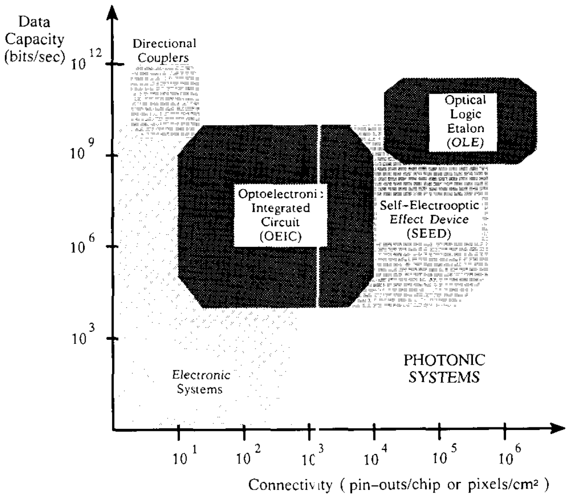

Fig. 8. Photonic device capabilities [25].

Alternatively, a number of networks have been proposed that use multiplexing schemes in which multiple signals share a common path. The most well known schemes are time-division multiplexing (TDM) and wave-length division multiplexing (WDM) [56]. In TDM, connections between communication units are established through assignment of time slots. The data pulses from multiple sources are interleaved and can boost the capacity of the common optical media, which can be either bus, free space, or integrated waveguides. In the WDM configuration, fixed wavelengths (frequencies) are assigned to transmitting and receiving units. The signals of various wavelengths share a common media without causing interference.

The objective of this section is to provide an assessment of the current state-of-art in optical interconnection technology and to speculate on the future trends in building large switching networks, which can be used in supercomputer systems.

A. Interconnection Networks in Supercomputer Systems

The architecture of many supercomputer systems is a complex combination of clusters of processors, memory modules, interconnection networks, and I/O subsystems [57], [58], [59]. As shown in Fig. 9, the general organization of these systems is hierarchical in the sense that processors are clustered to form system modules, which are in turn interconnected to communicate with each other through shared memory modules or by sending messages. Within each system module, processors and memory units are interconnected through a high-speed switching system. A processor itself is generally a combination of a number of VLSI chips such as CPU, memory, clock generator, etc., which are interconnected on a board. The characteristics and design requirements of these switching systems can vary considerably, depending upon where they are employed in the system. We present a characterization of this hierarchy of interconnection and discuss how emerging optical technology can provide solutions to switching problems at various levels of the hierarchy.

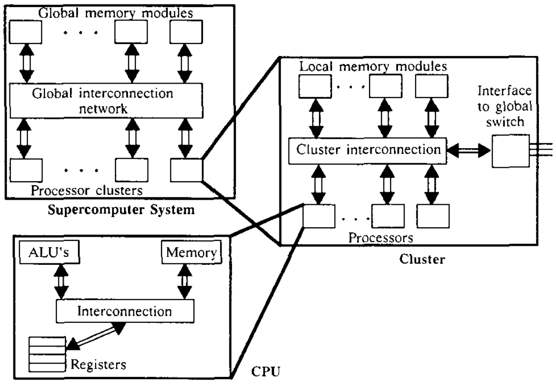

Fig. 9. Hierarchy of interconnections in supercomputers.

B. Module-to-Module and Board-to-Board Optical Switching

Inter-module and inter-processor communication for large supercomputer systems can require enormous bandwidth and a high degree of parallelism. In order to satisfy these requirements the most feasible, but rather costly, solution is to use crossbar or multistage interconnection networks (MINs) based on the space division configuration. With crossbar switches not economically feasible, many leading supercomputers including Cedar [57], GF11 [58], and RP3 [59] employ MINs at this level, which in turn are composed of small crossbar switches. In Cedar, for example, the global switch is a 2-stage shuffle network consisting of 16 8 × 8 crossbar switches. Although the cost of MINs is moderate in general, the control of these switches is much more complex than a crossbar switch. We now discuss current optical switching technology that can be used for inter-module communication.

1) Guided Waves for Space Division Switching: Optical guided wave technology has shown promising results in building large space-division switches, since an increase in the optical path does not reduce the bandwidth of the link. This is due to absence of stray capacitance which is experienced in electrical interconnections and negligible attenuation on optical fiber over short distances. The switching networks built with this technology use dense fabrication of optical waveguides on a single chip. This device can also be an optical shutter in the form of a logic gate [25]. Based on optical switches and gates, a number of prototype crossbar networks have been built.

A 2 × 2 optical switch is a four-port device such that two input ports can be connected directly or crossed to two output ports. The key operating parameters for such a switch are the insertion loss, crosstalk level, wavelength response, switching time, and precision. The connection in the device can be controlled with an electric voltage, an acoustic signal, or a beam of light. The most commonly used optical switch is the directional coupler which implements a 2 × 2 switching element and consists of two parallel optical channels (waveguides), as shown in Fig. 10(a) [60], [61].

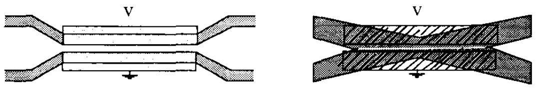

(a) (b)

Fig. 10. Directional optical couplers.

Under the normal mode of operation the light signal in one channel couples to the other, which results in the exchange connection of the switch. Under the application of control input (DC voltage) across the channels, the refractive index can be changed and the coupling between the channels becomes negligible. This forces the light to propagate within a single channel, which results in the straight connection of the switch. By changing the geometry of the channel layout, an improvement in switching capability can be obtained. One such switch, shown in Fig. 10(b), uses intersecting waveguides, in which the two channels intersect at a small angle, which is generally of the order of 1-8° [60]. An optical 8 × 8 crossbar interconnection network based on these switches is shown in Fig. 11 [62]. The advantage of this architecture is that the intersecting waveguides are straight, with slow bends at the edge of the switch, making fabrication simpler.

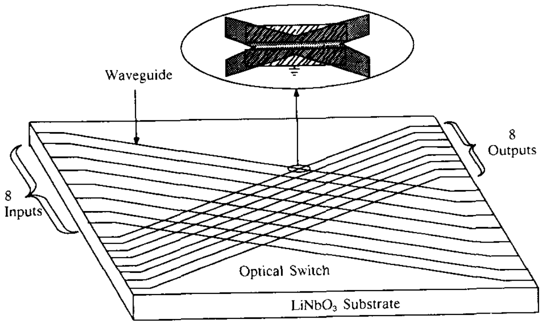

Fig. 11. An optical 8 × 8 crossbar network based on X-switches [62].

Many directional coupler-based interconnection networks have been implemented experimentally. The reported results show that low-loss switches can be easily fabricated using, for example, single-mode fiber. The largest experimental single-chip interconnection networks are 8 × 8 nonblocking networks [63]. One such network uses 64 directional couplers with DC voltage as the control input. The overall length of the chip is 60 mm. The data-handling capacities of these networks are in the range of 800 Mbits/s. Interesting experimental demonstrations have shown that self-routing of packets in these networks can be achieved using a scheme based on correlation and spread spectrum [64]. Each destination bit is encoded using a spread spectrum bandwidth expansion code. An optical controller, which is essentially an incoherent fiber-optic delay-line correlator, performs the correlation function on the spread spectrum sequence of the destination address. The output of the correlator is converted into electrical switching pulses which ultimately control the optical switching elements. The combination of this type of control and switch setting can then be used to realize the required permutation for routing.

Another promising technology used to build low-loss crossbar switches is based on polarizing gates, that can be dynamically reconfigured to control the input light [55], [65]. The operation of a crossbar network is viewed as a matrix-vector multiplier architecture, as shown in Fig. 12. N serial data input lines are connected to N serial data output lines (each one-bit wide) in form of a matrix-vector product $\vec{c} = \mathbf{A}\vec{b}$ [66]. The vector $\vec{b}$ is a binary column vector of length N of the data on the input lines, while the vector $\vec{c}$ represents output data. The $N \times N$ binary matrix $\mathbf{A}$, whose elements are 0's and 1's, is a generalized permutation matrix and represents the interconnection pattern provided by the switch. An entry of 1 in row i and column j means that input j is connected to output line i.

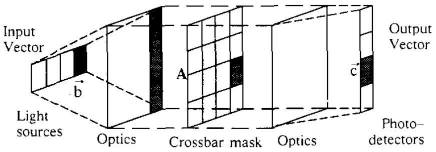

Fig. 12. $N \times N$ parallel matrix-vector inner product processor [66].

As shown in Fig. 12, the optical implementation of this network assumes free space communication. The N input lines drive an array of N light sources which are either light emitting diodes or laser diodes and can be turned on and off to represent binary input signals. This array is followed by a light focusing system which spreads the light from each input source vertically on a column of a crossbar mask. The crossbar mask provides the desired interconnection pattern. It consists of a set of optical shutters which can be dynamically controlled by using various types of control inputs such as electrooptical or acoustooptical signals or beams of light. Each gate can either be a spatial light modulator, a laser diode, or a photoconducting device. By controlling the gates at appropriate positions in the matrix, the desired interconnection can be achieved. A generalization to a matrix-matrix architecture, $\mathbf{C} = \mathbf{A} \times \mathbf{B}$, is possible. This allows interconnection of N parallel input lines to N parallel output lines, each M bits wide.

