High-Speed Optical Interconnections for Digital Systems

D.Z. Tsang

■ As an alternative to electrical interconnections, optical interconnections are being investigated for digital systems that require high speed, density, and connectivity. In particular, diode-laser-based optical interconnections with high optoelectronic conversion efficiencies promise simple interconnections that can operate even at gigahertz rates. Optical modules in which critical alignments have been performed during module assembly can reduce module placement requirements from micrometers to millimeters. Such modules can also eliminate the need for micropositioners in digital systems. This article describes the basic concepts of optical interconnections, the progress on experimental diode-laser-based free-space optical interconnections, and the future prospects for this technology.

SEVERAL TRENDS IN COMPUTING—the use of multiprocessor and parallel-processing systems (for example, see Reference 1), and the drive for higher operating speeds and larger as well as denser very large-scale integrated (VLSI) circuits—are pushing the development of advanced interconnection technology. Multiprocessor computers tend to require both a larger number and a higher density of interconnections with a higher degree of connectivity than single-processor machines. The combination of faster operating speeds, longer interconnection distances, and higher densities makes maintaining signal integrity difficult because of parasitic reactances, impedance mismatches, crosstalk, dispersion, and frequency-dependent skin-effect losses. Indeed, VLSI-circuit designers have become increasingly concerned with the pin-out problem, in which the number of input/output pins and pads is limited by the density of wire bonds as chips get larger and denser.

It is important to note that the clock rate of high-speed digital systems is limited by packaging and electrical-interconnection technology. For example, although the fastest transistors have switching times on the order of 10 psec, package parasitic effects limit the fastest commercial integrated circuits (gallium-arsenide logic) to risetimes and falltimes on the order of 100 psec. Signal propagation within circuit boards results in additional deterioration, and the electrical backplane, into which boards are inserted, can only transfer risetimes of approximately 1 nsec between boards. It should be pointed out that high-speed electrical backplanes often need to be custom engineered with computerized circuit-modeling programs—an expensive process for small production quantities. All told, the clock cycle times for a typical high-speed digital system are on the order of 10 nsec after adequate margins for timing errors and stabilization of logic levels are included.

Seeking an alternative to electrical interconnections, researchers have experimented with optical links for connecting mainframes, modules, boards [2], chips [3–5], and even points within a chip [6]. Optical interconnections can be viewed both as a technology to be used in conjunction with electrical interconnections and electronic logic circuits, and as an enabling technology for optical computers. The various concepts for optical computers share the key feature that optical waves carry information throughout the computer, even down to the gate level via optical interconnections. Optical interconnections are expected to provide advantages in those situations which require high speeds, high interconnection densities, high connectivities, and long distances. This article describes the basic characteristics and advantages of optical links, the requirements and current status of diode lasers for efficient optical transmission, an experimental 1-Gb/sec board-to-board optical interconnection, and the prospective performance and role of optical interconnections.

Optical Interconnections

Recent interest in the use of optical interconnections was stimulated by a 1984 paper by J.W. Goodman et al. [3], which contained ideas and considerations for the incorporation of optics into VLSI systems. Various implementations of optical interconnections have been subsequently proposed and studied. Figure 1 contains some of the various possibilities, which can all be partitioned into a source, a channel, and a receiver.

As an example of a source, a semiconductor laser or light-emitting diode can be connected to a digital integrated circuit either directly or through a special driver, and the source can be amplitude modulated. (The semiconductor laser source will be extensively described later in this article.) Optical modulators can also be directly connected to the output of a digital gate. A modulator requires an optical input from a diode laser or from a powerful laser beam that is split into multiple channels, one for each modulator. The modulator could be electrooptic, absorptive, or reflective, and an array of modulators such as spatial light modulators can accommodate one- and two-dimensional arrays of optical channels.

The optical channel(s) can consist of free space with lenses or holograms acting to direct the light. The channel might also be a guided-wave channel such as a single

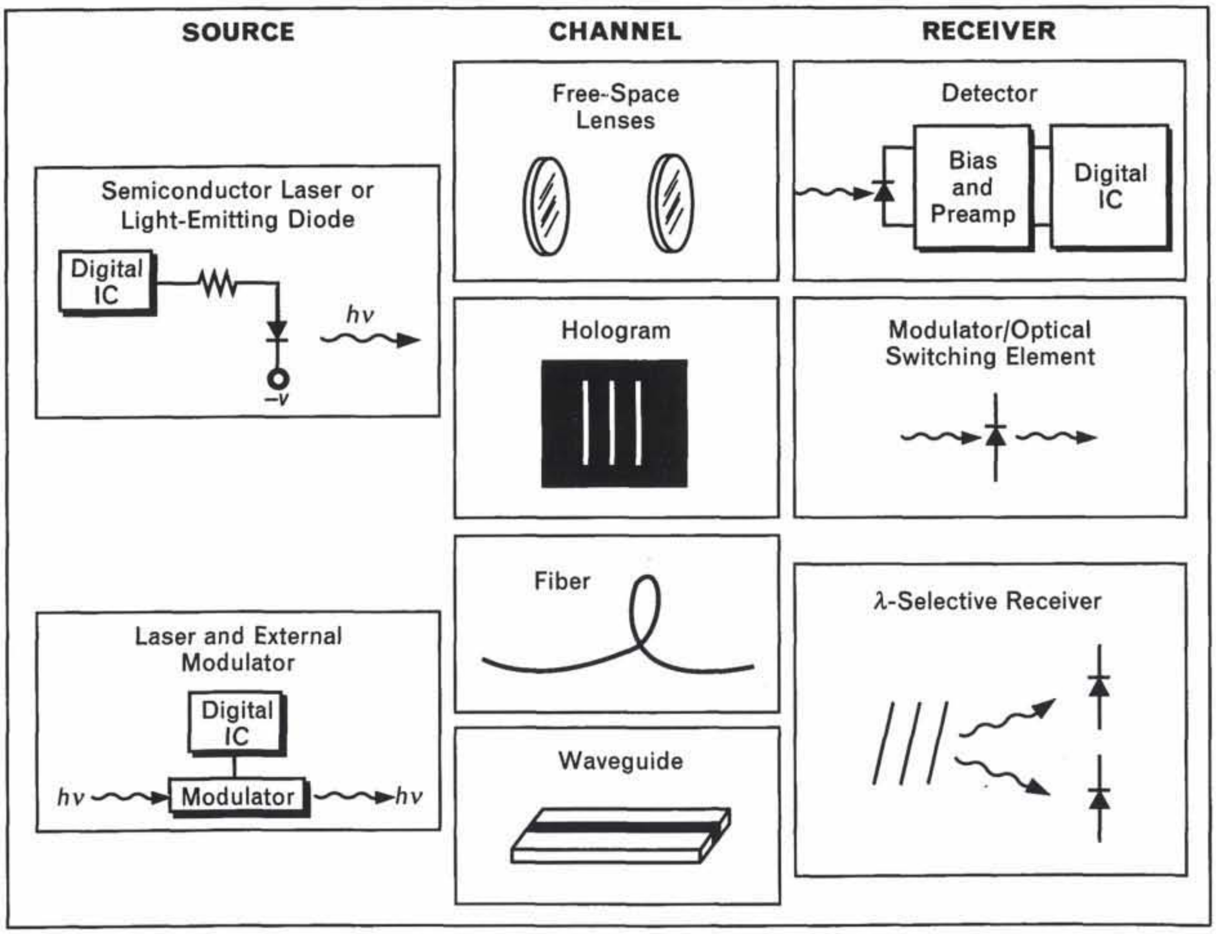

FIGURE 1. Optical interconnections can be partitioned into a source, a channel, and a receiver. Some of the various possibilities are shown.

