A Systems Perspective on Digital Interconnection Technology

Ronald A. Nordin, *Senior Member, IEEE,* A. F. J. Levi, Richard N. Nottenburg, *Member, IEEE,* J. O'Gorman, T. Tanbun-Ek, and Ralph A. Logan, *Fellow, IEEE*

Abstract— **Advances in packaging and interconnection technologies for future telecommunication switching systems are paramount to remain competitive while offering enhanced services. This paper describes, from a switching system designers perspective, present digital interconnection technologies, and needs. The discussion begins at the chip to chip interconnection level and proceeds to the frame to frame level. We evaluate the optimal packaging and interconnection technology that should be used in future system designs. The analysis suggests that parallel optical data links based on a laser array technology implemented at the board to board level is presently advantageous. This analysis, as well as a detailed description of laser gate arrays, is also included.**

I. Introduction

INCREASED interconnection density and stringent reliability specifications are emerging for advanced wide bandwidth switching and high throughput computer system architectures. This need (or desire) arises from the demand for enhanced digital telecommunication services (e.g., Broadband ISDN, High Definition TV, video conferencing, cellular, etc.) and advanced computer services (e.g., graphics, real time image processing, pattern recognition, modeling, speech processing, etc.). These unremitting demands for higher interconnection density and bandwidth at all levels of the interconnection hierarchy, dictates advances in traditional electronic packaging and/or introduces the opportunity for new optical techniques into system designs. End-to-end optical data link products (e.g., ODL40® & ODL50®) have been used successfully for some time in high performance computer and telecommunication switching systems (e.g., 5ESS®). The authors contend that the next application of optics, to enhance system performance, will be the implementation of a generic parallel optical component which will serve to communicate over distances between 1 m and 10 km.

The purpose of this paper is to describe from a system's perspective the future needs for digital interconnection technologies. The discussion begins by describing present electronic interconnection technologies throughout the interconnection hierarchy (i.e., from chip-to-chip, to printed circuit board (PCB)-to-PCB on a shelf, to PCB-to-PCB across shelfs, to PCB-to-PCB across frames, etc.), in order to find where optics gains strategic advantage for high performance systems. This paper will then strive to show how optical schemes can be used to implement parallel synchronous data links, for PCB-to-PCB interconnections. The results suggest that optics at this level of interconnection is a competitive alternative to electronic strategies.

Finally the paper concludes with a discussion on a new optical component called a laser gate array (LGA), which optimally implements parallel optical bus requirements. This device (i.e., LGA) uses a voltage controlled saturable intracavity absorber to digitally (ON-OFF) modulate lasing light output [1]. Direct compatibility with emitter coupled logic (ECL) and voltage control (allowing correct 50 Ω termination of Gb/s digital electrical data) are unique advantages of this approach. LGA transmitters are a practical solution to the present 1 m to 10 km communication bottleneck in high performance computer and switching systems.

Manuscript received March 20, 1991.

R. A. Nordin is with AT&T Bell Laboratories, Naperville, IL 60566-7050.

A. F. J. Levi, J. O'Gorman, T. Tanbun-Ek, and R. A. Logan are with AT&T Bell Laboratories, Murray Hill, NJ 07974-2070.

R. N. Nottenburg was with AT&T Bell Laboratories, Murray Hill, NJ 07974–2070. He is now with the Department of Electrical Engineering, University of Southern California, Los Angeles, CA.

IEEE Log Number 9107427.

II. Electronic Interconnection Technology

To quantify where a communication "bottleneck" within a computer or switching system might exist, the problem can be classified into three categories: 1) IC die-to-IC die (chip-to-chip) communication, 2) PCB-to-PCB (intrashelf) communication, and 3) PCB-to-PCB (intershelf) communication. In each category, parameters such as power, bandwidth, complexity, fanout reliability, and cost must be evaluated to determine whether optics can be a competitive alternative.

A. Chip-to-Chip

The intercommunication between integrated circuits (chip-to-chip) can be studied based on three trends in device technology, namely

1. transistor (or gate) density on a chip,
2. density of chip-to-substrate contacts (I/O density), and
3. substrate interconnection density.

These issues will determine if optics is a cost competitive alternative to electronics, at the chip-to-chip interconnection level. Each issue will now be subsequently addressed.

1) *Integrated Circuit Transistor Density:* Fig. 1(a) shows the transistor density trend for two architectural types of IC's: Dynamic Random Access Memory—DRAM (highly symmetrical layout) and Digital Signal Processing DSP (computational type chips similar to switching and computer chips) [2]. For example, Motorola's 68040 processor (clock rate = 33 Mb/s)

chip has 1.2M transistors. A digital gate is an architecture dependent grouping of transistors normally on the average between 3 and 4 transistors per gate. The trend for increasing density will continue, due to improvements in IC scaling. However, as the IC dimensions approach physical limitations, increases in process complexity are mandatory to gain further miniaturization. An alternative, or a supplement to IC scaling, is the vertical stacking [3] of devices to gain perhaps a five-fold increase in circuit density. The technique [4] of selective epitaxial overgrowth is critical to obtain a high quality vertical stacked CMOS inverter structure. It is possible therefore, that the transistor density trend, shown in Fig. 1(a), may increase at a more rapid rate as this technology emerges from research. The chip technology that is incorporated into a system design (i.e., silicon—CMOS, NMOS, ECL, ... or GaAs [5]—HFET, BFL, ED, ...) is strictly a function of the system design requirements. Fig. 1(b) and (c) show a current view of different chip technologies versus two important system design parameters (gate delay and power dissipation per gate). Improvements in these parameters are expected to continue with IC scaling. BiCMOS [6] is currently a very important silicon IC technology because it combines the high transistor density and low power of CMOS with the high speed capability of ECL [7]. Note that this technology does not supplant scaled CMOS (e.g., 0.25-μm CMOS will out perform 1.5-μm BiCMOS), but adds design flexibility in layout (i.e., a gate on one side of the chip can communicate to another gate on the other side of the chip without submitting to a new chip layout) and adds higher performance at a given gate length process technology. The design of the NPN transistor in BiCMOS is disproportionately more complex (with respect to IC fabrication) than the CMOS transistor design as one scales the gate length below 1 μm. This perhaps, will manifest itself as an increased processing cost over CMOS as one scales both technologies (i.e., 1.25-μm BiCMOS is approximately 25% higher in cost than 1.25-μm CMOS). As one scales to 0.25 μm BiCMOS, the cost would be greater than 25% higher than 0.25-μm CMOS. Hence, it is not clear that the cost-performance ratio of BiCMOS will be lower than CMOS at all gate lengths, so one must carefully evaluate the extra cost in scaling CMOS to ascertain which technology is more cost effective. However, in the near future, BiCMOS should offer higher performance chip designs using complex circuits in the hundreds of kilo gates at very high data rates (hundreds of Mb/s). Therefore, this chip technology will most likely be prevalent in the next generation switching systems. The gate density trending exhibited in Fig. 1(a), can also estimate the number of required I/O terminals that would be needed, as shown in Fig. 2. Rent's rule is an empirical formula that can relate the chip density to the expected terminal count. Rent's rule can be expressed as, [8]

$$\text{gate} = \left(\frac{I/O}{k}\right)^{c}$$

where gate = average number of gates supported by I/O terminals,

I/O = number of signal terminals (chip pinout),

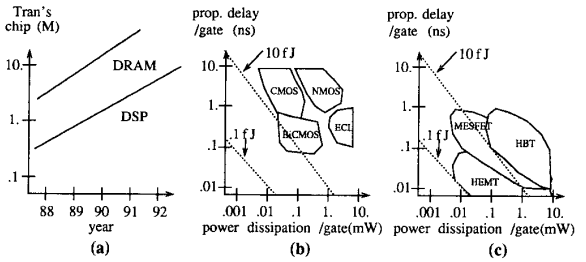

Fig. 1. Integrated circuit technology, (a) silicon transistor density trending and delay versus power in (b) silicon, and (c) GaAs material systems.

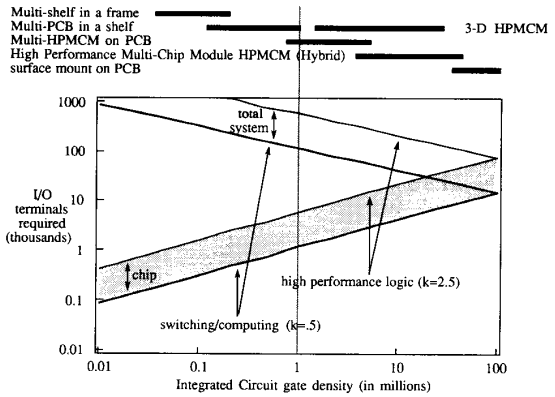

Fig. 2. Large system terminal estimation versus IC density (using Rents rule) and packaging technology options (assuming a required system complexity on the order of 100 million gates).

k = constant, the value of which depends on terminal sharing (i.e., $.15 < k < 0.2$ for memory chips, $0.4 < k < 0.6$ for switching and computer chips, $k > 1$ for high performance logic chips, and

c = constant in the 1.5 to 3.0 range ($c = 1.79$ was used).

Interpreting values from Fig. 2, a chip density on the order of 1 million gates would require approximately 1120 I/O terminals (pins). As further chip integration is obtained, the chip I/O demand increases. If given the problem of implementing a complex function (e.g., a 1024×1024 switch matrix) measured in the number of gates, the total system I/O terminal demand shown in Fig. 2 is created. In this case, a system complexity of 100 million gates is assumed. Here, as expected, the total system I/O terminal demand decreases with increasing chip density. For example, at a 1 million gate density (the dotted vertical line in the figure) the chip terminal demand is 1124 terminals, and to implement a 100 million gate function could require a minimum of 100 of these chips. Hence, the total system I/O terminal demand would be 100×1124 or 112, 400 terminals. This fact is highly informative, since it suggests the type of packaging that would be required to implement this complex function. The top portion of Fig. 2 shows the likely packaging technology options for a system complexity of 100 M gates versus the IC gate density.

Elaborating further, these 100 chips probably would not fit on one high performance multichip module (HPMCM),

therefore, multiple HPMCM's on a printed circuit board (PCB) would probably be the best choice of packaging. Additionally, this figure suggests what the bandwidth limitations would be as derived from the packaging choice. These interdependencies of packaging technology options and bandwidth are summarized in Fig. 3. Once again using this example problem, since multiple HPMCM's would need to be interconnected on a PCB, this interconnection length would dictate the maximum clock rate on the system (unless demultiplexing and remultiplexing occurs at the HPMCM's boundaries). Assuming an interconnection length of 25 cm on a glass epoxy board, the maximum clock rate would perhaps be between 400 and 600 Mb/s. (Note that these numbers are very sensitive to both architecture, PCB layout, and PCB material construction.) This limitation is based on a source terminated or a lossy unterminated transmission line structure. If one employs a terminated transmission line structure, the bandwidth could increase to 10 Gb/s. Note however, that in order to achieve this bit rate, the transmission line loss must be reduced by increasing the width and the thickness of the signal conductor. This will necessitate a decrease in the interconnection density, and hence tradeoffs between interconnection density and bit rate become apparent. The operating region shown in the shaded portion of Fig. 3 is where these trade offs exist. When moving vertically from the lossy line to the lossless line boundaries, the bit rate increases while the interconnection density (which is not shown in the figure) decreases for loss reasons as described above, and for potential architectural changes (i.e., demultiplexing and remultiplexing). As further chip integration is obtained, there is a higher chip I/O demand, a need for less total system I/O, fewer chips are required, and interconnection lengths are shorter (which would result in higher clock rates). The ramifications of increased chip density are advantageous to system designers. Note also that as chip density improves, the thermal management of the system is stressed, and new thermal strategies are needed. To summarize this subsection on transistor density, one can correctly conclude that further chip integration will occur, which will increase I/O demand off of the chip. However, this implies that for a given implementation of a complex function:

1. fewer (overall) chips will be required for a given circuit function which will lower the system's cost,
2. shorter chip-to-chip interconnection lengths will be required which results in a higher bandwidth,
3. and overall fewer chip I/O will be required which results in less system power.

