INTERCONNECTION PROCESSES AND MATERIALS

Greg E. Blonder, Richard A. Gottscho, and King L. Tai

Greg E. Blonder *is head of the Interconnection Research Department.* ***Richard A. Gottscho*** *is head of the Electronics Packaging Research Department, and* ***King L. Tai*** *is a supervisor in that department. They are with AT&T Bell Laboratories in Murray Hill, New Jersey.*

Mr. Blonder's research involves low-cost optical components and systems, optical integration, and the underlying basic physics of materials and devices. He joined the company in 1982 and has an M.S. and Ph.D. in physics from Harvard University (Cambridge, Massachusetts).

Mr. Gottscho's interests are plasma processing, real-time process monitoring, and plasma physics. His department is responsible for electronic packaging and interconnection research. He joined the company in 1980 and has a B.S. in chemistry
(continued on page 59)

The performance of today's electronic and photonic systems is largely limited by interconnection technology. Using silicon-based multichip modules and hybrid optical packaging as exemplary models, we show how high-performance interconnection technology can be achieved at low cost. The key is to choose materials and processes that are compatible with today's very-large-scale integrated-circuit manufacturing.

Introduction

The genius in Seymour Cray's invention of the supercomputer was in his exploitation of interconnection design and technology.[1] However, the performance of today's electronic systems is still limited by interconnections between components and subsystems and not by the high-speed, very-large-scale integrated (VSLI) circuits from which the systems are constructed. This paper concerns processes and materials for the next generation of high-performance interconnection technology. To achieve the overriding goal of high performance at low cost, it is imperative that we build on existing manufacturing and materials expertise. (Panel 1 defines acronyms and terms used in this paper.)

In any electronic or photonic system, there is a plethora of interconnection technologies. Figure 1 shows how the cost per connection increases dramatically from on-chip to system-to-system interconnections. It is beyond the scope of this article to discuss all these technologies. Instead, interested readers should refer to a recent issue of the *AT&T Technical Journal* that was devoted entirely to interconnection technology.[2] In this paper, we focus on chip-to-chip interconnection materials and the processes that are compatible with conventional VSLI manufacturing.

System Needs

At each level of interconnection, the driving force is to lower cost and improve performance. To achieve this goal, we must:

- *Reduce the cost per connection.* VSLI manufacturing has shown that we can achieve this most effectively by increasing the interconnection density. As the level of integration increases, the cost per

Panel 1. Acronyms, Elements, and Terms

Ag	silver
Al	aluminum
Au	gold
ASIC	application-specific integrated circuit
AVP	advanced VSLI packaging
Cr	chromium
Cu	copper
CVD	chemical-vapor deposition
DARPA	Defense Advanced Research Projects Agency
EMI	electromagnetic interference
HIC	hybrid integrated circuit
HOPS	hybrid optical packaging on silicon
I/O	input/output
IC	integrated circuit
M1	first-level metal
M2	second-level metal
M3	third-level metal
MCM	multichip module
MFLOPS	10^6 floating-point arithmetic operations per second
MI	microinterconnect
Ni	nickel
Pb/Sn	lead/tin
Pd	palladium
PdAg	palladium silver alloy
PIN	positive-intrinsic-negative
POLYHIC	polymer hybrid integrated circuit
PWB	printed-wiring board
Si	silicon
SiO_2	silicon dioxide
SINCAP	silicon nitride capping layer
SRAM	static random-access memory
Ta_2Si	silicon resistor
TCE	thermal coefficient of expansion
TFLOPS	10^{12} floating-point arithmetic operations per second
Ti	titanium
Ti/Pd/Cu/Ni/Au	layered titanium, palladium, copper, nickel, and gold
TV	television
VSLI	very-large-scale integration
VCR	video-cassette recorder
W	tungsten

connection decreases because more connections are made simultaneously. This results in increased yield and reduced use of materials and processing.

- *Shorten the interchip distance.* If we increase the level of integration, we also enhance system throughput, which improves performance.
- *Increase input/output (I/O) per chip.* This directly increases system throughput. To increase the I/O per chip, we use higher density interconnection technologies such as multichip packaging.
- *Improve power management.* Materials with high thermal conductivity, such as silicon (Si), are preferred over thermally insulating ceramics. These materials reduce the need for expensive heat sinks and liquid cooling in high-density packages.
- *Use commodity chips.* On-chip customization is prohibitively expensive unless large volumes are produced and sold. If we passively interconnect small, inexpensive commodity chips with self-contained functions, we can minimize the length of the design-test-manufacture cycle and provide cost-effective, complex, customized circuitry.
- *Use common technology for photonics and electronics.* Future systems will require increased integration of photonic components with electronic components. When common technology is used to package these components, material and processing costs are shared and, thereby, reduced.

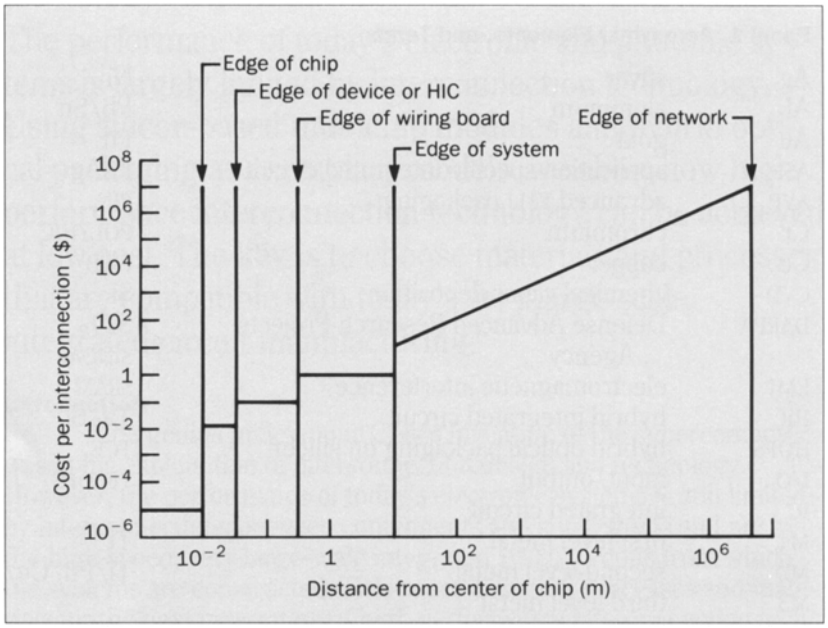

Figure 1. Approximate cost of interconnecting a conductor or fiber in a telecommunications system. The most effective way to reduce the cost per interconnection is to increase interconnection density. As the level of integration increases, more connections are made simultaneously. (Adapted from "Materials for Information and Communication," by John S. Mayo. Copyright© 1986 by Scientific American, Inc. All rights reserved.)