In another approach, input and output vectors are fully connected to a switching matrix by means of passive fiber-optic splitters and combiners, as shown in Fig. 13. In this switch, four incoming biconical fibers are connected to 4 vertical columns of a 2-dimensional optical switching array. The light from each fiber is coupled through its switching gate into a fiber on the opposite side of the array where rows are combined into four output fibers. Experimental systems based on this technology have been successfully demonstrated [55], [65]. Based on SLM technology the PLZT switch using electrooptic (EO) effect (or any magnetooptic (MO) device) can operate quite effectively. A number of 4 × 4 experimental switches based on both EO and MO effects have been demonstrated. In one such demonstration a network was operated at a channel rate of 139 Mbits/s, with average per-channel loss of 19.5 db and error rate of 10^{-9}. The average reconfiguration time for the switch was about 20 μs.

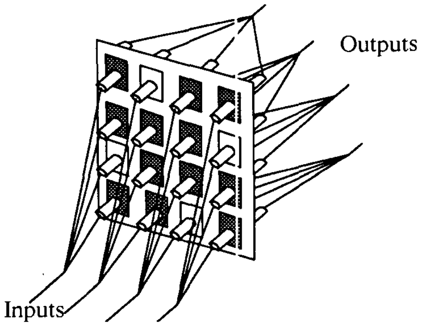

Fig. 13. A 4 × 4 fiber-optic crossbar switch [55].

In a second experimental switch a data rate of almost 400 Mbits/s was reported but at a cost of increased channel loss [65]. The high bandwidth was possible due to a smaller reconfiguration time which was about 5 μs. In this architecture the light is split from a single fiber to N fibers and combined from N fibers to a single fiber. These operations provide the network with the highly desirable property of full broadcast. On the basis of these experiments, it is predicted that 16 × 16 switches can be fabricated to provide a channel rate of 1 Gbits/s.

A further increase in data rate is possible by using wavelength division multiplexing, since the PLZT-based matrix switch has been shown to be relatively insensitive to the wavelength. These experimental systems reveal that low-loss optical switches with considerably high data rates are feasible, as compared to electronic switches for which metallic losses at high channel rates can be significantly higher.

The idea of building crossbar switches based on gate matrices has been further developed by using active optical-fiber splitters. The heart of this technology is the use of integrated laser-diodes as switching devices [65], [16], thus avoiding electrical-to-optical or optical-to-electrical conversion. A laser-diode is basically a nonlinear bistable optical device and has numerous advantages over other optical switches. Noted among them are low switching energy, smaller size for substrate length, and better switching dynamics in terms of time and precision. A number of experimental switching systems have been built based on this technology, using both SD and TDM configurations [67]. Using the well-known Benes MIN architecture, an interesting system has been proposed in [67], which allows self-routing of packets in the network. The destination bits in this scheme are used to control the injection current into the laser-diode which ultimately controls the voltage across a directional-coupler to route the message. The proposed system is quite economical, since a 2 × 2 switch has been fabricated in a 250 μm square, while a switch of size 16 × 16 can be fabricated on a 4 mm square chip.

The ultimate goal is the all-optical switching device, where an optical control signal causes another optical signal to switch between two paths. One such well known nonlinear optical device is the Fabry-Perot resonator [16], depicted in Fig. 6. It can have a switching speed of less than one picosecond. A successful fabrication of a 100 × 100 array of such a switch, with each device about 9 μm across, has been reported [68]. It is projected that optical NOR gates with a diameter of 1 μm can be fabricated based on Fabry-Perot etalons, which can be used to realize complex logic for switching systems.

It now appears feasible to build crossbar switches with rather small average loss incurred per channel. A number of efforts to build 32 × 32 and larger optical crossbar switches are underway. Also, small switches can be treated as the basic building blocks of larger switching systems for communication. The larger switches can provide extremely high bandwidth with considerably less power requirements than electronic switches. The interconnection of these crossbar building blocks can be in the form of Clos, Benes, Banyan, shuffle-exchange or Omega networks [57], [59], as is the case of Cedar and other leading supercomputer systems.

2) TDM and WDM for Optical Shared Bus Interconnection: As mentioned earlier, attenuation on optical fibers over short distances is negligible and bandwidth is not affected by the length of the fiber. Optical buses, therefore, offer an attractive choice for module-to-module and board-to-board communication. Currently, electronic buses are used at various levels in computer systems as backplane interconnection. However, in supercomputer systems the use of a single electronic bus for intermodular communication is not feasible due to its limited bandwidth. However, in the optical regime a fiber network can provide much larger bandwidth (order of Terabits/s) with moderate cost.

Within the last decade, a new generation of optical fiber based networks has been developed which use various multiplexing schemes, including TDM and WDM [69], [70], [71]. Typically, the processors are connected to a multiplexer which can concentrate the data from these processors and transmit it on optical fiber using TDM or WDM. Such a system generally uses a star coupler which provides a passive beam-splitting mechanism. Optical signals from each transmitter are multiplexed and combined by the coupler with the combined signal being equally split among the outputs. An important consideration for the operation of a TDM system is to maintain proper clock synchronization, which becomes a serious problem at high data rates.

The WDM configuration also uses a star coupler, however, this configuration does not encounter the problem of clock synchronization. In WDM, as shown in Fig. 14, a fixed optical wavelength is assigned to each transmitting module to broadcast its signal through the star coupler. Any receiving module connected to the coupler can tune to any desired wavelength to select an appropriate transmitting module.

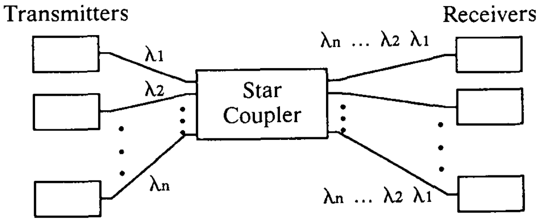

Fig. 14. Wave-division multiplexing of optical signals.

A number of experimental optical buses have been constructed. This technology is more popular for the local area network environment, but can also be used in computer systems. Multiple buses using TDM and WDM configurations can easily provide much higher data rates and are expected to serve as backplane interconnection for future supercomputer systems.

C. Chip-to-Chip Optical Switching

Large supercomputer systems are expected to have hundreds, if not thousands, of VLSI chips on each board. Each chip may contain as many as 100 000 gates, and require 1000 pins. There are a number of considerations when interconnecting these chips, the most important being the bandwidth. For example, interconnection of 100 chips supporting 32 bits in parallel at clock rates of about 20 Mbits/s may require a bandwidth exceeding 1 Gbits/s. Parallel data channels may require a large number of fan-outs, where a limited energy signal needs to drive multiple components. Also, these chips may be arbitrarily placed on the board. In general, interconnection at the chip level must support large data rates and must have a high degree of parallelism with no limit on fan-outs. Furthermore, it should be flexible enough to support arbitrary board layouts. These diverse requirements are difficult to meet with electronic technology. In this technology, the most flexible solution is to use a shared bus in the backplane. But, a large number of fan-outs can cause reflection and ringing of electrical signals on the bus. Also, an increase in fan-out increases the impedance of the bus which in turn requires VLSI chips to output a high power signal. An increase in power results in a larger area on the chip, which is not desirable [72]. On the other hand, interconnections based on metallic waveguides suffer from the problems of interference among channels and high stray losses. Optics can provide an efficient interconnection mechanism at this level.

1) Holographic Switching using Free-Space Communication: On a board the communication distances are small and hopefully no spiders are present; therefore, free-space can be effectively used for optical communication. This allows interconnection among various VLSI chips present on the board. Free-space communication allows exploitation of the third dimension and beams of light can cross each other without causing any appreciable interference. Furthermore, it allows the density of the interconnects to increase to the fundamental diffraction limit of optics with minimal crosstalk. Such interconnections can be implemented using holograms [53] and systems using noncoherent white light and bulk optics. These schemes are discussed in the following sections.

There are two types of interconnects based on free-space communication, unfocused and focused. The unfocused technique is used to broadcast optical signals from a single source to an entire area of sinks. These sinks are basically photodetectors, located on multiple chips. One important use of this form of communication is optical clock distribution. One such system is described in [53] where a source illuminates an entire chip. Due to the location of the source, the optical signal is received by the detectors with identical delays and therefore with no clock skew. The problem with this type of interconnect is its inefficiency because of the errors and possible improper operations, that may be caused by energy incident on portions of the chip where it is not wanted. Also, most of the optical energy is wasted since only part of it falls on the photodetectors.

To increase efficiency of optical signal distribution, the use of optical holographic techniques is promising. A hologram is used to change the direction of rays of light [73]. The heart of this system is a holographic optical element (HOE) which allows focusing of an optical source onto multiple detector sites simultaneously. A number of holographic imaging techniques based on HOE have been proposed for optical interconnection [53], [74], [75]. This mechanism can provide interconnection among multiple chips located on a board as well as among various functional units (CPU, memory, clock generator, etc.) within a chip.

In general, the interconnection requirements among multiple chips may vary drastically. In most of the cases multiple sources and sinks are arbitrarily positioned on their respective chips. Each signal source may need to be connected to several detectors for limited data broadcasting. Each detector may also receive inputs from several signal sources. The desired interconnection among them can be established in a very compact manner by arranging sources and detectors in a plane, with a reflection hologram suspended above them at a distance of about 5–10 mm, as shown in Fig. 15 [53]. The HOE used in this scheme is of the reflective type which can be fabricated as a surface-relief pattern etched into a silicon surface [74].

