A System Design Perspective on Optical Interconnection Technology

Ronald A. Nordin
Anthony F. J. Levi

Demands for increased interconnection density and higher bandwidth, coupled with constraints on the cost of advanced wide-bandwidth telecommunication switching and high-throughput computer architectures, are exhausting the capabilities of conventional electrical interconnection. Advances in integrated circuit technologies and enhanced digital services have created a "bottleneck" at the circuit-board-to-circuit-board level of the interconnection hierarchy. To alleviate this problem, AT&T is developing new techniques in parallel optical interconnection. Inserting parallel optical interconnection technology into advanced electronic processing systems not only meets projected performance requirements, but also potentially offers them at a competitive cost. This paper describes how future high-performance parallel optical interconnect technology might be packaged and used.

Introduction

Technological advances in integrated circuits (ICs) and demand for enhanced digital telecommunication services and advanced computer services are creating a need for increased interconnection density and higher bandwidths, while maintaining stringent reliability specifications. Enhanced digital telecommunications services, which include Broadband Integrated Services Digital Network (BISDN), high-definition television (HDTV), video conferencing, cellular services, and microcell interconnection, have spurred research and design in graphics, real-time image processing, pattern recognition, modeling, speech processing, and parallel processing. (See Panel 1 for definitions of abbreviations, acronyms, and terms.)

Increased interconnection density and bandwidth at all levels of the interconnection hierarchy are shaping system design advances in traditional electronic packaging and optical techniques.[1-4] End-to-end optical data link (ODL) products (e.g., ODL® 40 and ODL 50) have been used successfully in high-performance computer and telecommunication switching systems, such as the AT&T Network Systems 5ESS® switch. To meet future system needs, AT&T is developing new techniques in parallel optical interconnection.

Parallel optical interconnection will most likely be achieved using a generic component designed to communicate over distances between 1 meter (m) and 10 kilometers (km). These devices are known as one-dimensional ODLs (1D-ODL).

Figure 1a shows the conventional electronic packaging technology for large digital switching and computer systems. The length of the corresponding interconnection hierarchy is roughly bounded by its packaging technology, as shown in Figure 1b. Electrical interconnections are efficient and support the density and bandwidth needed to interconnect lengths of less than 1m. Discrete ODLs are efficient for interconnection lengths of more than 25m, but they cannot support the density or cost constraints imposed by high-performance system design in the 1- to 100m range. In this range, we expect the 1D-ODL to be a competitive solution.

To fulfill the needs of many small-volume users, which in total create a large demand for 1D-ODLs, the industry needs a standard, such as the IEEE 1596-1992 Scalable Coherent Interface (SCI), that will not need to be customized for different applications. Although it supplies computer-bus-like services, it uses a collection of fast point-to-point links (of size 16N + 2), instead of a

Panel 1. Abbreviations, Acronyms, and Terms

1D-ODL — one-dimensional optical data link
AlGaAs — aluminum gallium arsenide
Au — gold
BER — bit error rate
BH-GRINSCH — buried-heterostructure graded-index separate-confinement heterostructure
BISDN — broadband Integrated Services Digital Network
CW — continuous wave
dBm — decibels referenced to a milliwatt
DFB — distributed feedback
ECL — emitted coupler logic
EEL — edge-emitting laser
EMI — electromagnetic interference
FP — Fabry Perot
GaAs — gallium arsenide
GaInAsP — gallium indium arsenide phosphide
Gbit/s — gigabit(s) per second
IC — integrated circuit
InP — indium phosphide
km — kilometer
LED — light-emitting diode
LGA — laser gate array
$LiNbO_3$ — lithium niobate
m — meter
MAC — multi-fiber array connector
Mo — molybdenum
NRZ — nonreturn to zero
NPN — transmitter
OC — optical carrier
ODL — optical data link
OEIC — opto-electronic integrated circuit
PiN — diode
PN — diode
POLYFIC — thin-film polymer film integrated circuit
POLYHIC — thin-film polymer hybrid integrated circuit
PQFP — plastic-quad flat pack
ps — picosecond
RTV — room-temperature vulcanization
SCI — scalable coherent interface
SEED — self-electro-optic effect device
S-SEED — symmetric SEED
SEL — surface-emitting laser
Sn — tin
W — watt

physical bus, to reach data rates as high as 1 gigabit per second (Gbit/s) per link. This standard provides a convenient framework for developing the 1D-ODL.

The key technological developments needed by the 1D-ODL are laser diode arrays, detector diode arrays, solder bump bonding, electronic laser driver arrays, electronic receiver arrays, optical submount technologies (e.g., micro-machined silicon with waveguide capability), optical fiber ribbon cables, and optical fiber connectorization. Figure 2a shows a schematic diagram of AT&T's 1D-ODL technology. System design advantages of this technology include low-cost optical data links (package and manufacturing cost are shared over many channels); high bandwidth; high reliability; piece parts rugged enough for manufacturing environments; high optical interconnection density; laser-to-laser optical delay differences, or *skew*, low enough to be used in synchronous designs; a small module footprint (approximately 6.5 cm^2); and low electromagnetic interference (EMI). This represents a significant cost/performance advantage over a serial ODL scheme, which uses discrete devices, or an all-electronic interconnection scheme.

Figure 2b shows how 1D-ODLs may be incorporated in a system design. They can be mounted to the surface of a board, and then, at the appropriate point in the manufacturing process, the multi-fiber array connectors (MAC)-IIs can be mounted and connected to the 1D-ODL. Figures 2c and d shows 18 1D-ODL modules (18 wide) mounted on a circuit board. Each module has 18 fibers, and each fiber transmits at 1 Gbit/s, collectively yielding a bandwidth of 324 Gbits/s. For high-bandwidth connections, this represents a significant improvement over "traditional" connections between a circuit board and backplane. The slim profile and low power dissipation of 1D-ODLs allow them to support more modules per circuit board and more circuit boards per equipment shelf.

Critical to the development of the 1D-ODL program is the adoption of an appropriate packaging