These observations imply that the macroscopic (or system) interconnection complexity will be reduced but will be vastly increased at a microscopic level (i.e., chip I/O demand). This increase in chip I/O density is further explored in the next subsection.

2) *Chip-to-Substrate I/O Contact Density:* This section concerns itself with I/O density (i.e., now that a complex high density, high speed chip design has been devised, what are the limitations getting off the chip). The contact density from the chip-to-substrate (where the HPMCM used in the package is defined as a thin/thick film hybrid, silicon hybrid, or a

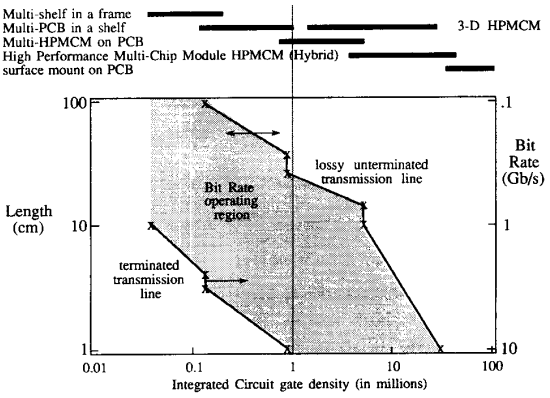

Fig. 3. Maximum system interconnection length and bit rate versus IC density. Packaging technology options also shown assuming a required system complexity on the order of 100 million gates.

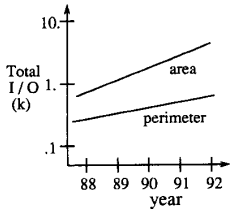

		year		
Parameter		88	90	92
perimeter (μm)	pad size	135	68	60
	pitch	200	100	90
area (μm)	pad size	100	65	25
	pitch	200	130	50

Fig. 4. Integrated Circuit I/O contact density trending.

PCB) may consist of a linear array of contacts along the chip perimeter, or a two-dimensional array of contacts over the chip's area. Fig. 4 shows how these technologies (perimeter and area) will trend in the short term. Presently, conventional wire bonding (ball and wedge) dominates low lead count IC's (i.e., < 200). This is an example of chip-to-substrate contact occurring along the perimeter of a chip. Obviously, for I/O starved designs, the array contact is preferred. However, this will demand a higher capital investment by the factory.

The reliability of a chip-to-chip interconnection is inversely proportional to the actual number of intermediate contacts required to provide the interconnection. The chip-to-substrate contact level is sometimes quantified as the number of actual contacts needed to connect the chip to the substrate. For example, Fig. 5 depicts a chip that is wirebonded to a lead frame, which is then solder connected to a substrate. Three actual connections are made, and is thereby referred to as a level-three interconnect. Tape automated bonding techniques (TAB) is a level-two interconnect since the inner lead solders to the chip, while the outer lead solders to the substrate. Both TAB, and wirebonded leadframes are perimeter contact techniques. TAB is generally preferred over wirebonding because of:

1. improved reliability (i.e., lower number of intermediate contacts per connection),
2. parallel bonding (leads are gang bonded instead of sequentially bonded),
3. higher density (narrow inner lead pitch),
4. better signal transmission [9] (lower inductance), and

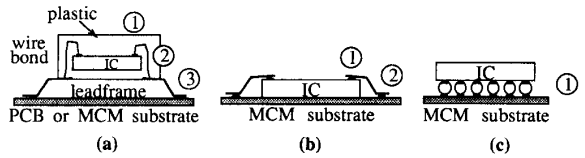

Fig. 5. Packaging interconnection levels, (a) level 3 (wirebond), (b) level 2 (TAB), and (c) level 1 (C4).

5. TAB provides a lower surface profile.

Present TAB technology [10] provides an inner lead pitch of under 100 μm, which results in greater than 500 leads per die. This density is expected to increase (perhaps to near 700 or 800 leads), but, in order to obtain factors of 2 to 4 times improvement, other technologies should be exploited. One such peripheral contact array technology is the metallized SiO_2 microcantilevers [11]. They are similar to AT&T's beam lead technology (with dimensions of 20 μm wide on a 40-μm pitch), and can yield more than 1000 I/O's around the perimeter of a 1 cm $\times$ 1 cm chip.

An increasingly popular area contact technique is the controlled collapse chip connection (C4) [12] process (sometimes referred to as flip chip or solder bump attach). The C4 technique is a level-one interconnect since the solder bump connects the chip directly to the substrate. A product example that uses C4 technology contains 762 solder bumps (a 29×29 array) [13]. The most dense area reported has been a 128×128 array of 25-μm bumps with a 60-μm pitch, which resulted in nearly 16 000 connections [14]. C4 is generally preferred over TAB and wire-bonding for the following reasons:

1. improved reliability due to its lower interconnection level,
2. higher density from 2D arrays, which can have a finer pitch than TAB,
3. better signal transmission due to its lower lead inductance,
4. provides a better thermal contact to the substrate for localized heating in proximity to the solder bump,
5. improved chip density on HPMCM, and
6. self alignment feature offers advantage for optical connections.

Another area contact technology that provides direct chip interconnect, is a polymer thick film [15] (PTF) process. With this method a polymer adhesive can be used for gold (on chip) to gold (on PCB) pressure contacts, sometimes referred to as "microbump." This technology has been exemplified through work done by Matsushita Electric Industrial Co., [16] by applying 54 chips to a substrate via microbump technology. Chips having I/O contacts at a pitch of 63.5 μm were separated by 10 μm from their nearest neighbors. An alternative contact technology uses anisotropic conductive polymer adhesives (ACPA) which can be employed to bond very fine pitched conductors (10 μm). Advantages derived through ACPA technology include: high density interconnections, increased mechanically flexible contacts, and provides a probable path to three-dimensional packaging technology.

Table I provides numerical information on the chip-to-substrate technology that is currently available. Note, that if one employs Rent's rule (Rent's rule usually over estimates I/O contact needs), then for 1 M and 10 M gate density chips, on the order of 1124 and 4070 I/O's would be required, respectively. This is good news for system designers since no new electrical contact technology needs to be developed in order to meet near term chip evolutions. One might suspect that with an increasing number of I/O pads on a chip, the contact pads will take up usable silicon surface area, and thereby subtract from the gate density capability. However, this is a not a major concern since the contact pads are on the chip's surface (on top of a dielectric spacer) with only small vias that route down to the I/O transistor driver (i.e., many transistors can be placed underneath an I/O contact pad). The only practical concerns are the actual size of the output transistor (which is much larger than the gate transistors), and the complexity of the driver circuit. For a fixed chip size this aspect will contribute to I/O limitations. For example, in order to drive an output at an electrical power level of say 10 mW (e.g., $V\text{out} = 0.7\text{V}$ and $I\text{out} = 15$ mA), the data in Table II depicts the circuit and transistor areas that would be required for each transistor technology (assume pad size of $30\ \mu\text{m} \times 30\ \mu\text{m} = 900\ \mu\text{m}^2$). The HBT transistor technology offers the highest return when regarding the minimal driver area usage. As time progresses, the other technologies will benefit from IC scaling, which leads one to conclude that the I/O density will eventually be limited by the chip surface contact technology utilized, and not by the driver surface area.

In summary, one concludes that present C4 technology is adequate for near term (i.e., less than 10 years) high speed and high density chip designs. In addition, C4 technology should be more reliable, have higher bandwidth, have higher density, provide a lower junction thermal resistance, and have a lower per-contact cost than alternative electronic contact technologies. One will have to supply additional capital expense to enter the technology, but the advantages gained should outweigh the cost. A final observation, with respect to contact density of a chip, is that there are no significant differences between optics and electronics at this stage of packaging. The differences, if any, will occur in the density of substrate interconnections.

3) *Substrate Interconnection Density:* The final issue to discuss, with regards to the intercommunication between IC chips, is the substrate interconnection density [17], [18]. Given that one has a potentially complex, high-speed chip, along with a high density I/O contact technology, the remaining question that arises is; can the substrate provide the interconnection density imposed by the chip? In actuality, there are two aspects that control this interconnection density. These aspects concern the interconnecting transmission line structures from chip-to-chip, and the routing from the contact pad site (underneath the chip) to the transmission line (outside the chip). As one would expect, the transmission line width is much greater than the routing line width, and hence different signal densities are obtained. Additionally, vias, which provide connection from the contact pads on the substrate to internal signal conductors within the substrate, block otherwise available routing tracks underneath the chip, and thereby reduce routing interconnection density. We define the substrate transmission

TABLE I
CHIP TO SUBSTRATE TECHNOLOGIES

Technique	Min Lead Pitch	Contacts 15 sq-mm	Bandwidth	Complexity*	reliability	level	ref
wirebond	100 μm	600	$<$ 5 GHz	Low	Medium	3	8
TAB	90 μm	650	5 GHz	Medium	Medium+	2	9, 10
C4	65 μm	50 000	$\gg$ 5 GHz	High	High	1	12
C4	200 μm	5600	$>$ 5 GHz	High	High	1	11
microbump	10 μm	2M	$\gg$ 5GHz	Medium	High	1	14, 15
microcantilevers	40 μm	1500	$>$ 5 GHz	High	Medium+	2	13

*complexity in terms of cost (initial and maintenance of capital equipment).

TABLE II
TRANSISTOR DRIVER CIRCUIT AREA COMPARED TO PAD AREA

Transistor Technology	# of transistors in driver circuit	Area of output transistor	total driver circuit area	total % of pad space
BJT	2	510 μm^2	675 μm^2	75
HBT	1	135 μm^2	135 μm^2	15
MOSFET	4	550 μm^2	810 μm^2	90
MESFET	4	550 μm^2	810 μm^2	90
HFET	2	410 μm^2	540 μm^2	60

line interconnection density, and the routing interconnection density as the number of signal lines per cross-sectional width. Numerous reports [19], [20] on both substrate types (silicon and ceramic) have achieved both high chip interconnection densities, high frequency, and improved cost-performance ratios. Due to possible thermal expansion coefficient mismatches between the chip and substrate (for cases of very high density I/O's), the silicon chip-to-silicon substrate is preferred. However, other options are available (silicon chip-to-ceramic substrate, gallium arsenide chip-to-silicon substrate, etc.) with reduced I/O and/or chip dimensions.