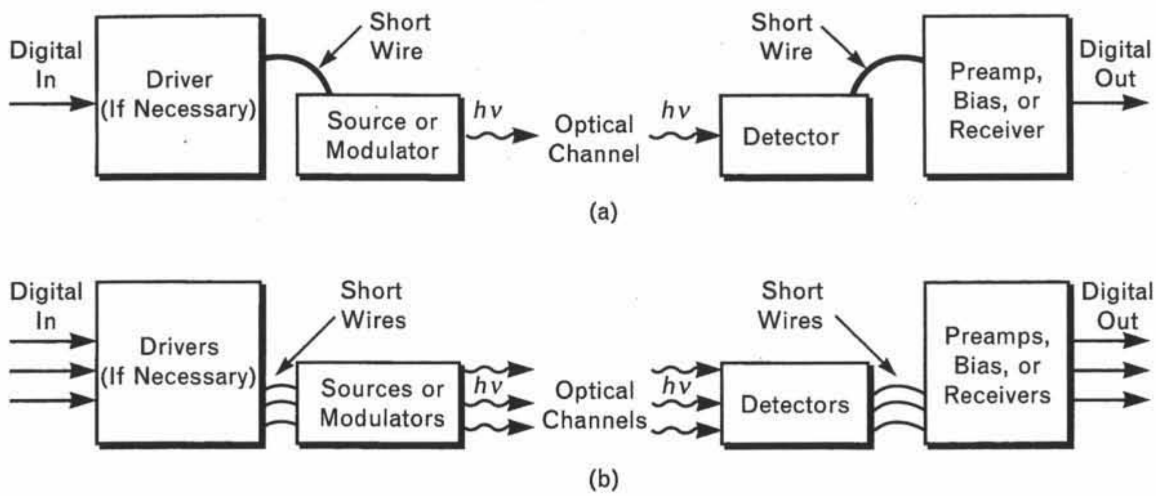

FIGURE 2. Optical interconnection systems: (a) single channel and (b) multiple channel. With optical interconnection systems, high-speed interconnections can be made with short electrical lines while the optical link maintains the integrity of the signal. When multiple channels are considered, signal channels can cross one another, optical imaging techniques can be applied, and holographic and reconfigurable holographic channels can be considered.

or multimode optical fiber, or it could be a waveguide such as polymer waveguides that can be photolithographically defined on a circuit board. The latter approach was pursued by C.T. Sullivan [7], and D.H. Hartman et al. [8].

The receiver is generally a photodiode or photoconductor connected to bias and preamplifier circuitry. Variations include the use of a modulator as a combined detector and switching element such as the self-electrooptic-effect devices (SEED) studied at AT&T, and the inclusion of a wavelength-selective filter in front of the detector as is under consideration at IBM.

We can take combinations of a source, a channel, and a receiver from each of the respective columns in Figure 1 and form the various optical interconnections that different groups have pursued. The central idea behind all the combinations is to locate the optical source physically close to the driver so that all electrical interconnections are short and a simple wire can be used between the driver and source. The signal is converted to light before significant signal degradation occurs. For a single optical channel, as shown in Figure 2(a), this design results in a number of advantages:

- Impedance matching, reflections, and ground planes are less of a problem than with high-speed electrical interconnections.
- Any signal loss is independent of the interconnection length and the dispersion is negligible for distances of interest.
- The optical channel itself has low crosstalk, although stray light at the detector can be a concern.
- With free-space interconnections, no physical contact is needed.
- The channel has a high bandwidth.

On the receiver side, the detector is again mounted very close to the receiver electronics so that only a short simple wire is needed. The electrical part of the interconnection is so short that it does not limit the interconnect in speed and the end-to-end response is independent of the distance of the interconnection. Disadvantages of optical interconnections include coupling and alignment problems, the cost of sources and detectors, and the loss of signal due to limited optoelectronic conversion efficiencies.

Multiple optical interconnections, as shown in Figure 2(b), also result in a number of advantages:

- With free-space channels, beams can cross with no interference and several channels can coexist in the same space. (With guided-wave optics, waveguides can also cross at an angle with little penalty.)
- The number of interconnections possible with holographic optical channels has been calculated to

be as high as the volume divided by the wavelength cubed, but limitations in the holographic process and difficulties in physical implementation will probably make such a high density very difficult to achieve.

- For arrays of highly parallel interconnections, two-dimensional imaging techniques can be applied.

It is interesting to note that computer architectures uniquely suited to optical implementations are possible. In fact, many researchers are studying techniques for reconfigurable optical interconnections in which the computer architecture can be optimized for specific applications.

Can Optical Telecommunications Technology Be Used for Interconnections?

Fiber optic telecommunications technology has been used for optical interconnections in some applications. In a recent demonstration of commercial fiber optic technology in a multiprocessor computer, 1024 wires were multiplexed onto two 400-Mb/sec optical fiber links to form one leg of a hypercube architecture [9]. But 1-Gb/sec fiber optic transmitter modules currently cost $5000 and receivers cost another $5000—far too expensive for most computer applications, particularly at the board or chip level, where the number of interconnections is large.

Much of the cost and complexity in telecommunications systems result because the received signals are at low levels, which is a consequence of the loss that results when repeaters are spaced at maximum distances. With the short-distance and small-signal-loss characteristics of most interconnection applications, the cost and complexity of typical telecommunication systems can be substantially reduced, especially if very efficient lasers and detectors are used. Moreover, in digital systems relatively high-level optical signals at the detector are necessary to avoid high-gain, and thus electromagnetic interference (EMI)-susceptible, receiver circuits.

High-speed optical fiber communications systems designed for long-distance transmission typically have an overall electrical differential-current efficiency of 1 to 3% for short links. (The overall electrical differential-current efficiency is defined as the ratio of differential current out of the detector to differential current into the laser when both devices are biased at the operating point). Attenuation in a long fiber will reduce the efficiency even further. Consequently, when components designed for long-distance fiber systems are employed in a short-distance optical interconnection, the detector delivers only 1 to 3% of the laser current. An optical interconnection based on such devices ordinarily requires a laser driver, preamplifier, and possibly an automatic gain control, timing recovery, and other interface circuitry. The use of efficient optoelectronic components can eliminate much of the interface circuitry.