Manufacturing Needs

Although backplane or circuit-board technology appears mundane, such disinterest is, in reality, a testament to the low cost and high quality that is routinely achieved in conventional, backplane manufacture. Indeed, new technologies such as those we discuss here must satisfy these high expectations and add enough value to justify the costs of the new technology. To achieve this goal, we must understand the interactions among the materials that make up packages, design processes and assembly techniques that are simple and insensitive to variations, develop low-cost and reliable testing methods, and create designs that permit repair or replacement of individual components.

Materials Interactions. Even in conventional printed-wiring board (PWB) designs, the interactions among materials are important. For example, solder must be designed not to dissolve copper PWB wiring, yet must assure long-term adhesion and electrical and thermal contact. Connectors must be of low resistance and function in the presence of corrosive atmospheres, high temperatures, and humidity.

When we go to higher interconnection densities or optical interconnections, these conventional problems are compounded. For example, dust and dirt often collect on wire leads and can short out devices.[3] The closer the wires, the greater the problem. During manufacture or in the field, plastics liberate high-vapor-pressure components that can coat optical surfaces or prevent good solder adhesion. Many optical components, such as lasers, require a hermetic environment for long-term reliability, but fabricating the hermetic environment may loosen the solder that holds the devices in alignment. The components, board, and backplane must work together through many thermal cycles, even though they consist of diverse materials with different expansion coefficients.

Assembly. Open the back of a computer, and you will see integrated-circuit (IC) packages attached to PWBs, and boards slid into edge-card connectors along the PWB backplane. Typically, the package leads are inserted into holes on 2.5-mm (millimeter) centers (pitch), and soldered into place. More recently, surface mounting of components has become popular to increase component density. Here, the ICs are soldered directly to copper traces on the board surface, with 0.6 mm between the centers of the leads. The assembly process relies on a special-purpose robot that picks up the IC (i.e., the chip), and, through a combination of dead reckoning and vision, puts the chip in its proper location. Because the leads are placed in holes or restrained by small dabs of adhesive, the chip is held in place during the soldering operation.

These operations are much too coarse-grained for electrical leads on 0.06-mm centers or for optical devices that require accurate and stable alignment to 0.001 mm. New methods are required. Among the possible solutions are: improved, active alignment and self-alignment.

Improved, active alignment can include high-performance vision systems that view the chip and the pad. Assembly methods also include *capacitive alignment*, where the bonding station measures the capacitance between the chip and the pads as assembly is taking place. When the leads are directly above the correct pads, the capacitance is highest; when they are misaligned, the capacitance drops.[4] The apparatus can grip optical devices, such as lasers or photodetectors, in a fixture that also provides power to the devices. Then, by monitoring the device's light output through the waveguide intended for attachment, highly accurate alignment is possible.

A jigsaw puzzle is an everyday example of *self-alignment*. You need to get the right piece near the right location. Then, the piece's shape and some gentle vibration will align and hold it in registration. Solder reflow can be used in the self-alignment of chips to bonding pads. When the solder melts, surface tension "sucks" the pad and chip leads into mutual alignment.

Because self-alignment fails when the chip and pad are offset by more than one-half the pad spacing, this technique is generally combined with another passive restraint. The added restraint could be a well with sloping sides to contain the leads. The chip is etched with a matching bevel and, like the puzzle pieces, is simply vibrated into place. Or, both the chip and the circuit board can be etched to accept a few small alignment pins, such as ball bearings.

In general, active alignment is more expensive than self-alignment, and the latter will be used wherever possible.

Testing. In a quality design, yields should be high enough to avoid almost all testing. Nevertheless, the final part must be tested and exercised before use. Of course, testing is also needed during early development. Thus, we may still need to probe the board's function at the finest wiring pitch, even though this pitch may be too fine for conventional probe tips.

Solutions to this problem take many forms. Some of the chips on a board can be designed only for testing the remaining chips. In this way, we can address a *proctor* chip with conventional probes and use it to execute a built-in program that exercises the higher interconnection density part of the board. Or, special test probes can be designed with the same interconnection technology as the circuit board. For example, a probe—a "dummy" chip with tiny (0.01 mm), lithographically defined spring fingers—can be pressed onto the board, taking advantage of the self-alignment techniques described above.

In an optical backplane, the original design can include extra waveguides that tap a small portion of the light.

Repair. Even the best products sometimes need repair. By using localized heating and easily cleaned fluxes, we can remove chips from a board by remelting

the solder. After we check the bond site for solder integrity, we can install a new chip. Again, self-assembly techniques help to prevent other chips from moving during this operation and simplify insertion of the replacement chip onto a crowded board.

However, hermetic or Si-encapsulated packages are impractical or too costly to repair. Often, it is cheaper to replace the part entirely. Yields and reliability must be high enough so that repairs are infrequent and can be done cost-effectively at the module level.

Multichip Modules

Because multichip modules (MCMs) provide increased integration, they offer the low cost, high performance, and compact packaging that system designers need. An MCM consists of several VSLI chips, each performing a different function, that are mounted on a base, called a substrate, that is patterned with thin-film resistors, capacitors, and conductors. Multichip packaging offers tenfold improvements in the ratio of system cost to system throughput compared to single-chip packaging.

The MCM concept is similar to on-chip, VSLI interconnection but is really a compromise. Ideally, one would prefer to gain further functional integration by making wafer-sized (150-mm diameter) circuits with the interconnection density of VSLI (0.001 mm). However, this approach currently is too expensive because of the number of defects found on Si wafers. Today's chip sizes (100 mm^2) are chosen because the probability of finding a defect on a chip this size with VSLI interconnection density is less than one. MCMs use interconnection structures with at least 1/100 the interconnection density of VSLI. Then, the module is constructed using individually tested VSLI chips that are attached to the MCM substrate, which prevents a single defect from destroying the module function. From the system perspective, MCMs look like subsystems on a giant, but inexpensive, customized chip.

