SYSTEM INTERCONNECTS

Due to the rapid advances in microprocessor, storage, and communication technologies, it has become feasible to develop increasingly more complex computing functions such as high performance computing, broadband multimedia (e.g., virtual reality, digital library, video on demand), massively multiparty online gaming, and business applications (e.g., data mining, data warehouse, content and knowledge management). However, the unrelenting demand for performance and quality in these applications creates pressure for further performance improvement on the Central processing unit (CPU) and the communication bandwidth within the system. Performance improvement thus far has been realized by advances in (1) processor architectures, such as the use of pipeline, superscalar, and parallel processing structures, and (2) Very Large Scale Integration (VLSI) technologies, such as the feature size reduction of CMOS. Yet speed of devices is not the only criterion that determines the acceptance of a technology. When these devices are used in a computing or communication system, problems of parasitics, circuit design, power consumption, and packaging often dominate over sheer speed. Therefore, it is insufficient to improve the performance of a system simply by choosing a faster device technology or a better system architecture. In particular, a number of factors such as the reflections due to electrical discontinuities of the transmission lines, crosstalk, skin effect, signal engineering, such as the use of a controlled-impedance transmission line, fully differential signaling (P1394, PCI express, InfiniBand), and limiting the interconnection structure to point-to-point (P1394, PCI express, InfiniBand) can push the metal interconnection technology to higher bandwidths than could be supported with existing technology. Nevertheless, the cost of the new technology is initially high, and there is no guarantee that the solutions can be scaled to clock speeds higher than their present limits. If any of these interconnection technologies is going to succeed, it must be available at low cost in high volume, and must be able to evolve to support future clock speeds of 3∼5 GHz and beyond.

Optical interconnections provide an alternative technology to solve the interconnection problem. The multigigabit bandwidth allowed by this technology is more than sufficient for applications such as communications within a multimedia system for the foreseeable future. Furthermore, it is easier to control reflections in this technology for both point-to-point and multidrop structures than in metal links (1). Optical links generally exhibit less ground-loop noise because fibers do not carry currents as do metal links. For these reasons, optical interconnections may be an attractive alternate technology for building high-speed board and backplane interconnections in future multimedia systems.

An optical link can be designed to be an almost one-to-one replacement for metal point-to-point or multidrop connections. The conventional line driver is replaced by a laser driver and an edge-emitting or surface-emitting laser diode/Light Emitting Diodes (LED), or a laser diode and an external modulator such as a Mach–Zehnder interferometer, directional coupler, total internal reflection (TIR) modulator spatial light modulator (SLM), self-electro-optic device (SEED), or vertical-to-surface transmission electrophotonic device (VSTEP) at the transmission end. The conventional line receiver is replaced by a light-sensitive device such as a PIN or metal–semiconductor–metal (MSM) photo detector and an amplifier at the receiving end. The light can be guided from the transmission end to the receiving end through single-mode or multimode fiber ribbon cable, polyimide, or silica-on-silicon channel waveguides, or free-space microlenses and/or holograms.

This article will survey a number of promising optical interconnect architectures and technologies. The goal is to investigate the potentials and limitations of optical interconnects. The organization of the rest of this article is as follows: The next section describes the existing metal interconnection hierarchy and potential impairments at high frequency. The advantages and drawbacks of optical interconnects are then discussed. Then a number of optical interconnect architectures and technologies are surveyed. Some of the research activities in this area during the past decade are then reported. The article concludes with a brief discussion.

EXISTING INTERCONNECTION HIERARCHY

Currently available packaging and interconnect technology at various packaging levels (as shown in Fig. 1) are chips, single-chip modules (SCM), multichip modules (MCM), cards, boards, backplanes, and cables between/among backplanes (2):

- *Chip-to-Package Interconnections.* These technologies include wirebonding, tape automated bonding (TAB), and flip-chip bonding using solder ball.
- *Ceramic and Plastic Chip Modules.* Each module made of ceramic or plastic encapsulation contains a single chip (single-chip module or SCM) or multiple chips (multichip module or MCM). The interconnections on these modules can have multiple signal layers using thin-film or thick-film processing techniques. (Thin-film packaging refers to packages in which the conductor and insulators are fabricated using deposition and patterning techniques similar to those used for fabricated integrated-circuit chips.)
- *Package-to-Board Interconnections.* Existing technologies can be categorized as pin-through-hole (PTH), leadless chip carrier (LLCC), and surface-mount technology (SMT). A through-hole on the printed circuit board is provided for each pin of a chip package in PTH. Both mechanical joint and solder joint are feasible for this technology. On the other hand, both LLCC and SMT, which are more area efficient and provide better signal quality, requires older joint between each lead of a package and the pad on a circuit board.
- *Printed-Circuit Board.* This technology has been around since before 1960. The progress over the past 30 years includes the decrease in the through-hole diameter (∼840 μm to ∼350 μm), the increase in the

J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*. Copyright © 2007 John Wiley & Sons, Inc.

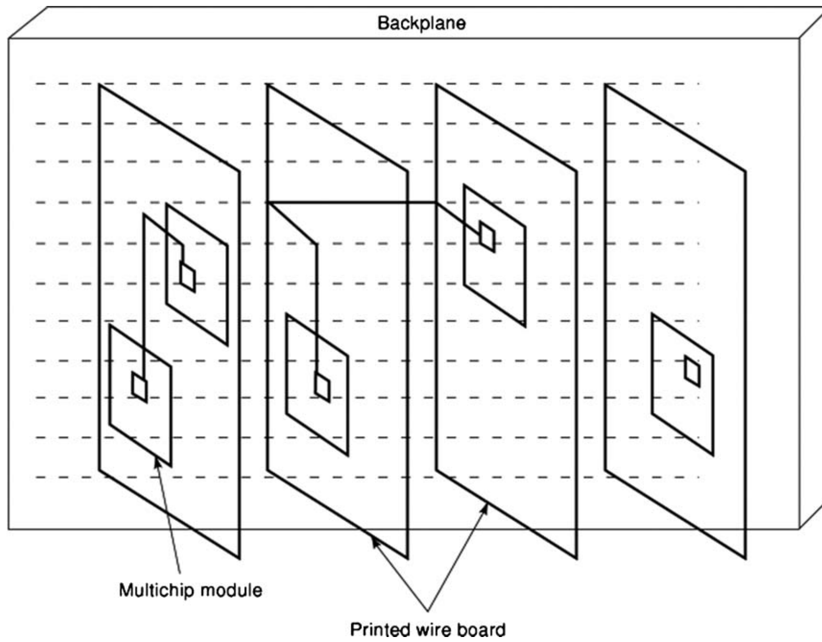

Figure 1. A typical packaging hierarchy, from bottom to top includes chips, multichip modules (MCM), cards, boards, and backplanes.

through-hole density (∼9/cm^2 to 64/cm^2), the increase in the number of signal planes (∼4 to 50), and the decrease of the interconnect width (∼250 μm to ∼50 μm).