Fig. 6 shows I/O signal density trends for the two types of hybrids (silicon and ceramic based). Fig. 6(a) shows the routing density, while Fig. 6(b) shows the transmission line signal density. The authors will now consider a two-dimensional packaging technique, possessing only a two-layer capability on one side (usually having a polymer as the dielectric), and transmission line rules at 10 μm width and 9 μm space (i.e., pitch = 19μm). This configuration will result in a signal density of 1575 signals/cm. The 9-μm spacing of the transmission line, is designed such that the crosstalk would be less than -30 dB. Once the signal arrives at and is propagating in the transmission line, any additional crosstalk from neighboring transmission lines would be inconsequential. Now, multiplying the above signal density by 4 (for the four sides of a chip) will allow the substrate to connect to 6300 contacts on a 10 mm $\times$ 10 mm chip. From Table I, less than a 200-μm pitch C4 style package would be required (on the order of 125 μm) to support this substrate signal density. Note, if one used the transmission line structure to connect to the substrate pad, then, assuming a pad-to-pad gap (routing track width) of 70 μm, only 12 transmission lines could fit through the gap, and hence would limit the pad matrix size to 24 $\times$ 24 pads. To increase this contact density, two methods could be pursued. One can increase the gap width between pads, and/or replace the transmission line with a metalized lead (routing interconnection). One must be concerned with the length of this routing interconnection because of crosstalk and bandwidth. However, by simply keeping this length short, the crosstalk will remain small and the bandwidth will remain high. The routing interconnection density (using the 1990 data from the table in Fig. 6), is approximately 1.21 per μm. Assuming a routing track width of 70 μm, there would be 85 routing tracks between contact pads. With a square contact pad matrix of $80 \times 80 = 6400$ contacts, there is an excess of routing tracks to interconnect the contact pad to the transmission line (e.g., $70 \times 1.21 = 85$ routing tracks, which theoretically would allow one to route a 170 $\times$ 170 pad matrix). Since the number of routing tracks greatly exceeds the required amount, one does not have to limit the layout to straight lines, but could entertain some random type interconnections.

An excellent example of the capabilities of electronic sub-

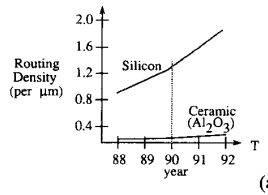

		year		
	metal line pitch	88	90	92
Silicon (μm)	metal level 1	2.6	1.9	1.3
	metal level 2	3.0	2.3	1.5
	metal level 3	5.2	3.9	2.6
Al_2O_3 (μm)	metal level 1	30	25	20
	metal level 2	40	30	25
	metal level 3		40	35

(a)

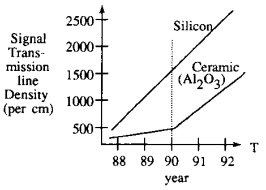

		year		
	Parameter	88	90	92
Silicon (μm)	line width	35	10	6
	dielectric depth	10	3	2
	# signal layers	3	3	3
Al_2O_3 (μm)	line width	102	51	25
	dielectric depth	38	19	9.7
	# signal layers	2	3	6

(b)

Fig. 6. Substrate signal density trending, (a) routing and (b) transmission line interconnections.

strate interconnection comes from Boeing Aerospace Electronics High Technology Center [21]. They achieved 800 lines per cm using a two layer 25-μm pitch and a bandwidth (for a 5 cm line) of approximately 2 GHz when uftilizing lossy transmission line technology. Additionally, they showed an excellent planarization technique in the presence of 6-μm-thick conductors with a 9-μm polyimide interlevel dielectric layer. For further I/O densities, one would have to go to double-sided substrates and/or four layer capabilities per side. This is not challenging substrate technologies, but only adding to the process complexity and cost.

A final comment on the limitation of electrical contact density concerning ground bounce will now be discussed. Ground bounce occurs during the transient switching time of an output driver, and increases when more outputs switch at the same time. One way to minimize ground bounce is to supply many ground and power contacts to the chip that are capacitively coupled to ground. The layout of the ground and power planes, as well as the signal interconnection distribution across these planes are critical in determining the ground/power plane inductance which ultimately leads to the calculation of ground bounce. Another approach to minimize ground bounce, is to interconnect chips with coupled differential transmission lines, such that the bounce associated with the noninverted output cancels the bounce associated with the inverted outputs. Although this technique does not eliminate ground bounce, it tends to minimize it. However, the interconnect density is adversely affected since there would be twice the number of signal conductors and the density will decrease by slightly more than a factor of two. Presently, there are a number of groups across the industry that are addressing this problem (with respect to CAD tools), so that one will be able to simulate different layouts in effort to minimize ground bounce.

The bandwidth of these transmission lines (defined by an interconnection length of 2 cm at a loss of 3 dB) is approximately 5 GHz since lossy unterminated transmission lines are generally delay limited. When terminated in its characteristic impedance, the bandwidth increases to approximately 15 GHz (assuming: width = 35μm, metal thickness = 1μm, and dielectric thickness = 10μm). Higher bandwidth transmission lines can be fabricated with increased processing complexity. For example, a coplanar strip transmission line in an air trench [22] has a measured rise time of 0.8 ps for a 2.8 mm length of propagation. The power advantage from utilizing substrates like silicon and ceramics over that of printed circuit boards can be utilized because the interconnection length can be made very small so that the line need not be treated as a transmission line, and therefore not need to terminate to a 50 Ω load resistor.

Considering a clock frequency of 200 MHz with a rise time of 1 ns, implies that the bandwidth of the transmission line must be greater than 1 GHz in order for the Fourier component contributing to the rise time to propagate through the line. For line lengths less than 2 cm (where this length is defined as $L = \nu_{\text{prop}} \cdot \frac{f}{10}$) one need not treat them as transmission lines. Hence, the power required to drive a 2 cm line $= C \cdot \frac{V^2}{\tau_r} = \frac{(1\text{ pF/cm})(2\text{ cm})(1\nu)^2}{1\text{ ns}} = 2$ mW, whereas for a line length greater than 2 cm, a transmission line (either terminated or unterminated-lossy) must be employed where the power would be nearer to 10 mW (e.g., $\frac{V^2}{R} = \frac{1^2}{100}$). These power considerations will be considered in more detail later in the paper.

The three-dimensional multichip packaging technology has emerged over recent years and is expected to be a competitive packaging technology for high performance systems. Advantages of this technology include: reduced size, weight, power, cost, delay and increased reliability.

From a packaging viewpoint, in order for optics to play a competitive role, it must at least provide the performance parameters as outlined in Table I. It is the authors' opinion, that an electronic "bottleneck" has not yet limited further system development, and the problem (or bottleneck) has been successfully pushed to the next level of packaging (i.e., PCB-to-PCB or intra-shelf communication). Note however, that optics (especially modulator based) can provide an advantage with respect to chip power dissipation (for the interconnection function only) over that of an electrical alternative, when the interconnection length is greater than 1 mm. This is one of many important system packaging concerns, as it suggests where optics may be successfully injected, and hence will be consequently discussed.

The comparison between electrical and optical (free-space) techniques, with respect to power and speed requirements, [23]–[25] will be presented in the problematic form of interconnecting two processor chips. (The authors are assuming that the ratio of processing electronic gates to electrical/optical I/O is high.) The parameter plot in Fig. 7 shows the results of this comparison. Part (a) of the figure shows the chip power dissipation required from the source versus the interconnection length with the bit rate and the source type (i.e., electrical (e.g., C4), laser (e.g., Surface Emitting Laser SEL), or modulator (e.g., Symmetric Self Electrooptic Effect Device S-SEED)) as parameters. Part (b) of the figure, simplifies these results in a plot of the bit rate versus the interconnection length (i.e., where the curves of Part (a) intersect). By increasing the interconnection length, the electrical and modulator inter-

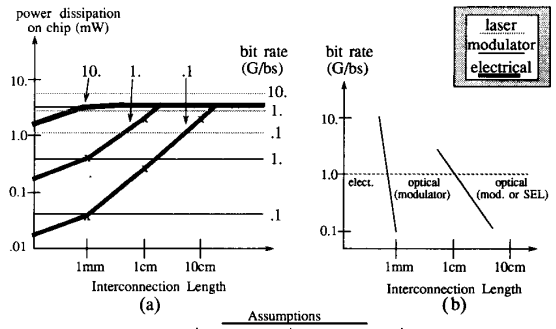

Assumptions			
parameter	electrical	optical - laser	optical - modulator
received detector voltage	1V	1V	1V
electrical driver efficiency	75%		
transmission line capacitance	1pF/cm		
input gate capacitance	20 fF	20 fF	20 fF
input pad capacitance	20 fF		
photo detector capacitance		5 fF	5 fF
load resistor capacitance		5 fF	5 fF
responsivity of photo diode		.8 A/W	.8 A/W
interconnect efficiency		90%	80%
modulator efficiency			50%
laser efficiency		30%	

Fig. 7. A comparison of Electrical versus Optical interconnection, (a) chip power dissipation per I/O (includes transmitter and receiver) and (b) efficient interconnection regimes w/r/t power dissipation.

connection schemes cross over near 1 mm, at a 1 Gb/s bit rate. For interconnection lengths greater than 1 mm, optics (based on modulators) is more efficient ($w/r/t$ power dissipation). For lengths less than 1 mm, electrical techniques are more efficient. Increasing the length even further, the electrical and laser based interconnection schemes cross over near 1 cm. Therefore, one may conclude that optics can reduce the total operating power of a digital system when interconnecting substrate-to-substrate as opposed to chip-to-chip and perhaps spread substrates out further to gain advantage in system thermal management. Hence, in systems where high density I/O's are needed with a high percentage of the I/O required to be active at one time, optics may provide an advantage ($w/r/t$ power dissipation) at the chip-to-chip interconnect level. The above assertion does not clearly describe where optics provides an advantage, since the class of applications that require high density I/O's with a required high percentage of them being active at one time is not an obvious one at this time. More importantly, it is simply the cost of the interconnection that dictates which technology is ultimately used in advanced systems.

It is interesting to note that when the chip-to-chip spacing is kept to a minimum, it does not matter whether the interconnect medium is optical or electrical since the chip power dissipation will limit the quantity of output connections. As example, assume the chip designer allocates 5 W to the output driver portion of the chip. If assuming an output dominated design, then at 1 mW of dissipation per output implies 5000 output contacts could be utilized. At a thermal resistance of 8°C/W, the chip will be 40°C above the ambient temperature, which probably exceeds the cooling capability of the system. It is in this chip thermal management research area, that advantages in large system performance can be exploited. Not only can one increase the number of connections but also enhance the reliability of the chip by improving the thermal design of the

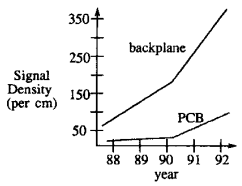

	year		
Parameter	88	90	92
line width (mils)	17.4	14	7
dielectric depth (mils)	10.	8	4
PCB (62.5 mils) # signal layers	3	4	6
backplane thickness (mils)	160	256	320

Fig. 8. PCB and backplane signal density trending.

chip. Irregardless of the progress in optical or electrical interconnection technology, before interconnection technology can be effectively exploited, enhancements in thermal management technology must take place.

In summary, the authors conclude that the electrically interconnected substrate can manage high density chip I/O contacts adequately without much further technology improvement. In addition, not only can it manage this density, but also provide a high bandwidth connection, lower the required source power, and lower the cost per connection below that of alternative technologies. Therefore, the insertion of optical interconnects as a direct replacement of point-to-point electrical interconnects (at the intra-substrate, chip-to-chip level) does not offer any major advantages for present day computing or switching systems. Advantages are visible when the interconnect length, speed, and fanout are pushed to the electrical realizability limits [26]. Since the packaging "bottleneck" seems to have been successfully pushed to the next level of interconnection (PCB-to-PCB—intra shelf or module-to-module communication), it will be discussed in the next section.

B. PCB-to-PCB Intrashelf Communication

In this level of intercommunication, the problem is to interconnect a PCB to another PCB that resides on the same shelf (i.e., in the electrical domain, the interconnection is implemented in the backplane). This interconnection level is sometimes referred to as the electrical or optical backplane technology. The interconnection between shelves is deferred to the next section. Ignoring, for the moment, the needs of the PCB to backplane interconnection, the authors will first evaluate the backplane needs separately. Fig. 8 shows I/O density trends for backplane boards and for PCB's. (Note that for the backplane interconnection densities shown in the figure, that only 70 to 90% will actually be routable due to staked pins or vias blocking routing tracks.) The point at which optics will become competitive with electronics is a function of cost, interconnection distance, signal density, bit rate, and fanout. In order to provide the comparison, an isocost curve is provided in Fig. 9 for point-to-point connections on a glass epoxy backplane. The interconnection distance is limited by conductor and dielectric loss, where the signal density dictates physical dimensions of the microstrip lines used.