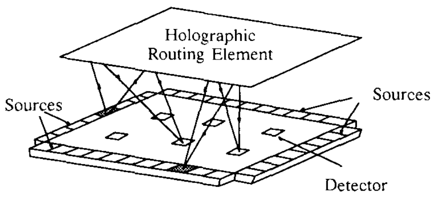

Fig. 15. Use of reflective type holograms for interconnection [53].

Since holograms are not easily changed, it is necessary that a holographic interconnection scheme be dynamically reconfigurable. One possible approach is to design programmable holograms so that they can allow dynamic coding of multiple interconnection patterns. Towards this objective, a novel hybrid approach based on a combination of optical switches and static transmitting holograms has been proposed and an experimental prototype switch (called Holoswitch) demonstrated [76]. The diagram of such a 4 × 4 switch is shown in Fig. 16. This scheme has been extended to build a system with 16 channels.

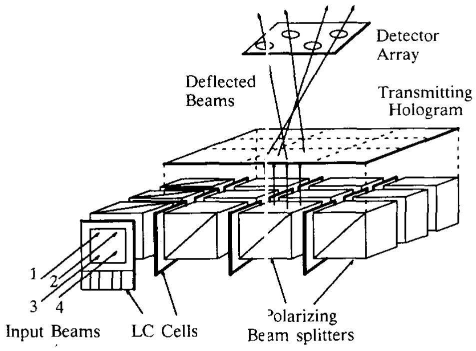

Fig. 16. Diagram of a 4 × 4 Holoswitch [76].

A Holoswitch consists of various optical components including liquid crystals (LC), polarizing beam splitters, and a transmitting hologram. The basic switching element is an electrooptical gate which is a combination of the twisted-nematic liquid crystal and a polarizing beam splitter. The electrical control input is applied to the liquid crystal to control its polarization, which in turn rotates the polarization direction of the input light signal. This light, when incident on the beam splitter, either passes straight through or is reflected at 90°, in order to select one of nine interconnection patterns. These patterns include a perfect shuffle, a linear string, a 2-D mesh, a binary tree, a fan-in to a single channel, and four types of row and column exchanges. The important performance parameters for the fully assembled switch include optical throughput, which is the designated output intensity relative to the input, number of channels supported, level of crosstalk, alignment accuracy, switching speed of the gates, and the compatibility with the other components in the computer system [76].

2) Switching Systems using Noncoherent White Light and Bulk Optics: Holographic interconnection schemes require the use of monochromatic light for input sources. The use of noncoherent white light illumination for free-space interconnection is also possible by using classical optics consisting of combinations of prisms, lenses, and optical masks. A number of important interconnection patterns using such an approach have been proposed. A few of these have also been implemented. These include crossbar switches; single stage and multistage perfect shuffle networks [77], [78]; the Clos network, which also allows one-to-many connection [79]; and a full broadcast interconnection [80].

The implementation of a crossbar network using this approach has already been shown in Fig. 12. Although, the design and implementation of a crossbar switch is rather straightforward and its control complexity is low, its major drawback is the number of switching elements needed, which is N^2. In its optical implementation it can require considerable precision in alignment if N is large. The reduced number of switching elements can be obtained by using reconfigurable multistage interconnection networks, such as the perfect shuffle or Clos networks, which only require $(N/2) \log_2 N$, 2 × 2 switching elements.

A perfect shuffle (PS) network can also serve as the basic interconnection network for a large number of MINs, due to its high flexibility and comparatively simple routing mechanism. Furthermore, its implementation requires a relatively small number of optical components. A number of parallel algorithms such as fast Fourier transforms (FFT), sorting, etc., can be executed using this network.

An optical implementation of the PS network is shown in Fig. 17 [77]. This system first stretches the inputs into two halves and then combines them using an exchange function [41]. The system consists of four prism wedges and two positive spherical lenses. The input light sources can be LEDs or SLMs. The first pair of prisms splits the inputs into two halves, with inputs stretched out in one direction so that the size of each half is equal to the size of the original input. Using a second pair of prisms, placed in the Fourier plane, a shift is applied to the inputs of the two halves, which allows interleaving of these inputs. This results in the desired shuffled output, which appears as an overall reversed sequence. Using other standard optical components, such as mirrors, lenses, and prisms, etc., the undesired reversal can be removed.

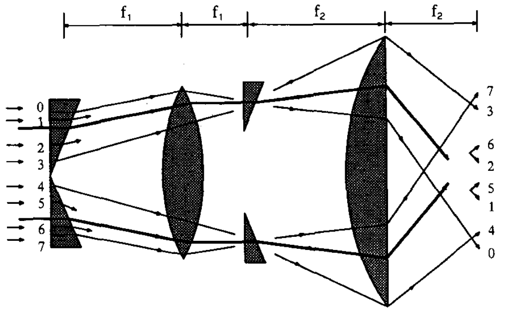

Fig. 17. Simple perfect shuffle set-up with two prism pairs [77].

Another important class of MINs is the Clos networks which have the same characteristics as that of a crossbar switch, that is, they are also nonblocking networks. A Clos network allows both one-to-one and one-to-many interconnections. A typical Clos network with the general set of parameters (p, q, r) is shown in Fig. 18. It connects $p \cdot r$ inputs to $p \cdot r$ outputs, using three types of switches arranged in three stages. The sizes of these switches are pxq, rxr and qxp. Each switch can be further decomposed such that it can be realized using small crossbar switches.

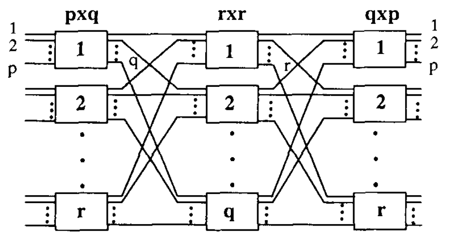

Fig. 18. A Clos (p, q, r) interconnection network.

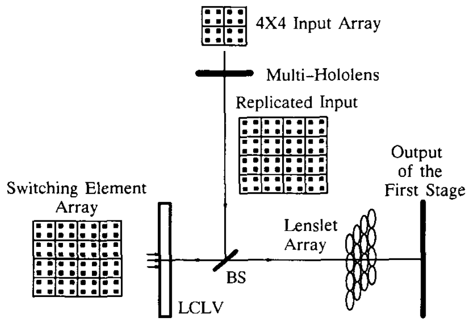

Fig. 19. Optical implementation of one stage of a (4, 4, 4) Clos network [79].

Fig. 19 shows an optical setup of the first stage of a 2-dimensional Clos (4,4,4) network, which has been experimentally built [79]. The complete network requires three such stages connected in a cascaded form. In this setup, a 4 × 4 input array is replicated 2 × 2 times using a combination of lenses (indicated as Multi-Hololens). The total of 64 lines, in the form of 4 quadrants, then shine on a beam splitter (BS). One of the BS outputs reflects back from a switching element array, consisting of liquid crystal light valves (LCLV). The second output, along with the reflected signals from the LCLV, is then collected by a 4 × 4 array of small lenses, which can feed the second stage of the network. By controlling the LCLV one-to-many connections are realizable. The proposed layout is quite flexible. One can form a system with a number of stages cascaded together, since the data paths are just straight parallel lines without any crossing. The layout has also been extended to carry out various single instruction multiple data (SIMD) processing involving matrix transposition, bivariate polynomial evaluation, 1-D and 2-D FFT computations [79]. The overall speed of the system is limited by the switching speed of the LCLV devices. The main limitation of this system is the precision in alignment of light beams. With the increase in the number of input lines, alignment becomes a major problem. Another limitation is fan-out and restoration of light energy as the signals flow through a series of stages, because restoration and amplification of light signals in free-space communication are more difficult than in a guided communication system.

For broadcasting data from a single memory system to an array of processors, a system called OPTIMUL [80] has been designed. The heart of this system, as shown in Fig. 20, is a lenslet array (called a fly's-eye lens), which allows a light signal coming from the main memory system to shine on an array of light shutters. The shutters are electrooptical gates which are connected to the local memory units of the processors. The role of these gates is to control the flow of data coming from the main memory to the local memory. These gates in turn are controlled by their respective processors. The main memory system is coated with a thin polymeric film, which when shined by a laser reflects the beam. The reflected beam is intensity-modulated by the electric field at each memory position and carries the complete bit map of the contents of the main memory. At the receiver end, after passing through the shutters, the signal is demodulated using photosensitive material. Although, no

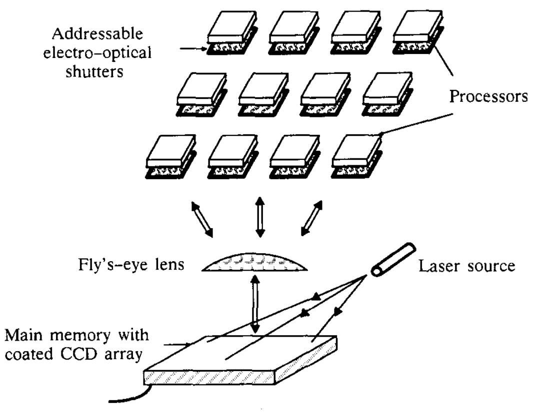

Fig. 20. The OPTIMUL interconnection configuration [80].

prototype exists, the proposed system is expected to provide a data rate on the order of 10^{14} bits/s.