- *Connectors and Cables.* A connector can be defined as a separable interface between the boundaries of two electronic elements. These elements may be unalike, similar, or identical. A connector usually provides connection inside component case, from component to PC board or wire, from PC board to wire or another PC board on a chassis, from internal chassis to another internal chassis in the same housing, case, or cabinet, or from one piece of equipment to another. An electric cable can be defined, in general, as a flexible (discrete or multiple) conductive path or link that has as its elemental building unit at least two conductors spaced uniformly apart by an internal dielectric, a single conductor, and a ground-return conductor. Cables can be further divided into open wire, where there is no shielding of the signal conductor, and closed wire, where there is shielding that provides enhanced noise immunity. The recent progress in this area includes smaller size, higher-reliability, zero-insertion force (ZIF) design, and various new materials which provide lower dielectric constants and low-dispersion signal propagation.

As the speed of devices increases, existing metallic interconnect technology is no longer adequate due to its performance degradation at high-frequency.

Sources of performance degradation include

- Reflections,
- Ground-loop noise,
- Crosstalk among adjacent interconnects,
- Frequency dependent signal distortion.

Reflections

A high performance system requires more than one level of packaging and interconnects to accommodate complicated logic functions. A typical packaging hierarchy includes chips, single-chip modules (SCM), multi-chip modules (MCM), cards, boards, and backplanes. However, electrical discontinuities exist between any two packaging levels. Discontinuities may be primarily inductive (such as electrical connectors) or capacitive (such as stubs in a multidrop net and 90-degree bends in a microstrip line). Depending upon the nature of the discontinuities and impedance changes, the resulting reflections may be either positive or negative.

Various methods exist to reduce the reflections resulting from impedance mismatch. For example, a termination resistor is usually placed at the receiving end of an interconnect in order to reduce the reflections. However, a perfect matching between the characteristic impedance of the interconnection and the impedance of the load is difficult to achieve because of the parasitic capacitance and inductance. If the round-trip propagation time between the source and the discontinuities is less than the rise time of the signal, these reflections can be absorbed by the interconnect driver with a net effect of an increased signal rise time. On the other hand, the waveform of the signal is severely degraded by the multiple reflections if the round-trip propagation time is longer than the rise time of the signal, resulting in a reduced noise margin or/and false switching.

Ground-Loop Noise

The ground plane of a packaging system usually cannot achieve zero resistance and inductance. Any local injection of current from the devices changes the electrical potential at that point. For a single-ended interconnection, the receiving side has to rely on the potential of the local ground plane as a reference to determine the amplitude of

the incoming signals. Any disturbance of the ground plane is therefore coupled in to the received signal.

One way of alleviating this problem is to transmit differential signals so that the signal can be interpreted unambiguously at the receiving end of an interconnect. However, the required interconnect density has to be doubled and thus more signal layers are necessary to accommodate the increased interconnect complexity. Some of the chips that are already pin-count-limited cannot afford this option either.

Crosstalk Among Adjacent Interconnects

For a given interconnect density, crosstalk between adjacent interconnects increases as the rise time of the signal decreases. Furthermore, crosstalk of the transmission lines with a TEM or near-TEM structure, such as slotted lines and microstrip lines, usually couples with switching noise and may consume the entire noise margin if they are not carefully controlled (3). TEM mode is the fundamental mode supported by a transmission line such as a coaxial cable. The transmission line structures that can support TEM mode are said to have TEM structures. Therefore, either the interconnects have to be spaced farther apart or additional shielding lines have to be inserted between signal lines to reduce crosstalk to an acceptable level. In both cases, the effective interconnect density is reduced.

Frequency-Dependent Signal Distortion

Packaging discontinuities introduce frequency-dependent signal distortion as a result of the inductive or capacitive nature of the discontinuities. Additional signal distortion is introduced by the dispersion and the skin effect of metal interconnects.

The microstrip lines on a printed circuit board are inherently dispersive, since they are incapable of supporting a pure TEM mode (4). The mode's effective dielectric constant is a function of frequency, causing different frequency components of the signal to travel at a different speed. This effect becomes significant when the rise time of the waveform is smaller than 100ps and the signal has to travel more than a few centimeters.

The skin effect also contributes to frequency-dependent signal distortion for metal interconnects when the thickness of the interconnects is large compared to the skin depth. Due to the skin effect, high-frequency components within the signal experience higher attenuation, yielding nonnegligible waveform distortion. In order to reduce the skin effect, the thickness of the metal has to be less than the skin depth of the metal. The skin depth of copper is 2 μm at 1 GHz, and becomes 0.7 μm at 10 GHz. A wider transmission line is thus required to accommodate signals with higher data rate while maintaining an acceptable DC and low-frequency loss, resulting in a net reduction of the interconnect density.

SYSTEM INTERCONNECT STANDARDS

The system level interconnects among processors, and between processors and I/O's that have been widely used commercially at various packaging levels include:

- PCI: The original PCI bus debuted more than a decade ago at 33MHz with 32-bit bus and peak bandwidth of 132MB/s.
- PCI-X: The PCI-X was an attempt to update PCI to allow it to extend its useful life for a few more years. It essentially doubles the bus width from 32 bits to 64 bits, and increases the clock rate from 66 MHz to 133 MHz and boost the bandwidth to 1GB/s. The latest version of PCI-X (PCI-X 266) also double-pumps the bus so that the data is transmitted on the rising and falling edges of the clock.
- PCI Express (PCIe): The PCIe spec was finalized in 2002. As opposed to the original PCI and PCI-X, PCIe adopts point-to-point topology, and a shared switch replaces the shared bus to route traffic among PCIe devices. PCIe links can be composed of from 1 lane (x 1 link at 2.5 Gb/s) to 32 lanes (x32 link at 80 Gb/s or 10 GB/s).
- InfiniBand (http://www.infinibandta.org): Similar to PCIe, Infiniband is also based on point-to-point bidirectional serial link. The serial connection's bandwidth is 2.5 Gb/s in each direction per connection. InfiniBand supports double and quad data speeds at 5 Gb/s and 10 Gb/s, respectively. The links can be aggregated in units of 4 or 12 (denoted as 4X and 12X, respectively). Most of the InfiniBand products in the market today are based on 4 X 2.5 GB/s connections.

OPTICAL INTERCONNECTS

Because of the bandwidth bottleneck associated with the existing interconnect and packaging technology, optical interconnect using free-space, optical waveguides, or optical fiber thus becomes a viable and attractive alternative to increase the total system throughput. In this section, issues associated with using optical interconnect for high-speed digital systems are investigated. In particular, we will examine the potential problems and solutions of using dense optical interconnects for high-performance multimedia systems. In such systems, serialization of data cannot be employed to increase the channel density if the data rate of each channel is very high before serialization is introduced. Therefore, an interconnect technology with the capability of providing high density and high bandwidth is necessary for acceptance in digital systems.

Potential Advantages

Dense optical interconnects have the potential of offering the following advantages:

- *More sophisticated interconnection pattern:* Light beams from different sources do not interfere with each other upon crossing. Very sophisticated 2-D and 3-D interconnect patterns based on planar optical waveguide and free space interconnect technologies, respectively, can thus be built from this principle, achieving a higher packaging density and shorter av-

erage signal propagation distance.