Fig. 9(c) provides cost information at a bit rate of 200 Mb/s and a signal density of 20 connections/cm. For interconnection distances less than 1m, the cost of providing a 200 Mb/s interconnection from one PCB to another PCB on the same

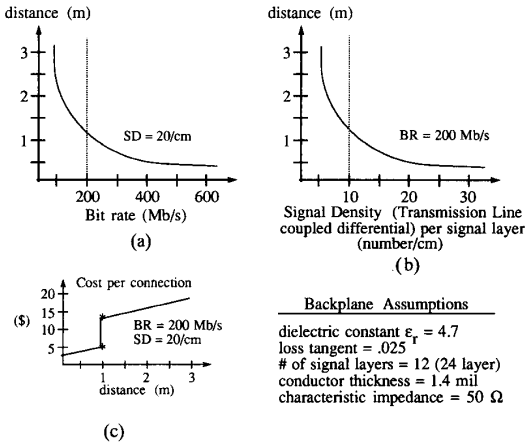

Fig. 9. Iso-cost curves for (a) distance versus bit rate (with the signal density held constant at 20 connections/cm), (b) distance versus signal density (with the bit rate held constant at 200 Mb/s), and (c) cost curve versus distance (with both the signal density and bit rate held constant) for electrical backplanes.

shelf (e.g., backplane ($2.00), driver ($1.), receiver ($1.), and 4 pins ($1.20)) is approximately $5.20. Beyond a meter, the cost jumps because alternative electrical media must be utilized (i.e., coaxial lines or twisted pair cable). Therefore, when operating below this isocost curve (approximately $5.20), optics must be more cost competitive than electrical backplane. Note however, if one includes fanout, the isocost curves shift downward, thereby making optics more competitive because the cost of providing a passive fanout in optics (e.g., a star coupler implemented with polymetric mixing rod [27], [28] techniques) could be relatively cheap. The magnitude of the downward shift depends on the following cost relationships: the cost of the optical power splitter, and the additional cost of supplying more optical power versus the cost of the "electrical power splitter" (i.e., either the additional cost of more drivers or the cost of FET buffers to provide active taps off of the transmission line).

Another issue that one must address is the number of I/O pins on a PCB edge that can connect it to the backplane. For a relatively large board (43 cm× 33 cm), 600 staked pin type connectors are readily available. However, the bandwidth is relatively low, on the order of 500 MHz. If one adds ground blades around the pins in such a way that no reduction in the number of signal I/O pins will occur for grounding purposes, one can increase this bandwidth to approximately 1 GHz. With other connector technologies (i.e., Gold Dot, Cinch—Synapse, Beta Phase—thermal shape memory) the density exceeds 2000 I/O's per PCB, and bandwidths in the range of 2 GHz can be obtained.

Guided wave optics can supply an advantage since it can provide a factor of 1.3 to 1.5 reduction in the connector area (assuming a multifiber array connector (MAC) type [29], compared to a staked pin electrical connector). Hence, for an equivalent electrical connector area of 600 contact pins, which could supply perhaps 120 signal I/O's and the remaining pins being held at ground, the optical ribbon connector could supply 216 I/O's (assuming the MAC type connector). In addition to an improved connector I/O density, the proposed optical package (see section V–III) could reduce by a factor of 4 the PCB area required to drive the I/O's, increase the connector as well as the interconnection bandwidth, lower I/O power requirements, lower electro-magnetic emission (and raise susceptibility), and increase fanout possibilities. All of these improvements are key items for switch system designers. (We note, that in the future, it is possible that free-space optical interconnection of PCB's might exceed the optical guided wave approach by an appreciable factor, and hence might be an attractive technology.)

The final issue to discuss on PCB-to-PCB communication is the PCB signal density. This concern addresses the problem of handling the high pinout of the substrate and the routing of these contacts to the PCB edge connector. First of all, the signal density for a standard PCB (with 4 high-speed layers with FR4 as the dielectric) is near 150/cm. The contact density need not be greater than this, since the additional contacts could not be routed away from the substrate. For perimeter contacts, this would mean that a contact pad would have a width of 33 μm with a 33 μm space. Obviously, this is pushing the limits of PCB manufacture, and a staggered perimeter contact array or a two-dimensional contact array would have to be utilized.

To summarize, for a substrate size of 6 cm × 6 cm, 3600 connections could theoretically be removed from the substrate. These connections could do one of two things, route to the board's connector, and/or route to another substrate on the PCB. Since these connections are long and could be at a high data rate, optics could play a role in providing substrate-to-substrate communication but perhaps a more likely role would be in the substrate on one PCB to a substrate on another PCB (intra-shelf communication). An optical solution becomes even more competitive in the PCB-to-PCB inter-shelf communication, basically because the interconnection length is longer. This packaging level will be discussed next.

C. PCB-to-PCB Intershelf Communication

In this level of interconnection an electronic backplane is not available to interconnect PCB's together. The following options exist: coaxial cables (single or ribboned), twisted pair (single or ribboned), flex microstrip circuits, and of course the optical alternatives. In these instances, optics [30] can compete quite well with the electrical alternatives, since the electrical technologies become expensive (i.e., with respect to cable cost and system cost) whereas the optical cost as a function of distance is relatively flat. As was shown earlier in this section, the cross over point where optics is cheaper than electrical is a function of bit rate, distance, density, skew, and fanout requirements. For a point-to-point interconnection, the crossover point is near a bit rate of 200 Mb/s, an interconnection length of 1 m, and a signal density of 20 connections/cm. Therefore, with the additional advantages as listed above, fiber optics presents an attractive alternative for use in advanced switching (or computer) system designs.

Note also, that optics (using fiber ribbon) has improved skew performance (delay variation) to 2–5 ps/m as compared to

electrical coaxial ribbon cable which is in the 300 ps/m range. This implies that a system can be designed synchronously without the cost burden of providing clock recovery circuits per channel. The data is simply re-timed at the destination by a simple digital skew cancellation technique. Additionally, optics provides nearly a factor of 100 times improvement in the signal density (i.e. signals/cross sectional area) in the backplane as compared to coaxial ribbon connections. Although a cost figure advantage is difficult to quantify, Optics is clearly advantageous with respect to assembly, maintenance and repair.

In summary, it has been shown that an optical solution is best utilized as a backplane technology. Optics will have a very difficult time replacing electronic interconnection technologies at the chip-to-chip or at the substrate-to-substrate (within a PCB) packaging levels in the near future. In order to penetrate this area, optics has to provide revolutionary functionality, and not compete in an evolutionary electronic arena. However, in the backplane environment, Optics does compete effectively as an evolutionary interconnection technology. The backplane provides interconnection for PCB's on the same shelf, PCB's on different shelves in the same frame, and/or PCB's in completely different frames. The advantages ultimately manifest themselves as a lower cost technology. The additional performance parameters outlined in the preceding section, make the optical performance-cost ratio far superior than the electrical alternatives. The remainder of this paper will be concerned with optical solutions for backplane interconnection technologies.

III. System Backplane Interconnection Requirements

The backplane interconnection system requirements are implementation independent (optical or electronic) and provide the requirements for intrashelf (0 to 1 m distances) as well as intershelf (1 m to 1 km distances) PCB interconnections. Fundamentally, the interconnection technology should look like a digital computer bus. This implies that the signal levels are quantized into two levels (digital) with sufficient noise margins such that optical, electrical or thermal feedback are not necessary to readjust threshold levels in the transmitter or the receiver. This is radically different from traditional long haul optical transmission systems where feedback is necessary in order to get the highest level of system performance. The problems that are being solved in switching are very different. They allow simplified transmitter and receiver designs at a low cost per interconnect.

The general system requirements (based on generic switching central office environments) are as follows:

1. Cost/Interconnection $<$ \$10 (includes: driver, receiver and connectors and back-plane media)
2. Reliability: 10^5 hours $< \rho$
3. Bit Error Rate (BER): $\text{BER}_{\text{DATA}} < 10^{-14}$ and $\text{BER}_{\text{VOICE}} < 10^{-12}$ (at 1 Gb/s)
4. Temperature (ambient): $0°\text{C} < T < 70°\text{C}$
5. Absolute Humidity < 0.026 (5% $<$ relative humidity $<$ 95%)
6. Distance: 1 $<$ Length $<$ 100 m
7. ECL I/O level compatible
8. Power Supplies: $+5$, ground, -2, and -4.8
9. Bit Rate (BR): $0 <$ BR $<$ 650 Mb/s
10. Delay Variation ($\Delta\tau_d$): $<$ 200 ps
11. SKEW: across the total interconnect skew $<$ 400 ps (driven by the desire not to have clock recovery circuitry nor allow for any jitter accumulation.)
12. Fanout: (e.g., 8) as much as possible (application dependent). The reason is that point-to-point connections (i.e., ring networks where the fanout is one) are not as desirable as a bus architecture (where the fanout > 1) because of the added functionality a bus offers.
13. Packaging Size: 2.5 cm$\times$ 2.5 cm (integrated 18 channels)
14. Receiver Sensitivity < -20 dBm
15. Dynamic Range $>$ 16 dB

IV. Optical Solutions

There are three competitive optical transmitter technologies which provide an optical backplane strategy [31]. They are

1. Light Emitting Diode (LED) arrays (one- or two-dimensional),
2. laser diode arrays (one- or two-dimensional),
3. optical modulator arrays (one- or two-dimensional), and
4. laser gate arrays (LGA).

Each will be subsequently discussed, with respect to system performance and cost.

A. LED Arrays

LED arrays [32] have been regarded as an attractive optical transmitter for four reasons: they have a low fabrication cost, high reliability, are easily extendible to large two-dimensional arrays (i.e., surface emitting), and are more temperature stable than other optical methods. These reasons can lead to a lower system cost solution.

One must keep in mind however, that the chip cost, especially compared with that of the optical device packaging, becomes a small percentage of the total packaged cost. Therefore, differences in chip fabrication cost may have less impact on the total cost. The extendibility to large two-dimensional arrays is an important advantage as the demand for optical interconnect density increases, and is especially important as a chip-to-chip interconnect vehicle, where the interconnecting electronics are in close proximity to the LED's, or are monolithically integrated on the LED array.

Unfortunately, there are three disadvantages with LED's. Their external power efficiency is low (a fiber coupled efficiency on the order of 0.36 mW/A (with integral lens) at a 200 Mb/s bit rate), and manifests itself in unwanted chip power dissipation, thereby limiting the size of the LED array. For example, assume that only 100 μW of optical source power is required from each LED. Note that this is very low power with respect to large digital system needs. This would require an LED drive current of 280 mA, leading to each LED dissipating approximately 0.42 W (assuming $v_f = 1.5$V). Now, assuming that with a thermal management system, a total dissipation of 10 W could be tolerated, then only 24

LED's could be integrated on a chip. This is probably the most fundamental drawback of utilizing LED arrays in a large optically interconnected switching (or computing) system. Additionally this higher power dissipation also leads to higher chip temperatures, which results in an optical output power decrease (approximately 0.5%/°C). This power fluctuation could lead to thermal crosstalk.

The other disadvantages (a long recombination lifetime and nonlinear termination) of LED configurations are somewhat interrelated, in that they both limit the LED's bandwidth. The modulation bandwidth of LED's (neglecting parasitic capacitance) is generally controlled by the long recombination lifetime of the carriers, in comparison to the reduced lifetime of the carriers in a laser which results from stimulated emission. Consequently, LED's are fundamentally slower than injection lasers, from carrier lifetime arguments, in addition to the nonlinear termination a LED provides.

