

Photonic Interconnections

Integrated circuits (ICs) are traditionally connected internally and externally using metallic wires. Electronic connection has many virtues, among them simplicity, compatibility with internal electronic devices, and the ability to make "tee" connections. However, electronic connection also suffers several drawbacks, particularly at high data rates. These include susceptibility to parasitic loading (capacitive and inductive) as well as to electromagnetic interference (EMI). These effects limit the overall usable bandwidth in electronic interconnects and can adversely affect system performance.

Microwave electromagnetic waveguide interconnects (microstrips), although operable at high frequencies, consume far too much real estate to be used in ICs. In contrast, the smaller wavelengths of photonic frequencies enable waveguides to be made much smaller. Thus, the alternative of connecting ICs photonically is not only attractive, but may become a necessity.

Ironically, it is the interactive nature of electrons through Coulombic forces that makes them ideal for information processing but poor candidates for information transmission. Processing of information requires the alteration of signals, but transmission of information requires the preservation of signals. Thus the virtually non-interactive nature of photons makes them ideal for communication. As such, optical interconnection is free of EMI and parasitic capacitive and inductive loading. The absence of loading problems also can improve the fan-out capability of optical interconnects.

A very common misconception is that the signal propagation time is shorter for an optical interconnect than for an electronic interconnect. This is not necessarily so. Propagation time depends on velocity and path length. For a free-space optical interconnect, an example of which is shown in Figure 1, the velocity is maximized, but for more common optical-fiber-based interconnection, the velocity might be comparable to that in an optimized coaxial transmission line. Furthermore, metallic wires can turn corners more sharply than optical fibers, and this can actually increase the path length for an optical interconnect. The performance advantage of optical interconnects in terms of speed lies in bandwidth potential rather than in sheer signal propagation time. Utilization of optical fiber bandwidth capability is limited by the source and detector electronics, so techniques to multiplex data economically in an interchip interconnect are currently of considerable interest.

Another potential advantage of using optical interconnects is communication with low-temperature electronics. The next generation of electronics may be able to operate at liquid nitrogen temperatures, including circuits based on new "high-temperature" superconducting materials. The cooling requirements of these circuits can be minimized if connections are made via poor thermal conductors such as glass, rather than excellent conductors such as conventional metal wires.

The development of optical interconnects may also help speed the development of electro-optical computing systems. In almost all E-O

computing systems under consideration, the conversion between optical and electronic signal representation is inevitable, and a viable technology for reliably implementing this conversion is essential.

Components

The taxonomy of photonically interconnected circuits consists of three primary components: an optical source, a transmission medium and an optical receiver. Optical sources include LEDs and laser diodes. Transmission media include free space, optical fibers and integrated photonic waveguides. Receivers can be photodiodes (PN, PIN, avalanche or Schottky) or photoconductors. The choice of components, as in all systems, is application-dependent. For example, the speed and length of the interconnect, as well as the materials of the integrated circuits, are important considerations.

LEDs and laser diodes are considered the dominant candidates for sources because of their compatibility with current or near-future circuit fabrication processes. Laser diodes generally offer speed and output-power advantages over LEDs, but for short-distance low-fan-out applications, LED sources may suffice. They are simpler to fabricate and require less real estate. Laser sources, while offering superior performance, require significant power to operate and, because of their temperature sensitivity, require special attention to heat sinking.

An interesting, but thus far unexplored, alternative to source design is the use of a more powerful but remote laser source that is locally modulated. The local modulator or light valve, perhaps a semiconductor multiple quantum well device, plays the role of source. Several such sources could be driven from a single laser diode, but an additional optical interconnect would be required to pipe in the external supply of photons.

At the other end of the interconnect lies the receiver. This device must translate the incoming photons back to an electronic signal. Several semiconductor devices are well established for this purpose. A PN junction has the advantage of simplicity, but unless the photons are absorbed in a space-charge region where they are collected rapidly by the electric field, the response time is




Figure 1. Free-space interconnect scheme pursued at the Jet Propulsion Laboratory. Light from sources is focused to points on the silicon chip by means of the holographic grating.

dominated by the slower diffusion process, and the detector speed is substantially reduced. The PIN structure avoids this problem by extending the electric-field collection region through the use of the lightly doped or intrinsic region. The structure, by increasing the effective "plate" spacing, also has lower capacitance, thus reducing circuit R-C time constraints. However, the PIN structure is harder to fabricate.

In gallium-arsenide-based devices, in which photon absorption lengths are substantially reduced, metal-semiconductor contacts that behave like Schottky diodes often are used as detectors. For example, lateral back-to-back diodes offer high speed and simple fabrication. However, the dark current can be dominated by surface effects and can introduce significant noise into low-light-level (i.e., high-fan-out) applications. More sophisticated devices, such as those produced by bandgap engineering, can provide amplification with significant noise reduction. These laboratory devices are difficult to fabricate and their widespread role in optical interconnects is uncertain.