- *Electrical reflection reduction:* The reflections due to electrical discontinuities of a packaging system does not seriously affect the signal waveform, as long as the round-trip propagation delay is less than the rise-time of the signal waveform. Therefore, multiple reflections due to impedance mismatch between different levels of packaging can be reduced or eliminated by replacing metal interconnects on higher packaging levels (such as the boards and the backplanes) with optical interconnects so that the round-trip propagation delay of any metal interconnects is shorter than the signal rise time.
- *Higher bandwidth:* The bandwidth of the optical interconnects is mainly limited by the interface electronics and has the potential to achieve a multi-gigabit data rate with very little signal distortion.
- *Higher spatial density:* The potential spatial density of either optical-waveguide or free-space interconnect technology is an order of magnitude higher than what can be achieved by the current metal interconnect technology (line spacing between two thin-film metal interconnects is $\geq$25 μm with a propagation distance less than 7 cm and increases to $\geq$100 μm for longer distance in order to avoid large crosstalk between adjacent interconnects (2)).
- *Freedom from electromagnetic interference (EMI).* The propagation of light does not generate EMI to interfere with the surrounding circuit, nor can it be affected by the EMI produced by the environment. As we will show in later chapters, optical crosstalk between adjacent optical interconnects is usually negligible.
- *Breaking of ground-loops:* By using optical interconnects, current is no longer transferred between the transmitters and receivers, and thus the disturbance on the ground plane is reduced. In addition, optical signals in an optical interconnect can not be disturbed by the noise of the ground plane and therefore the signal quality is improved.
- *Greater flexibility:* Currently, many printers, Personal Digital Assistants (PDA), digital cameras, and virtually all of the new notebook computers are equipped with the IrDA (infrared Data Association)-compliant transceivers. These low bit rate (usually less than 4 Mbps) free-space infrared links provide a very inexpensive and convenient alternative to those heavy cables.

Potential problems

On the other hand, we also have to be aware of the potential problems if optical interconnects are used to replace metal interconnects:

- *Modal noise (5):* When multimode waveguides or fibers are used in conjunction with highly coherent lasers, the coherent interference of different spatial waveguide or fiber modes give rise to a speckle pattern. Fluctuations of the speckle pattern due, for example, to fluctuations in the spectrum of the optical source, can lead to modal noise if a mode-selective loss (such as a bad connector) is present in the optical link. Modal noise can cause a bit-error-rate (BER) floor which might not be tolerable in applications which require extremely low BER. Modal noise problem can be solved by either using a laser diode with large linewidth or premodulating the laser at a frequency comparable to the relaxation oscillation frequency of the laser diode (6).
- *Optical reflections (7):* Index discontinuities are also unavoidable in waveguide or fiber interconnects. Reflections from the laser/waveguide interface might increase the linewidth as well as the relative intensity noise (RIN) of the laser. Other reflections due to the discontinuities along the optical path degrade the signals arriving at the receiver by reducing the eye opening and increasing the RIN.
- *Optical crosstalk:* Optical crosstalk can occur at the coupling between laser array and optical waveguide array, between adjacent waveguides, or between the waveguide array and the photodetector array as a result of the high packaging density required by the system.
- *Threshold uncertainty (8):* The large number of interconnects within a digital system require all of the receivers to be set at the same threshold. In practice, this threshold cannot be individually adjusted according to the characteristics of the source. This means there is no feedback between the driver and the receiver to adjust the laser output, which deteriorates with time. Local feedback might be able to correct for this problem, but the added logic circuitry would compete for chip area with other logic circuitry.
- *High density required for optoelectronic components:* Each typical single-chip module (SCM) may have over 100 signal-I/O's, while a multichip module (MCM) can have several hundreds to several thousands of signal-I/O's. In order to provide optical interconnect in this environment, we have to be able to fabricate equally dense optoelectronic devices such as LD/LED, PIN/APD arrays, driver arrays and receiver arrays.
- *Propagation delay (8):* The propagation delay of light in a waveguide is unlikely to reduce below the 5.0 ps/mm value currently available. This compares unfavorably with the 3.5–4.0 ps/mm for metal interconnects if suitable fabrication processes are developed to use expanded PTFE type material as an insulator in multi-chip modules and boards. This seems to be a fundamental limitation for waveguide optical interconnects. However, metal interconnects suffer additional delay at each discontinuity, as well as require longer settling time due to switching noise, crosstalk, and reflections. Therefore, propagation delay alone cannot be used to evaluate the performance of an interconnect technology.
- *Conversion delay:* Signals are useful only in their electrical forms. Therefore, electrical-to-optical (E/O) and optical-to-electric (O/E) conversions are necessary for every interconnect, which always involve nonnegligi-

ble conversion delay.

- *Sensitivity to noise during E/O and O/E conversion:* Existing optical interconnect technology has more loss than metal interconnection for such distances, due to the insertion loss of the connector and scattering loss of the surface defects of a waveguide. Therefore, more amplification and a higher sensitivity to both power supply noise and electrical crosstalk are experienced by the receiver. This problem is further aggravated by the high density required by a dense optical interconnect environment. In such an environment, there could be significant electrical interference either through the shared common power supply or through the parasitic inductance and capacitance.
- *Thermal interactions:* Laser characteristics, such as the wavelength, threshold current and differential quantum efficiency, are strongly affected by the operating temperature. Thermal interactions between adjacent lasers in a dense laser array could thus significantly degrade the system performance.

ARCHITECTURE OF OPTICAL INTERCONNECTS

Figure 2 shows the structure of a typical optical interconnect, which consists of a driver array, a laser diode or LED (LD/LED) array, a waveguide or fiber ribbon array, a photodetector array (p-i-n or MSM), and a receiver array. Using optical interconnects for high-bandwidth communication channels between boxes has been demonstrated, for example, in (9). It is conceivable that optical interconnect can also be used within a box (both at the board and backplane levels), shown in Fig. 3.

In this section, possible architectures for dense optical interconnects at the board and backplane levels are investigated. The constraints for designing the architecture of an optical interconnect systems are:

- Compatibility with existing packaging technology,
- Flexibility in fitting into the architectures of digital systems,
- Fault tolerance,
- Easy engineering change and fault diagnostics.

In what follows, first the available interconnect forms will be examined. Possible interconnect architectures at backplane, board, and multi-chip module level will then be investigated. Possible E/O and O/E conversion schemes will also be evaluated.

Interconnect Media

Possible media that can be used for optical interconnects include:

- *Free-space interconnect:* Light travels fastest in free space. In addition, free-space interconnects also offer the highest density and the most sophisticated interconnection patterns. Currently, implementations of free-space optical interconnects range from very low bit rate links such as the majority of TV/VCR/Cable remote controls to low bit rate links such as those IrDA-compliant infrared links between notebooks and printers. For high-speed optical links, unfortunately, bulk optical elements such as lenses, holograms, beam splitters, etc., are usually unavoidable in free-space optical interconnects and thus make the alignment of optical beams very difficult and unstable with respect to environmental disturbances.
- *Optical fiber ribbon:* Optical fiber has the least loss compared with the other two media, and most of the technologies used in fabrication are already mature. Fiber ribbon cable also has the potential of providing reasonable interconnect density with regular interconnection pattern. However, fibers are incompatible with the existing packaging technology at the board or MCM level and they are not suitable for interconnects with very short distance or complicated patterns due to the possibly excessive volume occupied by the fiber cable. A lot of research effort has thus been devoted to develop compatible connector and packaging technologies to interconnect fiber ribbon and optical transceiver array.
- *Planar optic waveguide:* Passive planar optic waveguides are emerging as a viable alternative to optical fiber for very short distance interconnects. It has a higher propagation loss than optical fiber ($0.01{\sim}0.5$ dB/cm as compared to 0.2 dB/km) but uses technologies that are compatible with existing PCB or MCM technology. Therefore, it is more suitable for short-distance dense interconnect applications. However, coupling of light into and from the waveguides is also difficult and careful alignment cannot be avoided.