Although a variety of MCM technologies are available (see Table I), all have in common a multilevel structure of insulators and conductors. However, the packaging technologies use different materials and methods for chip attachment. These differences determine not only the electrical performance of the subsystem but also the extent to which we can build on existing expertise with materials and packaging.

A polymer hybrid integrated circuit (POLYHIC) uses thin polymer films on ceramic substrates. In cofired ceramic technology, ceramic and metallic films are cocentered. Here, we focus on the microinterconnect (MI) technology that uses aluminum (Al) metalization on Si substrates (Figure 2). MI technology is similar to AT&T's advanced VSLI packaging (AVP) technology but uses Al instead of copper (Cu).[2] In AVP, the interconnect metalization actually consists of three metals: Cu, chromium (Cr), and nickel (Ni). Although Cu is a better conductor than Al, Cu processing is considerably more complex because multiple metal layers are needed.

MI technology offers both superior performance (see Table I) and an opportunity to exploit existing processing and design expertise. The two key elements in fabricating MI MCMs are:

- Fine-line interconnection on Si substrates
- Flip-chip solder attachment.

Fine-Line Interconnections on Si. On Si VSLI chips, for example, line dimensions are less than 0.001 mm and the cost per connection is only $\$10^{-5}$ (Figure 1). This low cost results primarily from the simultaneous, high-yield production of millions of interconnections. On MCMs, line dimensions are typically 0.01 to 0.02 mm and the cost per connection is $\$10^{-2}$. The larger, MCM line dimensions are tailored to provide critical damping of signal reflections without the need to use termination resistors that dissipate excessive power.[5] This performance advantage cannot be realized with either finer or coarser dimensions.

Although such lines are difficult to fabricate on ceramic substrates, they are routinely fabricated on Si for VSLI. Also, Si is preferred to ceramic as a substrate material because of Si's high thermal conductivity and perfect thermal expansion coefficient match with Si chips.

Table I. Current Multichip Module Packaging Technologies

Properties	POLYHIC (standard)	Cofired ceramic		MI (development)
		Standard	Development	
Conductors				
Metalization	Ti/Pd/Cu/Ni/Au	W	Au, PdAg, Ag	Al
Minimum line width (mm)	0.05	0.125	0.1	0.01
Minimum line pitch, P (mm)	0.125	0.25	0.2	0.02
Thickness (mm)	0.005 to 0.01	0.0075 to 0.025	0.0075 to 0.025	0.0025
Dielectric layers				
Material	Triazine/96% alumina	Glassy ceramic	Glassy ceramic	Polyimide/silicon
Maximum number	2	>30	>30	2
Maximum signal layers, N	~2.5	>20	>20	2
Minimum thickness (mm)	0.025	0.0925	0.0925	0.0025
Via diameter (mm)	0.15 by .015 or 0.1 by 0.2	0.25	0.125	0.0075
Glass transition temperature, T_g (°C)	~200	—	—	350 to 400
Interconnection density				
Signal lines only (mm/mm^2), N/P (excludes vias)	20	79	98	98
Electrical				
Dielectric constant	2.8	8.0	4.0 to 4.5	3.5
Propagation delay (ns/mm)	0.0055	0.009	0.0067 to 0.0071	0.0063
Resistance (Ω/mm)	0.098	0.039	0.020	0.98
Inductance (mH/mm)	0.28	0.43	0.39	0.24
Capacitance (pF/mm)	0.11	0.20	0.14	0.098
Bandwidth (GHz)	>1	N/A	N/A	>1
Attenuation at 3.5 cm, 20 GHz (%)	7	N/A	N/A	20
Thermal				
TCE matched to	Alumina	Alumina	Silicon	Silicon
Thermal conductivity, dielectric (W/cm°C)	0.35	0.03	≤ 0.03	1.6
Miscellaneous				
Flip-chip attachment	To be developed	Yes	To be developed	Yes
Termination or damping resistor	Yes	Yes	Yes	No

To keep the costs of materials and processing low, polymers such as polyimide are the preferred interlevel insulators. Here, processing is similar to that used in conventional photolithography. That is, films are spun onto the Si substrate, are cured by heating to 300°C to 400°C (degrees Celsius), and then are patterned.[6]

Interfacial adherence is a common problem in building multilayer structures. If the Al is free of water and organic residues, then adhesion is not a problem when polyimide is deposited on Al. But when Al is deposited on polyimide, the adhesion is usually poor.[7] Often, an intermediate layer such as titanium is used to provide

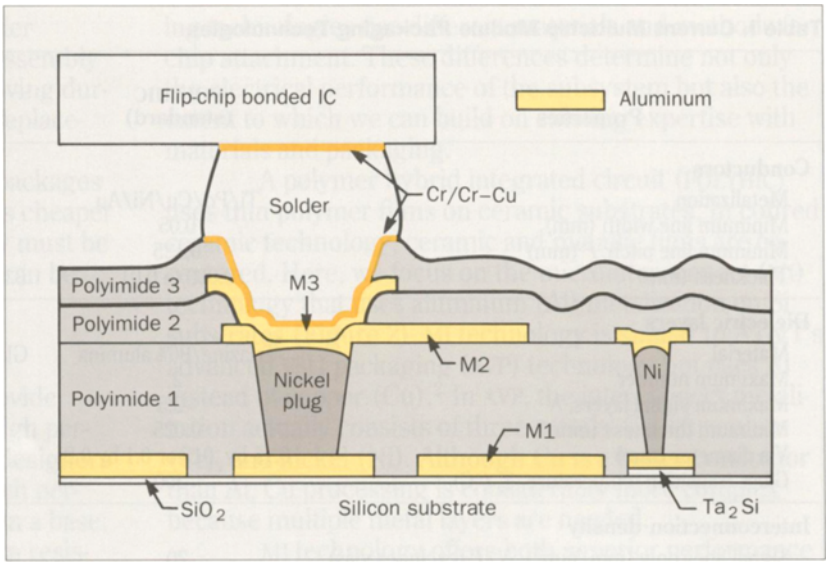

Figure 2. The microinterconnect (MI) structure used in fabricating multichip packages. A chip is bonded face down ("flip-chip") onto a Si substrate using Pb/Sn solder bumps anchored by Cr/Cu composite base metals. Polyimide is used as an interlevel dielectric, typically in a three-level Al metal—M1, M2, M3—scheme. One level is used for power and two levels are for signals.

the "glue" that bonds Al to polyimide. This approach is complicated and costly.