There are three fundamental contributions to the overall differential-current efficiency of a link. For a laser biased above threshold, the efficiency is proportional to the product of the laser differential quantum efficiency (the fraction of electrons injected into the laser that are converted to photons), the detector quantum efficiency (the fraction of photons that are converted into electrons), and the optical transfer efficiency (the fraction of the photons emitted by the laser that are incident on the detector).

Because the electrical power out of the detector is proportional to the square of the detector current, the electrical output power varies as the square of the quantum efficiencies. Thus efficient devices are very important. For example, if we consider a laser and detector with differential efficiency values of 20% and 50%, respectively, an optical link with matched source and load impedances would already have an insertion loss of 20 dB. Since a digital level must be recovered from the optical signal, 20 dB of amplification is required even with an optical efficiency of unity. The need to maintain high optical transfer efficiency dictates the use of lenses or holograms for free-space board-to-board applications to ensure that a large fraction of the emitted photons arrives on the desired detectors. Because the emission cone from the micron-size aperture of most diode lasers results in large beam spreading, or diffraction, the fraction of photons collected by a reasonable-size detector without intervening optics is small for interconnection distances of a few centimeters. Thus an optical interconnect cannot afford to use inefficient lasers and detectors or to spray the board with photons unless high-gain amplifiers in front of each digital input are used—a practice that is best avoided in a digital environment with nearby logic gates switching on and off.

Diode Lasers for Efficient Optical Interconnections

High-efficiency optical links require efficient lasers, optics, and detectors. In addition, other considerations are important for the laser. For interconnection applications, if the threshold of the diode laser is sufficiently low, high-speed operation (~100-psec turn-on delay, or risetime, and multigigahertz bandwidth) is possible with the current supplied directly from a digital circuit (~50 mA). Lasers with low threshold (<10 mA) also eliminate the laser-driver circuitry that would otherwise be required for higher-threshold devices. It should be noted that the differential quantum efficiency of most high-speed lasers is only about 20% per laser facet at low current levels, and the efficiency can be even lower at the current levels required for high-speed operation. Furthermore, reliability at high speed is an issue, particularly at the >100-mA current levels often required by such operation. The temperature dependence of the laser could also be a concern, especially at elevated ambient temperatures. This section discusses the development of high-speed diode lasers for interconnection applications.

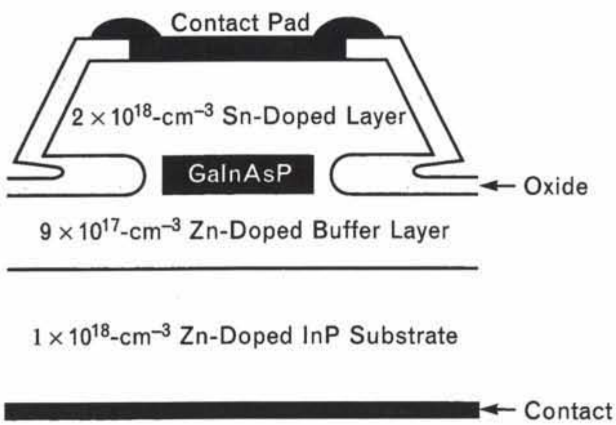

FIGURE 3. Cross section of p-substrate mass-transported laser. The contact pad is confined to the top of the laser mesa to minimize parasitic capacitance.

The 1.3-μm-wavelength mass-transported buried heterostructure laser fabricated from GaInAsP (Figure 3) had a low threshold, high speed, and high efficiency [10, 11]. In general, the 1.3-μm GaInAsP lasers have a longer lifetime than 0.8-μm AlGaAs lasers. To fabricate the 1.3-μm lasers, we used liquid phase epitaxy to grow wafers, starting with a 1×10^{18}-cm^{-3} Zn-doped InP substrate onto which a 9×10^{17}-cm^{-3} Zn-doped buffer layer was deposited. The next layer grown was a GaInAsP active layer that was nominally undoped and 0.2 μm thick. The final layer was Sn doped to 2×10^{18} cm^{-3}. The lasers were then fabricated in a fashion similar to the procedure described in Reference 10. The width of the transported region was about 0.2 μm with a nominal active-region width of 1.5 μm, and the laser wafer was thinned to a thickness of 75 μm. The ohmic contacts to the n-type and p-type regions were Au-Sn and Au-Zn, respectively. The contact metallization pads were confined to the top of the mesa (Figure 3) for the virtual elimination of parasitic capacitance between the contact pad and the substrate.

The best laser of this type without any facet coating has a CW (with a dc current) threshold of 4.5 mA [11], which is among the lowest thresholds reported for 1.3-μm lasers. Figure 4(a) shows the optical-power output versus current for another of these devices. The laser threshold without facet coatings is 8 mA and the differential quantum efficiency is 32% from each of the laser's two facets. The differential quantum efficiency decreases as the laser is driven with increased current well above threshold because a given increment of current well above threshold results in a lower output than for the same increment of current just above threshold. In an optical link, the reduction in quantum efficiency can result in increased incremental insertion loss as the laser is operated far above threshold for maximum speed.

For applications in which the output from only one facet is required, calculations for 1.3-μm buried-heterostructure lasers show that a high-reflectivity coating on one facet can produce differential quantum efficiencies greater than 60% with no increase in laser threshold. Link current efficiencies of 56% or higher might be possible in the future with such lasers and antireflection coatings on the optics and detector. Even higher efficiencies should be possible for other laser structures with lower internal loss.

Recent experimental results for the laser of Figure 4(a) have demonstrated such an increase in differential quantum efficiency. High reflectivity can be accomplished either with a reflective metal coating or a quarter-wave stack. The latter consists of dielectric layers of alternating low and high refractive indexes; each layer has a thickness equal to a quarter of a wavelength of light in the material.

Although a reflective metal coating is easier to deposit, the quarter-wave-stack approach was chosen because it forms a nonconductive coating, has low optical loss, and can have a reflectivity virtually equal to unity as the number of layers is increased.