Backplane Optical Interconnects

The function of a backplane is to provide a logical bus for all of boards connected to it. Free space, fibers, and planar waveguides are all suitable for backplane interconnects. An optical backplane can be achieved through using of star couplers, as shown in Fig. 4. Each board in the architecture occupies one input port and one output port from each of the star couplers so that signals input to any of the input port will be broadcast to all of the output ports. The total number of star couplers required can be greatly reduced by multiplexing several channels into a single waveguide with each channel using a different wavelength.

On the other hand, a topological bus can also be used to interconnect from one board to another, as shown in Fig. 5. The bus is either folded back at the end or two independent buses are used because a unidirectional optical bus structure is usually easier to implement.

Board and Multi-Chip-Module Optical Interconnects

Board-level optical interconnects have to provide interconnects between different SCM's or MCM's while MCM-level optical interconnects have to provide interconnects between unpackaged wire-bonded or solder-ball-bonded flipped chips. At the board level, the E/O and O/E conversion can be performed within an SCM/MCM, or through separate special-purpose E/O and O/E chips. Similarly, the

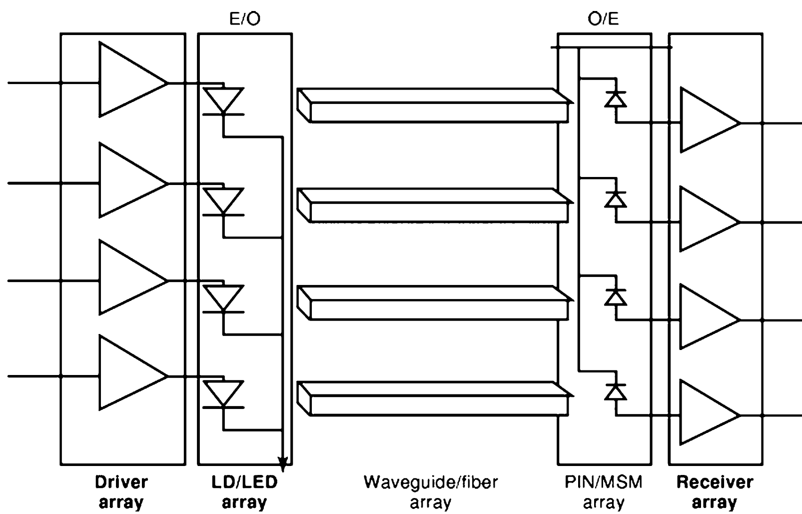

Figure 2. Structure of an optical interconnect system.

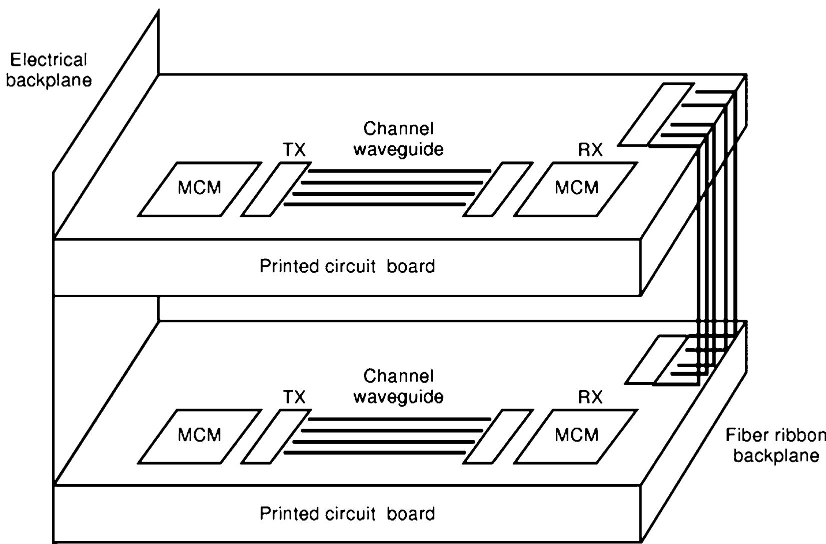

Figure 3. A packaging structure uses optical interconnect at board and backplane levels.

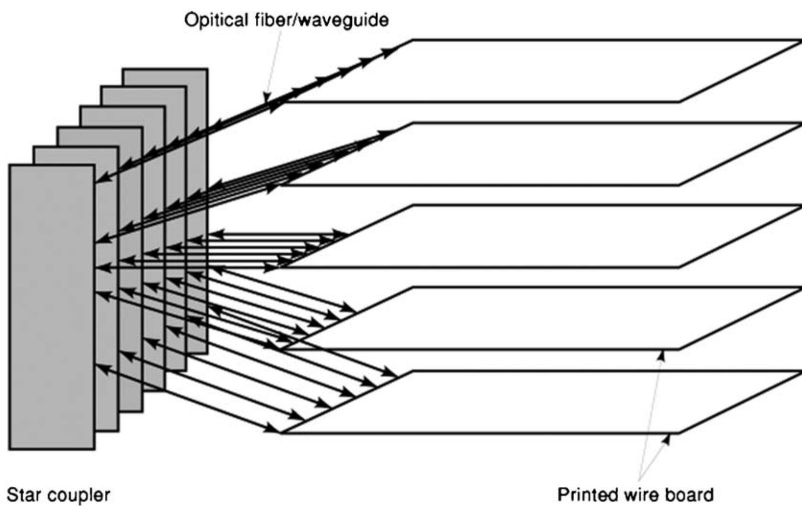

Figure 4. Optical backplane interconnects. Star couplers are used to combine and redistribute the data signals.

E/O and O/E conversion at the MCM level can be performed within the chip where the logical signals are generated or via separate special-purpose E/O and O/E chips on an MCM.

If the E/O and O/E conversion is performed before the package is connected to the next higher level, as shown in Fig. 6, the electrical discontinuity can be minimized but the optical alignment is more difficult. On the other hand, more

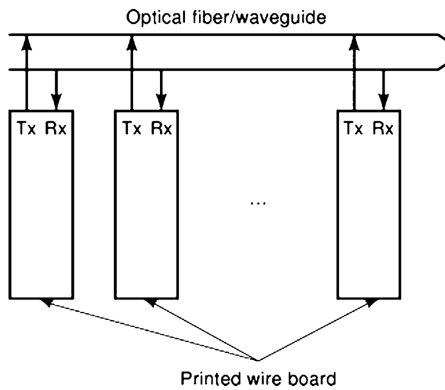

Figure 5. Optical backplane interconnects. A topological bus is used to provide communication path between any two boards connected to the backplane.