3) Interconnection Using Guided Waves: The primary problem with using free-space focused interconnection techniques is the precision of the alignment which is needed to assure that the focused beams of light strike the appropriate places on the chip or board area. However, in order to provide high-speed communication among chips, the most effective solution is to use switches and waveguides (optical fibers). Switches based on couplers/gates similar to the one used for module-to-module communication can also be employed, provided the number of chips is small. Otherwise, the overall size of the switch can pose a serious problem if hundreds of VLSI chips on the board need to be interconnected. Similarly, direct optical waveguide (fiber) connections among chips can also be used and wired on the board, provided the number of chips remains limited. Both approaches are efficient because optical paths do not have the problem of crosstalk.

The use of an optical bus can provide a simple and cost-effective solution as opposed to using free-space communication or guided-wave switching. As mentioned earlier, an optical bus can provide an enormous data rate and can be used as backplane interconnection for hundreds of chips. The important parameter in a bus-based system is the bus fan-out. This parameter greatly reduces the applicability of electronic buses for chip-to-chip communication. Optical bus fan-outs are greater than those of electrical buses, since they are only limited by the minimum detectable power at the detectors. A laser source can easily deliver up to 10 db output power, while detectors for −40 to −30 db sensitivity at 1 Gbits/s have been designed. The availability of these components can easily meet the fan-out requirements for interconnecting large numbers of VLSI chips. A system using a TDM-based optical bus has been implemented. It multiplexes 4 ECL chips using a GaAs-based 4:1 multiplexer and can operate at a data rate of almost 2 Gbits/s [72].

If the number of chips is extremely large, then the combination of free-space holographic switching and multibus architectures can provide the desired interconnection.

D. Intra-Chip Optical Switching

With the level of integration on a chip increasing drastically over the years, communication among various functional areas on a chip is becoming more challenging. The current trend towards meeting this challenge is to deposit material on the silicon or GaAs device as a waveguide. At the same time the use of microstrips and striplines which can provide a bandwith on the order of Gbits/s is the leading emerging electrical technology. However, the problem of fan-out is more critical in electrical technology, since the power to drive the chip is limited due to small substrate and thermal radiation area. Also, each electric stripline path requires a certain minimum area to operate at a high frequency, otherwise the attenuation can be excessively large. Similarly, there is a minimum spacing requirement (at least 20 μm) to avoid crosstalk. These problems are not present in optical waveguides. However, in optical interconnections, problems arise in stringent packaging and fabrication requirements. For example, mechanical coupling between electrooptical components and waveguides needs to have very small alignment tolerance (less than 1 μm). Due to thermal expansion, this problem becomes more severe [81].

A more promising scheme is to extend holographic interconnection proposed for inter-chip communication, to intra-chip interconnection. The concept is essentially the same, except a hologram can provide interconnection among various functional areas on the chip. In this approach, a heteroepitaxial growth of GaAs on silicon can allow fabrication of surface-emitting type sources in the GaAs. Using a combination of HOE and imaging components, free space can be utilized for providing efficient interconnection at this level [53].

VI. Optical Processing

The use of optics for performing high-speed digital computation has intrigued researchers for almost three decades. Optics has a number of advantages over electronics, which include high speed, inherent parallelism, large connectivity, and a low level of interference among data channels. Furthermore, optics can be very effective for pipeline and dataflow computations. In earlier efforts, optics was used to process only analog waveforms for various signal processing applications, such as estimation, correlation, etc. But with the advancement of nonlinear optical devices, the interest in digital processing has grown rapidly. However, even after three decades, the use of optical technology for designing high speed processors is still in its infancy.

A number of research activities are currently underway with the purpose of building general-purpose digital optical computers [82]–[88]. Most of these systems mimic existing electronic computers by using optical logic gates, switches, interconnections, and memory elements. Consequently, progress in nonlinear optical devices is expected to have a major impact on optical processor technology, since these devices can be used to design complex optical logic circuits. This logic can in turn lead to the implementation of various units of a conventional electronic computer, such as arithmetic logic units, instruction set decoders, register stacks, and counters.

As mentioned earlier, high density physical integration of optical switching devices is currently feasible. Therefore, future optical processing systems are expected to have massively parallel architectures with fine-grain computations. Optically controlled devices can have switching speeds on the order of 10^{-15} seconds with as little power requirements as 10^{-6} watts. The resulting processor architectures, therefore, have the potential of higher performance with small power requirements. However, the current limitations of alignment and precision in optics is a major drawback to building highly reliable, massively parallel optical processors.

The currently proposed optical processors can be broadly classified into two categories: numeric and nonnumeric. These processors can be further categorized as analog and digital [89]. Digital processors use binary or mixed-binary data for carrying out the desired computation, while analog processors operate on continuous time signals. Our discussion is confined to digital processors, since the mode of computation in a conventional electronic computer is digital (binary).

A. Numeric Processors

Optical numeric processors are being designed to carry out CPU intensive computations and involve various mathematical calculations. Such computations are generally algorithmic in nature, such as matrix operations [90]–[93], signal processing algorithms to perform Fourier transformations [94], eigenvalue and eigenvector determination [95], finite impulse response filtering [96], Kalman filtering [97], and image processing [98], [99]. Being algorithmic, such computations are easily rendered to systolic and pipeline processing.

Exploration in designing complex optical functional units of a general purpose computer has grown steadily [100], [101]. Part of these efforts are motivated by the fact that special-purpose processors have a rather limited scope and if optics is to play an important role in computing, then general-purpose optical computers must be realized. A number of approaches have been proposed, which range from building systems that mimic conventional electronic computer systems to developing new theories for optical logic operations such as symbolic substitution.

1) Matrix Multiplication: The approach shown in Fig. 12 can be generalized to perform matrix-matrix multiplication: $\boldsymbol{C} = \boldsymbol{A} \times \boldsymbol{B}$. Using the outer-product algorithm, matrix multiplication of two $N \times N$ matrices, $\boldsymbol{A}$ and $\boldsymbol{B}$, can be optically realized as shown in Fig. 21. The matrix $\boldsymbol{C}$ is obtained by adding N matrices, each of rank one. These matrices are generated by taking the outer product between column vectors of $\boldsymbol{A}$ and the corresponding row vectors of $\boldsymbol{B}$. In general, the complete computation can be obtained as follows:

$$C = \sum_{i=1}^{N} C^i = \begin{pmatrix} a_{1i}b_{i1} & a_{1i}b_{i2} & \cdots & a_{1i}b_{iN} \\ a_{2i}b_{i1} & a_{2i}b_{i2} & \cdots & a_{2i}b_{iN} \\ \vdots & \vdots & \vdots & \vdots \\ a_{Ni}b_{i1} & a_{Ni}b_{i2} & \cdots & a_{Ni}b_{iN} \end{pmatrix}$$

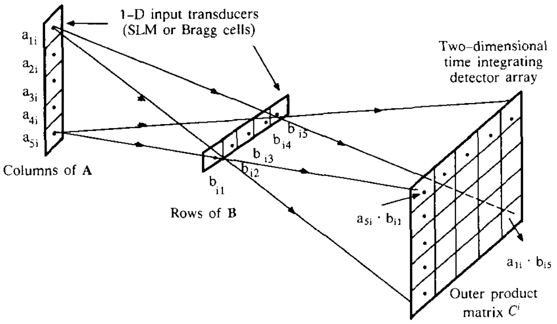

Fig. 21. An optical set-up for outer product [91].

If both $\boldsymbol{A}$ and $\boldsymbol{B}$ are $N \times N$ matrices, the above computation can be represented in a simple algorithmic form as follows:

```
   DO 10 $i = 1,N$
     DO 10 $j = 1,N$
       DO 10 $k = 1,N$
10 $C_{ij} = C_{ij} + A_{ik} \cdot B_{kj}$
```

Note that the operations for the outer two DO loops can be parallelized, such that N^2 multiplications and N^2 additions can be performed in parallel. The operations of the inner loop can be performed using time integration on the elements of C^i through a 2-D array of photodetectors. The optical realization to calculate the C^is is in Fig. 21.

In this set-up, the cells of the transducers are SLMs. A column and a row can be fed to these SLMs, which can be EO or AO devices. Performing multiplication by shining data values from one transducer onto the other is possible due to the spatial parallelism of optics. The result of the product is received by the 2-D array. The detectors keep on integrating the outputs received for all the outer products. The final result is the total time integration of all the outer products.

The performance of this simple set-up depends on a number of factors. The most important are the response time of the SLMs to store new input data, the noise characteristics, and the response time of the photodetectors used to collect the result of the outer product.

The set-up shown in Fig. 21 can serve as a functional unit to implement various iterative algorithms for which matrix-multiplication is the fundamental computation. Noted among these algorithms is LU decomposition, which is used to solve a set of linear algebraic equations. The generation of the $\boldsymbol{L}$ and $\boldsymbol{U}$ triangular matrices from a matrix, say $\boldsymbol{A}$, uses Gaussian elimination, which requires matrix multiplication to perform elementary row operations.

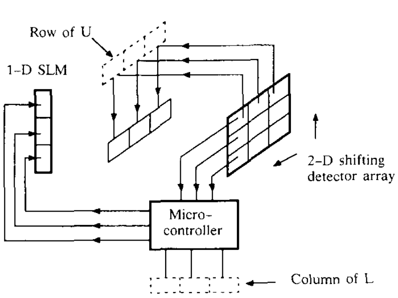

Fig. 22. Optical iterative computation for LU decomposition [91].

The schematic diagram for the LU decomposition algorithm using optical components is shown in Fig. 22 [91]. In this set-up, the matrix $\boldsymbol{A}$ is stored in a 2-D photodetector array with the capability that its elements can be shifted simultaneously vertically (row movement) and horizontally (column movement) across the array. The rows and columns of ***A*** are fed to the input devices (SLMs in this case), to perform the desired outer products. The elements of ***L*** and ***U*** ultimately appear in the array of detectors.