The last aspect to consider involves the driver to LED connection. For high data rates, where the interconnection length (l) is long (i.e., $1 > \lambda/10$) the interconnecting media must be treated as a transmission line. In this case, either a source termination or a load termination (or a lossy transmission line concept) must be employed. Since the diode presents itself as a nonlinear load (in the digital on/off modulation scheme), reflections will occur with either termination technique. Matching networks can be used to make the load appear to be 50 Ω, but this adds complexity and increases power dissipation. However, with no matching network, the bandwidth will be lower when compared with a characteristic impedance load terminated transmission line. Note that the bit rate is dependent on the delay that occurs in the transmission line (propagation delay) for nonpipelined architectures. This delay or transmission line settling time is usually greater than twice the propagation delay, hence limiting the bit rate to approximately one over the delay. For linear devices, (e.g., LED's and lasers) these oscillations will couple into the optical output, and thereby place further design constraints on the receiver circuitry.

In summary, the efficiency and modulation bandwidth are fundamentally limited in the LED configuration. A lasing configuration would be advantageous in both of the above respects, while inherently providing smaller beam divergence. Note that some of these limitations could be overcome with an edge emitting LED. However, as will be shown in subsequent sections, a laser would still be preferable.

B. Laser Arrays

In recent years, laser arrays have attracted much attention as candidates for the optical transmitter function. One-dimensional arrays of edge emitting lasers (EEL's) and two-dimensional arrays of surface emitting lasers (SEL's) in both GaAs and InP Material systems have stimulated the search for new nonlinear devices [33] with beneficial characteristics for computing/switching architectures. Lasers are preferred over LED's for the following reasons:

1. higher external power efficiency (usually mW's of optical power are available with less than 20 mW of input electrical power),
2. better spatial coherence (allows simple butt coupling of fibers), and
3. modulation capabilities that extend into the GHz [34] regime (resulting from a reduced carrier lifetime arising from 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, as discussed in the next section), or intra-cavity modulation (as will be discussed in the laser gate array section). Different system specifications (or performance specifications) generally dictate the type of modulation that can be used. Research on digital applications of laser arrays to date is focused on direct drive lasers [35] (directly driven from a logic family like ECL without current pre-bias). This technique is highly desirable since it allows high contrast switching, and obviates the need for bias monitoring and feedback control.

With a low threshold current (on the order of a mA) the laser can be directly driven by an ECL output buffer. This is advantageous for system designers since a custom high power current driver is not required and high switching currents are minimized. System designers will also benefit from a laser that is temperature stable (i.e., a design that does not require a Thermal Electric Cooler—TEC). The benefits include lower cost and higher reliability. Two general techniques can be pursued in order to obtain temperature stability: drive the threshold currents lower, and change the material system to a wider bandgap material. Temperature stability is improved with GaAs lasers over that of InP lasers due to the improved T_0 constant (i.e. $I_{\text{th}} \propto \exp(T/T_0)$). Note also, that the gain of a GaAs laser is larger than InP lasers. One might conclude then that GaAs laser arrays would be advantageous over InP laser arrays. This is basically true (neglecting wavelength differences which lead to more/less dispersion in fiber). However, an important concern is that of the reliability differences between GaAs and InP lasers. InP lasers are presently more reliable than GaAs lasers [36] (less facet deterioration with time, less dark line defects, and less nonradiative recombination in the cavity). But, as GaAs laser technology improves, the reliability difference should disappear. Even lower wavelength materials (higher bandgaps) may be advantageous in future switching systems arising from lower cost optical connectors and plastic fibers. Since the optical interconnection applications being described here are for short lengths (e.g., less than 100 m) and modest bit rates (1 Gb/s) the wavelength used need not be dictated by minimum loss or dispersion in the fiber. The fundamental concerns for interconnect designers, are low cost and minimum skew (matched delay) in the fiber ribbon. These concerns do not necessarily control the wavelength. Wavelength does become an issue above 1 Gb/s and or greater distance (1 km). In this situation InGaAsP/InP lasers operating at the 1.5-μm fiber dispersion minimum have a distinct advantage.

A good example of a high bandwidth direct drive laser array has been produced by AT&T Bell Laboratories. They have reported an InP/GaInAsP laser array [37] (12 lasers) whose lasing threshold was controlled to 9.8 ± 0.9 mA, an optical

power to 8.6 ± 0.4 mW per facet (at 100 mA), bandwidth greater than 4 GHz (at 3 mW of optical power), and crosstalk less than -26 dB. Similar reports on laser arrays have been made by IBM [38] and ORTEL [35] in GaAs, and by NTT [39] and Lincoln Labs [40] in InP. Generally as the amplitude of the current pulse increases, the turn-on delay decreases. Turn-on delays less than 100 ps are obtained with 20 mA current pulse. Note that this delay together with relaxation oscillations dictate the laser bandwidth. Additionally, the uniformity of turn-on delay across the array (skew) will also limit the usable bandwidth assuming synchronous operation. By decreasing the laser's cavity length, and improving facet reflectivities, the turn-on delay may be decreased. The disadvantages of present laser arrays are summarized below:

1. *Inter-Symbol Interference (Pattern Effect):* The starting condition for each pulse is affected by the charge left over from the previous pulse. This charge decays (for a laser biased below threshold) with a time constant equal to the carrier life-time. This effect is found to be proportional to the following ratio:

$$\frac{I_{\text{pulse}}}{I_{\text{pulse}} - I_{\text{threshold}}}.$$

   Therefore, in order to keep intersymbol interference small, the current pulse must be large compared to the threshold.
2. *Driver to Laser Connection:* Here again, as with the LED, the laser diode presents itself as a nonlinear termination. To obtain high speed operation (assuming direct drive technique), the driver must be in close proximity to the laser.
3. *Voltage Level Shift:* The direct drive laser still requires a level shift function or an additional power supply to be ECL compatible. This could be overcome by the integration of transistors with the lasers [41] (OEIC technology), however, the added fabrication complexity might degrade reliability, performance and increase cost [42]. However, in the future, it is expected that these issues will be overcome, and OEIC technology will be viable.
4. *No Regeneration:* Since the laser is a linear analog electrical-to-optical conversion device, the input current pulse is not "cleaned up" by the laser, so that anomalies in the current pulse are passed on to the receiver.
5. *Thermal Crosstalk:* There is laser to laser thermal crosstalk caused by the high current drive in neighboring lasers that are relatively close to each other.
6. *Skew:* There are laser to laser optical delay differences (skew), arising from the driver circuitry (ac and dc components) and laser threshold variation.

When very low threshold laser arrays become prolific, all of the above disadvantages become minor in magnitude except for the cost of the electronic driver IC, its power dissipation, and the added system skew. But as one would expect in time and volume production, the cost of the electronic driver IC's will decrease, and the system performance (power and skew) improve. Therefore, the difference between a very low threshold laser, and gated laser array technology (see Section V) would be minor, suggesting that the technology choice will simply be a cost ($w/r/t$ manufacturability, reliability, yield, alternative outside second source, etc.) decision. Presently the gated laser configuration, to be discussed, provides the advantages a laser affords, and overcomes the above outlined disadvantages. However, this is not necessarily a long term technology choice. This device should make significant penetration into "backplane" applications where higher speeds (> 100 Mb/s) and longer distances (> 1 m) are required.

C. Optical Modulator Arrays

The last conventional optical candidate to be discussed in this paper is the class of optical devices called modulators. The advantage of this type of device is that they are easily extendible into large two-dimensional arrays and that their power requirements are small. Generally, these devices are interconnected in free space (as opposed to a guided wave technique) due to beam combination requirement complexities for large two-dimensional arrays. A modulator can be either active or passive. Examples of passive optical modulators include:

1. Deformable Mirror Devices (DMD)
2. Liquid Crystal Devices (Electrooptic Devices)
3. Magneto-Optic Devices.

The characteristics of these passive modulators are such that they: do not regenerate the optical signal, resynchronization can only occur at the expense of pulse width narrowing, only regular interconnection patterns can be exploited, and generally no fanout can occur. These disadvantages can be overcome by utilizing an active optical modulator. Presently, there are two approaches as to where the modulators can physically reside. They can be integrated with the laser source [43], or they can reside far from the source. An example of this latter type of active modulator is the self electrooptic effect device (SEED). The principle of operation of this device is based on the quantum confined stark effect where the optical absorption depends on changes in the electric field applied normal to the device. When two SEED's are connected in series, the circuit behaves as a set-reset latch with differential optical inputs and complementary optical outputs. This circuit configuration has been named the Symmetric SEED (S-SEED) [44], [45]. Since this structure has memory, resynchronization can occur without pulse width narrowing. This circuit also has time sequential gain (device state set by low power beams and sampled with high power beams), which "reconditions" the input pulse shape. Also, since the circuit is differential, it is not sensitive to optical power supply fluctuations. From the above, active modulators have advantages over passive modulators, however the following improvements must be achieved before large arrays of active modulators can become competitive with alternative techniques:

1. improve optical modulation intensity (> 1 mW),
2. reduced power dissipation per gate,
3. increased interconnection density ($> 128 \times 128$ array),
4. improved signal functionality (i.e., digital signal processing),
5. irregular and bussed interconnection architectures,

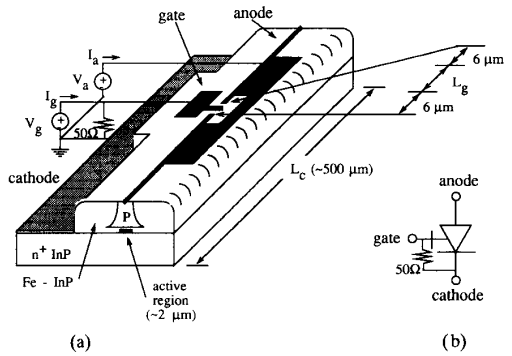

Fig. 10. Intra-cavity saturable absorber laser, (a) schematic sketch and (b) schematic representation.

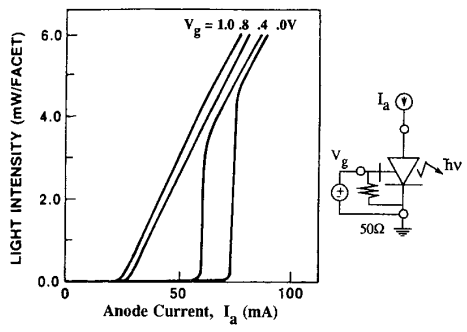

Fig. 11. Measured light intensity per facet as a function of anode current, I_a, for various gate voltages.

6. lower optical hardware cost and complexity (i.e., lens, gratings, beam splitters ...),
7. reduced loss budget from input light source to modulator,
8. improve reliability on the optical power source.

While research on these and additional related topics are under way, a free-space interconnection technology based on active modulators is unlikely to be exploited in switch systems in the near future, but may offer future advantages with respect to relieving I/O congestion.

D. Laser Gate Array

The laser gate consists of a small intracavity electroabsorptive loss modulator [46], integrated into a semiconductor laser diode cavity as shown schematically in Fig. 10(a). This paper will focus on short haul switching/computing applications (length < 1 km), where digital (ON-OFF) Fabry–Perot (FP) devices are advantageous from a cost perspective. For longer distance communications (e.g., local area networks) distributed feedback (DFB) lasers are necessary, due to dispersion limitations in the fiber. DFB lasers with intra-cavity absorbers can be an effective alternative to conventional laser techniques [47], [48].