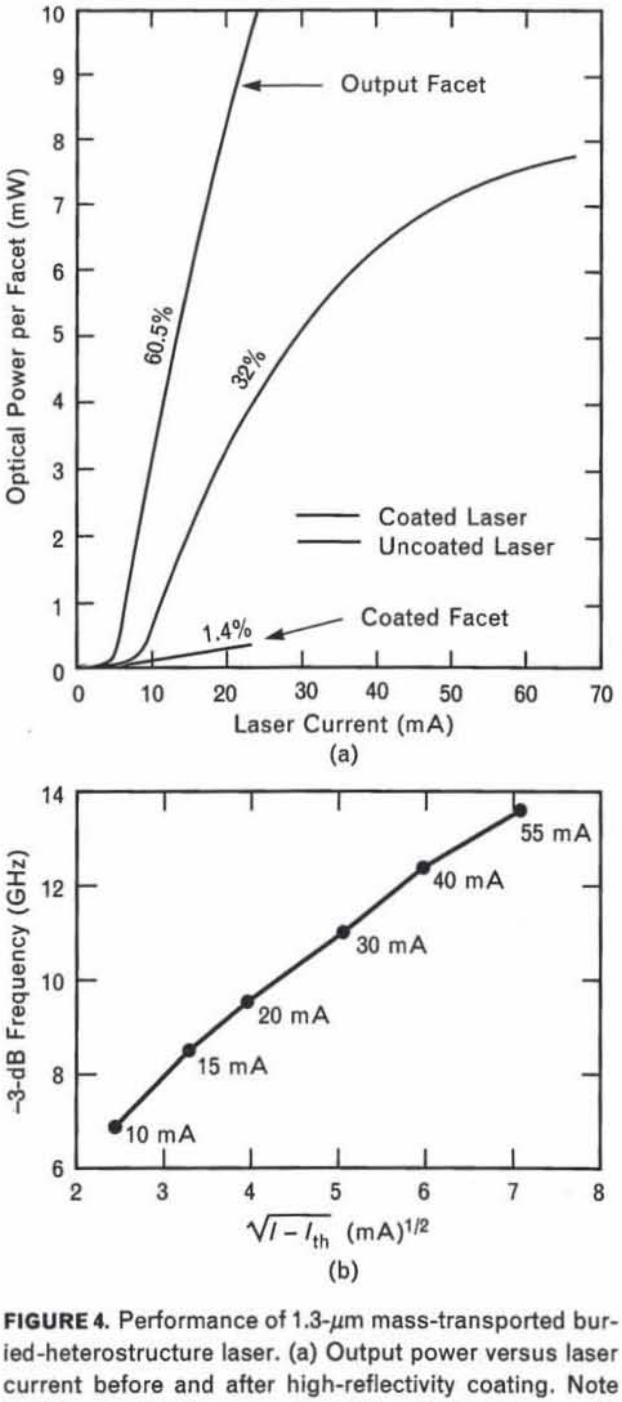

FIGURE 4. Performance of 1.3-μm mass-transported buried-heterostructure laser. (a) Output power versus laser current before and after high-reflectivity coating. Note that the coating improves the differential quantum efficiency to greater than 60%. (b) The –3-dB frequency for the laser of part *a* is more than 12 GHz with only 40 mA of current.

Using electron-beam evaporation, we deposited a quarter-wave dielectric stack (three pairs of SiO_2/Si layers) onto the back facet of a laser, which was removed and tested after the deposition of each pair of SiO_2/Si layers. After the first pair of coatings, the nominal reflectivity was 80%, the laser threshold had dropped to 5.5 mA, and the differential quantum efficiencies of the front and back facets were 44% and 18%, respectively. After two pairs of coatings, the nominal reflectivity was 96%, the laser threshold was 4 mA, and the efficiencies had changed to 58% and 6%. And, after the third and final pair of coatings, the nominal reflectivity of the back facet was 99%, the threshold relative to the uncoated laser had actually dropped by a factor of two to 3.8 mA, and the efficiencies had changed to 60.5% and 1.4%. Note that the sum of the differential quantum efficiencies from the two facets remained roughly constant, and the linearity of the laser improved markedly, ostensibly because the desired output was attained before nonlinear mechanisms became an issue. With three pairs of coatings this laser had 10 mW of output power with only 25 mA of current.

The –3-dB frequency of the laser of Figure 4(a) is about 10 GHz at 25-mA current levels (Figure 4[b]). With only about 40 mA, the –3-dB frequency increases to approximately 12 GHz. Commercial GaAs integrated circuits, which can provide drive currents of 60 to 70 mA, can easily supply the current levels of Figure 4(b). Mass-transported buried-heterostructure lasers of the type that is described in this article have operated at 16 GHz and at digital rates of 16 Gb/sec [11], but the total current level of about 100 mA required for those speeds exceeds the drive capability of most digital circuits.

Experimental 1-Gb/Sec Free-Space Optical Interconnection

Electrical interconnections are currently arranged in two-dimensional planes in chips, on boards, and in backplanes. For interboard communications, a signal from a chip on one board must be sent to the edge of the board, through a backplane, and then from the edge of the next board to the desired chip on that board. The backplane often constrains the speed of such interconnections. With an optical interconnection, a light beam could transmit the

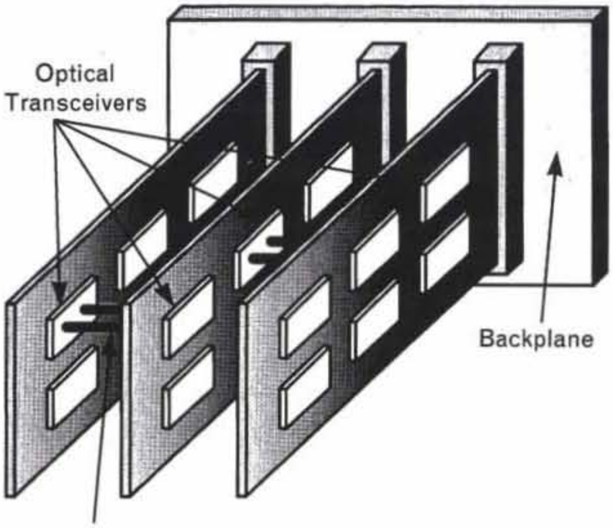

FIGURE 5. Free-space optical interconnections allow interconnections to be made in the third dimension in direct paths perpendicular (or at angles) to the boards.

signal to the next board without any physical contact between the boards, as shown in Figure 5. Optical interconnections can also be located anywhere on the board, not just on the edges. This flexibility literally adds a third dimension to the layout of logic systems.

At Lincoln Laboratory, we have concentrated on implementing an optical interconnection between digital gates without the use of a complicated fiber optic type of transmitter and receiver [12, 13]. Figure 6(a) is a schematic of the optical interconnection that was implemented by connecting a diode laser directly to a digital circuit on one board, and a detector directly to another circuit on another board. The circuits used were both GaAs integrated circuits with emitter-coupled-logic (ECL) compatible voltage levels. A diode laser was chosen over a light-emitting diode because of the former's inherent higher efficiency, narrower emission cone, and higher speed.