Optical set-ups for various other numeric processors can be found in [83], [87], [88], [98], [100]. All these schemes, like the ones described above, exploit the spatial and temporal parallelism of optics in one form or another. However, their performance greatly depends upon the performance of the optical devices as well as interfacing with other computer components, clocking of optical data, and control of the flow of data.

2) Optical Arithmetic Logic Unit: Various optical nonlinear and bistable devices can be used to realize important logic operations such as AND, OR, and EX-OR. Thus, these devices can be interconnected to realize complex boolean expressions to support arithmetic and logic functions. One such approach [101] directly transforms the design of a general-purpose electronic arithmetic logic unit (ALU) into the optical domain. The system uses binary data representation and only optical components for its operation. The proposed architecture implements four fundamental binary arithmetic functions, namely, addition, subtraction, multiplication, and division, which are generally provided by the conventional electronic ALUs. These functions can be represented by the following set of boolean equations:

$$S = (A \oplus (B \oplus X) \oplus C_i)F + A\bar{F}$$

$$D = B$$

$$C_{i+1} = (B \oplus X)(A + C_i) + AC_i$$

where A, B, C_i, are the input data lines and S, D, and C_{i+1} are the desired outputs. The fundamental logical operation in these equations is the EX-OR (half-adder) operation, which can be realized using a suitable combination of nonlinear optical devices, a set of mirrors, and beam splitters. One such layout is shown in Fig. 23 [101].

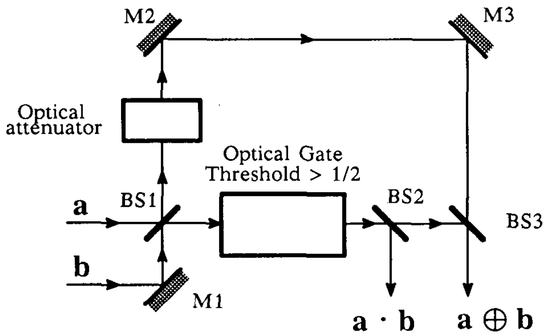

Fig. 23. Optical realization of a half adder [101].

The gate shown in this figure can be any bistable device which can be implemented using an optooptical switch, such as the Fabry-Perot etalon described earlier. This scheme is representative of many other schemes which are based on mapping functions of electronic components directly into the optical regime. However, such schemes also suffer from the same limitations as encountered by special-purpose optical processing approaches.

A number of other approaches have been proposed for building general-purpose optical processors. These include *symbolic substitution logic* (SSL) [102]–[104] *and residue arithmetic* [105]. The former approach exploits the third dimension available in the optics domain to generate a superset of boolean logic. SSL can be functionally divided into two phases: recognition and substitution. Binary patterns of 1's and 0's are recognized in parallel and substituted by other output patterns. These two phases can be used to define various boolean operations such as AND, OR, EX-OR, and NOT.

In the residue arithmetic approach, a set of relatively prime numbers are selected, which are termed moduli and used to represent integers. The main advantage of this scheme is that calculations for the different moduli can be performed in parallel, since there is no interaction among them. Due to this inherent parallelism, the approach is easily implementable in optics. A number of implementations have been proposed, based on this concept, with a few of them being analog. These include diffraction gratings and interferometric conversion [106], which are basically used to encode and decode data. The use of the conventional truth table can be extended beyond logic functions and general arithmetic functions based on residue arithmetic can also be tabulated by using precalculated results. Since table look-up is very attractive using optics, a number of approaches based on this general concept have been reported [105].

3) Optical Correlator: The computation of correlation is a lengthy and time-consuming process for electronic computers but it can be carried out extremely fast using optics [107]. A hybrid optoelectronic system could be realized in which correlation and other similar operations (convolution, Fourier transforms, matched filtering, etc.) are performed using an optical set-up connected to an electronic computer via arrays of laser diodes and photodetectors.

Figure 24 depicts an optical accelerator that performs digital correlation between an input two-dimensional pattern and a reference pattern. The input pattern is stored on the first spatial light modulator SLM_1, while an appropriate reference filter is loaded on SLM_2. Both of these modulators are two-dimensional, electrically-addressed devices operating in the transmissive mode. The correlator can have thousands of digital filters available for read-in on SLM_2. The projected throughput for such a system is 50 000 frames/s, which equates to 100,000 MOPS for a 128 × 128 SLM [40]. A similar design can be used effectively for digital pattern recognition.

B. Nonnumeric Processors

1) Optical Text Processor: One of the major problem areas in the information storage and retrieval field concerns full text search. When the text database contains newspaper articles, case histories, or legal briefs, one must often resort to full text search of the data. Even with the most sophisticated computing equipment, this process may be time-consuming and expensive. Because of its inherent speed and bandwidth we believe that optics may offer some future solutions to these problems.

Almost every full text search operation requires the comparison of the input data to a reference word or words. Following an approach in [85], [108], optical comparisons can be performed based on the exclusive-or (EX-OR) primitive using dual-rail logic as shown in Fig. 25. Two *n*-bit words

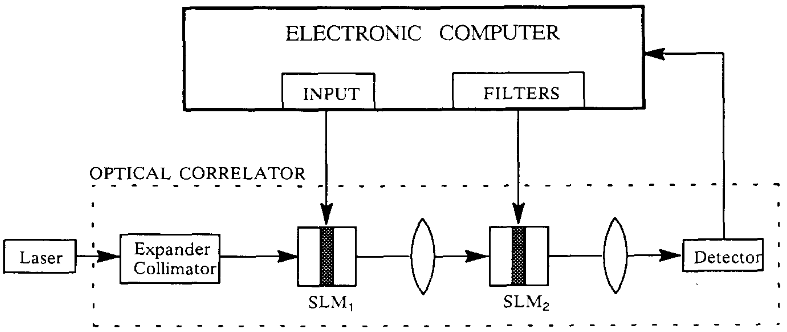

Fig. 24. An optical correlator connected to an electronic host.

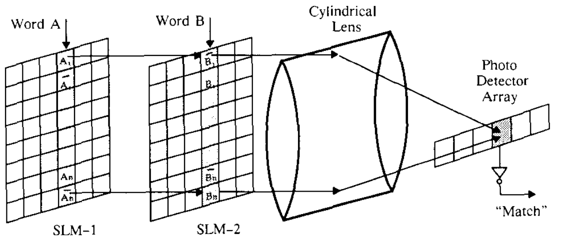

Fig. 25. Two-array configuration for multiple optical comparisons.

A and B are equal if $A_i\bar{B}_i + \bar{A}_iB_i = 0$ for each pair of corresponding bits A_i and B_i. For the comparison, both the value of the bit and its complement are needed, thus each n-bit word will be represented by $2n$ light beams (i.e., the 4-bit word 1011 will become 10-01-10-10). Using this method a 00 combination corresponds to a "don't-care" character while the 11 combination always produces a "Not-Equal" result. The coding scheme for the two logical values, 1 and 0, can be either light and no light, or horizontal and vertical polarization, respectively. The light beams are superimposed bitwise and are focused by means of a cylindrical lens on a single photodetector, which performs the logical OR (or summation) of all the beams. If no light is detected, the two words are equal, while any level of light intensity other than zero indicates that the two words differ in at least one bit.

Extending these ideas, multiple word comparisons can be made using the arrays of Fig. 25. In this case, we assume that the search arguments are latched in SLM-2 while data to be searched are stepped through SLM-1. In [109] a full range of text operations is developed based on this technique.

2) Optical Data/Knowledge Base Machines: A common approach in dealing with very large knowledge bases is the incorporation of a Data/Knowledge Base Machine (D/KBM). A D/KBM with multiple storage units, multiple processors, and the appropriate interconnection network, operates as a back-end to a host computer. D/KBMs must have very large storage capacity, a high degree of parallelism, and specialized processing units. Many electronic D/KBMs have been proposed but only a few have been implemented. While electronic processing may be acceptably fast, the Input/Output process remains a severe bottleneck, particularly when a transaction requires a global search through the entire database.

Optics can offer an alternative solution [110]–[112]. The massive data rates potentially available from modified optical disks or POHM will have a profound effect on the I/O problem. Information exchange in distributed databases as well as intracommunication within data/knowledge base machines will benefit by the speed and parallelism of optical interconnections. Furthermore, a number of relational database operations can be performed on the optical data, providing a much richer input to the electronic host. In [111] an optical configuration based on the scheme depicted in Fig. 25 is used to perform selection, projection, and equijoin.

VII. CONCLUSION

In this paper we have attempted to assess the impact of optics on supercomputing. We have presented our remarks in three major categories: storage, interconnection, and processing.

Storage has benefited greatly from progress in optical technology, mostly manifested as optical disks. With large storage densities and the apparent solution to the read/write problem, optical disk technology will play an expanding role in mass storage for supercomputer systems. Because of their unique characteristics, massive data rates (100's of megabytes per second) from optical disks are feasible and have been demonstrated in the laboratory. When these devices become a commercial reality, they will have a far-reaching impact on supercomputing. In the last few years there has been a rekindling of interest in holographic memories. These devices offer the promise of even greater data rates than optical disks. However, considerable research still needs to be conducted before holographic memories become a commercial reality.