The understanding of adherence phenomena such as these remains a fundamental materials-science problem whose solution will lead to simpler and cheaper structures.

The choices of substrate, metalization, and interlayer dielectric just discussed provide an important benefit: We can fabricate MCM substrates in existing IC facilities. In building customized MCM substrates, we can effectively use our vast knowledge of IC processing and reliability, as well as the convenience and low cost of commercial equipment.

Flip-Chip Solder Attachment. The second critical MCM building block is flip-chip solder attachment. In ordinary IC packaging, the chip is wire-bonded face-up with respect to the mounting surface. The term *flip-chip* is used because the chip is bonded face down toward the substrate (Figure 2). (IBM uses the term C4, or *controlled collapse chip connection*, for this technology.[8])

Thin metal films anchor the solder joint to both the chip and the substrate. Solder may be applied to either the chip, the substrate, or both. But to make MCM repair simpler and cheaper, it is preferable that solder be applied only to the chip.

Also important is to have solders with different melting points. For example, the solder used to bond chips to MCM substrates must have a higher melting point than the solder used to connect the MCM to the PWB. We can easily achieve this by varying the tin concentration in lead-tin solders. However, the melting point is not the only issue, and solder metallurgy is more complicated than first meets the eye.

It is important to make reliable solder joints. Again, this means we must use multiple layers to obtain strong, reliable adhesion between the solder bump and the chip and between the solder bump and the substrate. The solder must easily wet the metal layers but, at the same time, repeated solder reflow must not dissolve the metal layers.

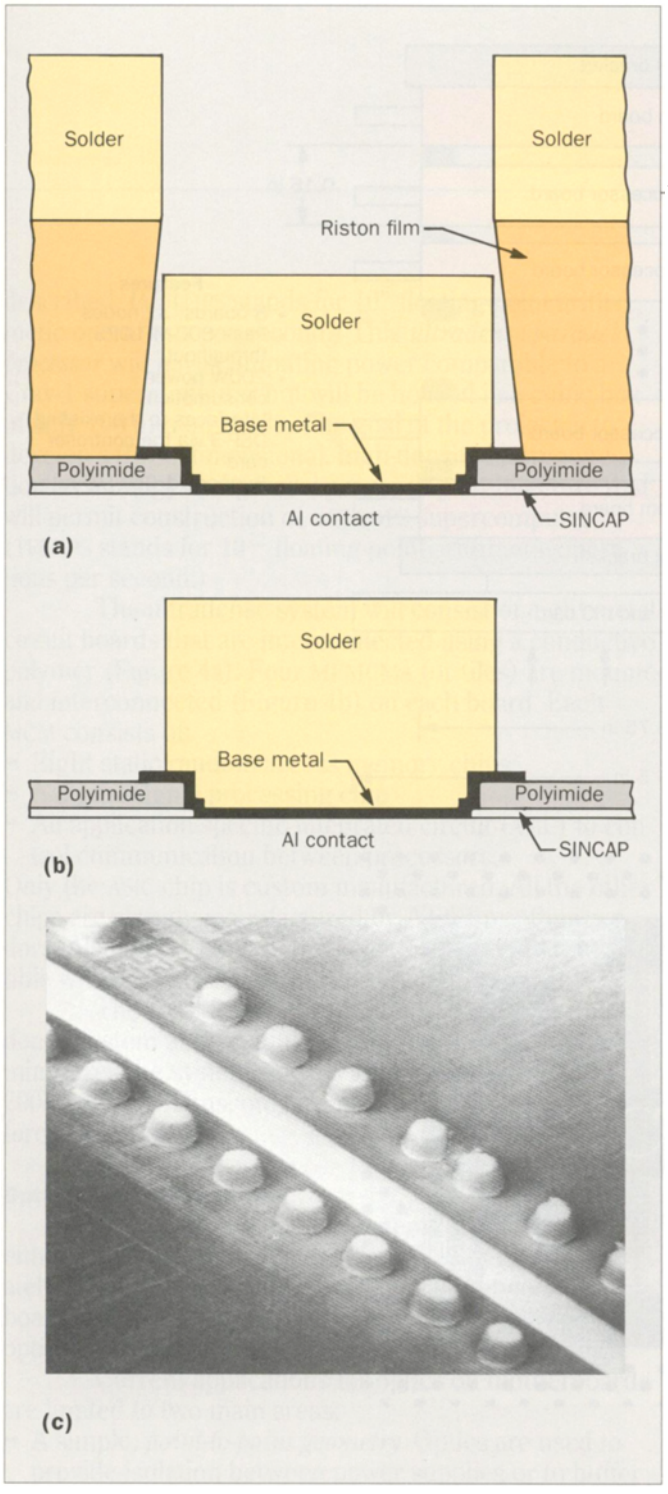

Figure 3. Solder-bump structure. (a) Solder and base metals are deposited into the holes in and on top of the dry-film resist (Riston photopolymer film), which is then (b) dissolved during the lift-off process. The device is protected by a silicon nitride coating (SINCAP). This technique produces well-controlled, uniform, high solder bumps. (c) This scanning-electron micrograph shows the row of solder bumps on a device wafer before solder reflow.

To satisfy these contradictory requirements, we use a base-metal system of two interleaved composites of Cr and Cu.[8] During reflow, the Cu absorbs molten solder into the Cu/Cr composite but the Cr remains inert and is not dissolved by the solder. After the solder cools, it is locked into the Cr structure. Even after repeated reflows, the structure remains stable.