To test the interconnection, we connected the laser to a GaAs code generator that generated a fixed word sequence. The 1.3-μm GaInAsP laser used had a threshold of 5 mA and an external differential quantum efficiency of 35% per facet. (The facets of this laser were uncoated.) The low threshold of the laser allowed the logic gate to drive the interconnect directly, while the high differential efficiency enabled the detector to generate signal levels high enough so that the detector could be connected directly to another digital gate. An inexpensive aspheric 5-mm-diameter compact-disc lens with a numerical aperture of 0.55 collected and collimated the laser light. To demonstrate the capability of the interconnection, we sent the light over a 24-cm path, many times the distance required for a typical board-to-board interconnect. The light traveled to an identical lens that focused the light onto a commercial 100-μm-diameter GaInAs pin photodiode with a nominal 70% quantum efficiency. The output of the photodiode was connected directly to the input of another GaAs digital circuit—a D-type flip-flop that had an internal amplifier, or comparator (Figure 6[a]). The D-type flip-flop regenerated the digital output at rates from 100 Mb/sec up to 1 Gb/sec, as shown in Figure 6(b). The top trace of the figure shows both the output from the D-type flip-flop and the zero level when the optical beam is blocked. The lower trace shows the pattern of the electrical current into the laser. The limit of the GaAs code generator was 1 Gb/sec.

As stated earlier, the laser was connected directly to one digital circuit while the detector was connected directly to another digital circuit. No external laser drivers or preamplifiers were used. This experiment demonstrated that with efficient optoelectronic components, digital circuits can be interconnected directly. Indeed, efficiency allows a reduction in the complexity of optical interconnections. It should be noted that the experimental interconnection has the potential for a fan-out of two because the light from the other facet could also be collected and sent to another gate.

In a separate measurement involving a laser with 30%/facet efficiency, we determined the overall differential-electrical-current transfer efficiency of the interconnection to be 12.5%. This transfer efficiency was high compared to the 1 to 3% common in high-speed fiber optic systems. Even higher differential current efficiencies of 19% were measured with a 61%-efficiency laser.

The power consumption of the final stage in the GaAs code generator and the diode laser was 2 V multiplied by the diode laser current. For the 5-mA-threshold laser and the 1-Gb/sec interconnection, the estimated power consumption was 20 to 25 mW for a bit pattern with an equal probability of ones and zeroes. Although the resistor-network interface between the photodiode and the input to the GaAs D-type flip-flop helped to supply any input bias current to the gate in addition to providing an interface to the photodiode, the interface was probably not optimal. The calculated power dissipation in the receiver bias circuit was about 20 mW with up to 2 mA of photocurrent.

Micropositioner stages currently position the transmitter and receiver boards, and both the lenses and stages are mounted on an optical table. Because the setup is not practical for electronic systems, it is desirable to eliminate the optical-translation stages. Thus the question arises as to the kinds of performance and mechanical-alignment tolerances that will accompany such modifications. The following section addresses this issue in greater detail.

Prospective Performance of Optical Interconnections

Successes in the development of hybrid integration, uniform laser arrays, surface-emitting lasers, monolithic optoelectronic integrated circuits, microoptics, fiber-coupling techniques, and the heteroepitaxy of GaAs on Si will play a key role in the evolution of optical interconnections. Although these technologies are not reviewed here, prospects for optical interconnections can be examined from more basic considerations.

Power Dissipation

Several research groups have investigated the power dissipation of electrical interconnections in CMOS cir-

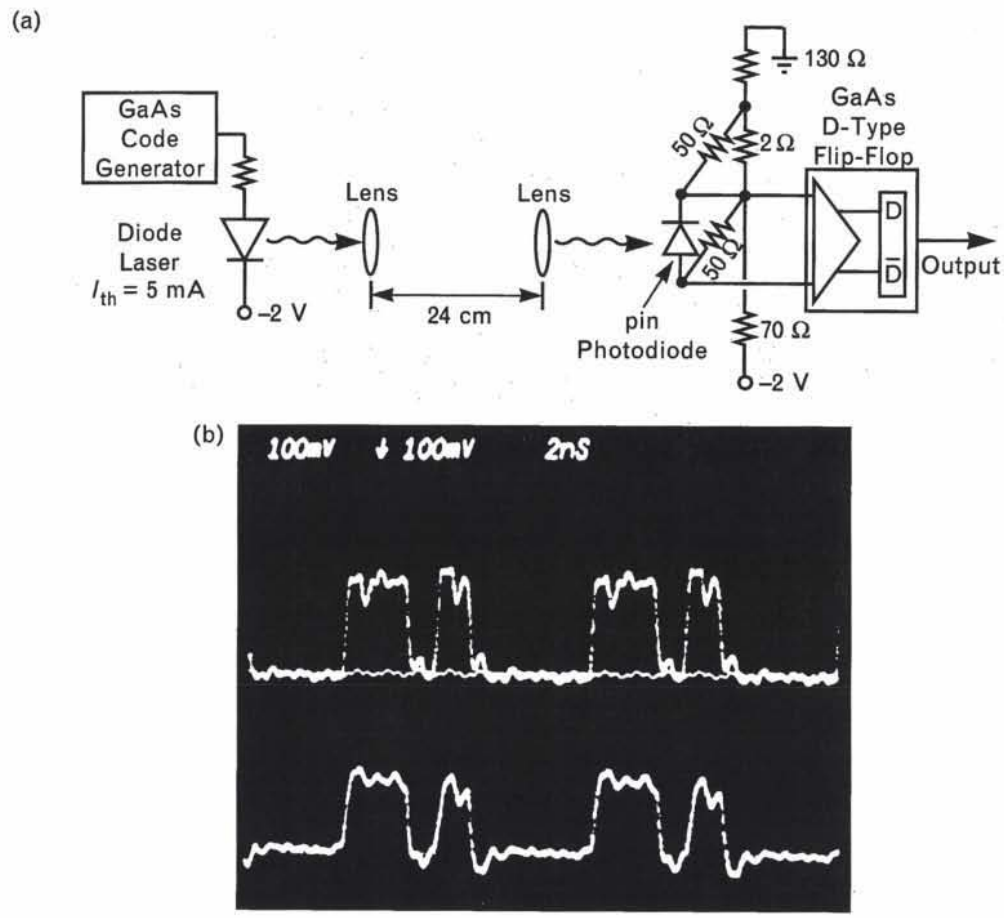

FIGURE 6. Experimental free-space optical interconnection. (a) High-efficiency optoelectronic components allow interconnections to be made directly to digital circuits without the extensive use of interface circuitry. (b) The simple interconnection system shown in part *a* operates at clock rates of 1 Gb/sec. The top trace shows the output of the D-type flip-flop and the zero level when the optical beam is blocked. The lower trace shows the desired digital pattern. The vertical scale is 1 V/division and the horizontal scale is 2 nsec/division.

cuits, and have compared the dissipation with that of optical interconnections. For clock rates of 1 GHz or faster and 3-μm CMOS design rules, the somewhat optimistic calculations by M.R. Feldman et al. [6] predict that diode-laser-based optical interconnections can dissipate less power than electrical wires at distances of 1 mm or longer. The comparison is based on a calculation of the energy required to charge the capacitance of an integrated-circuit trace (interconnection line) and gate. For a link efficiency of 45%, a similar calculation for external modulators indicates that optical interconnections will dissipate less power at distances of the order of 1 mm or longer for risetimes of 1 nsec or less.