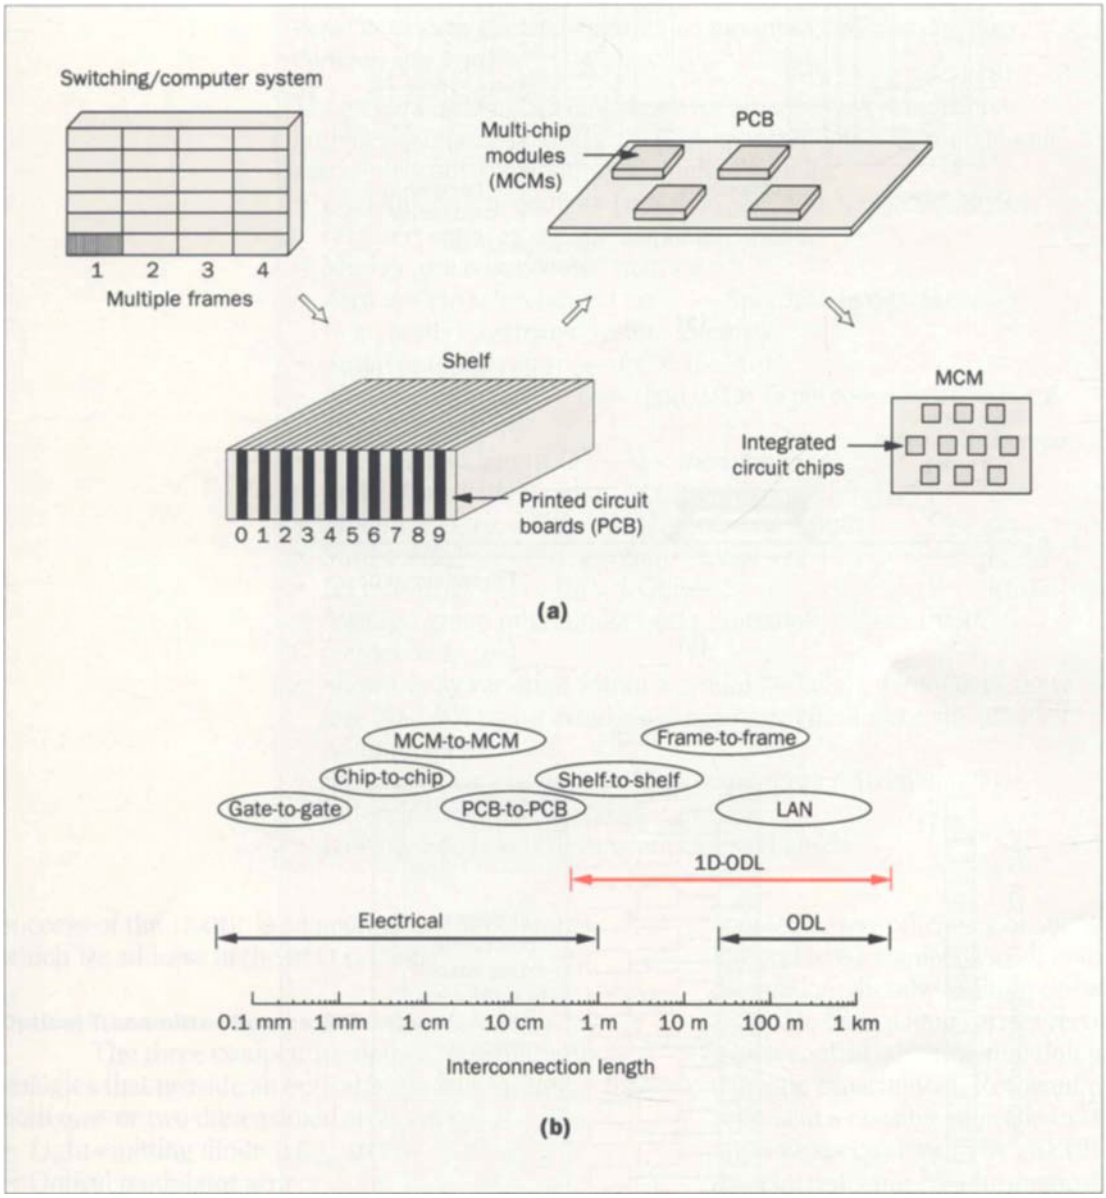

Figure 1. Conventional large system (a) packaging and (b) interconnection hierarchies and their typical lengths.

scheme, in which the transmitter and receiver modules are compatible with an existing fiber ribbon connector. One such fiber ribbon connector is AT&T's MAC-II, which consists of 12 or 18 multimode (or single-mode) glass fibers placed in accurately machined silicon v-grooves on 250-micrometer (μm) centers. Guide pins ensure that the polished ends of the two connector halves can be accurately attached with less than a 0.5-dB coupling loss. For the MAC-II, the standard deviation for fiber misalignment is less than 2 μm. Bonding the laser array (or detector array) on a silicon submount that incorporates v-grooves and guide pins makes it compatible with the MAC-II connector, and usable in a manufacturable, self-aligned packaging scheme.

Backplane Interconnection Requirements

Panel 2 lists general system specifications for an advanced synchronous high-performance system based