Two important considerations in receiver design, in addition to speed and detectivity, are packing density and crosstalk. For high-density interconnects, detectors that are spaced too closely are susceptible to crosstalk. The crosstalk arises from a combination of photonic and electronic effects. Since photon absorption is a statistical process, some photons may stray into adjacent detector regions. This effect can be alleviated by proper photonic interface design and increased detector spacing. Electronically, carriers generated by the absorption of photons may also stray into adjacent detector carrier collection regions. This can be managed by increased detector spacing, improved detector isolation, and by strong confinement of carriers by electric fields.

Between the source and the receiver is the transmission medium. The transmission of photons can be guided or unguided. In one sense, the choice between guided and unguided interconnects is analogous to choosing between microwave links or fiber optic links in long-distance telephony. In the unguided or free-space mode, photons emitted by the source are focused to a point or several points on an integrated circuit by means of lenses or holographic gratings.

The approach of free-space interconnects has some interesting possibilities, including splitting a light source to several receivers. With

light valves, the destination might be dynamically reconfigured, allowing for adaptive interconnect architectures. Since photon beams can cross each other in free space without appreciable interaction, the density of interconnects is potentially much higher for free-space interconnects than for guided transmission. Near-term applications for free-space interconnects generally involve high-fan-out intrachip connections. Examples are dynamic random access memory addressing and VLSI-circuit skewless clock distribution.

Although viewed as perhaps the ultimate interconnect methodology, there are several technological hurdles to be overcome before free-space interconnects can be realized practically. For example, to optimize detector real estate consumption, the focus and directionality of the beam must be tightly controlled through the design of the grating. Also required are the exact alignment of source, holographic grating, and integrated circuit during manufacturing and the configuration's continued mechanical stability. Stray photon diffraction from the grating must be minimized to avoid crosstalk and electronic effects.

Guided-mode transmission offers the advantage that the signal is guided directly from the source to the receiver without requiring precise geometrical alignment or "line-of-sight" topology. For interboard and interchip communication, optical fibers are the most readily apparent technology for guiding the lightwave. However, for some interchip and especially intrachip applications, waveguides fabricated directly in the semiconductor by material modification, or on the semiconductor by thin-film techniques, offer the advantage of monolithic fabrication. Yield and reliability would be expected to be significantly higher for the latter approach, as was the case when transistors were first integrated through the use of thin-film metallic wires.

Choice of materials

Today's ICs are fabricated primarily in silicon, yet silicon is not the optimal material in which to implement optical interconnects. There are several reasons for this, first among which is that silicon is an indirect bandgap semiconductor. A consequence of this is that radiative recombination is so unlikely that fabrication of light-emitting sources in silicon would be very difficult.

Some optical interconnect applications, though, do not rely on integrated source fabrication, such as in

a free-space interconnect for RAM addressing or clock distribution. These applications do rely on the proper detection of the photonic signal, and this is a second area in which silicon is a nonoptimal material. Silicon photoreceivers tend to consume more real estate than gallium arsenide receivers because it is harder to isolate silicon devices than GaAs devices.

Since silicon devices are generally minority-carrier devices, it is critical to control the generation and diffusion of these carriers outside the detector region. Stray photons from imperfect grating structures or lossy optical fibers can cause generation of carriers in undesired regions. For silicon CMOS, the consequences could be dire. These circuits are particularly prone to minority-carrier-triggered latch-up, which invariably leads to circuit meltdown. In DRAM applications, these minority carriers could lead to random bit errors.

Photon absorption in semiconductors is a statistical process; therefore minority carriers can be generated photonicly many absorption lengths beneath the detector surface. Isolation of these generated carriers is difficult in silicon. Silicon devices are isolated in today's technology by trenches several microns deep that are filled with an insulator. The depth of these trenches is usually small compared to minority-carrier diffusion lengths, so they can be ineffective for photonic injection. The electronic consequences can be the same as those described above, or can be manifested as channel crosstalk.

Gallium arsenide and other non-silicon semiconductors offer many advantages over silicon, such as speed, ease of source fabrication, reduced absorption length, ease of device isolation, higher operating temperatures and lower power consumption. However, note that the most mature of these alternate material technologies, GaAs ICs, is only now beginning to make an impact in the electronics market. The dilemma that faces systems designers who are contemplating using optical interconnects is how to make the decision between committing to GaAs or making do with silicon.