Base metal and solder bumps are deposited by sputtering and evaporation, respectively, into the holes of a film resist (i.e., Riston® photopolymer film) that is defined using a lithographic process (Figure 3a), as well as onto the resist surface. (Riston is a registered trademark of E. I. Du Pont de Nemours and Company.) The holes provide access to underlying Al contact pads. The Riston film is then dissolved (during a lift-off process, Figure 3b) leaving behind well-controlled, uniform, high solder bumps on the chip's I/O pads (Figure 3c).

First, chips are attached to the substrate through the application of heat and pressure. Then, the solder is reflowed by heating in the presence of flux. This process forms good solder joints, and the surface tension of the molten solder aligns the chip to the substrate. Thus, the placement tool needs to align the chip to the substrate to within a distance of only half the size of a bonding pad—typically, 0.05 mm—and chips need not be perfectly attached to the substrate initially.

Ultradense Parallel Processor. Under contract to the U.S. Defense Advanced Research Projects Agency (DARPA), AT&T Federal Systems is constructing an 800-MFLOPS computer using the MI technology just

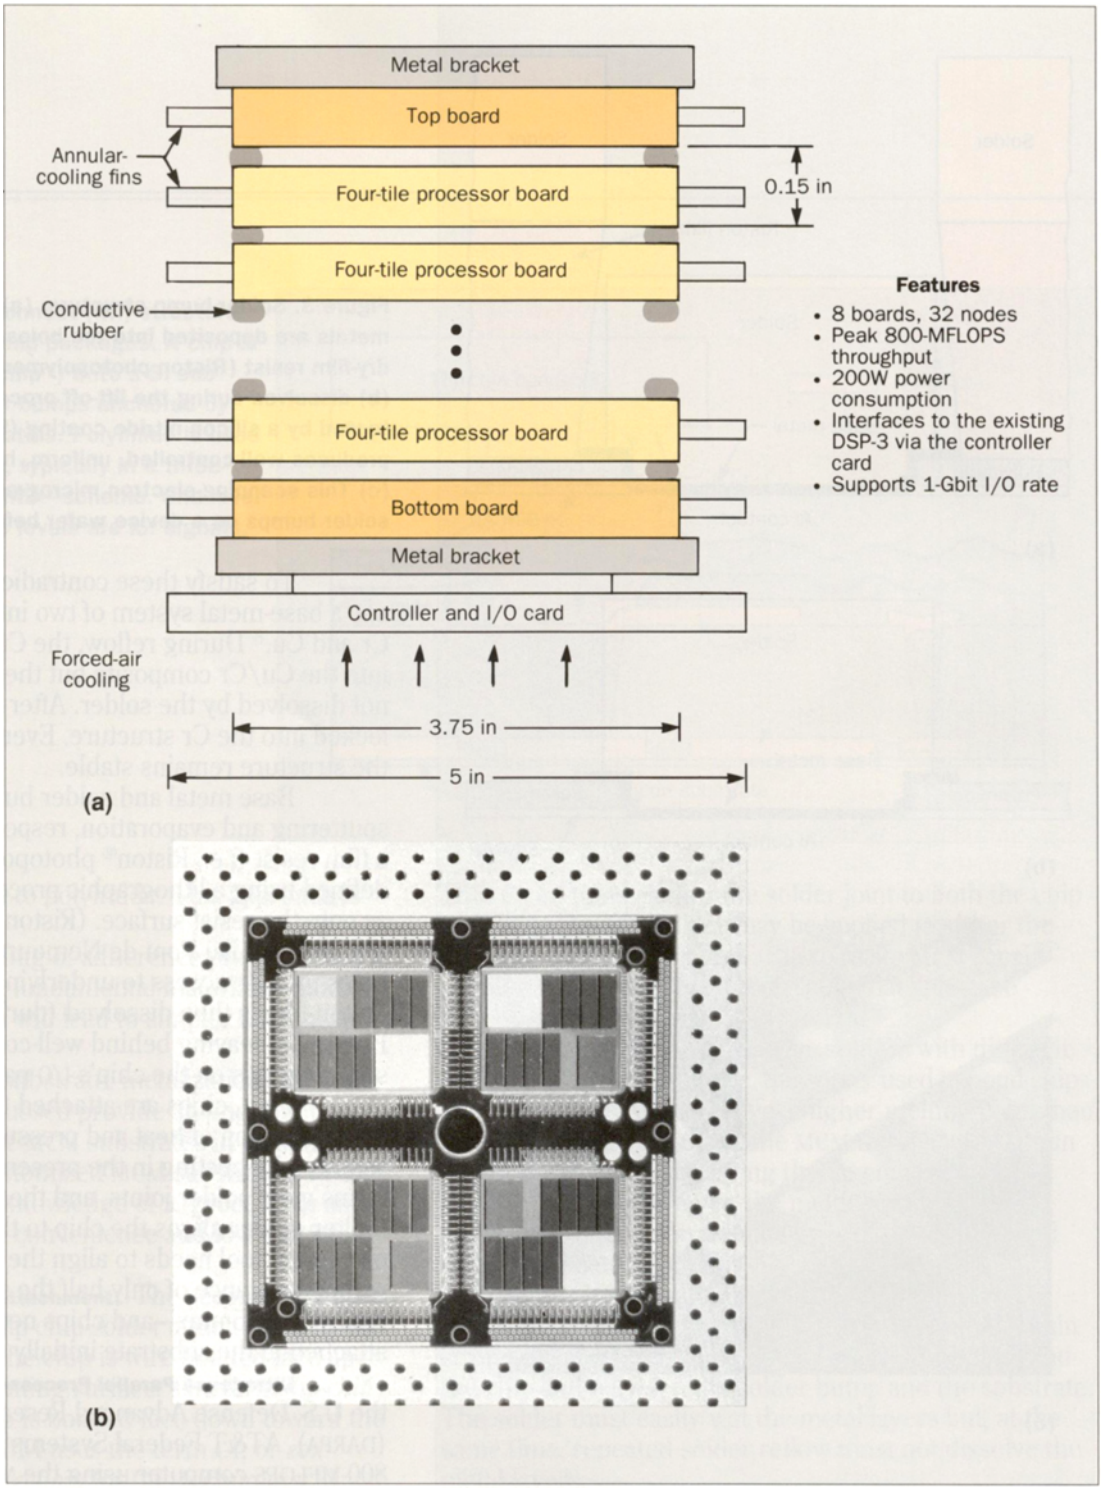

Figure 4. AT&T's ultradense parallel processor is a building block for a TFLOPS computer. (a) By using multichip module (MCM) technology, the ultradense processor will provide 800-MFLOPS throughput in a cubic box of only 100 mm on a side. The system's power consumption will be less than 200W, so a small fan easily cools the device. (b) Each of the eight metal boards in the ultradense processor uses four MCMs, called tiles. Each MCM consists of a DSP, an ASIC, and eight SRAMs that are flip-chip bonded onto a fine-line patterned, Si substrate.