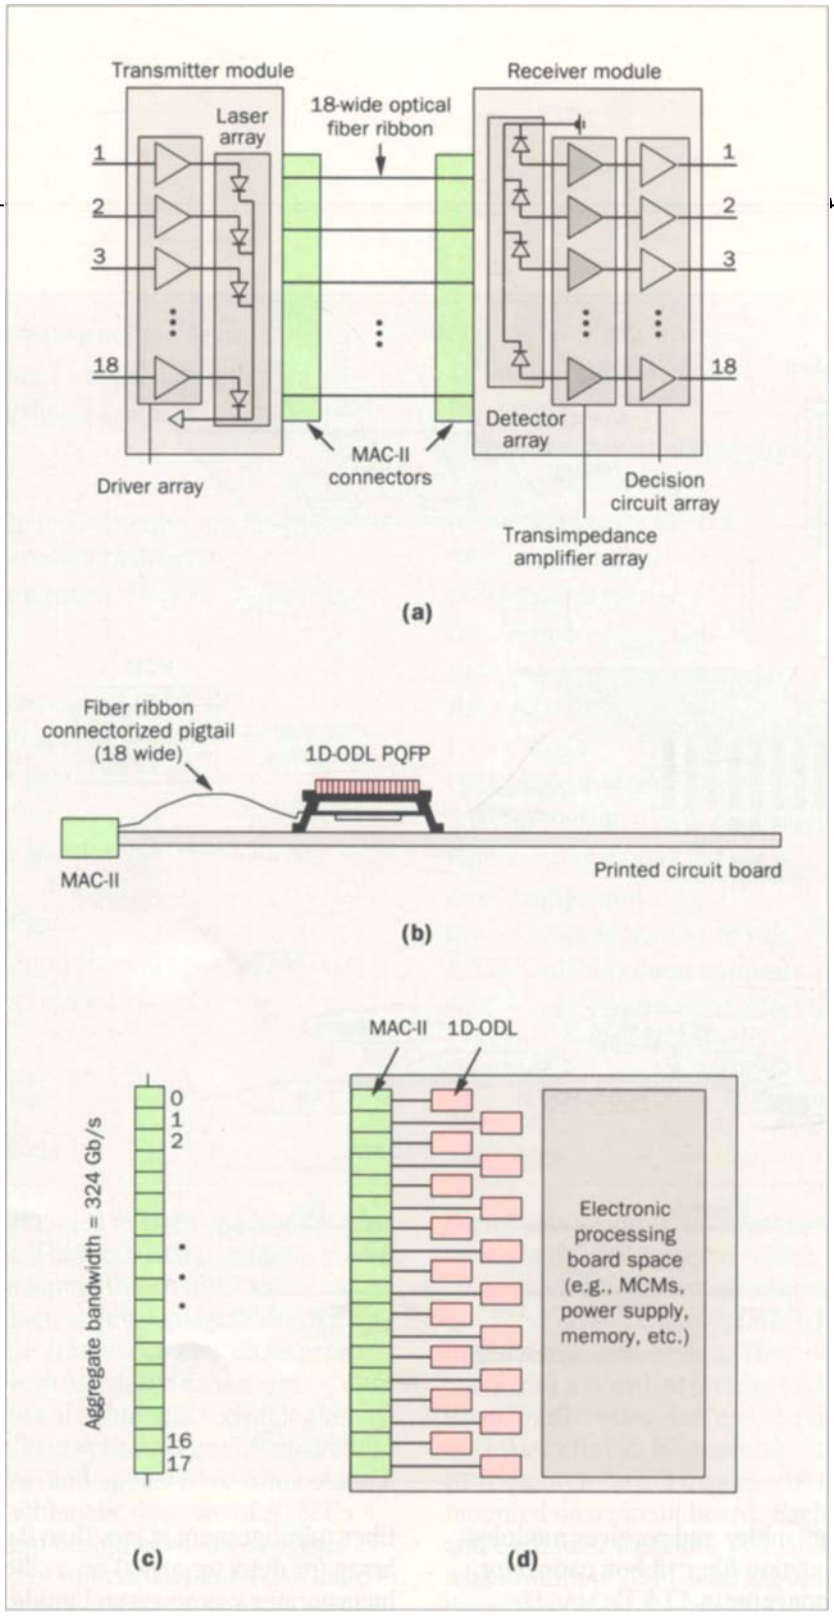

Figure 2. 1D-ODL technology. (a) Schematic diagram of a transmitter connected to a receiver by 18-wide optical fiber ribbon. The physical design of the system is shown in (a) a side view, (b) an end view, and (c) and (d) a top view.

on environments that contain generic central office switching equipment. To attain the required performance within acceptable constraints, the electronic driver and receiver must be integrated into an array form. Although integrating very high gain/bandwidth amplifiers is difficult, because of crosstalk and power/ground "noise," integrating low-gain transimpedance amplifiers (e.g., 300Ω) is much more feasible. The even distribution of power dissipation between the transmitter and receiver is also desirable, from a system perspective. Key to the

Panel 2. System Specifications for an Advanced Synchronous High-Performance System

The general system specifications for an advanced synchronous high-performance system based on environments that contain generic central office switching equipment include:

- Cost/interconnection — Less than $50, which includes driver, receiver, connectors, and backplane media
- Module reliability — 10^5 hours < σ
- Zero system interconnect error — Appropriate digital coding (e.g., parity) to ensure system integrity
- Ambient temperature — 0°C < T < 70°C
- Absolute humidity — Less than 0.026 (5 percent < relative humidity < 95 percent)
- Interconnect length (L) — 1 < L < 100m
- Compatible with standard ECL input/output levels
- Module electrical differential inputs and outputs
- Power supplies — +5, ground, –2, and –4.8
- Bit rate (BR) — 0 < BR < 1 Gbit/s
- Average group interconnect delay variation — Less than 200 picoseconds (ps)
- Skew (delay variation within a group) — Total interconnect skew less than 400 ps (to avoid clock recovery circuitry or any jitter accumulation)
- Small package size — 1" × 1" (e.g., integrated 18 channels)
- External to module dc coupling
- Low module power dissipation of less than 2W.

success of the 1D-ODL is an appropriate light source, which we address in the next section.

Optical Transmitter Source Selection

The three competitive optical transmitter technologies that provide an optical backplane strategy[5] for both one- or two-dimensional arrays are:

- Light-emitting diode (LED) arrays
- Optical modulator arrays
- Laser diode arrays.

The system performance and cost of each is discussed in the sections that follow.

LED Arrays. LED arrays[6] have been regarded as attractive, economical optical transmitters because they are inexpensive to manufacture; highly reliable; easily extendible to large, two-dimensional (i.e., surface-emitting) arrays; and more stable in temperature than other optical methods. Unfortunately, LEDs have a few disadvantages. Their low external power efficiency (a coupled power efficiency of about 0.2 percent, with an integral lens) manifests itself in unwanted chip power dissipation, thereby limiting the width of the LED array. LEDs also have a long carrier recombination lifetime, which controls their modulation bandwidth (neglecting parasitic capacitance). Resonant cavity ("super") LEDs represent a possible upgrade in LED performance. They improve spectral width by an order of magnitude, thereby reducing the chromatic dispersion in fiber. Quantum wells also improve LED bandwidth. Despite all this, LEDs have a fundamentally longer carrier lifetime than injection lasers, and, more significantly, cannot supply the light power necessary for the 1D-ODL.

Optical Modulator Arrays. Another of the optical transmitter technologies is the class of optical devices called modulators. These devices have an input light power source directed to each modulator and an output connection that directs light to the receiver. The light is modulated either with an electrical signal (e.g., $LiNbO_3$

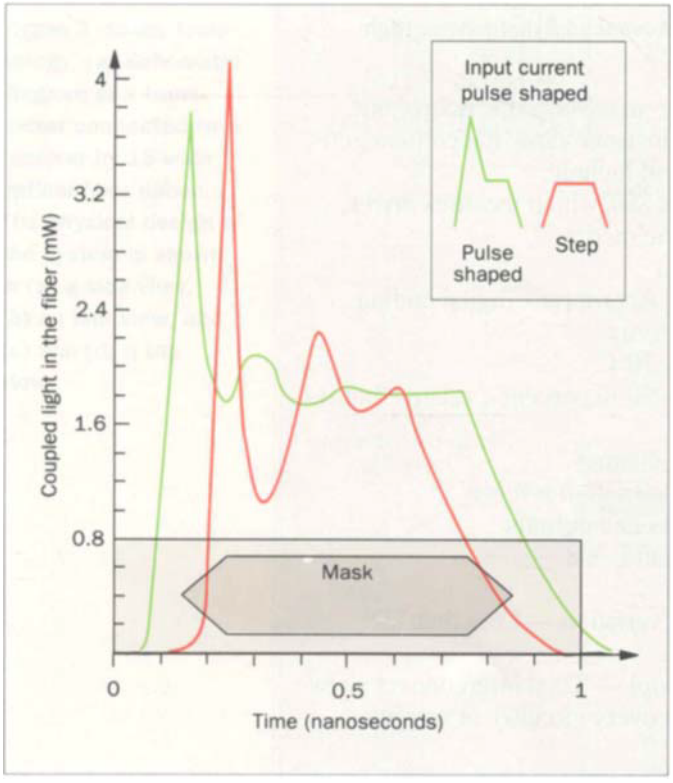

Figure 3. A simulated "optical eye diagram" using two different current pulse shapes to support a 1-Gbit/s data rate. The bright red line represents a step function current pulse with $t_r = t_f$ = 250 ps and a 20-mA magnitude. The other line is a shaped current pulse with a 40-mA initial spike settling back down to the 20-mA level. Pulse shaping relaxes oscillation peaks and reduces settling time. This delay, together with relaxation oscillations, dictates the laser bandwidth.

or InP directional coupler), an optical signal, a magnetic signal, a mechanical signal, or a combination of all of these. Usually, the light power source is carried in a fiber-optic cable. Once inside the 1D-ODL, the light is power divided and routed to each modulator. This allows the continuous wave (CW) high-power light source to be physically separate, unconstrained by size or power, and capable of providing reliability through redundancy. The incident CW light is directed on the optical device, which can then reflect or transmit light. The modulated output light is carried out of the 1D-ODL in waveguides. These modulators can either interact with, or latch, the data signal, or they can operate transparently. Examples of data-transparent optical modulators include deformable mirror devices,[7] $LiNbO_3$[8] liquid crystal devices (electro-optic effect),[9] and magneto-optic devices.[10] The self-electro-optic effect device (SEED)[11] is an example of the latter type of latchable modulator.