The device essentially is a three-terminal FP laser shown in Fig. 10(b). The anode (the gain section) is dc biased above its lasing threshold. The cathode is at ground potential and the gate (the absorber section) is voltage modulated. The high input impedance gate is then 50 Ω terminated. A change in the gate voltage changes the magnitude of the loss (or gain) in the absorber section which modulates the output light intensity. Basically, the light output is either in the spontaneous state (gate voltage low), or the stimulated emission state (gate voltage high). When the anode is dc biased (e.g., I_a = 60 mA), there is a large nonlinear regime in which a large change in light output intensity occurs for small changes in the gate voltage, and hence a large contrast ratio is obtained as shown in Fig. 11. Note that the absorber section saturates (electrically), and near 100% modulation can occur. This aspect is very important since the input pulse (electrical or optical (with integrated P-I-N)) can be reconditioned and can support a wider temperature change.

Additionally, these devices exhibit power gain, and as such may be considered as having transistor-like amplification

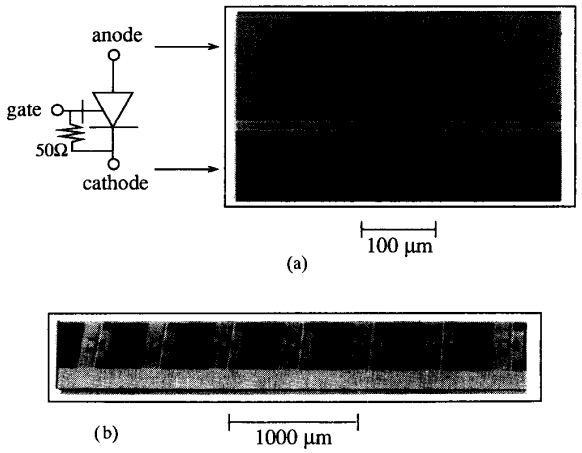

Fig. 12. Photomicrograph of (a) a laser gate and (b) an array of laser gates.

characteristics (i.e., 30 μW change of electrical input power can deliver 7 mW of optical output power).

The laser, described in this memo, is an InGaAs/InP buried-heterostructure graded-index separate-confinement heterostructure (BH-GRINSCH) laser diode with four quantum wells. The physical construction of the device is detailed in [1]. The laser gate array's optical outputs (edge emitting) are at a 508-μm pitch. A photomicrograph of a laser gate and an array of laser gates are shown in Fig. 12.

As outlined in the preceding sections, optics can provide a better performance-cost ratio than an electrical equivalent for high density backplane applications (see Fig. 9). Conventional lasers outperform LED and modulator techniques, and at a better system cost (when one includes the cost of the receiver function). LGA have the following advantages over direct drive, low threshold laser arrays:

1. wider bandwidth due to 50-Ω device termination
2. lower power dissipation due to the lower effective drive current, as well as, no power dissipation from an IC driver
3. directly ECL compatible ⇒ lower cost (no extra power supply or level shift)
4. reconditions pulse shape (absorber section saturates) ⇒

better noise margin (NM) and improved bit error rate (BER)
5. less intersymbol interference
6. wider temperature margins (no thermal crosstalk)
7. lower delay times (minimizes skew).

Currently, the only identified disadvantage of LGA compared to direct drive, low threshold lasers is a marginally increased process complexity. The only additional processing required is that needed to isolate absorber regions and to define gate contact pads. As with direct drive, low threshold lasers, relaxation oscillations are observed and tend to limit the device bandwidth.

The result is an attractive candidate for an optical backplane interconnection technology. In the following subsections, the discussion of the laser gate device characteristics will be separated into its dc and ac performance.

1) *DC Device Characteristics:* Fig. 11 shows the typical light output per facet for various gate (absorber section) voltages, versus anode (gain section) current I_a. When V_g is logically high (approximately equal to the anode voltage V_a), the laser has a threshold current of $I_{\text{th}} = 20$ mA. By decreasing V_g to near 0.4 V (a logic zero), the absorber section increases its loss so that the threshold current shifts out to 60 mA. Note the nonlinear behavior of the device near the gate voltage of 0.4 V. This occurs due to the exponential nature of the absorber loss and the saturation effect of the absorber with respect to the applied gate voltage. This nonlinear behavior is advantageous for the following reasons:

1. it reconditions the input digital waveform,
2. it increases the noise margin of the device,
3. it increases the contrast ratio, and
4. it increases the temperature range of operation.

These attributes are helpful for achieving error free data transmission. Fig. 13 shows the transfer curve of the device for various I_a (dc) operating currents. For a 0.6 V logic swing (like the logic family ECL), the noise margins are nearly identical at approximately: NM0 = 135 mV and NM1 = 175 mV at $I_a = 60$ mA and the temperature was controlled at 15°C (where the noise margins are defined at the 10 and 90% points). Note that the threshold gate voltage can be controlled with process variations to recenter the transfer curve. These values are quite good, compared to ECL noise margins (NM0 = 155 mV and NM1 = 125 mV). Note that the actual system noise margins should be determined with respect to the input voltage and the output voltage. The output voltage is obtained through an optical receiver circuit. The above calculation serves only as an instructive model for relative comparisons of optical transmitters. Figs. 11 and 13 also demonstrate that milliwatts of optical power can be switched using relatively low input electrical power. This power efficiency is approximately 30 to 50% (slope efficiency for a coated laser) [49], which could be improved when coupled to a lensed fiber. This is an important parameter when integrating many lasers together on a die for thermal management reasons. With this large switched optical power, system design benefits by:

1. relaxing requirements on the optical receiver (perhaps to −10-dBm sensitivity which eases amplifier and decision circuit design),
2. relaxing connector loss specifications,
3. provide for large fanout of the optical signal,
4. relaxing fanout splitter efficiency, and
5. lowering packaging costs (e.g., through the use of non-hermetically packaging, the excess optical source power may offset facet degradation that may occur at a higher rate).

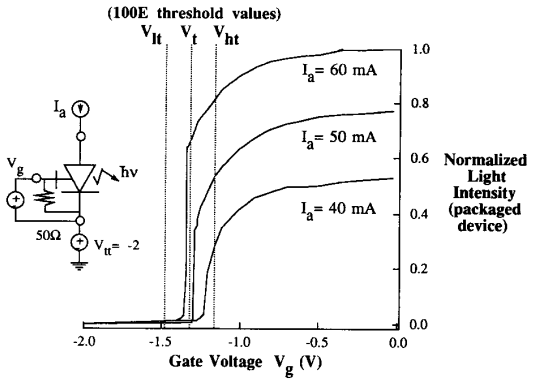

Fig. 13. Transfer characteristics of the packaged laser gate.

Overall the system design benefits from a more flexible architectural design, thereby dramatically lowering the cost of the optical interconnection.

The LGA is directly compatible to the ECL logic family. This is a very important concept since the voltage level shift circuitry is not required, nor is an extra power supply needed. The ECL compatibility is obtained by connecting the cathode of the device to a −2 V power supply (which should already be present in ECL logic designs). Since the LGA's gate voltage threshold is near 0.7 V, when the cathode is biased to −2 V, the gate voltage threshold shifts to −1.3 V. This new threshold is exactly centered in the ECL logic family transfer curve, and hence with a simple power supply connection, the LGA becomes ECL level compatible.

Another important system design concern is the uniformity (with respect to anode threshold current and gate voltage) of the gated lasers along the array and from die to die. Preliminary measured results of the threshold current (at $V_g = 1$ V), and the gate voltage threshold (at $I_a = 60$ mA) along the array, are very encouraging and even tighter distributions are expected with process refinements.

The temperature dependence of the threshold current and the gate voltage for the InP laser gate was also measured. Assuming an ECL logic swing of 0.6 V, implies that a temperature-safe operating window of 10°C is possible. This suggests, that initially a thermal electric cooler (TEC) must be incorporated to control the device operating temperature. Effort is underway to eliminate the TEC to relax packaging requirements and lower the device cost. This may be possible if two steps are accomplished: first, change laser technology from InP to a wider bandgap material (e.g., GaAs) and secondly lower the threshold current to near 1 mA. By evolving the

gated laser technology from InGaAs ($\lambda = 1.55$ μm) to a wider bandgap material, one benefits by the increase of the threshold temperature coefficient (T_o) by nearly a factor of two for InGaAsP ($\lambda = 1.3$ μm) and between a factor of 3 to 4 for GaAs [50] ($\lambda = 0.85$ μm) and higher factors for InGaAlP ($\lambda = 0.67$ μm) lasers. Note that the threshold current varies exponentially with temperature (phenomenological relationalship) as,

$$I_{th} \propto \exp(T/T_0)$$

and with such a large increase in T_0, the temperature stability should become greater. This stability will also improve by decreasing the magnitude of the threshold current. When biased at a dc current (see Fig. 11), as the threshold current decreases, the operating window is widened and the device can tolerate a larger variation in temperature. Only the contrast ratio is affected. Further temperature stability can be attained if the dc anode current is increased with increasing temperature (e.g., using a negative temperature coefficient resistor bias scheme). Attainment of a 70°C safe operating temperature window obviously relaxes the thermal management constraints.

These dc characteristics outlined above, suggest that the LGA optical interconnection technology will be a superior product with respect to performance and cost.

2) *AC Device Modeling and Measurements:* The model [46] for the gated laser structure, begins with the single mode rate equations for a laser structure of cavity length L_c, absorber section length L_s, stripe width W and active region thickness d (see Fig. 10).

$$\frac{dn}{dt} = \frac{I_a}{q\nu} - R \quad \text{and} \tag{1}$$

$$\frac{dp}{dt} = \left[\Gamma G - \frac{1}{\tau_p}\right]p + R_{sp} \tag{2}$$

where n: carrier density, I_a: anode current, ν: active region volume ($= (L_c - L_s)Wd$), R: charge carrier recombination rate, p: photon density, Γ: confinement factor, G: gain function, τ_p: photon lifetime $\left(= \left(\frac{c}{n_g}\alpha\right)^{-1}\right)$, n_g: group velocity index of refraction, α: lumped absorption (loss) per unit length of the laser, R_{sp}: spontaneous emission coupling into the laser, mode ($= \beta\Gamma B_r n^2$), and β: spontaneous emission factor, B_r: radiative recombination parameter.

For a conventional semiconductor laser, one generally assumes that the carrier and photon lifetimes are relatively constant. These assumptions are not valid with multielectrode lasers, where large changes in the carrier density occur, and the cavity losses are determined in part by a voltage controlled loss modulator. Hence, the carrier recombination rate must be expressed in terms of carrier density,

$$R = (A_{nr} + B_r n + C_r n^2)n + Gp \tag{3}$$

where A_{nr}, B_r, and C_r describe the nonradiative and radiative recombination processes in the semiconductor [51]. The laser gain is assumed to be linearly proportional to the carrier density, but is subject to gain saturation [52], $G = G_o(n - n_o)(1 - \epsilon p)$ where G_o is the gain constant, n_o is the carrier density required for transparency and $\epsilon(\ll 1.0)$ is a gain saturation parameter.

To describe absorption in the laser, one must include both the saturable and nonsaturable components (which are clearly apparent from Fig. 11). In the above-threshold regime, an increase in the applied gate (absorber) voltage, shifts the light-current curve to higher anode (gain section) currents. In addition, the laser displays an absorption saturation (with respect to anode current) at a threshold which is also a function of the applied gate voltage. We treat these saturable and nonsaturable elements as additive terms and write the total absorption for the laser as

$$\alpha = \alpha_o + \alpha_m + \alpha_s(V) + \alpha_{ns}(V)$$

where α_o and α_m are, respectively, the internal and mirror losses which determine laser threshold and differential quantum efficiency when the absorber section is forward biased, $\alpha_s(V)$ is the saturable loss and $\alpha_{ns}(V)$ is the nonsaturable loss. The nonsaturable loss term functions as a voltage dependent photon extractor which varies approximately linearly as $\alpha_{ns}(V) = \alpha_1 + \alpha_2 V$. The saturable term varies as $\alpha_s(V) = \alpha_{so}/(1+I/I_{so}(1+\kappa V))$, where I_{so} is the voltage independent saturation intensity.