There is another difficulty arising from the relative immaturity of GaAs, namely the inability to fabricate large-scale ICs with integrated optoelectronic components such as lasers and receivers with an economical manufacturing yield, especially with the added complexity of coupling the optoelectronic portions to optical fibers. Although optical inter-

connects may become technologically viable, their acceptance in other than high-performance DoD hardware will be extremely limited until manufacturing costs become manageable.

A compromise that decouples the optoelectronic IC (OEIC) manufacturing yield from the non-optoelectronic IC, and allows the latter to remain realized in silicon is to adopt a hybrid approach. Here, the OEIC contains the sources, receivers and driving and amplification circuitry. The OEIC is coupled to the silicon (or GaAs) circuit using established hybridization techniques such as solder ball bumping ("flip chip" technology) or wire bonding. The silicon IC designer need not be concerned with the optical interconnect, and the OEIC, being a standard interface for optical interconnects, need not be redesigned for each IC. Manufacturing efficiency will probably improve considerably with this approach.

Coupling technology

The physical connection between the IC source or receiver and the transmission medium is perhaps the most difficult technological challenge that faces optical interconnects. There are several constraints on the design of the coupler. It must consume a minimal amount of chip real estate, it must facilitate alignment between the source/receiver and transmission medium, it must have low loss, and it must be mechanically stable. Several schemes have been proposed recently for achieving these requirements, but up until now only two have demonstrated their potential.

In a horizontally coupled GaAs receiver scheme pursued by Honey-

well (Figure 2), an optical fiber's output end is polished at an angle permitting the reflection of the light vertically down into the IC receiver. The receiver is an interdigitated back-to-back Schottky diode structure. The alignment between the fiber and the IC is maintained by using an alignment fixture, a piece of silicon that has been anisotropically etched to form V-grooves. The fibers are aligned by the V-grooves and glued in place. The fibers and the end of the silicon fixture are bevelled by simultaneous polishing.

The V-groove alignment fixture method allows a high density of optical interconnects to be made to the perimeter of an IC. However, IC trends indicate that in order to satisfy the number of I/O connections required for VLSI circuits, connections must be made to the interior regions of the chip as well. Thus, while horizontal coupling schemes maintain the planarity that has so far characterized IC development, they may not be suitable for interconnected systems of very high density.

Vertical coupling of optical fibers allows the highest possible density of interconnects. In a scheme developed at Columbia University, optical fibers are directly connected to ICs using a vertical coupling technique. Since the detector footprint need be only as large as the fiber core, the cladding footprint can be utilized for amplification or driving circuits. In this coupler, a vertical cavity is etched into silicon using an anisotropic etching process. Cylindrical cavities are typically 20µm deep and 20µm in diameter for front-side interconnects. The inside of the cavity is doped to form a PN junction on the interior surface. An optical fiber is

conically tapered using wet etching and inserted into the cavity. Thus, the cavity not only provides alignment and mechanical stability, but the detector surrounds the fiber tip and acts as an integrating sphere, yielding high quantum efficiency. However, the PN junction detector structure, due to carrier diffusion, is relatively slow, particularly at longer wavelengths.

The cavity concept is being extended to high-speed GaAs circuits using a backside coupler approach. The cavity, the diameter of which is matched to the optical fiber diameter, is etched from the back side, terminating close to the front surface of the wafer. Detector performance can be dramatically improved using this strategy. Such 3-D interconnects may be particularly useful in optical computing applications.

Another challenge for any coupling scheme is efficient coupling of photons from a source to an optical fiber. Conventional sources, including most laser diodes, are poorly suited for this. However, there has recently been considerable research on vertically emitting laser sources. The alignment of fibers to sources is particularly critical, and means of manufacturing such interconnects reliably and economically will be critical to the realization of optical interconnects.

Future prospects

The future for optical interconnects looks promising. Although the widespread replacement of electronic interconnects in electronic systems with optical interconnects is unlikely, there are many niches in which optical interconnects will play an important role. Certainly, inter-board connections will benefit from advances in optical interconnects and will probably be implemented with optical fibers. Optical interchip communications will initially find use in harsh environments or high-data-rate/high-fan-out applications, but widespread adoption will not occur until packaging economics justifies it. Optical printed circuit boards fabricated with silicon may be an interesting alternative to connecting chips with optical fibers. Intrachip photonic communication remains a difficult challenge, but free-space transmission may improve these prospects. Overall, it seems likely that the path followed by photonic communications, from long-distance telephony to LANs, will lead to the very short-distance communications realm.

*Eric R. Fossum
Columbia University*




Figure 2. Coupling approach pursued by Honeywell. The fibers, which are aligned by the silicon V-groove fixture, are polished at an angle so that the optical signal is directed vertically into the photoreceiver.