Although a modulator-based 1D-ODL is unlikely to be available soon, research on modulators and additional related topics is under way, and may offer future advantages.

Laser Arrays. In recent years, laser arrays have been considered for use as optical transmitters. One-dimensional arrays of edge-emitting lasers (EELs)[12,13] and two-dimensional arrays of surface-emitting lasers (SELs),[14,15] in both GaAs and InP material systems, have stimulated the search for new nonlinear devices with beneficial characteristics for 1D-ODL applications. Lasers are preferred over LEDs because they have:

- Higher external power efficiency
- A narrow output radiation cone that allows simple butt coupling of fibers
- Modulation capabilities that extend into the gigahertz[16] range (reduction of carrier lifetime with stimulated emission).

There are three fundamental ways of modulating a laser: direct current drive (with or without pre-bias), external cavity modulation (e.g., $LiNbO_3$ or other electro-optic devices), or intra-cavity modulation. Different system specifications (or performance specifications) and device limitations dictate the type of modulation that can be used. If designers use pre-bias modulation, or an external modulation scheme in which the contrast is low, they must "code" the transported data to set the proper threshold adjustment for the receiver. This results in an ac–coupled system. For an ac–coupled system to work, the threshold level at the receiver must be adjusted automatically using automatic gain control. A back-facet monitor at the transmitter controls the source power. To reduce the complexity and associated cost, a more digital approach (i.e., dc-coupled systems) is desirable. Direct modulation produces a dc-coupled system (i.e., logic 0 equals no transmitted photons, which results in a high contrast).

Current research on digital applications of laser arrays focuses on direct-drive lasers,[17] driven from a logic family such as emitted coupler logic (ECL), without current pre-bias. The philosophy behind the transmitter module electronic design centers on the direct digital

drive of the laser. The electronic current driver delivers a step of current to the laser, in which the magnitude of the current pulse in the "on" state is usually two to ten times greater than the lasing threshold. This is done without feedback (i.e., back-facet monitors).

Two design constraints immediately arise when considering this strategy. First, the magnitude of the current pulse must be large enough to deliver the necessary amount of optical power in the worst-case conditions (i.e., high temperature and near-end of life), and large enough to minimize the turn-on delay. Second, because the laser is not pre-biased above threshold, the turn-on delay and jitter can be excessive. This obstacle may be averted using appropriately shaped current pulses and engineering the laser properties to achieve the optimal coating design, cavity length, active area design, etc. Maintaining a low laser threshold current (while still maintaining high output power levels) maximizes signal fidelity and minimizes transmitter power dissipation.

Figure 3 shows a simulated "optical eye diagram" using two different current pulse shapes to support a 1-Gbit/s data rate. Pulse shaping reduces oscillation peaks and settling time. This delay, together with relaxation oscillations, dictates the laser's useful bandwidth. The uniformity of turn-on delay across the array (i.e., skew) also limits the usable bandwidth in synchronous applications. Decreasing the laser's cavity length and improving facet reflectivities decrease the turn-on delay. A high-bandwidth, direct-drive laser array developed by AT&T Bell Laboratories consists of a 12-laser InP/GaInAsP array[18] in which the lasing threshold is controlled to 9.9±0.9 mA, the optical power is 8.6±0.4 mW per facet, the bandwidth is greater than 4 gigahertz (GHz) (at 3 mW of optical power), and the crosstalk is less than –26 dB. Similar reports on InP-based laser arrays have been made by IBM[19] and ORTEL in GaAs, and by NTT[20] and Lincoln Labs.[21] The conclusion drawn from this simulation is that direct-current modulation, with pulse shaping, can result in an excellent eye diagram at 1 Gbit/s, in which a digital logic zero corresponds to no emitted photons.

A temperature-stable laser — one that does not require a thermal electric cooler — costs less and is more reliable. To increase the temperature stability, designers can lower the threshold currents and/or change the material system to one with a wider bandgap. GaAs lasers are more temperature stable than InP lasers, owing to the improved T_o constant, that is $I_{th} \propto \exp(T/T_o)$, and typical values for InGaAsP and AlGaAs are 50°C and 120°C, respectively. In addition, the gain of a GaAs laser is larger than that of an InP laser, indicating that GaAs laser arrays should be superior to InP laser arrays. However, InP lasers are currently more reliable than GaAs lasers,[22] because they suffer less facet deterioration with time, fewer dark line defects, and less nonradiative recombination in the cavity. As GaAs laser technology improves, the reliability difference should disappear. Future switching systems, which will have lower-cost optical connectors and plastic fibers, will use even shorter-wavelength optical sources. Although the optical interconnection applications being described in this paper are for distances of less than 100m, and for modest bit rates (1 Gbit/s), the wavelength used need not be dictated by minimum loss or dispersion in the fiber. The fundamental concerns for interconnect designers include low cost and minimum skew (matched delay) in the fiber ribbon. Even though minimum loss and dispersion in the fiber do not control the wavelength, wavelength does become an issue at bit rates greater than 1 Gbit/s, and/or at distances greater than 1 km. Here, InGaAsP/InP lasers operating at a minimum 1.3-mm fiber dispersion have a distinct advantage.

Device and circuit designers are currently debating the use of n-substrate versus p-substrate lasers. Since there are no differences between them in factors such as device processing costs, yields, or characteristics, p-substrate lasers would be preferable, from a circuit design perspective, because the electronic driver:

- Uses NPN transistors, which have a higher f_t
- Has a less complex circuit design
- Has a lower power dissipation.

Recent research in SELs for optical interconnection applications also suggest that they are an attractive candidate for the 1D-ODL light source. First, they can be tested at the wafer level, which lowers the device cost and provides a natural way to evaluate defect maps. Contrary to popular opinion, however, this does not represent a significant advantage for high-volume array production.

The second advantage of SELs is their potentially high coupling efficiency to fiber, the result of a single-mode gaussian output profile. The coupling efficiency can be increased because of the single-mode

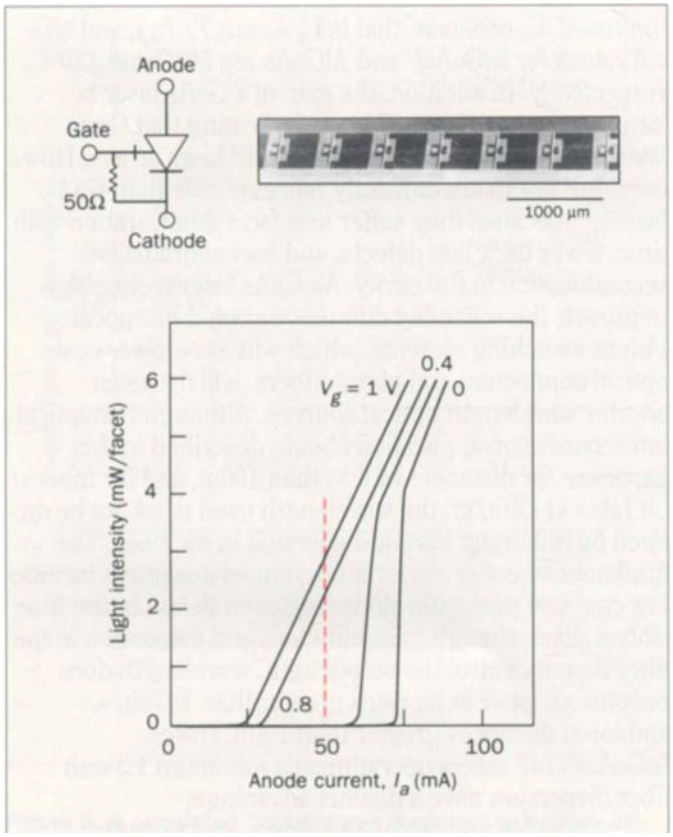

Figure 4. A schematic and a photomicrograph of a laser gate array.