In transmission-line electrical interconnections and ECL circuits, the power consumption for a 50-Ω transmission line and termination is 48 mW for a logic high and 8 mW for a logic low. If high and low levels are equally likely, the average consumption is 28 mW. A 1.3-μm diode-laser-based interconnection can have lower power consumption on the transmitter side. With the laser characteristics shown in Figure 4(b), only 10 mA is necessary for a 7-GHz response in which the average consumption is 10 mW. It should be noted that power consumption is not the same as power dissipation: the useful light output is not dissipated as heat in the transmitter but is transmitted to the detector.

We should also point out that the total power consumption of the diode-laser-based interconnection including the receiver can never be as low as the total power consumption of the electrical interconnection if the receiver is required to produce an ECL level for the next digital stage. A receiver circuit that translates optical signals into ECL levels with a 50-Ω termination must have the same current flow as the ECL driver for a 50-Ω electrical interconnection. Thus, if the optical interconnection is to have a lower power dissipation, the input structure of the logic gate must be adapted for optical inputs.

Fan-Out

Fan-out is one area in which optical interconnections present intriguing possibilities. With electrical interconnections and short links, we can connect several gates to the output of one gate provided that the fan-out does not exceed the rating of the driving gate. The difficulty with electrical interconnections, however, occurs at higher speeds and longer length lines. When a line is longer than an appreciable fraction of the wavelength of the signal, the line does not act like a simple wire. Controlled impedance or coaxial lines become necessary, and impedance matching and ground planes become important. With coaxial lines, one technique for implementing fan-out is to break the coaxial line at each gate and connect the gate. The last gate in the chain is terminated, although the best choice of the termination resistor is determined by the input capacitances of the gates. The major problem of this technique is the multiple reflections between the various input gates.

With optical interconnections, the fan-out is independent of the length of the interconnections for distances of interest. The fan-out can be performed electrically or optically, and each has its merits.

Electrical splitting can be accomplished for diode-laser-based optical interconnections if the diode lasers are connected in series. With several lasers connected in series to the output of an ECL gate, we can achieve fan-out simply by returning the lasers to a more negative supply, as shown in Figure 7. Note that although electrical interconnections cannot maintain standardized voltage levels when the gates are connected in series, the optical isolation property of optical interconnections permits electrical fan-out through series-connected lasers. In fact, two columns of diode lasers (Figure 7) can be connected in parallel to get fan-outs of up to 20 with a modest driver if the lasers have a low threshold. For example, speeds of 5 GHz (or even 10 GHz, as shown in Figure 4[b]) should be possible with only 20 to 30 mA of laser current if the lasers have a threshold of about 5 mA.

Fan-out can also be accomplished optically: one light beam can be split into N beams either through optics, fiber optic splitters, or waveguide splitters. With optical splitting, the splitting process is easy and very large fan-out ratios are possible. However, even for ideal splitting, the optical power drops as $1/N$, the electrical power out of the detector drops as $1/N^2$, and the signal level must be restored through amplification.

The work on optical amplifiers, including both diode laser amplifiers [14] and fiber amplifiers [15], is very promising for lossless optical splitting. Indeed, laser amplifiers with a gain of at least 20 dB could make

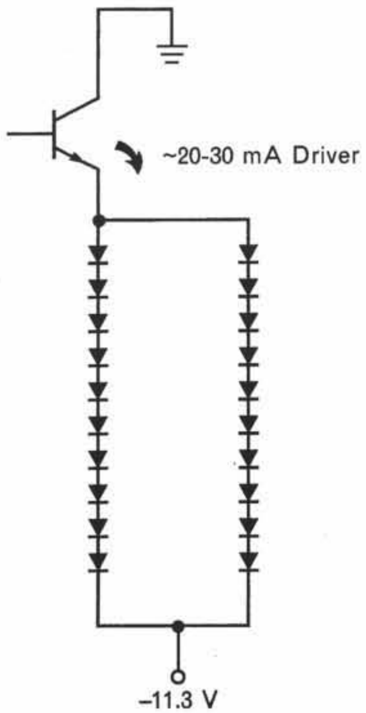

FIGURE 7. Fan-out can be achieved with no optical-splitting loss if the lasers are connected in series and parallel. Fan-outs of 20 should be possible with a modest drive current and 1-to-3-GHz operation. The schematic diagram assumes a laser series resistance ≤10 Ω, threshold current ~5 mA, and wavelength ~1.3 μm.

splitting ratios of 100:1 possible with no loss. Diode laser amplifiers have been demonstrated at 0.8 μm, 1.3 μm, and 1.5 μm. The work on fiber amplifiers is concentrated on 1.5 μm with Er-doped fibers that can be pumped with diode lasers. Such amplifiers have demonstrated net gains of 25 to 46 dB.

Alignment

High-speed optoelectronic components are generally small and difficult to align. There are several approaches to solving the alignment problems created by the use of free-space optics. One approach is simply to machine fixtures to tolerances as small as 6 μm where necessary, another approach is to borrow from technology developed for optical fiber alignment, and a third approach is to use active alignment with feedback. An alignment-invariant design in which software determines the alignment is also under consideration at Delft University [16]. Our approach was to design the optical system for reduced sensitivity to the misalignment. The following paragraphs briefly explain our work.

The use of optical modules with a source or detector and prealigned optics can greatly simplify the alignment [12]. A transmitter module can be designed to contain a diode laser and optics that allow the beam to expand to a ~1-mm-diameter collimated beam. The optics in the receiver module serve to focus the expanded beam onto a detector. All critical alignments between the lens and the source, and between the lens and detector, are performed during module assembly. Thus, during circuit assembly, the modules need only be aligned to the much less critical tolerances (on the order of millimeters) of the expanded beam input/output optics. This technique relaxes lateral-positioning requirements at the expense of tighter angular-alignment tolerances. For low insertion loss in the optical link, the receiver lens must be located in the near field of the transmitter lens so that most of the signal can be collected.

Using simplified thin-lens approximations of the optics and a uniform distribution of light, we considered the effects of misalignments one at a time to gain insight into the order of magnitude of the alignment problem. For a board-to-board interconnection with a 1-mm-diameter transmitter lens and a 2-mm-diameter receiver lens, a lateral misalignment of either the transmitter or receiver of up to ± 0.7 mm permits an optical transfer efficiency of 80%. For these lenses, board separations of 3 cm with an 80% optical transfer efficiency are possible if the transmitter module is aligned to within ± 2°. Also, risetimes of 50 psec are possible if the receiver module is aligned to within the same angular tolerance. The lateral and angular tolerances are compatible with modest mechanical accuracies.