Communication has long been a strength of optics. With the speed and bandwidth of optical communication via fibers, waveguides, and free space, massive amounts of information can be communicated among computing systems. Next logical steps are the use of optics for interconnection at the module, board, chip, and gate levels. We believe that progress will be made by moving down this hierarchy, thus capitalizing on current communication technology. However, considerable research and development must be conducted on optical device technology. Although optical devices have been built in a laboratory with extremely fast switching capability (femtoseconds) large switching systems are yet to be realized.

The least well developed aspect of optics applied to supercomputing is processing. It is quite clear that, at least in theory, optical processing can have a major impact on computing in general and supercomputing in specific. This is supported by various studies that indicate massive performance improvement based on speed and parallelism. However, optical device technology remains undeveloped. Considerable research is being conducted with the objective of developing optical devices for processing as well as switching with expected payoffs in the 1990s. Only time will tell whether the potential performance improvements of optical processing will become a reality.

In the next decade the major impact of optics will be felt through electrooptical systems. The speed and parallelism of optical interconnection, the massive data rates from optical storage media, and the superior processing capabilities of electronic technology will combine to yield new classes of supercomputers.

REFERENCES

[1] T. S. Perry and G. Zorpette, "Supercomputer experts predict expansive growth," *IEEE Spectrum*, vol. 26, pp. 26-33, Feb. 1989.

[2] Y. S. Abu-Mostafa, D. Psaltis, "Optical neural computers," *Scientifc American*, pp. 88-95, Mar. 1987.

[3] D. Psaltis, A. A. Yamamura, M. A. Neifeld, and S. Kobayashi, "Parallel readout of optical disks," in *Proc. of OSA Top. Meeting on Optical Computing*, vol. 9, (Salt Lake City, UT), pp. 58-61, Feb. 1989.

[4] P. P. Chen, "The compact disk ROM: How it works," *IEEE Spectrum*, vol. 23, pp. 44-49, Apr. 1986.

[5] G. Bouwhuis, J. Braat, A. Huijser, J. Pasman, G. V. Rosmalen, and K. Schouhamer Immink, *Principles of Optical Disc Systems*. Bristol, UK: Adam Hilger Ltd., 1985.

[6] J. Isailovic, *Videodisks and Optical Memory Systems*. Englewood Cliffs, NJ: Prentice-Hall, 1985.

[7] R. Barton, C. R. Davis, K. Rubin, and G. Lim, "New phase-change material for optical recording with short erase time," *Appl. Phys. Lett.*, vol. 48, pp. 1255-1257, May 1986.

[8] H. Oba, M. Abe, M. Umehara, T. Sato, Y. Ueda, and M. Kunikane, "Organic dye materials for optical recording media," *Appl. Opt.*, vol. 25, no. 22, pp. 4023-4026, Nov. 1986.

[9] J. S. Hartman and M. A. Lind, "Erasable bilayer dye-polymer optical recording medium," in *Proc. OSA Top. Meeting on Optical Data Storage*, vol. 10, (Stateline, NV), pp. 155-158, Mar. 1987.

[10] T. Murakami, K. Taira, and M. Mori, "Magnetooptic erasable disk memory with two optical heads," *Appl. Opt.*, vol. 25, pp. 3986-3989, Nov. 1986.

[11] R. P. Freese, "Optical disks become erasable," *IEEE Spectrum*, vol. 25, pp. 41-45, Feb. 1988.

[12] G. J. Ammon, J. A. Calabria, and D. T. Thomas, "A high-speed, large-capacity, 'jukebox' optical disk system," *IEEE Comput.*, vol. 18, pp. 36-45, July 1985.

[13] W. P. Altman, G. M. Claffie, and M. L. Levene, "Optical storage for high performance applications in the late 1980's and beyond," *RCA Engineer Mag.*, vol. 31-1, Jan./Feb. 1986.

[14] W. Carpenter and G. Herman, *Comparative Analysis of Online Data Storage Technologies*, Tech. Report MTR-86W117, MITRE Corp., Oct. 1986.

[15] P. B. Berra and S. J. Marcinkowski, "Optical content addressable memories for managing an index to a very large data knowledge base," *Data Engineering Bull.*, vol. 11, no. 1, Mar. 1988.

[16] H. Gibbs, *Optical Bistability: Controlling Light with Light*. London: Academic Press, 1986.

[17] See Special Issue on Photonic Switching, *IEEE Commun. Mag.*, vol. 25, no. 5, May 1987.

[18] H. Gibbs and N. Peyghambarian, "Nonlinear etalons and optical computing," *SPIE* 634, Optical and Hybrid Computing, pp. 142-148, Mar. 1986.

[19] M. Warren, S. Koch, and H. Gibbs, "Optical bistability logic gating and waveguide operation in semiconductor etalons," *IEEE Computer*, vol. 20, no. 12, pp. 68-81, Dec. 1987.

[20] S. D. Smith, J. G. H. Mathew, M. R. Taghizadeh, A. C. Walker, B. S. Wherret, and A. Hendry, "Room Temperature, visible wavelength optical bistability in ZnSe interference filters," *Opt. Commun.*, vol. 51, pp. 357-362, Oct. 1984.

[21] D. A. B. Miller, D. S. Chemla, T. C. Damen, A. C. Gossard, W. Wiegmann, T. H. Wood, and C. A. Burrus, "The quantum well self-electro-optic effect device: Optoelectronic bistability and oscillation, and self-linearized modulation," *IEEE J. Quantum Electron.*, vol. QE-21, pp. 1462-1476, Sep. 1985.

[22] D. A. B. Miller, "Quantum wells for information processing," *Opt. Engineer.*, vol. 26, pp. 368-372, May 1987.

[23] T. E. Bell, "Optical computing: A field in flux," *IEEE Spectrum*, vol. 23, pp. 34-57, Aug. 1986.

[24] D. A. B. Miller, J. E. Henry, A. C. Gossard, and J. H. English, "Array of optically bistable integrated self-electrooptic effect devices," in *Proc. of Conf.* on Lasers Electro-Optics, (San Francisco, CA) pp. 32-33, 1985.

[25] H. S. Hinton, "Architectural considerations for photonic switching networks," *IEEE J. on Selected Areas in Commun.*, vol. 6, pp. 1209-1226, Aug. 1988.

[26] C. Warde and A. Fisher, "Spatial light modulators: Applications and functional capabilities," in *Optical Signal Processing*, J. Horner, Ed. New York, NY: Academic Press, 1987, pp. 478-524.

[27] W. Penn, "Liquid crystals and spatial light modulators," Tut. T24, *SPIE* (30-th ITS on Optical and O-E ASE) Aug. 1986.

[28] G. R. Knight, "Interface devices and memory materials," in *Optical Information Processing—Fundamentals*, Topics in Applied Physics, vol. 48, New York, NY: Springer-Verlag, 1981, pp. 111-180.

[29] J. A. Grinberg, W. P. Bleha, L. Miller, L. Fraas, D. Boswell, and G. Myer, "A new real-time non-coherent to coherent light image converter—The Hybrid Field Effect Liquid Crystal Light Valve," *Opt. Eng.*, vol. 14, pp. 217-225, May/June 1975.

[30] P. Niselson and S. Iwasa, "Real-time optical processing with $Bi_{12}SiO_{20}$ PROM," *Appl. Opt*, vol. 11, pp. 2760-2767, Dec. 1972.

[31] C. Warde *et al.*, "Optical information characteristics of the microchannel spatial light modulator," *Appl. Opt.*, vol. 20, pp. 2066-2074, 1981.

[32] A. Fisher, L. C. Ling, J. Lee, and R. Fukuda, "Photoemitter membrane light modulator," *Opt. Eng.*, vol. 25, pp. 261-268, Feb. 1986.

[33] N. Sheridon, "The Ruticon family of erasable image recording devices," *IEEE Trans. Electron Dev.*, vol. ED-19, pp. 1003-1010, Sep. 1972.

[34] W. S. Colburn and B. J. Chang, "Photoconductor-thermoplastic image transducer,' *Opt. Eng.*, vol. 17, pp. 334-343, July/Aug. 1978.

[35] S. H. Lee, S. Esener, M. Title, and T. Drabik, "Two-dimensional silicon/PLZT spatial light modulators: Design considerations and technology," *Opt. Engineer.*, vol. 25, pp. 250-260, Feb. 1986.

[36] D. Armitage *et al.*, "High-speed spatial light modulator," *IEEE J. Quantum Electron.*, vol. QE-21, pp. 1241-1248, Aug. 1985.

[37] W. E. Ross, D. Psaltis, and R. H. Anderson, "Two-dimensional magneto-optic spatial light modulator for signal processing," *Opt. Engineer.*, vol. 22, pp. 485-490, July/Aug. 1983.

[38] R. E. Brooks, "Micromechanical light modulators in silicon," *Opt. Engineer.*, vol. 24, pp. 101-106, Jan./Feb. 1985.

[39] J. N. Lee and A. D. Fisher, "Development issues for MCP-based spatial light modulators," in *Proc. of OSA Top. Meeting on Spatial Light Modulators and Applications*, vol. 8, (S. Lake Tahoe, NV), pp. 60-63, June 1988.

[40] J. A. Neff, "Future roles of spatial light modulators in advanced computing architectures," in *Proc. of OSA Top. Meeting on Spatial Light Modulators and Applications*, vol. 8, (S. Lake Tahoe, NV), pp. 70-73, June 1988.

[41] K. Hwang and F. A. Briggs, *Computer Architecture and Parallel Processing*. New York, NY: McGraw-Hill, Inc., 1984.