nature of the output power. SEL technology has progressed enormously to date, especially in certain performance specifications. But for SELs to play an active role in 1D-ODLs, they must demonstrate:

- Direct, dc-coupled modulation
- Higher single-mode output power
- Higher "wall plug efficiency"
- Proven reliability performance
- High performance over a wider temperature range

than that of a comparable edge emitter. Although research on SELs and related topics continues, a SEL-based 1D-ODL is technologically premature.

Figure 4 shows another type of edge emitter, one that uses an intra-cavity loss modulator[26] to control light output. This FP laser gate is a three-terminal device. The anode (gain section) is dc biased above its lasing threshold. The cathode is at ground potential, and the gate (absorber section) is voltage modulated. The high-input impedance gate is terminated at 50Ω. A change in the gate voltage affects the magnitude of the loss (or gain) in the absorber section, which modulates the output light intensity. The light output is either in the spontaneous state, where the gate voltage is low, or the stimulated emission state, where the gate voltage is high. When the anode is dc-biased (e.g., I_a = 50 mA), there is a large nonlinear regime. Figure 4 shows how small changes in the gate voltage can cause a large change in light output intensity, producing a large contrast ratio. The absorber section becomes electrically saturated, and almost 100-percent modulation can occur. This enables the electrical or optical (with integrated PiN) input pulse to be reconditioned, provide dc coupling, and support a wider temperature change.

These devices also exhibit power gain, and, as such, may be considered to have transistor-like amplification characteristics. For example, a 30-μW change of electrical input power can deliver 7 mW of optical output power. The laser described is an InGaAs/InP buried-heterostructure graded-index separate-confinement heterostructure (BH-GRINSCH) laser diode with four quantum wells.

Conventional lasers outperform LED and modulator devices and cost less (if the cost of the receiver function is included). Compared with direct-drive, low-threshold laser arrays, laser gate arrays (LGAs) have:

- Wider bandwidth, because of the high input impedance of the gate
- Lower power dissipation, owing to the lower effective drive current and lack of power dissipation from an IC driver
- Direct ECL compatibility
- Reconditioned pulse shape, resulting in an improved bit error rate (BER)
- Less intersymbol interference
- Wider temperature margins (no thermal crosstalk)
- Lower delay times, to minimize skew.

Currently, the only disadvantage that LGAs have compared with direct-drive low-threshold lasers is a marginally increased process complexity. Isolating absorber regions and defining gate contact pads are the only additional processing required.

Receiver Design. The receiver module design philosophy contradicts that of traditional telecommunication receiver design strategies. Here, high-incident optical

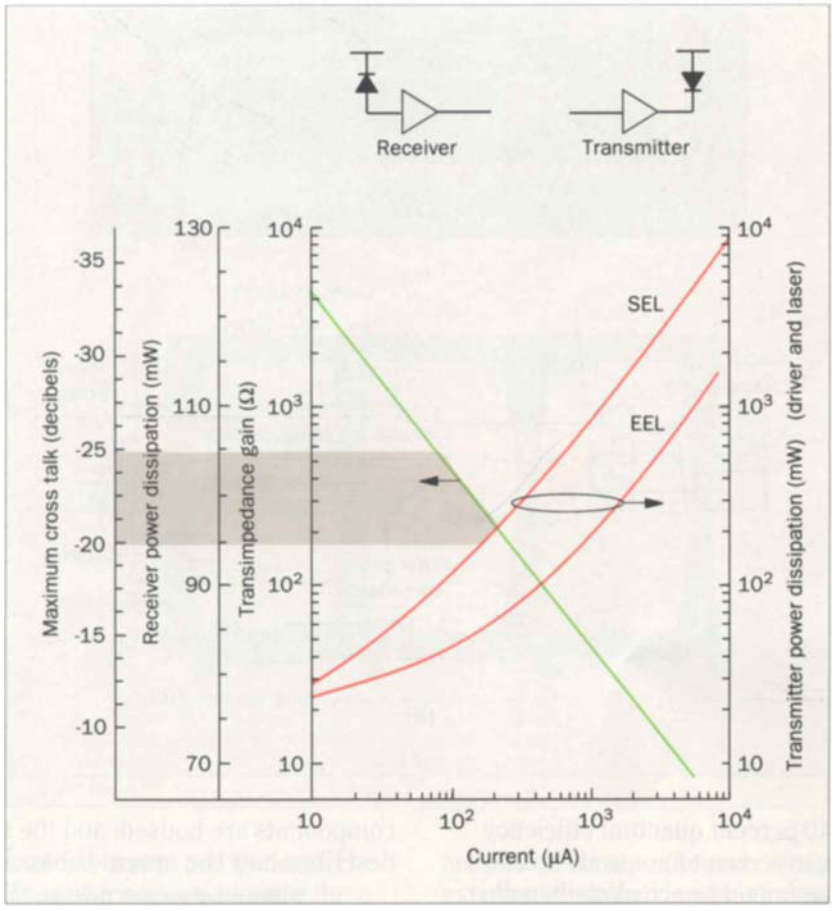

Figure 5. A comparison of light source technologies and receiver characteristics for an optimal data-link design. The receiver performance indices shown on the left are transimpedance gain requirements to produce a 50-mV output signal, the transimpedance amplifier power dissipation, and the minimum power/ground supply crosstalk allowed to prevent oscillation. The receiver power dissipation was determined by simply assuming 20 mW per gain stage (at a gain of 10 per stage).

power on the detector array produces a large current signal to the transimpedance amplifier array IC. Such a large signal significantly relaxes the gain and dynamic range requirements on the transimpedance amplifier. The transimpedance gain can be relaxed (e.g., to 300Ω), to integrate multiple amplifiers onto a single IC. This, in turn, will reduce signal delay variations and minimize cost. Otherwise, interchannel crosstalk resulting from the use of higher-gain amplifiers would prevent this level of integration. For example, 200 μW of light power absorbed in the detector at an 85-percent quantum efficiency will produce an output current of 170 μA. A transimpedance gain of 300Ω will produce an output voltage of 51 millivolts (mV). Generally, the output of a transimpedance amplifier is connected to a decision circuit, where the input signal is "regenerated." The 51-mV signal in this example is enough to drive the decision circuit. We estimate a per-channel power dissipation of 100 mW on the module, which generates a total power dissipation of 1.8W on the receiver module.

Figure 5 compares plots of different light source technologies. The receiver power dissipation was determined by simply assuming 20 mW per gain stage (at a transimpedance gain of 10 per stage). The crosstalk was scaled linearly with gain, and a maximum crosstalk of –30 dB was used as a reference at a transimpedance gain of 2000Ω (as per private discussions with Dr. R. N. Nottenburg on transimpedance amplifier array work at the University of Southern California).