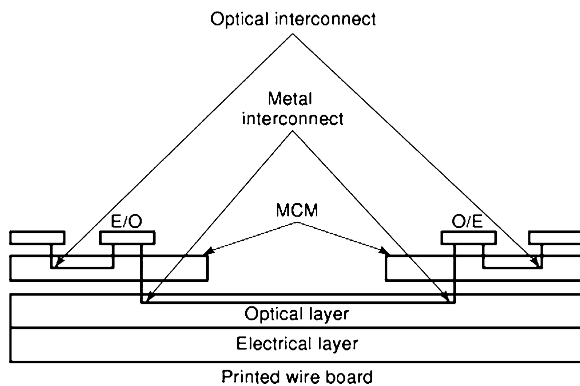

Figure 6. E/O and O/E conversion. Conversion is performed at the same packaging level as the electric signal is generated.

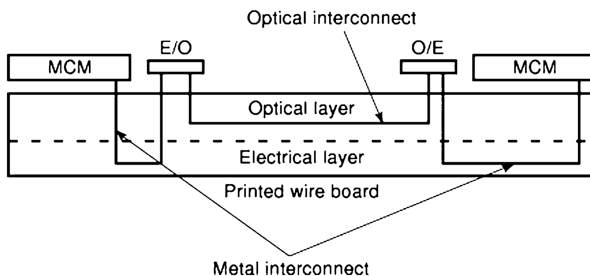

Figure 7. E/O and O/E conversion. Conversion is performed at the next higher packaging level.

electrical discontinuity and thus more signal distortion is introduced if the E/O and O/E conversion is performed after the package is connected to the next level, as shown in Fig. 7. However, this is acceptable for applications that require only moderate data rates.

In both cases, there already exist multiple layers of metal interconnect to provide signal lines as well as power and ground plane. Optical interconnects can be developed on top of these metal interconnect layers to allow optical signals to propagate from one chip/module to another chip/module. In some cases, more than one optical layer may be necessary in order to provide sufficient interconnect density (such as at the MCM level) just similar to its electrical counterpart.

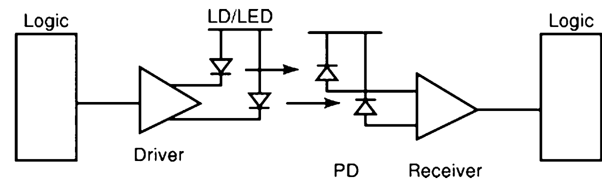

Figure 8. A fully differential optical interconnect architecture.

Fully Differential Optical Interconnect

As discussed earlier, optical interconnects for a digital system has the following potential problems:

- Threshold uncertainty,
- Latency due to serialization/deserialization, encoding/decoding,
- Sensitivity to the switching noise and power supply noise.
- Sensitivity to the thermal interactions.
- Sensitivity to the DC level of the data at the receiver.

A fully differential optical interconnect architecture, as shown in Fig. 8, was proposed in Refs. 10–12 to minimize the detrimental effects arising from these potential problems. In this architecture, complementary optical signals are generated, transmitted, and received along two independent optical channels.

In a fully differential optical interconnect, the threshold voltage at the output of the receiver is always located at differential zero, which is half way between two signal voltage of approximately equal amplitude but opposite sign, assuming two lasers at the differential transmitter have approximately the same average power and the attenuation along the differential path is similar. The threshold voltage is then independent of the actual power output of the lasers and the attenuation of the channel. Since the laser drivers and receivers are both fully balanced, the fluctuation of the current demands from the power supply is minimized, and thus the switching noise is reduced. An offset voltage is incurred when the laser output power or the attenuation is not balanced along the differential path. Furthermore, the differential structure for both drivers and receivers increases the common-mode rejection and thus reduces the sensitivity to the power supply noise. Therefore, differential optical interconnect is very attractive in a dense optical environment.

TECHNOLOGIES FOR OPTICAL INTERCONNECT SYSTEMS

In a dense optical interconnect, electrical interactions among elements in a transmitter or receiver array due to high density requirement might limit the system performance, and will be the subject of this section.

Transmitter Array

A transmitter array usually consists of a driver array and a laser diode array (or LED array). These two components might be monolithically or hybrid integrated on the same

substrate. Among possible interactions in a dense transmitter array are

- Electrical crosstalk between laser diodes due to the sharing of a common substrate,
- Electrical crosstalk due to parasitic capacitance and mutual inductance between adjacent channels,
- Switching noise due the sharing of a common power supply and ground.

These interactions increase with the increase of channel density, modulation speed, and modulation current of the transmitter.

Crosstalk among laser array elements has been a subject of continuous interest. Fabrication and characterization of one-dimensional individually addressable laser or LED array has been reported in Refs. 13–16. During the past 5 years, two-dimensional vertical cavity surface emitting laser (VCSEL) diode arrays or surface emitting LED arrays have already reached commercialization and have emerged as a very promising light source for two-dimensional optical interconnects. The performance of these LED's and laser diodes is reported in Refs 17–21. Modulation speed up to 10 Gb/s has already become commercially available for 10 Gb/s Ethernet and Fiber Channel for both directly modulated laser arrays (22) and laser arrays integrated with external modulators (23).

Most of the laser or LED driver circuits were published in the late 70's and early 80's (24–30). Recent laser driver circuit designs usually include monitoring circuits which calculate the peak and average of the laser output power in order to maintain a constant extinction ratio. Based on these designs, a driver array can be built by replicating the same design N times. Both monolithic integration (31) and hybrid integration (32, 33) of the driver array with the laser array have been exploited. Companies such as Vitesse and Broadcom are offering laser driver up at data rates up to 10 Gb/s. Crosstalk in these works is usually determined through experiments or simulations, but a systematic study of the crosstalk due to switching noise is yet to be addressed. However, this issue is important for choosing suitable driver architecture to minimize overall interference.

Waveguide

It has been shown that passive waveguides based on silicon nitride (35) and polymers (36, 37) are attractive for very short distance interconnections, such as those between chips on a multi-chip module or on a printed-circuit-board, or as backplane interconnections. Although suffering more loss than fiber, waveguides have the potential of providing much closer spacing and planar crossover geometries and can integrate modulators, optical amplifiers and receivers on the same substrate as well (38). The density of a waveguide array is limited mainly by the coupling-induced crosstalk between adjacent waveguides. In order to achieve the maximum density allowed by the required bit-error-rate (BER), it is necessary to determine the power coupling among waveguides in an array structure and thus the incurred system penalty.

Receiver Array

Both hybrid integration and monolithic integration technology can be used to package a photodetector array with an amplifier array. Hybrid integration allows separate optimization of the processing technology for the photodetectors and the amplifiers. This technology usually gives better device performance, at the expense of greater adjacent channel crosstalk and signal distortion introduced by the bonding wires. The photodetector array in a hybrid receiver array usually has a p-i-n structure and is made of Si, GaAs, or InGaAs/InP, depending on the wavelength of the light signals (32, 39). The amplifier array is made of Si bipolar (40) or GaAs MESFET. In a monolithic integration environment, both the photodetector array and the receiver array are integrated on the same semiconductor substrate. A planar process for the photodetectors is usually preferable for easier monolithic integration with other electronic circuits. A p-i-n structure usually has a vertical structure, which requires growing of a thick epitaxial layer in order to accommodate the intrinsic region of the p-i-n structure. The thickness of the intrinsic region is at least 2 μm in GaAs and 10 μm in silicon for efficient absorption of the light signals at $\lambda = 0.8$ μm. This process is usually incompatible with the processing steps used for electronic circuits that usually only require a thin epitaxial layer (< 2 μm). Metal-semiconductor-metal (MSM) with its planar structure has thus far emerged as the most popular structure for the photodetector array (41–56).