The total losses within the laser may now be conveniently expressed as

$$\frac{1}{\tau_p} = \frac{1}{\tau_p}\left[1 + r_1 + r_2 V + \frac{r_3}{(1 + I/I_{so}(1 + \kappa V))}\right] \tag{4}$$

where

$$\frac{1}{\tau_p} = \frac{c}{n_g}(\alpha_o + \alpha_m), r_1 = \frac{\alpha_1}{(\alpha_o + \alpha_m)}, r_2 = \frac{\alpha_2}{\alpha_o + \alpha_m},$$

and

$$r_3 = \frac{\alpha_{so}}{(\alpha_o + \alpha_m)}.$$

Equations (1) through (4) constitute a model for this device. Parameters α_o and α_m are obtained from the static light current curve of the device with the absorber section forward biased, and r_1, r_2, and r_3 are treated as fitting parameters for devices of differing absorber and gain dimensions. Fig. 14 shows a simulated set of light current curves obtained by integrating the rate equations, using the parameter values shown in the figure, and using a set of gate voltages in the range $0\ V < V_g < 1\ V$. The laser structure was a buried heterostructure graded-index separate confinement heterostructure (BH-GRINSCH) with four quantum wells and a cavity length $L_c = 500$ μm. The model accurately reproduces both the observed Q-switched response and the variation in laser threshold with the applied gate voltage. It is clear that saturable absorption plays a crucial role in determining the light-current characteristic of the device.

The model for this laser structure, can also be used to explore the ac behavior of the device. In Fig. 15(a) an experimentally observed ac response of a gated laser is shown when a 1 Gb/s nonreturn to zero (NRZ) fixed data pattern is applied to the gate (with $I_a = 60$ mA). The laser output is modulated by the applied gate voltage with approximately 100 ps turn-on delay. The output response contains laser relaxation oscillations which ride upon the pulse stream. In Fig. 15(b) the simulated ac response (using the same parameter values used

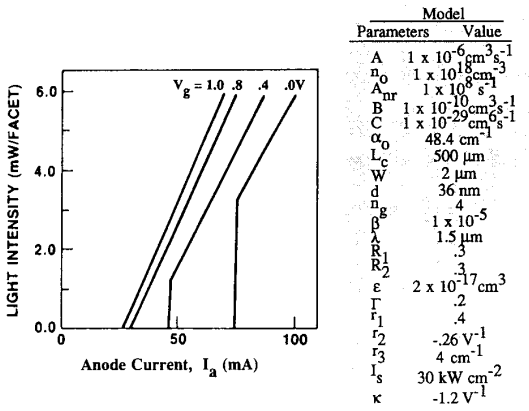

Model Parameters	Value
A	1×10^{-6} cm^{3}s^{-1}
n_o	1×10^{18}cm^{-3}
A_{nr}	1×10^{8} s^{-1}
B	1×10^{-10}cm^{3}s^{-1}
C	1×10^{-29}cm^{6}s^{-1}
α_o	48.4 cm^{-1}
L_c	500 μm
W	2 μm
d	36 nm
n_g	4
β	1×10^{-5}
λ	1.5 μm
R_1	.3
R_2	.3
ε	2×10^{-17}cm^3
Γ	.2
r_1	.4
r_2	-.26 V^{-1}
r_3	4 cm^{-1}
I_s	30 kW cm^{-2}
κ	-1.2 V^{-1}

Fig. 14. Simulated light intensity per facet as a function of anode current I_a, for various gate voltages.

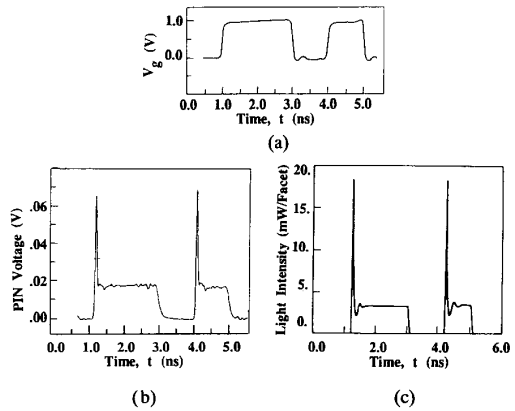

Fig. 15. Time domain response of an InP BH-GRINSCH gated laser, (a) fixed input NRZ pattern, (b) measured output response, and (c) simulated output response.

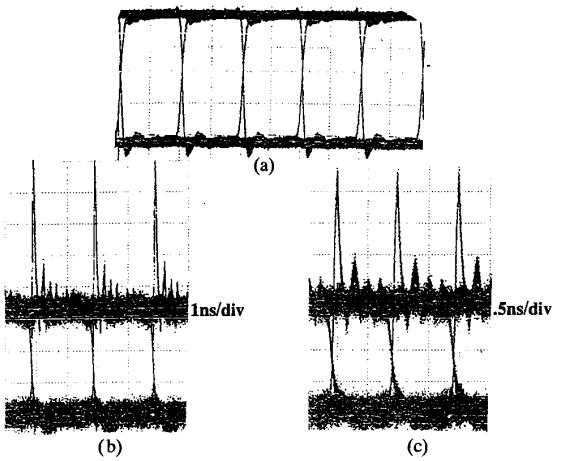

Fig. 16. Output eye diagrams from an input pseudo random NRZ data stream, (a) input, (b) measured output at 500 Mb/s, and (c) measured output at 1 Gb/s.

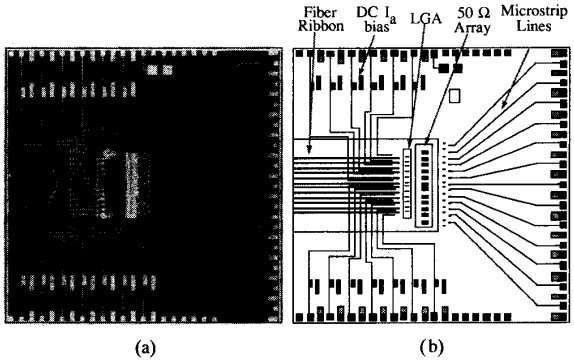

Fig. 17. Packaged laser gate array, (a) photograph and (b) schematic.

in Fig. 14(b)) to the same NRZ pattern is shown. Clearly, the modeled results are in good agreement with the experimental results. A 2^{20}-1 pseudo-random NRZ data pattern was applied to the LGA with the results at 500 Mb/s and 1 Gb/s shown in Fig. 16. A bit error rate (BER) of less than 10^{-14} was obtained.

3) *Packaging:* Fig. 17 shows a photograph of a packaged array (a) and a schematic of the packaged array (b). The currently implemented laser gate array consists of 12 individually addressable lasers monolithically integrated on a InP chip and packaged on a silicon hybrid. The hybrid has a dimension of 2.5 cm× 2.5 cm. High-speed microstrip transmission lines were fabricated using a polymer as a dielectric to obtain a 50-Ω impedance. The microstrips were then properly terminated to a 50-Ω resistance (wirebonded to a ceramic resistor array mounted on the hybrid). The termination resistor was also wirebonded to the gate of the LGA. Optical alignment (both lateral and height) of the multimode fiber ribbon (62.5-μm graded-index core) was obtained by fabricating v-grooves in the silicon hybrid, and by recessing the LGA die into the hybrid (accomplished with a deep chemical etch). The fiber ribbon was then simply butt-coupled to the laser, and epoxyed to the hybrid. A copper molybdenum heat spreader was attached to the back side of the hybrid for mechanical rigidity and to lower the package thermal resistance. The hybrid assembly was pressure contacted to a PCB, as shown in Fig. 18(a) the assembled view and (b) the dissembled view. Gate contacts on the hybrid connects to PCB microstrips which terminate to SMA connectors for external equipment modulation. Individual anode contacts (for DC current bias) come off the hybrid and terminate on a PCB connector for external current source connection. A TEC was thermally connected to the heat spreader as shown in Fig. 18.

The evolution of the device packaging is currently being planned. Emphasis in manufacturing, ease of assembly, and increased array size are expected to reduce the per interconnection cost. In the next version of the transmitter module, the LGA die will be bump bonded onto a silicon submount substrate with v-groove alignment features. This provides two accomplishments; better thermal contact to the substrate, and better alignment of the laser array outputs to the etched v-grooves. Multimode fiber ribbon [53] can then have its terminating ends chemically etched to form lenses to ease butt coupled alignment to the laser array. The ribbon will then

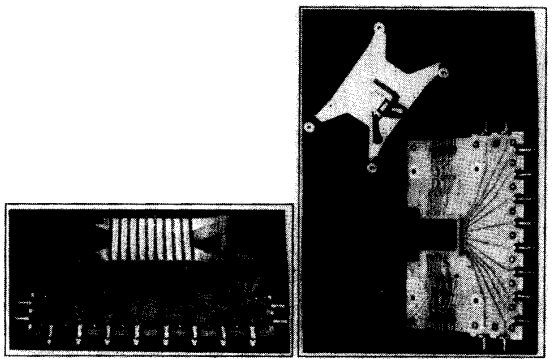

(a) (b)

Fig. 18. Test fixture photograph of the packaged laser gate array, (a) assembled view and (b) dis-assembled view.

be attached to the silicon submount. This sub-assembly will then be mounted on a PolyHIC (PHIC), which is a ceramic thin film multilayer (based on a polymer dielectric) hybrid. The high-speed interconnect and the dc biased connections on the silicon subassembly will then be wirebonded to the PHIC, and the fiber ribbon will be strain relieved to the PHIC. The 2.5 cm× 2.5 cm PHIC contains integrated 50 Ω resistor terminations, decoupling and bypass capacitors, high speed microstrip transmission line interconnections, power distribution resistors to dc bias the lasers, and a high-speed surface mountable lead frame that complies with JEDEC standards. The entire hybrid will then be coated with a clear RTV for environmental protection (note, not hermetically sealed).

In the longer term, we envision the transmitter module (or perhaps a transceiver module) being merged into the optical MAC, to form a smart optical connector. This would again reduce the cost per interconnection. The fiber ribbon pigtail from the hybrid to the optical MAC would be eliminated. The silicon submount and the silicon piece part in the MAC would be merged together, reducing the number of optical connections and alignments. Additionally the footprint required on a PCB for optical interconnection will be reduced even further with this scheme. Hence a circuit designer only interfaces to an ECL compatible smart optical connector.

V. Conclusions

Advanced switching (and computer) system architecture designs, require higher performance and lower cost packaging technologies to implement future new switch capabilities and enhanced services. Presently, electronic interconnection techniques (e.g., C4 and HPMCM's) offer superior performance, and at a lower cost for the chip-to-chip packaging level over that of alternative methods. Optical interconnection strategies offer advantages at higher levels of packaging (i.e., board-to-board, shelf-to-shelf and frame-to-frame interconnects). Furthermore, high data rate parallel optical data links can presently offer greater flexibility and lower cost than time multiplexed serial data links. This suggests that transmitter array technology research is clearly warranted.

There are numerous optical techniques for the implementation of a parallel data link, and the LGA solution has the potential of offering, in the near term, high performance at a low cost. The key concepts offered by LGA technology that designers can take advantage of include: no external driver required (i.e., ECL compatible), high optical output power (vastly simplifying the receiver circuit design), improved skew, lower total power dissipation, small PCB footprint, high temperature performance, and high contrast ratio.

Acknowledgment

The authors would like to thank and acknowledge the work done by the following people; D. B. Buchholz (thin film packaging), W. M. MacDonald (silicon submount packaging work), and R. F. Huisman (laser array measurements).