The EEL power dissipation was calculated with 85-percent detector efficiency, 30-percent quantum efficiency, 3-dB coupling efficiency at the detector and transmitter, 5-mA threshold current, twice calculated current pulse height (to minimize turn-on delay), and current derived from a 5V power supply. For the calculated SEL power dissipation, the assumptions varied from the

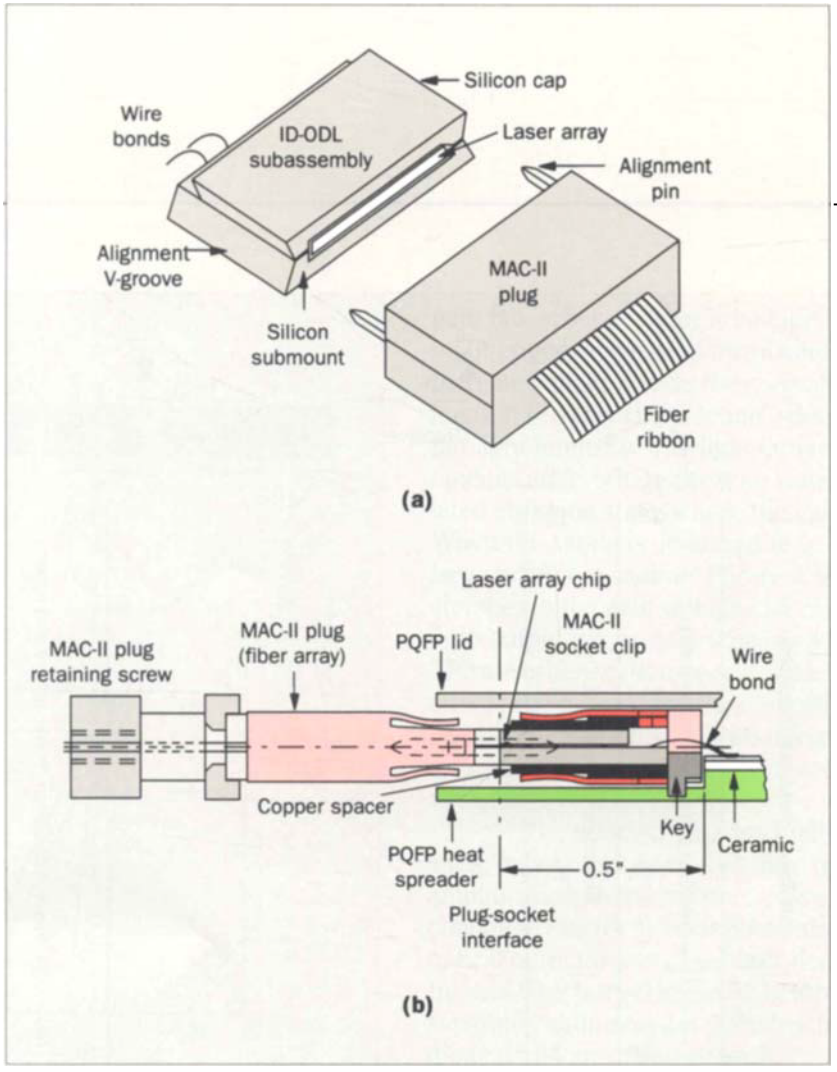

Figure 6. 1D-ODL transmitter. (a) Optical silicon subassembly connecting to MAC-II connector plug. (b) Module packaging scheme.

EEL only in two areas: a 10-percent quantum efficiency and a 6-dB coupling loss, to account for spacial hole-burning effects. These parameters were plotted against the output current exiting a detector array. With a detector current close to 500 μA, the receiver and transmitter (EEL) power dissipation balance to approximately 90 mW.

Clearly, it is desirable to lower the amount of detector current, even though this results in a slight impairment of the power balance with respect to the EEL. The evidence suggests that EELs should be the predominant light source technology for 1D-ODL applications. They provide a lower power dissipation and an overwhelmingly wider tolerance of power/ground crosstalk than do other technologies at this time.

Package Design for Parallel Optics. To successfully introduce new concepts into production-grade products, it is essential to design and develop an appropriate packaging scheme. It is advantageous to blend components currently in production (or extensions thereof) with a minimum of technological development, so products can be developed faster and with less risk. The 1D-ODL package design is described in the next two sections. The first section details the subassembly, where the optical components are housed, and the second section describes how the optical subassembly is packaged.

Silicon Submount Design. The 1D-ODL package was designed[24] as a connectorized transmitter (or receiver) module. The package was assembled within the tolerances needed to couple light from a laser array into a multimode fiber array. A strictly passive alignment scheme[25,26] was used, one in which the module was designed to fit inside a standard plastic-quad flat pack (PQFP). To make the design compatible with existing ribbon connectors, designers chose AT&T's MAC-II,[27] an 18-wide fiber ribbon connector. Consistent with the fiber array connector, the laser (and PiN) subassemblies used silicon. Silicon also provides a good thermal path for heat dissipation and can be manufactured to lithographic accuracies.[28] Figures 6a and b, respectively, show the alignment methodology and the 1D-ODL transmitter module packaging scheme. The alignment between the multimode fiber array and the laser was obtained by spring loading the alignment pins in the fiber array (MAC-II plug) into the alignment v-grooves in the transmitter module. The alignment v-grooves, located along the edges of the transmitter module, were formed by