[42] "Butterfly Parallel Processor Overview, Version 1," BBN Tech. Report No. 6148, Bolt, Baranek and Newman Inc., MA, March 6, 1986.

[43] "Multimax Multiprocessor Systems; The Power of Parallel Thinking," Encore Computer Corp., Marlboro, MA, 1987.

[44] T. Cheung and J. Smith, "A simulation study of the Cray X-MP Memory System," *IEEE Trans. on Comput.*, vol. C-36, no. 7, pp. 613-622, July 1986.

[45] D. Bailey, "Vector computer memory bank contention," *IEEE Trans. Comput.*, vol. C-36, pp. 293-298, Mar. 1987.

[46] G. F. Pfister and V. A. Norton, "Hot spot contention and combining in multistage interconnection networks," in *Proc. 15th Inter. Conf. on Parallel Processing*, pp. 790-797, Aug. 1985.

[47] A. Ghafoor, P. R. Austin, D. L. Bray, and D. Wong, "The effect of high performance processors in butterfly multiprocessor system," *ACM Performance Eval. Rev.*, vol. 17, no. 1, 1989.

[48] Z. Cvetanovic, "The effect of problem partitioning, allocation and granularity on the performance of multiprocessor systems," *IEEE Trans. Comput.*, vol. C-36, no. 4, pp. 421-432, Apr. 1984.

[49] A. Silbey, V. Milutinovic, and V. Mendoza-Grado, "A survey of advanced microprocessors and HLL computer architectures," *IEEE Comp.*, vol. 19, no. 8, pp. 72-85, Aug. 1986.

[50] B. Lazzerini, "Effective VLSI processor architectures for HLL computers: The RISC approach," *IEEE Micro*, no. 2, pp. 57-64, Feb. 1989.

[51] DARPA, "Strategic computing: New generation computing technology," Defense Advanced Research Projects Agency, Arlington, VA, Oct. 1983.

[52] M. R. Feldman, S. C. Esener, C. C. Guest, and S. H. Lee, "Comparison between optical and electrical interconnects based on power and speed considerations," *Appl. Opt.*, vol. 27, no. 9, pp. 1742-1751, May 1988.

[53] J. W. Goodman, F. Leonberger, S. Y. Kung, and R. A. Athale, "Optical interconnections for VLSI systems," *Proc. IEEE*, vol. 72, no. 7, pp. 850-866, July 1984.

[54] J. W. Goodman, "Fan-in and Fan-out with Optical Interconnections," *Optica Acta*, vol. 32, no. 12, pp. 1489-1496, 1985.

[55] A. R. Dias, R. F. Kalman, J. W. Goodman, and A. A. Sawchuk, "Fiber-optic Crossbar Switch and Broadcast Capability," *Opt. Eng.*, vol. 27, no. 11, pp. 955-960, Nov. 1988.

[56] R. I. MacDonald, "Terminology for phontonic matrix switches," *IEEE J. Selected Areas in Commun.*, vol. 6, no. 7, pp. 1141-1151, Aug. 1988.

[57] E. Davidson, D. Kuck, D. Lawrie, and A. Sameh, "Supercomputing tradeoffs and the cedar system," CSRD Report no. 57, May 1986, University of Illinois.

[58] J. Beetem, M. Denneau, and D. Weingarten, "The GF11 supercomputer," *Proc. IEEE Int. Symp. Computer Architecture*, 1985.

[59] G. F. Pfister, W. C. Brantley, D. A. George, S. L. Harvey, W. J. Kleinfelder, K. P. McAuliffe, E. A. Melton, V. A. Norton, and J. Weiss, "The IBM research parallel processor prototype (RP3): Introduction and architecture," *Proc. 15th Inter. Conf. on Parallel Processing*, pp. 764-771, Aug. 1985.

[60] R. C. Alferness, "Waveguide electrooptic switch arays," *IEEE J. Selected Areas in Commun.*, vol. 6, pp. 1117-1130, Aug. 1988.

[61] K. Iizuka, *Engineering Optics*, 2nd Edition, New York, NY: Springer-Verlag, 1987.

[62] T. Yasui and H. Goto, "Overview of optical switching technologies in Japan," *IEEE Commun.*, vol. 25, no. 5, pp. 10-15, May 1987.

[63] P. Granestrand, B. Stoltz, L. Thylen, K. Bergvall, W. Doldissen, H. Heinrich, and D. Hoffmann, "Strictly nonblocking 8 × 8 integrated optical switch matrix," *Electron. Lett.*, vol. 22, no. 15, pp. 816-818, July 1986.

[64] P. R. Prucnal, D. J. Blumenthal, and P. A. Perrier, "Self-routing photonic switching demonstration with optical control," *Opt. Engineer.*, vol. 26, no. 5, pp. 473-477, May 1987.

[65] A. Himeno and M. Kobayashi, "4 × 4 optical-gate matrix switch," *IEEE J. Lightwave Tech.*, vol. LT-3, no. 2, pp. 230-235, Apr. 1985.

[66] A. A. Sawchuk, B. K. Jenkins, C. S. Raghavendra, and A. Varma, "Optical crossbar networks," *IEEE Comput.*, vol. 20, pp. 50-60, June 1987.

[67] R. Kishimoto and M. Ikeda, "Optical self-routing switch using integrated laser diode optical switch," *IEEE J. Selected Areas in Commun.*, vol. 6, no. 7, pp. 1248-1254, Aug. 1988.

[68] T. Venkatesan, B. Wilkens, Y. H. Lee, M. Warren, G. Olbright, H. M. Gibbs, N. Peyghambarian, J. S. Smith, and A. Yariv, "Fabrication of arrays of GaAs optical bistable devices," *Appl. Phys. Lett.*, vol. 48, no. 2, pp. 145-147, Jan. 1986.

[69] J. R. Erickson, R. A. Nordin, W. A. Payne, and M. T. Ratajack, "A 1.7 gigabit-per-second, time-multiplexed photonic switching experiment," *IEEE Commun.*, vol. 25, no. 5, pp. 56-58, May 1987.

[70] See Special Issue of *IEEE Network* on Optical Multiaccess Networks, vol. 3, no. 2, Mar. 1989.

[71] S. Suzuki, T. Terahado, K. Komatsu, K. Hagashima, A. Suzuki, and M. Kondo, "An experiment on high-speed optical time-division switching," *IEEE J. Lightwave Tech.*, vol. LT-4, no. 7, pp. 894-899, July 1986.

[72] L. D. Hutcheson, "Optical interconnect technology development," in *Proc. IEEE Int. Conf. on Parallel Processing*, pp. 448-456, 1986.

[73] P. R. Haugen, H. Bartelt, and S. K. Case, "Image formation by multifacet holograms," *Appl. Opt.*, vol. 22, no. 18, pp. 2822-2829, Sep. 1983.

[74] L. A. Bergman, W. H. Wu, A. R. Johnston, R. Nixon, S. C. Esener, C. C. Guest, P. Yu, T. J. Drabik, M. Feldman, and S. H. Lee, "Holographic optical interconnects for VLSI," *Opt. Engineer.*, vol. 25, pp. 1109-1118, Oct. 1986.

[75] R. K. Kostuk, J. W. Goodman, and L. Hesselink, "Design considerations for holographic optical interconnects," *Appl. Opt.*, vol. 26, no. 18, pp. 3947-3953, Sep. 1987.

[76] J. B. McManus, R. S. Putman, and H. J. Caulfield, "Switched holograms for reconfigurable optical interconnection: Demonstration of a prototype device," *Appl. Opt.*, vol. 27, no. 20, pp. 4244-4250, Oct. 1988.

[77] A. W. Lohman, W. Stork, and G. Stuck, "Optical perfect shuffle," *Appl. Opt.*, vol. 25, no. 10, pp. 1530-1531, May 1986.

[78] G. Eichmann and Y. Lai, "Compact Optical Generalized Perfect Shuffle," *Appl. Opt.*, vol. 26, no. 7, pp. 1167-1169, Apr. 1987.

[79] S. H. Lin, T. F. Krile, and J. F. Walkup, "Two-dimensional optical Clos interconnection network and its uses," *Appl. Opt.*, vol. 27, no. 9, pp. 1734-1741, May 1988.

[80] N. Matloff, S. Kowel, and C. Eldering, "OPTIMUL: An optical interconnect for multiprocessor systems," *ACM Inter. Conference on Supercomputing*, 1988.

[81] L. D. Hutcheson, P. R. Haugen, and A. Husain, "Optical interconnects replace hardwire," *IEEE Spectrum*, vol. 24, pp. 30-35, Mar. 1987.

[82] H. Jordan, "Digital optical computing with fibers and directional couplers," in *Proc. of OSA Top. Meeting on Optical Computing*, vol. 9, (Salt Lake City, UT), pp. 352-355, Feb. 1989.

[83] V. Chadran, T. E. Krile, and J. F. Walkup, "Optical techniques for real-time binary multiplication," *Appl. Opt.*, vol. 25, pp. 2272-2276, July 1986.

[84] M. Hashimoto, K. Kitayama, and N. Mukohzaka, "Sequential logic operation using optical parallel processor based upon polarization encoding," in *Proc. of OSA Top. Meeting on Optical Computing*, vol. 9, (Salt Lake City, UT), pp. 406-409, Feb. 1989.

[85] P. S. Guilfoyle and W. Jackson Wiley, "Combinatorial logic based optical computing architectures," *Appl. Opt.*, vol. 27, pp. 1661-1673, May 1988.