There have been a number of receiver array designs using either hybrid integration (32,40,57) or monolithic integration (33,58–62) technology. Up to 32 and 12 channels/chip have been achieved thus far with monolithic (63) and hybrid technology (64), respectively.

Electrical crosstalk between photodetectors in a p-i-n array has been previously examined in Refs. 39 and 65. It was concluded in Ref. 65 that the common substrate of a p-i-n array introduces negligible DC crosstalk. A majority of the crosstalk came from the parasitic coupling between the bonding wires connecting between photodetectors and receivers.

RESEARCH ACTIVITIES ON OPTICAL INTERCONNECTS

Using optics for interconnections between VLSI systems was first suggested in Refs. 66 and 67. Early systems are mostly based on free-space interconnects with the use of holographic optical elements (HOE) and spatial light modulators (SLM) to establish interconnect patterns. More recent systems have begun to use both optical fibers and planar waveguides (68). An interprocessor optical link has been demonstrated between processor blocks in the Thinking Machines CM-2 at 400 Mbps (69). The feasibility of board-level optical interconnect using polymer (70, 71) and silica (72) have also been demonstrated. Both of these prototypes can demonstrate a bit rate higher than 300 Mbps. A DC-coupled, fully differential optical interconnect system was proposed, analyzed and simulated in Refs. 10 and 11 for connections within high-speed digital systems, specifically for board and backplane level interconnections. A chip set consisting of a 2.5 Gb/s bipolar differential laser driver,

an 800 Mb/s GaAs MSM-preamp array, an 800-Mb/s GaAs MSM-preamp-postamp array, and a GaAs MSM-preamp array with each preamp having a different bandwidth varying from 300 Mbps to 2 Gbps has been designed, fabricated, and tested to serve as a vehicle for verifying the concept (12).

Between 1991 and 1994, ARPA has supported the Optoelectronic Technology Consortium (OETC), which consists of Martin Marietta, AT&T, Honeywell, and IBM, to develop a 32-channel bus with a data transfer rate of 500 Mb/s (63). This system uses 850 nm VSCEL for the transmitter array, MSM for the photodetector array, and GaAs E/D MESFET for the AC-coupled receiver and driver array. The data is Manchester coded due to the ac-coupled design at the receiver. A 32×1 multi-mode fiber ribbon cable with 62.5 μm core diameter, 125 μm cladding diameter, and 140 μm pitch is used between the transmitter and the receiver. A NIST-funded consortium consisting of IBM, 3M, and Lexmark, has developed a 20-channel parallel optical interconnect, with a total throughput of 1 GB/s.

An ARPA-sponsored collaborative effort by the parallel optical link organization (POLO), consisting of HP, AMP, DU Pont, SDL, and the University of Southern California, operated between 1994 and 1997. The objective of this program is to provide a 10–20 Gb/s parallel channel with a manufacturing cost of $10/channel (73). A 10-channel DC-coupled parallel optical link, with each link operating at more than 500 MHz, using VCSEL polymer waveguide, and bipolar receiver has already been demonstrated (74).

Another ARPA-sponsored consortium consisting of GE, Honeywell, and Allied Signals is the POINT project (75). This project focuses on the batch processing and passive alignment between optical waveguides and traceivers for board and backplane applications.

In Europe, the European Strategic Programme for Information Technology (ESPRIT) has sponsored the OLIVES (Optical Interconnections for VLSI and Electronic Systems) between 1989 and 1992, and HOLICS (Hierarchical Optical Interconnects for Computer Systems) immediately after OLIVES. In HOLICS, 4-channel edge emitting laser diode arrays, InGaAs/InP p-i-n photodetector array, and 8-channel and 12-channel 1 Gb/s receiver arrays have been developed to be used in conjunction with 250 μm-pitch fiber ribbon cable (64).

SUMMARY

In this chapter, we have surveyed a number of promising optical interconnect architectures and technologies that can significantly improve the performance of high-throughput multimedia systems. These optical interconnect systems have already reached very high density, high data rate low insertion loss, as well as low optical and electrical crosstalk. Some of these technologies have already reached commercial maturity (such as fiber ribbon cable, MACII connectors, monolithically integrated photodetector and receiver array). Other technologies, such as packaging and transmitter array, still have a long way to go before they can reach the same level of reliability.

BIBLIOGRAPHY

1. H. S. Stone J. Cocke Computer architecture in the 1990s, *IEEE Comput. Mag.*, **24** (9): 30–38, 1991.
2. R. R. Tummala *et al.* Thin film packaging, in *Microelectronics Packaging Handbook*, R. R. Tummala and E. J. Rymaszewski (eds.), New York: Van Nostrand-Reinhold, 1989, pp. 673–723.
3. G. A. Katopis Delta-I noise specification for a high-performance computing machine, Proc. IEEE, 73: 1405–1415, 1985.
4. R. L. Veghte C. A. Balanis Dispersion of transient signals in microstrip transmission lines, *IEEE Trans. Microw. Theory Tech.*, **34**: 1427–1436, 1986.
5. R. E. Epworth The phenomenon of modal noise in analog and digital optical fiber systems, Proc. 4th Eur. Conf. Optical Commun., Genova, 1978, pp. 492–501.
6. R. W. Huegli C. M. Olsen Elimination of Modal Noise Bit-Error-Rate Floors by Strong Stimulation of the On-Level Relaxation Oscillations, IBM Res. Rep. RC 16675, Yorktown Heights, IBM T. J. Watson Research Center, 1991.
7. J. Gimlett N. K. Cheung Effects of phase-to-intensity noise conversion by multiple reflections on gigabit-per-second DFB laser transmission systems, *IEEE J. Lightw. Technol.*, **7**: 888–895, 1989.
8. G. Arjavalingam B. Rubin Electrical considerations for interconnections inside a computer, Proc. SPIE, 991: 12–21, 1988.
9. J. W. Lockwood *et al.* Scalable optoelectronic ATM networks: The iPOINT fully functional testbed, *IEEE J. Lightw. Technol.*, **13**: 1093–1103, 1995.
10. C.-S. Li H. S. Stone Differential board/backplane optical interconnects for high-speed digital systems. Part I: Theory, *IEEE J. Lightw. Technol.*, **11**: 1234–1249, 1993.
11. C. M. Olsen C.-S. Li Differential board/backplane optical interconnects for high-speed digital systems. Part II: Simulation results, *IEEE J. Lightw. Technol.*, **11**: 1250–1262, 1993.
12. C.-S. Li *et al.* Fully differential optical interconnections for high-speed digital systems, *IEEE Trans. VLSI Syst.*, **1**: 151–163, 1993.
13. J. P. Van Der Ziel R. A. Logan R. M. Mikulyak A closely spaced (50 μm) array of 16 individually addressable buried heterostructure GaAs lasers, *Appl. Phys. Lett.*, **41**: 9, 1982.
14. P. P. Deimel *et al.* Electrical and optical integration of a monolithic 1×12 array of InGaAsP/InP ($\lambda = 1.3$ μm) light emitting diodes, *IEEE J. Lightw. Technol.*, **3**: 1985.
15. D. Botez *et al.* High-power individually addressable monolithic array of constricted double heterojunction large-optical-cavity lasers, *Appl. Phys. Lett.*, **41**: 1040–1042, 1982.
16. L. A. Koszi *et al.* 1.5 μm InP/GaInAsP linear laser array with twelve individually addressable elements, *Electron. Lett.*, **4**: 217–219, 1988.
17. L. A. Koszi *et al.* Fabrication and performance of an InP/InGaAsP monolithic 12×12 element matrixed LED array, *Electron. Lett.*, **23**: 284–286, 1987.
18. D. L. McDaniel, Jr. *et al.* Vertical cavity surface-emitting semiconductor laser with CW injection laser pumping, *IEEE Photon. Technol. Lett.*, **2**: 156–158, 1990.
19. R. S. Geels *et al.* Low threshold planarized vertical-cavity surface emitting lasers, *IEEE Photon. Technol. Lett.*, **2**: 234–236, 1990.
20. Y. H. Lee *et al.* Characteristics of top-surface-emitting GaAs quantum-well lasers, *IEEE Photon. Technol. Lett.*, **2**: 686–688, 1990.