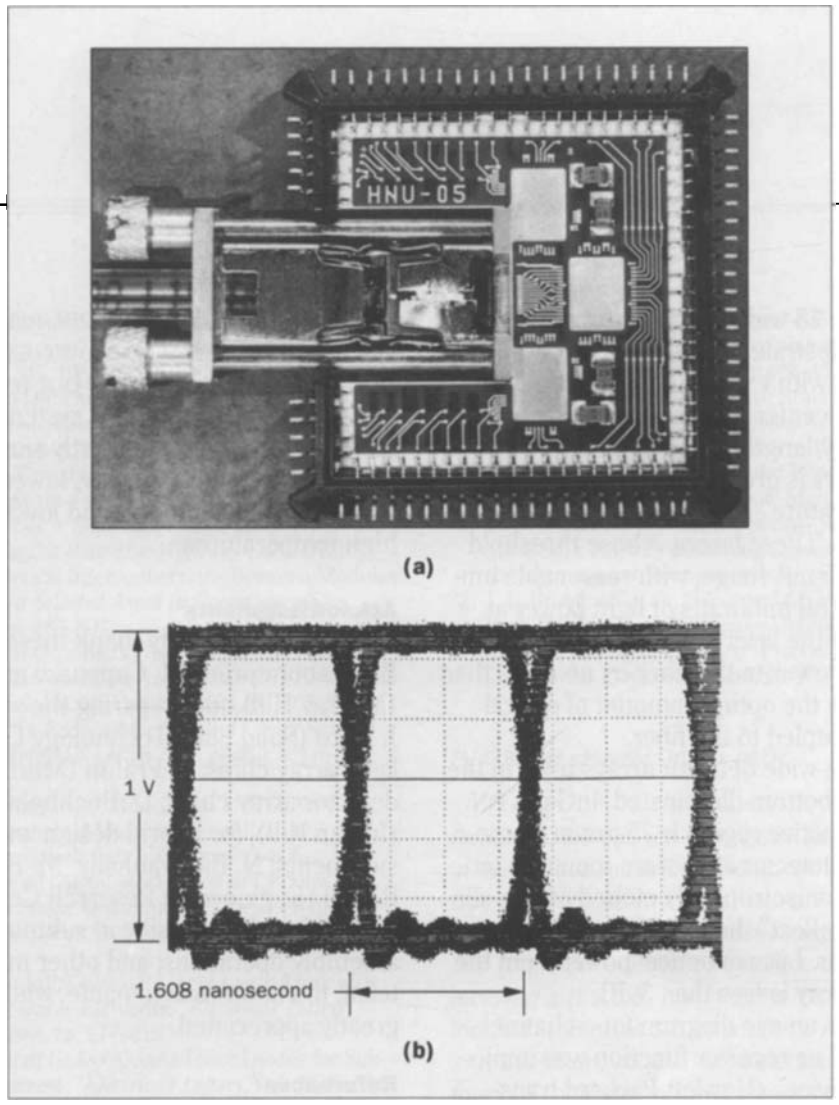

Figure 7. (a) Assembled transmitter module without its RTV or lid attachment. (b) Eye diagram of a pseudo-random pattern ($2^{15} - 1$) NRZ input signal transmitted through the 1D-ODL at 622 Mbits/s.

precisely etching the silicon submount and the cap. The bonding pads on the submount (corresponding to the p-contacts of the laser array) are located within lithographic accuracies of the etched v-grooves.

The laser array chip consists of 18 lasers located on 250-μm centers. It is bonded onto the silicon submount using Au/Sn solder, with the epitaxial side facing the submount. The solder bonding pads on the silicon submount in Figure 6a, located at the end of electrical interconnection lines, are used for driving the lasers. The solder bump height was optimized to provide a good mechanical joint, to dissipate the heat generated by the lasers, and to allow the laser array to self-align during the solder reflow process. A residue-free active atmosphere solder reflow process,[29] which does not require any post-cleaning, was used to keep the laser's surface free of contamination. Figure 7a shows a photograph of the completely assembled transmitter module.

Hybrid Design. The package for the 1D-ODL is based on a family of multilayer hybrid integrated circuits, called POLYHICs,[30,31] developed by AT&T. Precision thin-film resistors, fabricated on the ceramic substrate, are placed close to the ends of signal traces to provide excellent termination of high-data-rate lines. The lead frame supports data rates greater than 1 Gbit/s. Much of the POLYHIC package was designed to minimize development cost and interval, which was controlled by incorporating pre-existing, premolded lead frame and a quad-flat-pack geometry. For an initial predevelopment product, the package configuration was satisfactory.

1D-ODL Performance. How well any interconnection technology performs in a large digital system design is a function of:

- Signal density (on/off of the PCB and its footprint)
- Signal bandwidth
- Power dissipation
- Interconnection length limitations
- Cost per interconnection
- Skew
- Signal group delay variations
- Bit error rate
- Reliability.

The next two sections describe the performance of the transmitter and receiver modules in the 1D-ODL.

Transmitter. The 18-wide laser diode arrays used in the 1D-ODL were N-substrate, FP, InGaAs/InP (λ = 1.3 μm), bulk-active arrays, with PN junction-blocking layers. Laser facets, on 250-μm centers (to match the fiber-ribbon connector), have a cavity length of 250 μm. The –3 dB bandwidth of these lasers is greater than 2 GHz. Since the arrays are "high temperature" lasers, there is no need for thermal electric coolers. These lasers, whose threshold current is in the 10- to 15-mA range, with reasonable uniformity, can supply several milliwatts of light power at 85°C. "Butt" coupling of the lasers to the MAC-II connector reduces the amount of power to the fiber by no more than 3 dB. About 1 mW is the the optimal amount of optical power that should be coupled to the fiber.

Receiver. The 18-wide detector arrays used in the 1D-ODL are N-substrate, bottom-illuminated, InGaAs PiN diodes, whose optically active region is 75 μm in diameter on 250-μm centers. Detector arrays are mounted on top of a turning mirror (anisotropically etched on the silicon submount), to "up-reflect" the light from the butt-coupled MAC-II connector. Loss of optical power from the MAC-II to the detector array is less than 3 dB.

Figure 7b shows an eye diagram for a channel using 1D-ODL modules. The receiver function was implemented with discrete devices (Hewlett Packard transimpedance amplifiers and decision circuits), using an input synchronous optical network data rate of 622 Mbits/s (OC-12). The pattern generator produces a pseudo-random pattern length of $2^{15} - 1$ of non-return to zero (NRZ) data.

Conclusions

Advanced switching (and computer) system architecture designs require higher performance and lower-cost packaging technologies to produce new products and enhanced services. Electronic interconnection techniques — using technologies such as controlled collapse chip connect (C4) and multichip modules (MCMs) — offer performance superior to other methods at a cost that is lower for short-length interconnections (e.g., chip-to-chip packaging). Optical interconnection strategies are useful at higher levels of packaging (i.e., board-to-board, shelf-to-shelf, and frame-to-frame level of interconnections). Furthermore, parallel optical data links with high data rates now offer greater flexibility and lower cost than time-multiplexed serial data links.

Of the numerous techniques used to implement a parallel data link, the 1D-ODL may be the first to provide high performance at low cost. Designers agree that the key concepts of 1D-ODL technology include an ECL-compatible dc-coupled system, high optical output power and contrast that vastly simplify the receiver circuit design, improved skew, lower total power dissipation, small PCB footprint, and low thermal impedance at high temperatures.

Acknowledgments

We especially thank these people from AT&T Bell Laboratories: M. Cappuzzo and J. Shmulovich (Murray Hill), for preparing the silicon submounts; J. Zilko (Solid State Technology Center), for providing laser array chips; R. Frahm (Murray Hill), for providing detector array chips; D. Buchholz and R. Huisman (Indian Hill), for hybrid design and experimental measurements; N. Basavanhally, M. Brady, H. Nguyen, and R. Roll (Engineering Research Center), for designing and performing the silicon submount bonding and assembly operations; and other members of the 1D-ODL team, too numerous to name, whose efforts were greatly appreciated.