described. (MFLOPS stands for 10^6 floating-point arithmetic operations per second.) This *ultradense parallel processor* will have computing power comparable to a Cray-1 supercomputer but will be housed in a cubic box of only 100 mm on a side. The goal of the project is to develop a three-dimensional, high-density interconnection technology and parallel-processing architecture that will permit construction of a TFLOPS supercomputer. (TFLOPS stands for 10^{12} floating-point arithmetic operations per second.)

The ultradense system will consist of eight metal circuit boards that are interconnected using a conductive polymer (Figure 4a). Four MI MCMs (or tiles) are mounted and interconnected (Figure 4b) on each board. Each MCM consists of:

- Eight static random-access memory chips
- A digital-signal processing chip
- An application-specific integrated circuit (ASIC) to control communication between processors.

Only the ASIC chip is custom manufactured. All the other chips are already manufactured by AT&T or other vendors. Also, processing of the MCM Si substrate is compatible with existing AT&T VSLI production capability.

The MCM technology not only permits the ultradense system to be contained in a small area but also minimizes the system's power requirements to under 200W (watts). Thus, only a small fan will be needed for forced-air cooling.

Optical Interconnections

Most interconnection systems are terminated at either end with electronic devices. Thus, it is not immediately obvious why an optical connection between chips or boards (that require multiple electrical-to-optical and optical-to-electrical conversions) has any advantages.

Current applications for optics on motherboards are limited to two main areas:

- A simple, *point-to-point geometry*. Optics are used to provide isolation between power supplies or to buffer an analog board from a noisy digital board.
- The *free-space backplane*, which most of us own in the guise of a remote control for our VCR and TV. Here, we use light to control more than one board, and the light itself passes through the open air or bounces off the room's walls.

New applications for optics build on the isolation and transmission abilities already in use, as well as fundamental changes now occurring in the electronics industry. These changes include high-speed bottleneck and I/O bottleneck.

High-Speed Bottleneck. Conventional backplanes are limited to about 200 MHz (megahertz) because of capacitive and inductive coupling between lines. Capacitive coupling occurs when two metal lines are close enough to share their electric potential. Inductance (inductive coupling) represents the penalty paid when currents are forced to go around corners. These so-called parasitics cause fast pulses to spread out in time, and increase crosstalk between lines. Light, on the other hand, can pass freely through multiple beams or reflect off corners with no apparent interactions.

Many backplanes serve only to distribute (or broadcast) one signal to many destinations. It is particularly straightforward to split light into multiple paths without distorting the signal shape.

I/O Bottleneck. As integrated circuits increase in complexity and speed, more and more leads are required to satisfy the increased bandwidth and functionality. Even with advanced, high-density electrical interconnects, such as those discussed above, power dissipation and speed are an issue.

For some uses, optical I/Os offer increased levels of performance. This point is particularly pertinent as we learn to grow optically active devices directly on the surface of an IC. In one scenario, off-chip light from a common source might be broadcast to many chips, where local reflectance modulators would then impose their signal on the beam.

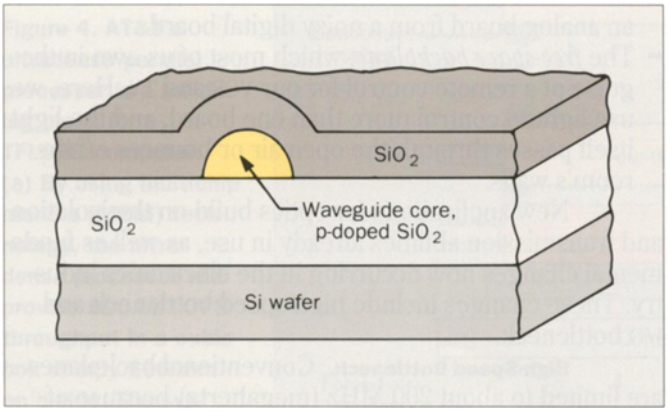

Figure 5. Schematic cross section of glass waveguide on Si wafer for optical interconnection. The transparent central core uses a high-index material and is surrounded by lower index cladding. This combination of indexes confines light to the central core. Thus, the light follows the core that bends and winds around the substrate between chips.

Waveguides. An optical waveguide is a planar version of an optical fiber.[9] As the cross section in Figure 5 shows, the optical waveguide has a transparent central core of a high-index material, surrounded by a lower index cladding. This combination of indexes confines light to the central core. Thus, the light follows the core, as it bends and winds around the substrate between chips.

Researchers at AT&T Bell Laboratories and elsewhere have demonstrated waveguides made from plastics,[10] glass,[11] and semiconductors.[12] We have worked mostly with glasses because of their excellent optical and mechanical properties and their compatibility with conventional semiconductor processes.

To begin, we oxidize a Si wafer in 20 atmospheres of steam to grow 0.015 mm of glass (SiO_2). On top of the silicon oxide (i.e., pure glass), we deposit 0.005 mm of glass using chemical-vapor deposition (CVD). To raise its index by 0.3 percent over that of pure glass, this glass is doped with phosphorus. Then, we photodefine the glass into a waveguide of rectangular cross section. Next, a thermal anneal at 900°C softens the glass, and surface tension reshapes it into a quonset-hut cross section. This reflow step removes small (and optically lossy) roughness on the sides of the waveguide. Finally, we use CVD to deposit undoped glass with the same index as the oxidized Si. This process forms a high-index core surrounded by a low-index cladding, in a way that is functionally similar to an optical fiber.