References

[1] A. F. J. Levi, R. N. Nottenburg, R. A. Nordin, T. Tanbun-Ek, and R. A. Logan, "Multielectrode quantum well laser for digital switching," *Appl. Phys. Lett.*, vol. 56, no. 12, Mar. 1990.
[2] K. Niwa, T. Araseki, and T. Nishitani, "Digital signal processing for video," *IEEE Circuits and Dev. Mag.*, vol. 6, no. 1, pp. 27–33, Jan. 1990.
[3] G. W. Neudeck, "Three dimensional CMOS integration," *IEEE Circuits and Dev. Mag.*, vol. 6, no. 5, pp. 32–38, Sept. 1990.
[4] R. P. Zingg, B. Hofflinger, and G. W. Neudeck, "High-quality stacked CMOS inverter," *IEEE Electron. Dev. Lett.*, vol. 11, no. 1, Jan. 1990.
[5] M. Shur, *GaAs Devices and Circuits.* New York: Plenum, 1987.
[6] A. R. Alvarez, *BiCMOS Technology and Applications.* The Netherlands: Kluwer Academic, 1989.
[7] J. D. Gallia *et al.*, "High performance BiCMOS 100k-gate array," *IEEE J. S. S. Circuits,* vol. 25, no. 1, Feb. 1990.
[8] R. R. Tummala and E. J. Rymaszewski, *Microelectronics Packaging Handbook.* New York: Van Nostrand Reinhold, 1989.
[9] S. M. Wentworth, D. P. Neikirk, and C. R. Brahce, "The high frequency characteristics of tape automated bonding (TAB) interconnects," *IEEE CHMT,* vol. 12, no. 3, Sept. 1989.
[10] K. Atsumi *et al.*, "Inner lead bonding technique for 500 lead dies having a 90 um lead pitch," *IEEE Trans. Comp., Hybrids, Manuf., Technol.*, vol. 13, no. 1, Mar. 1990.
[11] R. F. W. Pease, B. W. Langley, P. F. Marella, H. Blennemann, and O. K. Kwon, "High density interconnection in multichip packaging," Private Communication by the Semiconductor Research Corp., Apr. 1989.
[12] L. S. Goldmann and P. A. Totta, "Area array solder interconnections for VLSI," *Solid State Tech.*, June 1983.
[13] J. P. Dietz, "East fishkill advanced technology in new 9370 series," *IBM Press Release,* Oct. 1986.
[14] W. Weston, "High density 128 × 128 area arrays of vertical electrical interconnections," *Fourth Annual Microelectronic Interconnection Conf.,* July 1985.
[15] K. Gilleo, "Direct chip interconnect using polymer bonding," *IEEE Trans. Comp., Hybrids, Manuf., Technol.*, vol. 13, no. 1, Mar. 1990.
[16] K. Hatada, H. Fujimoto, T. Ochi, and Y. Ishida, "LED array modules by new technology microbump bonding method," *IEEE Trans Comp., Hybrids, Manuf. Technol.*, vol. 13, no. 3, Sept. 1990.
[17] C. A. Neugebauer and R. O. Carlson, "Comparison of wafer scale integration with VLSI packaging approaches," *IEEE Trans. Comp., Hybrids, Manuf. Technol.*, vol. 10, no. 2, June 1987.
[18] J. F. McDonald, H. T. Lin, H. J. Greub, R. A. Philhower, and S. Dabral, "Techniques for fabrication of wafer scale interconnections in multichip packages," *IEEE Trans. Comp., Hybrids, Manuf. Technol.*, vol. 12, no. 2, June 1989.
[19] J. L. Brandner *et al.*, "Electrical characterization of POLYHIC, a high density, high frequency, interconnection and packaging medium for digital circuits," in *Proc. 39th Electron. Components Conf.,* May 1989, pp. 759–764.
[20] R. R. Johnson, "Multichip modules: next generation packages," *IEEE Spectrum,* Mar. 1990.
[21] K. K. Chakravorty, C. P. Chien, J. M. Cech, M. H. Tanielian, and P. L. Young, "High density interconnection using photosensitive polyimide

and electroplate copper conductor lines," *IEEE Trans. Comp., Hybrids, Manuf. Technol., CHMT* vol. 13, no. 1, Mar. 1990.
[22] D. R. Dykaar, A. F. J. Levi, and M. Anzlowar, "Ultrafast coplanar air transmission lines," *Appl. Phys. Lett.,* vol. 57, p. 1123, 1990.
[23] M. R. Feldman, S. C. Esener, C. C. Guest, and S. H. Lee, "Comparison between optical and electrical interconnects based on power and speed considerations," *Appl. Opt.,* vol. 27, no. 9, May 1988.
[24] H. H. Arsenault, T. Szoplik, and B. Macukow, *Optical Processing and Computing.* New York: Academic, 1989, pp. 1–31.
[25] D. A. B. Miller, "Optics for low-energy communication inside digital processors: Quantum detectors, sources, and modulators as efficient impedance converters," *Opt. Lett.,* vol. 14, no. 2, Jan. 1989.
[26] L. L. Moresco, "Electronic system packaging: The search for manufacturing the optimum in a sea of constraints," *IEEE Trans. Comp. Hybrid. Manuf. Technol.,* vol. 13, no. 3, Sept. 1990.
[27] L. L. Blyler and G. J. Grimes, "Polymeric optical mixing rod coupler," *IEEE Trans. Comp., Hybrids, and Manuf. Technol.,* vol. 13, no. 1, Mar. 1990.
[28] J. Xu, R. D. Driver, and W. F. Gaslo, "Low cost high quality plastic fiber couplers," *Proc. SPIE,* vol. 988, pp. 51–55, 1988.
[29] R. E. Weiss, "A family of connectors for circuit pack to backplane optical interconnection," in *Proc. 9th Internat. Electron. Packaging Conf.,* Sept. 1989, vol. 2, pp. 1033–1042.
[30] G. J. Grimes and L. J. Haas, "An optical backplane for high performance switches," *Proc. Internat. Switching Symp.,* May 1990, vol. 1, pp. 85–89.
[31] M. K. Kilcoyne, S. Beccue, K. D. Pedrotti, R. Asatourian, and R. Anderson, "Optoelectronic integrated circuits for high speed signal processing," *Opt. Eng.,* vol. 25, no. 10, Oct. 1986.
[32] K. Kaede *et al.,* "12-channel parallel optical fiber transmission using a low drive current 1.3 μm LED array and a PiN PD Array," *J. Lightwave Technol.,* vol. 8, no. 6, June 1990.
[33] G. W. Taylor and P. Cooke, "Double-heterostructure optoelectronic switch as a single quantum well laser," *Appl. Phys. Lett.,* vol. 56, no. 14, Apr. 1990.
[34] H. Z. Chen, J. Paslaski, A. Yariv, and H. Morkoc, "High frequency modulation of AlGaAs/GaAs lasers grown on Si substrate by molecular beam epitaxy," *Appl. Phys. Lett.,* vol. 52, no. 8, Feb. 1988.
[35] K. Y. Lau, N. Bar-Chaim, P. L. Derry, and A. Yariv, "High speed digital modulation of ultralow threshold ($<$ 1 mA) GaAs single quantum well lasers without bias," *Appl. Phys. Lett.,* vol. 51, no. 2, July 1987.
[36] F. R. Nash and R. L. Hartman, "Accelerated facet erosion formation and degradation of (Al,Ga)As double-heterostructure lasers," *IEEE J. Quantum Electron.,* vol. QE-16, no. 10, Oct. 1980, pp. 1022–33.
[37] L. A. Koszi, B. P. Segner, H. Temkin, W. C. Dautremont-Smith, and D. T. C. Huo, "1.5 μm InP/GaInAsP linear laser array with twelve individually addressable elements," *Electron. Lett.,* vol. 24, no. 4, Feb. 1988.
[38] E. Marclay *et al.,* "Scaling of GaAs/AlGaAs laser diodes for submilliampere threshold current," *Electron. Lett.,* vol. 25, no. 14, July 1989.
[39] M. Nakao *et al.,* "1.55 μm DFB laser array with $\lambda/4$ shifted first order gratings fabricated by X-ray lithography," *Electron. Lett.,* vol. 25, no. 2, Jan. 1989.
[40] D. Yap, Z. L. Liau, D. Z. Tsang, and J. N. Walpole, "High performance InGaAsP/InP buried heterostructure lasers and arrays defined by ion beam assisted etching," *Appl. Phys. Lett.,* vol. 52, no. 18, May 1988.
[41] T. R. Chen, K. Utaka, Y. H. Zhuang, Y. Y. Liu, and A. Yariv, "Vertical integration of an InGaAsP/InP heterojunction bipolar transistor and a double heterostructure laser," *Appl. Phys. Lett.,* vol. 50, no. 14, Apr. 1987.
[42] M. Makiuchi *et al.* "GaInAs PIN photodiode/GaAs preamplifier photoreceiver for gigabit-rate communication systems using flip-chip bonding techniques," *Electron. Lett.*
[43] H. Ishikawa, H. Soda, T. Watanabe, H. Sudo, and K. Sato, "Low-threshold current and high-efficiency operation of electro-absorption modulator/DFB laser light source with AR-HR coating," *12th IEEE Internat. Semiconduct. Laser Conf.*
[44] A. L. Lentine *et al.,* "Symmetric self-electro-optic effect device: Optical set-reset latch," *Appl. Phys. Lett.,* vol. 52, no. 17, Apr. 1988.
[45] A. L. Lentine *et al.,* "Optical logic using electrically connected quantum well PIN diode modulators and detectors," *Appl. Opt.,* vol. 29, no. 14, May 1990.
[46] J. O'Gorman, A. F. J. Levi, R. N. Nottenburg, T. Tanbun-Ek, and R. A. Logan, "Dynamic and static response of multielectrode lasers," *Appl. Phys. Lett.,* vol. 57, p. 968, 1990.
[47] J. O'Gorman, A. F. J. Levi, T. Tanbun-Ek, and R. A. Logan, Asymmetric line broadening in intracavity loss modulated quantum well DFB lasers," *Appl. Phys. Lett.*
[48] J. O'Gorman and A. F. J. Levi, "Calculated chirp in intracavity loss modulated DFB lasers," *Electron Lett,* 1990.
[49] T. Tanbun-Ek *et al.,* "Very low threshold InGaAs/InGaAsP graded index separate confinement heterostructure quantum well lasers grown by atmospheric pressure metalorganic vapor phase epitaxy," *Appl. Phys. Lett.,* vol. 55, no. 22, p. 2283, Nov. 1989.
[50] P. Blood, S. Colak, and A. I. Kucharska, "Temperature dependence of threshold current in GaAs/AlGaAs quantum well lasers," *Appl. Phys. Lett.,* vol. 52, no. 8, Feb. 1988.
[51] G. P. Agrawal and N. K. Dutta, *Long Wavelength Semiconductor Lasers.* New York, Van Nostrand Reinhold, 1986.
[52] D. C. Channin, *J. Appl. Phys.,* vol. 50, p. 3858, 1979.
[53] K. W. Jackson *et al.,* "An enhanced ribbon structure for high fiber count cables in the loop," *Internat. Wire and Cable Symp. Proc.,* 1989.

Ronald A. Nordin (S'75–M'77–SM'87), photograph and biography not available at the time of publication.

A. F. J. Levi, photograph and biography not available at the time of publication.

R. N. Nottenburg (S'79–M'80) photograph and biography not available at the time of publication.

J. O'Gorman, photograph and biography not available at the time of publication.

T. Tanbun-Ek, photograph and biography not available at the time of publication.

Ralph A. Logan (M'73–SM'73–F'86), photograph and biography not available at the time of publication.