References

1. R. A. Nordin, et al.,"A Systems Perspective on Digital Interconnection Technology," *IEEE J. Lightwave Technology*, Vol. 10, No. 6, June 1992, pp. 811–827.
2. Y. Sakurai, et al, "Integrated Experimental System for BISDN," *ISS '92, Proceedings of the XIV International Switching Symposium*, Yokohama, Japan, October 25-30, 1992, Vol. 1, pp. 209–213.
3. J. F. Ewen, et al., "GaAs Fiber-Optic Modules for Optical Data Processing Networks," *IEEE J. Lightwave Technology*, Vol. 9, No.. 5, December 1991, pp. 1755–1764.
4. J. W. Parker, "Optical Interconnection for Advanced Processor Systems: A Review of the ESPRIT II OLIVES Program," *IEEE J. Lightwave Technology*, Vol. 9, No. 5, December 1991, pp. 1765–1773.
5. M. K. Kilcoyne, et al., "Optoelectronic Integrated Circuits for High Speed Signal Processing," *Optical Engineering*, Vol. 25, No. 10, October 1986, pp. 1119–1126.
6. K. Kaede, et al., "12-Channel Parallel Optical Fiber Transmission Using a Low Drive Current 1.3 mm LED Array and a PiN PD Array," *IEEE J. Lightwave Technology*, Vol. 8, No. 6, June 1990, pp. 883–887.
7. D. R. Pape and L. J. Hornbech, "Characteristics of the Deformable Mirror Device for Optical Information Processing," *Optical Engineering*, Vol. 22, 1983, pp. 675–681.
8. W. A. Payne and H. S. Hinton, "System Considerations for the Lithium Niobate Photonic Switching Technology," *Proceedings of the First Topical Meeting of Photonic Switching*, Incline Village, Nevada, March 18–20, 1987, pp. 196–199.
9. A. R. Tanguay, "Materials Requirements for Optical Processing and Computer Devices," *Optical Engineering*, January/February 1985, pp. 2–18.

10. W. E. Ross, et al., "Two-Dimensional Magneto-Optic Spacial Light Modulator for Signal Processing," *Society of Photo-Optical Instrumentation Engineers*, Vol. 341, 1982, pp. 191–198.
11. R. A. Novotny, et al., "Two-Dimensional Optical Data Lines," *Proceedings of the 43rd Electronic Components and Technology Conference*, Orlando, Florida, June 1–4, 1993, pp. 790–794.
12. S. Yamashita, et al., "Low-Threshold (3.2 mA per element) 1.3 mm InGaAsP MQW Laser Array on a p-type Substrate," *IEEE Photonic Technology Letters*, Vol. 4, No. 9, September 1992, pp. 954–957.
13. N. Yamanaka, et al., "A Gigabit-Rate Five Highway GaAs OE-LSI chip set for High-Speed Optical Interconnections Between Modules or VLSIs," *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 5, June 1991, pp. 689–697.
14. M. W. Maeda, et al., "Multi-Giga Bit/s Operation of 16 - Wavelength Vertical Cavity Surface Emitting Laser Array," *IEEE Photonics Technology Letters*, Vol. 3, No. 10, October 1991, pp. 863–865.
15. M. Orenstein, et al., "Matrix Addressable Vertical Cavity Surface Emitting Laser Array," *Electronics Letters*, Vol. 27, No. 5, January 1991, pp. 437–438.
16. H. Z. Chen, et al., "High Frequency Modulation of AlGaAs/GaAs Lasers Grown on Si Substrate by Molecular Beam Epitaxy," *Applied Physics Letters*, Vol. 52, No. 8, February 1988, pp. 605–606.
17. K. Y. Lau, et al., "High Speed Digital Modulation of Ultralow Threshold (< 1 mA) GaAs Single Quantum Well Lasers Without Bias," *Applied Physics Letters*, Vol. 51, No. 2, July 13, 1987, pp. 69–71.
18. L. A. Koszi, et al., "1.5 mm InP/GaInAsP Linear Laser Array with Twelve Individually Addressable Elements," *Electronic Letters*, Vol. 24, No. 4, February 1988, pp. 217–219.
19. E. Marclay, et al., "Scaling of GaAs/AlGaAs Laser Diodes for Submilliampere Threshold Current," *Electronic Letters*, Vol. 25, No. 14, July 1989, pp. 892–894.
20. M. Nakao, et al., "1.55 mm DFB Laser Array with 1/4 Shifted First Order Gratings Fabricated by X-ray Lithography," *Electronic Letters*, Vol. 25, No. 2, January 1989, pp. 148–149.
21. D. Yap, et al., "High Performance InGaAsP/InP Buried Heterostructure Lasers and Arrays Defined by Ion Beam Assisted Etching," *Applied Physics Letters*, Vol. 52, No. 18, May 1988, pp. 1464–1466.
22. F. R. Nash and R. L. Hartman, "Accelerated Facet Erosion Formation and Degradation of (Al,Ga)As Double-Heterostructure Lasers," *IEEE J. Quantum Electronics*, Vol. QE-16, No. 10, October 1980, pp. 1022–33.
23. J. O'Gorman, et al., "Dynamic and Static Response of Multielectrode Lasers," *Applied Physics Letters*, Vol. 57, No. 10, September 3, 1990, p. 968.
24. N. R. Basavanhally, et al., "Laser Optical Fiber Connector," U.S. Patent No. 104892, 1992.
25. M. S. Cohen, et al., "Passive Laser Fiber Alignment by Index Method," *IEEE Transactions in Photonic Technology Letters*, 1991, p. 985.
26. C. A. Armrento, et al., "Passive coupling of InGaAsP/InP Laser Array and Single Mode Fibers Using Silicon Wafer Board," *Electronic Letters*, Vol. 27, No. 12, June 6, 1991, p. 1109.
27. R. Bonanni, "Optical Connector," U.S. Patent No. 4818058, 1989.
28. P. E. Petersen, "Silicon as Mechanical Material," *Proceedings of the IEEE*, Vol. 70, No. 5, 1982, pp. 422–457.
29. R. D. Deshmukh, et al., "Active Atmosphere Solder Self Alignment and Bonding of Optical Components," *Proceedings of the International Electronic Packaging Society*, 1992, pp. 1037–1051.
30. C. Edge, et al., "Flip-Chip Solder Bond Mounting of Laser Diodes," *Electronic Letters*, Vol. 27, No. 6, March 4, 1991, pp. 499–501.
31. C. C. Shiflett, et al., "High Density Multilayer Hybrid Circuits Made with Polymer Insulating Layers (POLYHICs)," *Proceedings of the International Symposium on Microelectronics*, 1986, pp. 481–486.
32. J. L. Brandner, et al., "Electrical Characterization of POLYHIC, a High Density, High Frequency, Interconnection and Packaging Medium for Digital Circuits," *Proceedings of the 39th Electronic Components Conference*, 1989, pp. 759–764.

(Manuscript approved October 1993)

Ronald A. Nordin *is a distinguished member of technical staff in the Broadband and Photonic Switching Department at AT&T Bell Laboratories in Indian Hill Park (Naperville), Illinois. He develops optical interconnection equipment within the switch and to the subscriber for broadband systems. Mr. Nordin received a B.S.E.E. from Purdue University, West Lafayette, Indiana, and an M.S.E.E. and Ph.D. from Northwestern University, Evanston, Illinois, all in electrical engineering. He holds four U.S. patents, in the areas of optics and high-data-rate packaging. Mr. Nordin joined AT&T in 1977.*

Anthony F. J. Levi *was appointed Professor of Electrical Engineering at the University of Southern California in Los Angeles in 1993. From January 1984 to mid-1993, Mr. Levi worked in the Physics Research Division of AT&T Bell Laboratories in Murray Hill, New Jersey, first as a member of technical staff, and then, in 1988, as a distinguished member of technical staff. Mr. Levi's research interests include nonequilibrium phenomena in semiconductors, the scaling of electronic and photonic devices to submicron dimensions, and integration of electronic and photonic devices. He holds nine U.S. patents in these research areas. Mr. Levi received a B.Sc. from the University of Sussex, England, and a Ph.D. from Cambridge University, England, both in physics.*