[86] S. H. Lee, "Mathematical operations by optical processing," *Opt. Engineer.*, vol. 13, no. 3, pp. 196-207, May/June 1974.

[87] Y-Z Liang and H. K. Liu, "Optical matrix-matrix multiplication method demonstrated by the use of a multifocus hololens," *Opt. Lett.*, vol. 9, no. 8, pp. 322-324, Aug. 1984.

[88] A. W. Lohman and J. Weight, "Digital optical adder based on spatial filtering," *Appl. Opt.*, vol. 25, no. 18, pp. 3047-3053, Sep. 1986.

[89] E. Abraham, C. T. Seaton, and S. D. Smith, "The optical digital computer," *Scientific American*, pp. 85-93, Feb. 1983.

[90] A. Huang, "Parallel algorithms for optical digital computers," in *Tech. Dig. IEEE 10th Int. Optical Computing Conf.*, pp. 13-17, 1983.

[91] R. A. Athale and J. N. Lee, "Optical processing using outer-product concepts," *Proc. IEEE*, vol. 72, no. 7, pp. 931-941, July 1984.

[92] D. Casasent, "Acoustooptic linear algebra processors: Architectures, algorithms and applications," *Proc. IEEE*, vol. 72, no. 7, pp. 831-849, July 1984.

[93] H. Rajbenbach, Y. Fainman, and S. H. Lee, "Optical implementation of an iterative algorithm for matrix inversion," *Appl. Opt.*, vol. 26, no. 6, pp. 1024-1031, Mar. 1987.

[94] J. W. Goodman, *Introduction to Fourier Optics*, New York: NY: McGraw-Hill, 1968.

[95] H. J. Caulfield, D. Dvore, J. W. Goodman, and W. Rhodes, "Eigen vector determination by noncoherent optical methods," *Appl. Opt.*, vol. 20, pp. 2263-2265, July 1981.

[96] P. A. Ramamoorthy, G. Govind, and S. Antony, "Optical implementation of systolic finite impulse response filters," *Opt. Engineer.*, vol. 26, no. 11, pp. 1169-1174, Nov. 1987.

[97] J. L. Fisher, D. P. Casasent, and C. P. Neuman, "Factorized extended Kalman filter for optical processing," *Appl. Opt.*, vol. 25, no. 10, pp. 1615-1621, May 1986.

[98] S. D. Goodman and W. T. Rhodes, "Symbolic substitution applications to image processing," *Appl. Opt.*, vol. 27, no. 9, pp. 1708-1714, May 1988.

[99] K. M. Johnson, M. A. Handschy, and L. A. Pagano-Stauffer, "Optical computing and image processing with ferroelectric liquid crystals," *Opt. Engineer.*, vol. 26, pp. 385-391, May 1987.

[100] T. Yatagai, "Cellular logic architectures for optical computers," *Appl. Opt.*, vol. 25, no. 10, pp. 1571-1577, May 1986.

[101] P. P. Banerjee and A. Ghafoor, "Design of a pipelined optical binary processor," *Appl. Opt.*, vol. 27, no. 22, pp. 4766-4770, Nov. 1988.

[102] N. Streibl, K-H. Brenner, A. Huang, J. Jahns, J. Jewell, A. W. Lohman, D. A. B. Miller, M. Murdocca, M. E. Prise, and T. Sizer, "Digital Optics," *Proc. IEEE*, this issue.

[103] F. T. S. Yu and S. Jutamulia, "Implementation of symbolic substitution logic using optical associative memories," *Appl. Opt.*, vol. 26, no. 12, pp. 2293-2294, June 1987.

[104] K-H. Brenner, A. Huang, and N. Streibl, "Digital optical computing with symbolic substitution," *Appl. Opt.*, vol. 25, no. 18, pp. 3054-3060, Sep. 1986.

[105] T. K. Gaylord, M. M. Mirsalehi, and C. C. Guest, "Optical digital truth-table look-up processing," *Opt. Engineer.*, vol. 24, pp. 48-58, Jan./Feb. 1985.

[106] C. Hannel, D. Schuocker, and E. Klement, "Interferometric converter: Core of an optical residue computer," *Appl. Opt.*, vol. 26, no. 6, pp. 1019-1023, Mar. 1987.

[107] D. G. Feitelson, *Optical Computing: A Survey for Computer Scientists*, Cambridge, MA: MIT Press, 1988.

[108] P. S. Guilfoyle and W. Jackson Wiley, "Combinatorial logic based optical computing," *Proc. SPIE*, vol. 639-617, Apr. 1986.

[109] P. A. Mitkas, P. B. Berra, and P. S. Guilfoyle, "An optical system for full text search," in *Proc. 12th Annu. Int. ACMSIGIR Conf. on Res. and Dev. in Information Retrieval*, (Cambridge, MA), pp. 98-107, June 1989.

[110] P. B. Berra and N. B. Troullinos, "Optical techniques and data knowlege base machines," *IEEE Comput.*, vol. 20, pp. 59-70, Oct. 1987.

[111] P. B. Berra, A. Ghafoor, P. A. Mitkas, S. J. Marcinkowski, and M. Guizani, "The impact of optics on data and knowledge base systems," *IEEE Trans. Knowledge Data Engineer.*, vol. 1, no. 1, pp. 111-132, Mar. 1989.

[112] P. B. Berra, K-H. Brenner, W. T. Cathey, H. J. Caulfield, S. H. Lee, and H. H. Szu, "Optical database knowledgebase machines," to appear in *Appl. Opt.*, 1990.

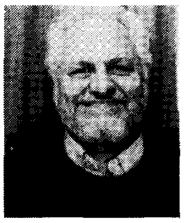

P. Bruce Berra (Fellow, IEEE) received the B.S. and M.S. degrees from the University of Michigan, Ann Arbor, and the Ph.D. degree from Purdue University, West Lafayette, IN. His industrial experience has included periods of service with Hughes, Bendix, and IBM. He is currently a professor of electrical and computer engineering and a member of the faculty of computer and information science at Syracuse University. He is also president of PBB Systems, an AI/database consulting firm. He previously held the position of professor and chairman of industrial engineering and operations research at Syracuse University, and his teaching experience includes posts at the University of Michigan's Dearborn Campus, Boston University, and Purdue University. His IEEE activities have included serving as general chair and program chair of the International Conference on Data Engineering, vice chair of the Publications Board, and member of the Distinguished Visitor Program. He has held a number of publications posts, including editor-in-chief of the Computer Society Press and editor of the *IEEE Transactions on Software Engineering*.

Dr. Berra is currently serving as a member of the Computer Society Governing Board, Chairman of the Transactions Advisory Committee, and editor of the IEEE Transactions on Computers.

Arif Ghafoor (Senior Member, IEEE) received the B.S. degree in electrical engineering from the University of Engineering and Technology, Lahore, Pakistan, in 1976, and the M.S., M.Phil., and Ph.D. degrees from Columbia University, New York City, in 1977, 1980, and 1985, respectively.

He has been a consultant to many companies, including Bell Labs and General Electric, in the area of telecommunications. He joined Syracuse University in September 1984 as an Assistant Professor. His current research interests include design and analysis of parallel and distributed systems, and optical engineering.

Dr. Ghafoor is a member of Eta Kappa Nu.

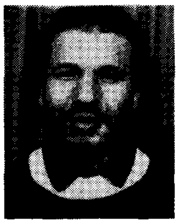

Mohsen Guizani received the B.S. and M.S. degrees in electrial engineering, and the M.S. degree in computer engineering, from Syracuse University, Syracuse, NY, in 1984, 1986, and 1988, respectively.

At present he is pursuing the Ph.D. degree at Syracuse, where he has been a Teaching and Research Assistant in the Department of Electrical and Computer Engineering since 1984. His research interests include solutions of linear systems using the conjugate gradient method (CGM), optical interconnection networks, and optical processing.

Mr. Guizani is a student member of IEEE and a member of Tau Beta Pi and Eta Kappa Nu.

Slawomir J. Marcinkowski received the B.S. and M.S. degrees in computer engineering from Syracuse University, Syracuse, NY, in 1986 and 1988, respectively.

At present he is pursuing the Ph.D. degree in computer engineering at Syracuse University, where he has been a Research and Teaching Assistant since 1986. He has also worked for Delco Products as a General Motors Scholar. His current research interests are in applying optics to database management, optical computing, computer architectures for very large data and knowledge bases, and parallel processing.

Mr. Marcinkowski is a student member of IEEE and the IEEE Computer Society, Tau Beta Pi, Eta Kappa Nu, and the Theta Tau National Engineering Fraternity.

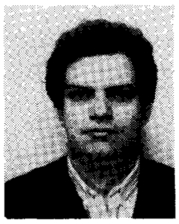

Pericles A. Mitkas received the Diploma in electrical engineering from the Aristotelian University of Thessaloniki, Thessaloniki, Greece, in 1985 and the M.Sc. degree in computer engineering from Syracuse University, Syracuse, NY, in 1987.

At present he is pursuing the Ph.D. degree at Syracuse University, where he was a Teaching and Research Assistant from 1985 to 1988. He is the recipient of a Syracuse University Fellowship for the 1988–1989 academic year. His current research interests include computer architectures for very large data and knowledge bases, optical data processing, and the application of optical computing to data, knowledge, and text bases.

Mr. Mitkas is a student member of IEEE and the IEEE Computer Society, ACM, and SPIE.