21. J. L. Jewell *et al.* Surface-emitting microlasers for photonic switching and interchip connections, *Opt. Eng.*, **29**: 210–214, 1990.

22. C.-E. Zah *et al.* Multiwavelength DFB laser arrays with integrated combiner and optical amplifier for WDM optical networks, *IEEE J. Sel. Top. Quantum Electron.*, **3**: 584–597, 1997.

23. G. Raybon *et al.* Reconfigurable optoelectronic wavelength translation based on an integrated electroabsorption modulated laser array, *IEEE Photon. Technol. Lett.*, **10**: 215–217, 1998.

24. M. Uhle The influence of source impedance on the electrooptical switching behavior of LED's, *IEEE Trans. Electron Devices*, **23**: 438–441, 1976.

25. R. Olshansky D. Fye Reduction of dynamic linewidth in single-frequency semiconductor lasers, *Electron. Lett.*, **20**: 928–929, 1984.

26. L. Bickers L. D. Westbrook Reduction of laser chirp in 1.5 μm DFB lasers by modulation pulse shaping, *Electron. Lett.*, **21**: 103–104, 1985.

27. M. A. Karr F. S. Chen P. W. Shumate Output power stability of GaAlAs laser transmitter using an optical tap for feedback control, *Appl. Opt.*, **18**: 1262–1265, 1979.

28. P. W. Shumate F. S. Chen P. W. Dorman GaAlAs laser transmitter for lightwave transmission systems, *Bell Syst. Tech. J.*, **57**: 1823–1836, 1978.

29. D. W. Smith M. R. Matthews Laser transmitter design for optical fiber systems, *IEEE J. Sel. Area Commun.*, pp. 515–523, 1983.

30. R. G. Schwatz B. A. Wooley Stabilized biasing of semiconductor lasers, *Bell Syst. Tech. J.*, **62**: 1923–1936, 1983.

31. O. Wada *et al.* Optoelectronic integrated four channel transmitter array incorporating AlGaAs/GaAs quantum well lasers, *IEEE J. Lightw. Technol.*, **7** (1): 186–197, 1989.

32. K. Kaede *et al.* 12-channel parallel optical-fiber transmission using a low-drive current 1.3-μm LED array and a p-i-n PD array, *IEEE J. Lightw. Technol.*, **8**: 883–888, June, 1990.

33. H. Yamanaka *et al.* A gigabit-rate five-highway GaAs OE-LSI chipset for high-speed optical interconnections between modules or VLSI's, *IEEE J. Sel. Areas Commun.*, **9**: 689–697, 1991.

34. R. Yu *et al.* HBT devices and circuits for signal and data processing, *Solid-State Electron.*, **41** (10): 1419–1431, 1997.

35. W. Stutius W. Streifer Silicon nitride films on silicon for optical waveguides, *Appl. Opt.*, **6**: 3218–3222, 1977.

36. B. L. Booth Low loss channel waveguides in polymers, *IEEE J. Lightw. Technol.*, **7**: 1445–1453, 1989.

37. J. M. Trewhella Polymeric optical waveguides, Proc. SPIE, 1177: 379–386, 1989.

38. M. Dagenais *et al.* Application and challenges of OEIC technology: A report on the 1989 Hilton Head workshop, *IEEE J. Lightw. Technol.*, **8**: 846–862, 1990.

39. M. G. Brown *et al.* Monolithically integrated 1×12 array of planar InGaAs/InP photodiodes, *IEEE J. Lightw. Technol.*, **4**: 283–287, 1986.

40. J. Wieland H. Melchior Optical receivers in ECL for 1 GHz parallel links, SPIE Proc. 1389: 659–664, 1990.

41. L. Figueroa C. W. Slayman A novel heterostructure interdigital photodetector (HIP) with picosecond optical response, *IEEE Electron Device Lett.*, **2**: 208–210, 1981.

42. W. Roth *et al.* The DSI diode-A fast large-area optoelectronic detector, *IEEE Trans. Electron Devices*, **6**: 1034–1036, 1985.

43. M. Ito O. Wada Low dark current GaAs metal-semiconductor-metal (MSM) photodiodes using WSIx contacts, *IEEE J. Quantum Electron.*, **22**: 1073–1077, 1986.

44. W. C. Koscielniak R. M. Kolbas M. A. Littlejohn Peformance of a near-infrared GaAs metal-semiconductor-metal (MSM) photodetector with islands, *IEEE Electron Device Lett.*, **9**: 485–487, 1988.

45. B. J. Van Zeghbroeck *et al.* 105 GHz bandwidth metal-semiconductor-metal photodiodes, *IEEE Electron Device Lett.*, **9**: 527–529, 1988.

46. H. Schumacher *et al.* An investigation of the optoelectronic response of GaAs/InGaAs MSM photodetectors, *IEEE Electron Device Lett.*, **9**: 607–609, 1988.

47. O. Wada *et al.* Noise characteristics of GaAs metal-semiconductor-metal photodiodes, *Electron. Lett.*, **24**: 1574–1575, 1988.

48. T. Kukuchi H. Ohno H. Hasegawa $Ga_{0.47}In_{0.53}As$ metal-semiconductor-metal photodiodes using a lattice mismatched $Al_{0.4}Ga_{0.6}As$ Schottky assist layer, *Electron. Lett.*, **24**: 1208–1210, 1988.

49. D. L. Rogers *et al.* High-speed 1.3 μm GaInAs detectors fabricated on GaAs, *IEEE Electron. Device Lett.*, **9**: 515–517, 1988.

50. W.-P. Hong G.-K. Chang R. Bhat High performance $Al_{0.15}Ga_{0.85}AsIn_{0.53}Ga_{0.47}As$ MSM photodetectors grown by OMCVD, *IEEE Trans. Electron Devices*, **36** (4): 659–662, 1989.

51. L. Yang *et al.* High performance of Fe:InP/InGaAs metal/semiconductor/metal photodetectors grown by metalorganic vapor phase epitaxy, *IEEE Photon. Technol. Lett.*, **2**: 56–58, 1990.