In preliminary experiments, modules with 1.8-mm-diameter graded-index (GRIN) lenses (Figure 8[a]) were mounted on boards. (The GRIN lens is a cylindrical glass rod with a graded refractive-index profile that is highest at the center of the lens. The optical surfaces of the glass are not curved as with conventional refractive optics.) The boards were then inserted into a card cage with precision-milled slots (Figure 8[b]), which were necessary to maintain the alignment and reduce the effects of warped boards. Preliminary results showed that boards could be removed and inserted while a 24% differential current efficiency was achieved with a

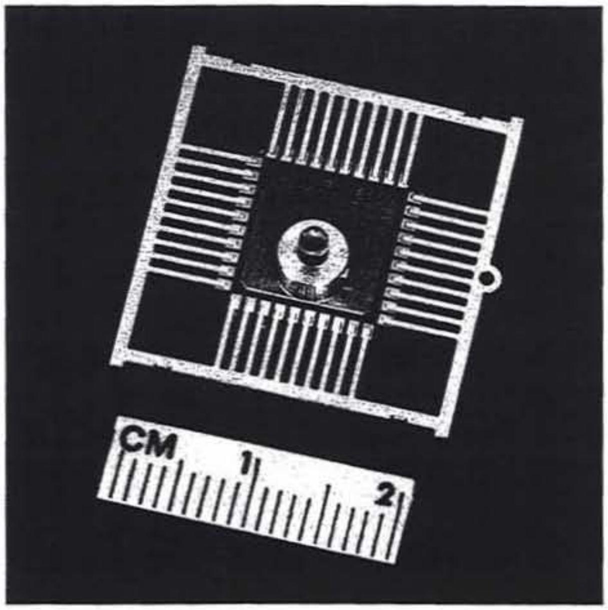

(a)

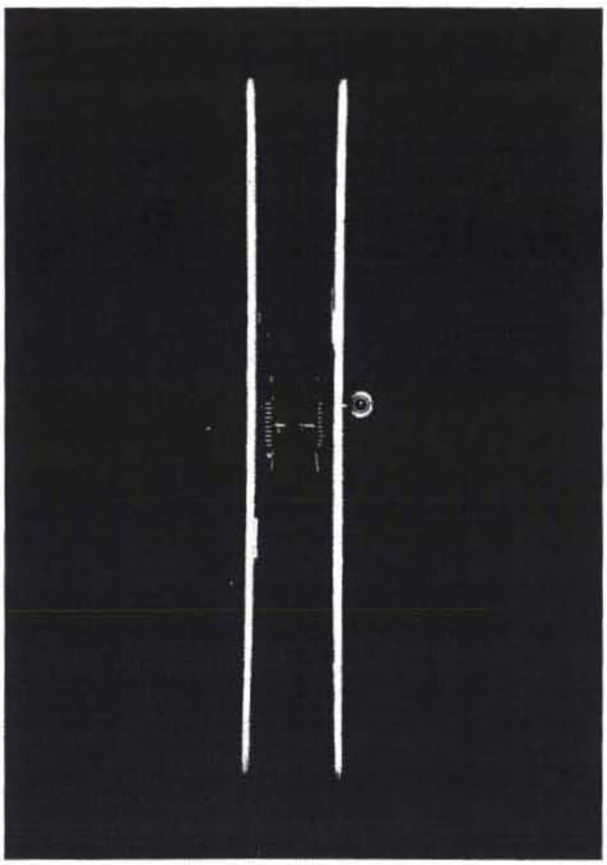

(b)

FIGURE 8. (a) Optical interconnection module that houses a laser or detector and a 2-mm-diameter GRIN (graded-index) lens in a high-speed flat-pack package. (b) Transmitter and receiver modules in a card cage.

58%-efficiency laser. No micropositioners were required to align the boards.

Physical Limits of Electrical Interconnections and the Role of Optics

The areas in which we believe optical and electrical interconnections will be important are shown in Figure 9. At low speeds and short distances (lower left region of the figure), simple wires will dominate because they are cheap and effective. The only reason to consider optics in this region is to overcome the difficulties associated with physical layouts and cross coupling of signals when high-density or high-connectivity interconnections are implemented.

As frequencies and distances increase, electromagnetic pickup becomes more of a problem and the reactance (capacitive or inductive) of a wire begins to induce signal distortions. In Figure 9, the blue line, which corresponds to an inductance of 5 Ω, gives a rough indication of the transition to the operating region where other interconnection techniques must be considered. The blue line is for a 50-μm-diameter wire. Depending on the impedances and current-handling capabilities of the digital circuits (e.g., CMOS versus ECL) and the geometry and resistance of the wire (e.g., flat versus round), wires might have limitations to the lower left of the 5-Ω line, or they might remain useful to the upper right of the line.

The physical limit on the use of a wire occurs when the length of the interconnection approaches a wavelength of the highest frequency being transmitted. An approximate limit is shown in Figure 9 by the red line, which represents an interconnection that is $\lambda/10$ long.

When simple wires induce unacceptable signal distortions or cross coupling, alternative interconnection technologies—electromagnetic transmission lines and optics—must be employed. Electromagnetic transmission lines include twisted pairs, microstrips, striplines, and coaxial lines. As the speed of transmission increases, transmission lines become dispersive and suffer losses from skin effect. The approximate limit for the use of transmission-line techniques is shown in Figure 9 by the green line, which represents a 3-mm-diameter 77-Ω copper airline (a coaxial line with air as the dielectric) at room temperature. The sloped portion of the line corresponds to a 3-dB propagation loss; the horizontal portion is at the frequency at which higher-order