Such waveguides lose less than 20 percent of the light in a meter of length, and are well matched to conventional optical fibers. Many passive devices (such as optical splitters and combiners, and filters that sort light by wavelength) have been built, tested, and applied in systems by researchers throughout the world.

HOPS (Hybrid Optical Packaging on Silicon)

Almost any material can serve as a support for optical waveguides and circuits. These materials include plastics, glass, ceramics, and semiconductors. Because planar optical-waveguide technology is a new field, no clear winner has yet emerged.

We have elected to use Si as the basis for our work on hybrid optical packaging for several reasons:

- *Good thermal properties, combined with environmental stability and low toxicity.* Solvents do not affect the properties of glass waveguide as they do with plastics. Also, we are free to solder devices to the optical waveguides at high temperatures (about 400°C), which permits us to use a hierarchy of temperatures to seal the package. This high-temperature ability is especially important to help freeze the laser in position with a long-term stability of less than 0.001 mm.
- *Good electrical properties.* By controlling Si conductivity, we have made packages that have bandwidths greater than 10 GHz (gigahertz). Simple device packages have excellent EMI (electromagnetic interference) performance.
- *Optical bench.* The Si-crystal planes have long been used to align optical components.[13] Anisotropic etches

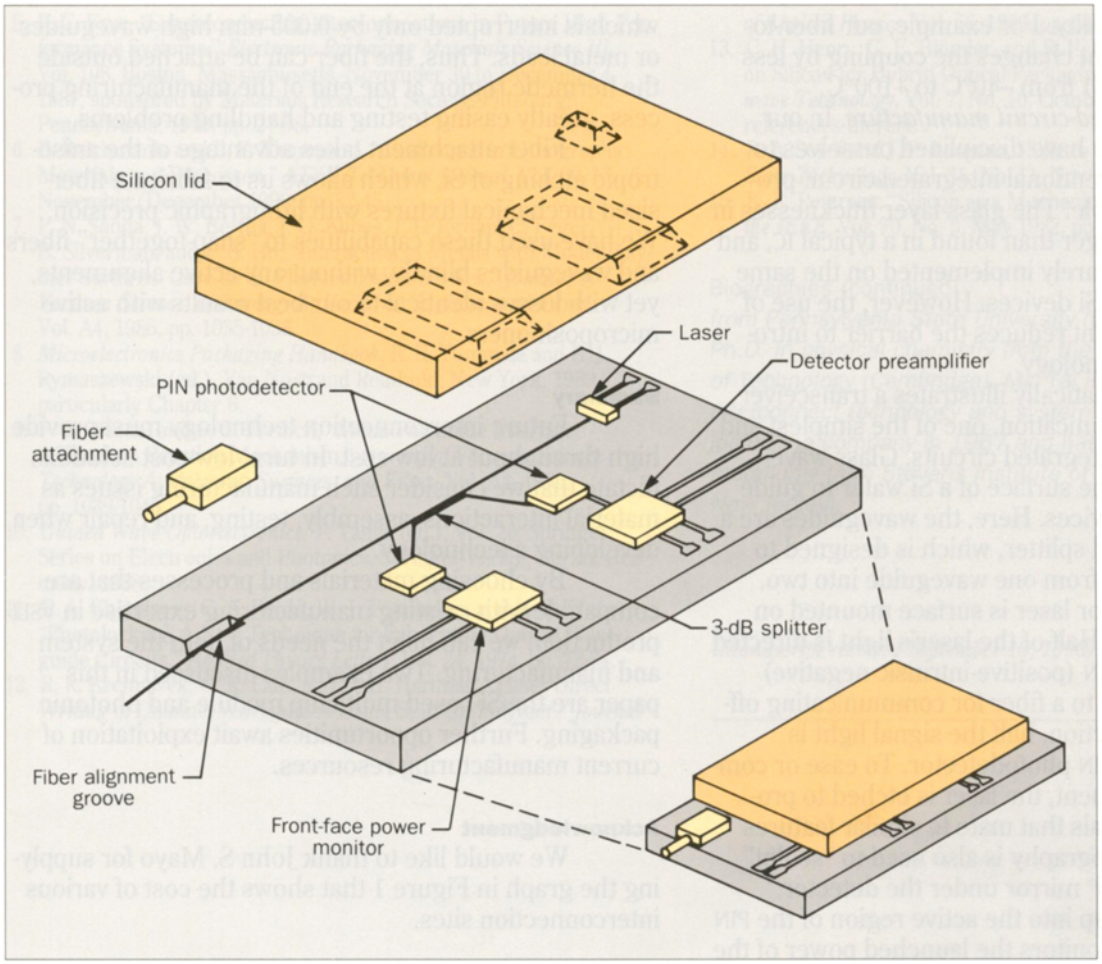

Figure 6. Exploded view of a transceiver made by hybrid assembly on silicon. The glass waveguides formed on the wafer's surface serve as an evanescent field splitter that divides the light equally. Half the light from the semiconductor laser mounted on one arm of the splitter is directed to a surface-mounted PIN photodetector, and half to a fiber for communicating off-chip. The other PIN detector detects half the incoming signal light. A Si lid mounted to the base isolates the surface-mounted devices optically and electrically, hermetically seals the devices, and reduces outside EMI.

allow us to use conventional semiconductor processing to penetrate deeply into a silicon wafer and stop on specific crystallographic planes. Lithographic precision of about ±0.0005 mm is routinely obtained. This length scale compares favorably with the size of single-mode waveguides (about 0.005 mm), photodetectors (about 0.05-mm active regions), and lasers (about 0.0001-mm active layers). Si can be chemically micromachined to hold fibers, lenses, devices, and lids in excellent registration without the need for active-alignment schemes.

- *Good mechanical stability.* Traditionally, optical packages contain many different materials bonded together with many different adhesives. As a result, creep and thermal expansion cause shifts in alignment. With proper design, the Si wafer can align optical compo-

nents with great stability. For example, our fiber-to-waveguide attachment changes the coupling by less than 0.03 dB (decibel) from –40°C to +100°C.

- *Synergy with integrated-circuit manufacture.* In our approach to HOPS, we have disciplined ourselves to leverage off the conventional integrated-circuit process wherever possible. The glass-layer thicknesses in a HOPS device are larger than found in a typical IC, and gold metalization is rarely implemented on the same production line with Si devices. However, the use of commercial equipment reduces the barrier to introducing this new technology.