52. L. Yang *et al.* Monolithically integrated InGaAs/InP MSM-FET photoreceiver prepared by chemical beam epitaxy, *IEEE Photon. Technol. Lett.*, **2**: 59–62, 1990.

53. K. Nakajima *et al.* Properties and design theory of ultrafast GaAs metal-semiconductor-metal photodetector with symmetrical Schottky contacts, *IEEE Trans. Electron Devices*, **37**: 31–35, 1990.

54. J. B. D. Soole H. Schumacher Transit-time limited frequency response of InGaAs MSM photodetectors, *IEEE Trans. Electron Devices*, **37**: 2285–2291, 1990.

55. D. Kuhl *et al.* High-speed metal-semiconductor-metal photodetectors on InP/Fe, *IEEE Photon. Technol. Lett.*, **2**: 574–576, 1990.

56. E. Sano A device model for metal-semiconductor-metal photodetectors and its applications to optoelectronic integrated circuit simulation, *IEEE Trans. Electron Devices*, **37**: 1964–1968, 1990.

57. Y. Ota *et al.* Twelve-channel individually addressable INGaAs/InP p-i-n photodiode and InGaAsP/InP LED arrays in a compact package, *IEEE J. Lightw. Technol.*, **5**: 1118–1122, 1987.

58. W.-P. Hong *et al.* High-functionality waveguide/MSM/HEMT integrated receiver prepared by one-step OMCVD grown on patterned InP substrates, Proc. Optical Fiber Commun., 1991, p. 5.

59. J. D. Crow Optical interconnects for high-performance data processing systems, Proc. IOOC'89, 1989.

60. J. D. Crow *et al.* A GaAs MESFET IC for optical multiprocessor network, *IEEE Trans. Electron Devices*, **36**: 263, 1989.

61. N. Yamanaka T. Takada A 1.5 Gbit/s GaAs four-channel selector LSI with monolithically integrated newly structured GaAs

ohmic contact MSM photodetector and laser driver, *IEEE Photon. Technol. Lett.*, **1**: 310–312, 1989.

62. M. Makiuchi *et al.* A monolithic four-channel photoreceiver integrated on GaAs substrate using metal-semiconductor-metal photodiodes and FET's, *IEEE Electron. Device Lett.*, **6**: 634–635, 1985.
63. Y.-M. Wong *et al.* Technology development of a high-density 32-channel 16 Gb/s optical data link for optical interconnection applications for the Optoelectronic Technology Consortium (OETC), *IEEE J. Lightw. Technol.*, **13**: 995–1016, 1995.
64. H. Karstensen *et al.* Parallel optical interconnection for uncoded data transmission with 1 Gb/s-per-channel capacity, high dynamic range, and low power consumption, *IEEE J. Lightw. Technol.*, **13**: 1017–1030, 1995.
65. D. R. Kaplan S. R. Forrest Electrical crosstalk in p-i-n arrays Part I: Theory, *IEEE J. Lightw. Technol.*, **4**: 1460–1469, 1986.
66. J. W. Goodman *et al.* Optical interconnections for VLSI systems, Proc. IEEE, 72: 850–866, 1984.
67. R. K. Kostuk J. W. Goodman L. Hesselink Optical imaging applied to microelectronic chip-to-chip interconnections, *Appl. Opt.*, **24**: 2851–2858, 1985.
68. L. D. Hutcheson P. Haugen A. Husain Optical interconnects replace hardwire, *IEEE Spectrum*, **24** (3): 30–35, 1987.
69. B. O. Kahle E. C. Parish Optical interconnects for interprocessor communications in the connection machine, Proc. ICCD, 1989, pp. 58–61.
70. D. H. Hartman G. R. Lalk T. C. Banwell Board level high speed photonic interconnections: Recent technology developments, Proc. SPIE, 994: 57–64, 1988.
71. D. H. Hartman *et al.* Radiant cured polymer optical waveguides on printed circuit boards for photonic interconnection use, *Appl. Opt.*, **28**: 40–47, 1989.
72. Y. Yamada *et al.* Optical interconnections using silica-based waveguide on Si substrate, Proc. SPIE, 991: 4–11, 1988.
73. W. S. Ishak *et al.* Optical interconnects: The POLO approach, SPIE Proc., 2400: 214–221, 1995.
74. K. H. Hahn *et al.* Gigabyte per second optical interconnect modules for data communications, Hewlett Packard J., 48: 53–61, 1997.
75. Y. S. Liu *et al.* High-density optical interconnect using polymer waveguides interfaced to a VCSEL array in molded plastic packaging, Proc. SPIE, 3288: 60–72, 1998.

Reading List

W. R. Blood MECL System Design Handbook, Motorola, Phoenix, 1988.

E. E. Davidson G. A. Katopis Package electrical design, in *Microelectronics Packaging Handbook*, New York: Van Nostrand-Reinhold, 1989, Chap. 3, pp. 111–165.

S. L. Diamond Micro standards: IEEE Standard P1394, *IEEE Micro Mag.*, **15** (2): 81–83, 1994.

M. P. Farmwald D. Mooring A fast path to one memory, *IEEE Spectrum*, **29** (10): 50–51, 1992.

S. Gjessing *et al.* A RAM link for high speed, *IEEE Spectrum*, **29** (10): 52–53, 1992.

K. Hwang F. A. Briggs Multiprocessor architecture and programming, in *Computer Architecture and Parallel Processing*, New York: McGraw-Hill, 1984, Chap. 7.

IEEE P1594 Working Group, *Scalable Coherent Interface: Logical, Physical, and Cache Coherence Specifications*, New York: IEEE Standard Department, 1992.

K. P. Jackson *et al.* High-speed characterization of a monolithically integrated GaAs-AlGaAs quantum-well laser-detector, *IEEE Photon. Technol. Lett.*, **2**: 832–834, 1990.

R. L. Khalil L. R. McAdams J. W. Goodman Optical clock distribution for high speed computers, Proc. SPIE, 991: 32–41, 1988.

W. C. Koscielniak J.-L. Pelouard M. A. Littlejohn Intrinsic and extrinsic response of GaAs metal-semiconductor-metal photodetectors, *IEEE Photon. Technol. Lett.*, **2**: 125–127, 1990.

A. W. Lohmann Optical bus network, *Optik*, **74**: 30–35, 1986.

M. W. Sachs A. Varma Fiber channel and related standards, *IEEE Commun. Mag.*, **34**: 40–51, 1996.

J. Shamir Three-dimensional optical interconnection gate array, *Appl. Opt.*, **26**: 3455–3457, 1987.

H. S. Stone Pipeline design techniques, in *High-Performance Computer Architecture*, New York: Addison-Wesley, 1987, Chap. 3.

R. Walker *et al.* A 2-chip 1.5 Gb/s bus-oriented serial link interface, Proc. Int. Symp. Solid-State Circuit, 1992, pp. 226–227.

R. C. Walker *et al.* A 1.5 Gb/s link interface chipset for computer data transmission, *IEEE J. Sel. Areas Commun.*, **9**: 698–710, 1991.

CHUNG-SHENG LI
IBM T. J. Watson Research Center, Yorktown Heights, NY