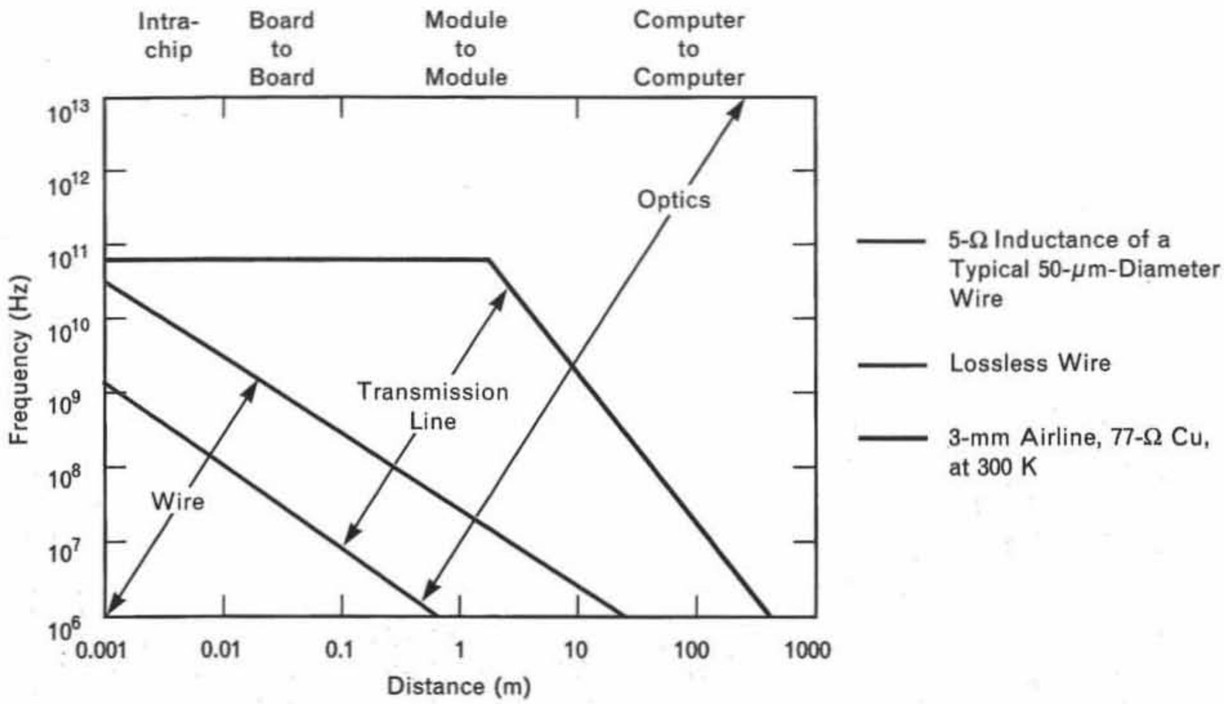

FIGURE 9. Regions of interest for both electrical and optical interconnections are determined by physical limits. At short distances and low speeds, wires are cheap and effective. At somewhat higher speeds and longer distances, both electrical transmission lines and optical interconnections are of interest. At the longest distances and high speeds, the physical limits suggest the use of optical interconnections. Although the optical channels for the highest speeds exist, the sources and receivers have not yet been developed.

modes become important.

The region where optical interconnection techniques begin to offer advantages in terms of power, isolation, and signal distortion begins roughly at the boundary where wires start to have difficulties. There exists a large and currently important range of speed and interconnection distance over which the trade-off between optical and electromagnetic-transmission-line techniques must be critically evaluated. At high speeds and long distances, optical interconnections offer the best solution because of the very high bandwidth and the low-loss characteristic of both optical fibers and free space. It is important to note, however, that at very high speeds the optical sources, drivers, detectors, and receiver circuits do not yet exist.

Summary

Optical interconnections can provide a number of benefits over electrical interconnections, particularly at long interconnect distances and high transmission speeds. Efficient high-speed low-threshold diode lasers can simplify optical interconnections. A diode laser with 3.8-mA threshold current, a –3-dB frequency of greater than 10 GHz, and 10 mW of output power at only 25 mA has been developed. The device makes possible optical interconnections with only a laser, two lenses, and a detector. Such a free-space board-to-board interconnection has been demonstrated at 1 Gb/sec, and the average power dissipation was only about 40 mW. Optical interconnections have prospects for even better performance with improved devices, and they should facilitate the development of advanced computers, particularly if the technology can be made cost effective.

Acknowledgments

The author thanks R.C. Williamson, D.L. Spears, C.E. DeVoe, J.R. Leger, and W.E. Courtney for helpful discussions. The author is greatful to Z.L. Liau and J.N. Walpole for their collaborations, and M.J. Corcoran for technical assistance. This work was sponsored by the Defense Advanced Research Projects Agency.

REFERENCES

1. I.H. Gilbert, "The Synchronous Processor," *MIT Lincoln Lab. J.* **1**, 19 (1988).
2. W.T. Cathey and B.J. Smith, "High Concurrency Data Bus Using Arrays of Optical Emitters and Detectors," *Appl. Opt.* **18**, 1687 (1979).
3. J.W. Goodman, F.J. Leonberger, S.-Y. Kung, and R.A. Athale, "Optical Interconnections for VLSI Systems," *Proc. IEEE* **72**, 850 (1984).
4. R.K. Kostuk, J.W. Goodman, and L. Hesselink, "Optical Imaging Applied to Microelectronic Chip-to-Chip Interconnection," *Appl. Opt.* **24**, 2851 (1985).
5. W.H. Wu, L.A. Bergman, A.R. Johnston, C.C. Guest, S.C. Esener, P.K.L. Yu, M.R. Feldman, and S.H. Lee, "Implementation of Optical Interconnections for VLSI," *IEEE Trans. Electron Devices* **34**, 706 (1987).
6. 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.* **27**, 1742 (1988).
7. C.T. Sullivan, "Optical Waveguide Circuits for Printed Wire-Board Interconnections" *Proc. SPIE* **994**, 92 (1988).
8. D.H. Hartman, G.R. Lalk, J.W. Howse, and R.R. Krchnavek, "Radiant Cured Polymer Optical Waveguides on Printed Circuit Boards for Photonic Interconnection Use," *Appl. Opt.* **28**, 40 (1989).
9. T.A. Lane, J.A. Quam, B.O. Kahle, and E.C. Parish, "Gigabit Optical Interconnects for the Connection Machine," *Proc. SPIE* **1178**, 24 (1989).
10. Z.L. Liau and J.N. Walpole, "Mass-Transported GaInAsP/InP Lasers," *MIT Lincoln Lab. J.* **2**, 77 (1989).
11. D.Z. Tsang and Z.L. Liau, "Sinusoidal and Digital High Speed Modulation of *p*-Type Substrate Mass-Transported Diode Lasers," *J. Lightwave Technol.* **5**, 300 (1987).
12. D.Z. Tsang, "Alignment and Performance Tradeoffs for Free-Space Optical Interconnections," *Optical Computing, 1989 Technical Digest Series* **9** (Optical Society of America, Washington, DC, 1989), pp. 146–149.
13. D.Z. Tsang, "One-Gigabit per Second Free-Space Optical Interconnection," *Appl. Opt.* **29**, 2034 (1990).
14. T. Saitoh and T. Mukai, "Recent Progress in Semiconductor Laser Amplifiers," *J. Lightwave Technol.* **6**, 1656 (1988).
15. P. Urquhart, "Review of Rare Earth Doped Fibre Lasers and Amplifiers," *IEE Proc.* **135** Pt. J, 385 (1988).
16. E.E. Freitman, private communication.

DEAN Z. TSANG
was born in Detroit and received the following degrees in electrical engineering: an S.B. from MIT, an M.S. from the University of Illinois, and an Sc.D. from MIT. As a graduate student in 1975, he joined Lincoln Laboratory, where he was involved with lead-salt diode lasers, integrated optics, and Q-switched diode lasers. He is currently a staff member in the Electrooptics Group; his present research interests include diode lasers and their applications to optical interconnections. Dean is a member of the Optical Society of America and Sigma Xi, and a senior member of the IEEE.