Figure 6 schematically illustrates a transceiver for bidirectional communication, one of the simplest and most generic, optical integrated circuits. Glass waveguides are formed on the surface of a Si wafer to guide light between active devices. Here, the waveguides are a form of evanescent field splitter, which is designed to divide the light equally from one waveguide into two.

A semiconductor laser is surface mounted on one arm of the splitter. Half of the laser's light is directed to a surface-mounted PIN (positive-intrinsic-negative) photodetector, and half to a fiber for communicating off-chip. In the return direction, half the signal light is detected by the other PIN photodetector. To ease or completely automate alignment, the laser is etched to produce mechanical fiducials that mate to similar features on the waveguide. Lithography is also used to "sculpt" the waveguide into a 45° mirror under the detector. Thus, light is directed up into the active region of the PIN detector, which then monitors the launched power of the laser and corrects for power drift.

All these devices, as well as possible preamplifier and driver circuits, are surface mounted to a Si wafer. Crosstalk can be a problem in a transceiver, so a Si lid is mounted to the base to isolate the devices optically and electrically. This lid also hermetically seals the devices and reduces outside electromagnetic interference.

Unlike the seal in conventional ceramic or metal packages, the hermetic seal here is to a planar surface, which is interrupted only by 0.005-mm high waveguides or metal leads. Thus, the fiber can be attached outside the hermetic region at the end of the manufacturing process, greatly easing testing and handling problems.

Fiber attachment takes advantage of the anisotropic etching of Si, which allows us to fabricate fiber-sized mechanical fixtures with lithographic precision. We have used these capabilities to "snap-together" fibers and waveguides blindly, without any active alignments, yet with losses identical to our best results with active micropositioners.

Summary

Future interconnection technology must provide high throughput at low cost. In turn, low-cost solutions dictate that we consider such manufacturing issues as material interactions, assembly, testing, and repair when developing a technology.

By choosing materials and processes that are compatible with existing manufacturing expertise in VSLI production, we can meet the needs of both the system and manufacturing. Two examples discussed in this paper are the Si-based multichip module and photonic packaging. Further opportunities await exploitation of current manufacturing resources.

Acknowledgment

We would like to thank John S. Mayo for supplying the graph in Figure 1 that shows the cost of various interconnection sites.

References

1. R. Mitchell, "The Genius," *Business Week*, No. 3157, April 30, 1990, pp. 80-88.
2. *AT&T Technical Journal*, Vol. 66, No. 4, July/August 1987, pp. 2-95.
3. R. B. Comizzoli, J. M. Landwehr, and J. D. Sinclair, "Robust Materials and Processes: Key to Reliability," *AT&T Technical Journal*, Vol. 69, No. 6, November/December 1990, pp. 113-128.
4. R. A. Boie, E. R. Wagner, and R. M. Richman, "Capacitance Guided Assembly," *Proceedings of the 1989 International Conference on Robotics and Automation*, Scottsdale, Arizona, May 14 to 19, 1989, IEEE, New York, 1989, pp. 496-502.

5. R. C. Frye, "Superconducting Interconnections in Future High Performance Systems," *Electronic Packaging Materials Science III*, Vol. 108, Boston, Massachusetts, November 30 to December 4, 1987, sponsored by Materials Research Society, Pittsburgh, Pennsylvania, 1988, pp. 27-38.
6. E. Reichmanis and L. F. Thompson, "Challenges in Lithographic Materials and Processes," *AT&T Technical Journal*, Vol. 69, No. 6, November/December 1990, pp. 32-45.
7. P. N. Sanda, J. W. Bertha, J. G. Abbes, J. L. Jordan, C. Feger, B. Silverman, and P. S. Ho, "Interaction of Metals with Model Polymer Surfaces: Core Level Photoemission Studies," *Journal of Vacuum Science Technology A: Vacuum, Surfaces, and Films*, Vol. A4, 1986, pp. 1035-1038.
8. *Microelectronics Packaging Handbook*, R. R. Tummala and E. J. Rymaszewski (ed.), Van Nostrand Reinhold, New York, 1989, particularly Chapter 6.
9. F. S. Ohuchi and S. C. Freilich, "Metal Polyimide Interface: A Titanium Reaction Mechanism," *Journal of Vacuum Science Technology A: Vacuum, Surfaces, and Films*, Vol. A4, 1986, pp. 1039-1045.
10. *Guided Wave Optoelectronics*, T. Tamir (ed.), Vol. 26, Springer Series on Electronics and Photonics, Springer-Verlag, Berlin, Germany, 1988.
11. E. A. Chandross, C. A. Pryde, W. J. Tomlinson, and H. P. Weber, "Photolocking, A New Technique for Fabricating Optical Waveguide Circuits," *Applied Physics Letters*, Vol. 24, 1974, pp. 72-74.
12. R. R. Krchnavek, G. R. Lalk, and D. H. Hartman, "Laser Direct Writing of Channel Waveguides Using Spin-On Polymer," *Journal of Applied Physics*, Vol. 66, 1989, pp. 5156-5160.
13. C. H. Henry, G. E. Blonder, and R. F. Kaarinov, "Glass Waveguides on Silicon for Hybrid Optical Packaging," *IEEE Journal of Lightwave Technology*, Vol. 7, No. 10, October 1989, pp. 1530-1538; and references therein.
14. R. A. Soref and J. P. Lorenzo, "Silicon Guided Wave Optics," *Solid State Technology*, Vol. 31, No. 11, November 1988, pp. 95-98.
15. K. E. Petersen, "Silicon as a Mechanical Material," *Proceedings of the IEEE*, Vol. 70, No. 5, May 1982, pp. 420-457.

Biographies (continued)

from Pennsylvania State University (University Park) and a Ph.D. in physical chemistry from the Massachusetts Institute of Technology (Cambridge). Mr. Tai is responsible for microinterconnect technology and system prototyping research. He joined the company in 1967 and has an M.S. and a Ph.D. in metallurgy from Stevens Institute of Technology (Hoboken, New Jersey).

(Manuscript received September 10, 1990)