INTERCONNECTION SYSTEM REQUIREMENTS AND MODELING

M. Robert Pinnel and Wulf H. Knausenberger

__M. Robert Pinnel__ is head of the Systems Analysis and Technology Planning Department at AT&T Bell Laboratories in Whippany, New Jersey, and __Wulf H. Knausenberger__ is a supervisor in that department. Mr. Pinnel is responsible for trend analysis, long-range planning, and development of advanced applications for interconnection technology and hardware. He joined the company in 1970 and has a B.S. in electrical engineering, an M.S. in metallurgy, and a Ph.D. in materials engineering, all from Drexel University. Mr. Knausenberger is responsible for interconnection system analysis and project planning. He joined the company in 1970 and has a B.S. in engineering science and a Ph.D. in solid-state science, both from Pennsylvania State University.

Rapid technological advances in electronic systems technologies are placing increasingly severe demands on interconnection media. One primary driving force is the evolutionary advance in the scale of integration in silicon with its inherent cost and performance advantages. Another is the revolutionary development of photonics, which is rapidly integrating into most levels of the interconnection hierarchy. Increasingly, the interconnection environment dominates and limits the performance of large-scale electronic systems. This paper explores traditional levels of interconnection from IC chip packages to frames. It shows how the interconnection levels interrelate and must be improved simultaneously to achieve full performance and cost benefits. The first major step in this evolution is well underway with the rapid transition to surface mounting of devices. We describe advanced modeling tools and relate the results of modeling analysis to the challenges faced by interconnection technology hardware and materials.

Perspective

Driven by rapid evolutionary and revolutionary technological advances, the electronics industry is entering a period of significant growth and change. For many years, the dual in-line package (DIP) served the needs for packaging silicon chips. Conventional printed-wiring substrates, connectors, and backplanes have successfully interconnected the various silicon and other components that do the system functions.

However, many believe that the DIP has run out of steam[1] and new integrated circuit (IC) packaging technologies are required. This has implications on each succeeding level of interconnection up to the full-system level, and a large-scale electronic system's performance will be increasingly dominated and limited by its interconnection environment. To sustain the current rate of growth in the performance of future systems, new technology directions will be needed.

This paper examines the various levels of interconnection from IC packages to frames. It shows how the various levels interrelate and how all levels must be improved simultaneously to achieve the full performance and cost benefits that advancing technologies offer. We describe technology trends and system requirements and—using analytical modeling tools—relate them to the challenges placed on interconnection technology hardware and materials in terms of performance characteristics, such as density, speed, and heat dissipation.

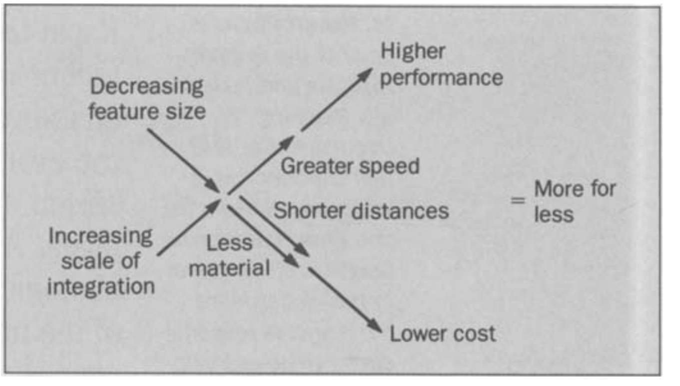

Figure 1. Trends in electronic packaging.

Trends and Technology Drivers

The major trends in electronic systems evolution are clear. The demand is for systems with ever higher performance that do their required functions in less physical space.[2] At the same time, international competition requires that this be done in the most cost-effective way possible. Figure 1 summarizes these key factors and trends.

Practical feature sizes (design rules and grid spacings at all levels) are continually decreasing, coincident with a rapidly increasing scale of integration. The progression of major system-level functionality—from full equipment bays, through single frames or shelves, down to single-board level—exemplifies this trend.[3] Single-wafer or chip-level systems of modest complexity are already being developed. The obvious by-products of this technology direction are greatly increased speed, shorter path lengths, and less material use. The result is a higher performance and lower cost system.

There are three primary technology drivers for electronic packaging: very-large-scale integration (VLSI), photonics, and software control. The most readily apparent is VLSI in silicon, because this driving force has been most influential for the longest time. But photonic signal transmission and the promise of optical signal switching and processing, coupled with ever-increasing use of software control, present two more technologies that are having a major effect. Table I summarizes the impact of these technologies on system-level products in terms of density, system capability (performance), and system complexity. The bottom line is the potential for lower cost per system function, and the historical record of VLSI capability (Figure 2) vividly demonstrates this potential.[4]

Over the last 15 years, the design rules on silicon have decreased by nearly an order of magnitude and are now about 1 μm (micron). Concurrent with this decrease in feature size, the typical density (gates per chip) has increased almost three orders of magnitude, while the cost (cents per gate) has declined almost two orders of magnitude. We do not foresee this trend changing for at least another 10 years and expect similar progress as photonics and software control experience innovation and maturity.

The chip-level trends scale to similar trends at the system or product level. As feature size decreases and memory capacity increases, system complexity and power density increase dramatically. With years of maturity, some portion of the advanced capability trickles down to consumer-grade products as the cost benefits are realized. But for these benefits to be realized, similar innovation and improvements must occur in the entire interconnection hierarchy.

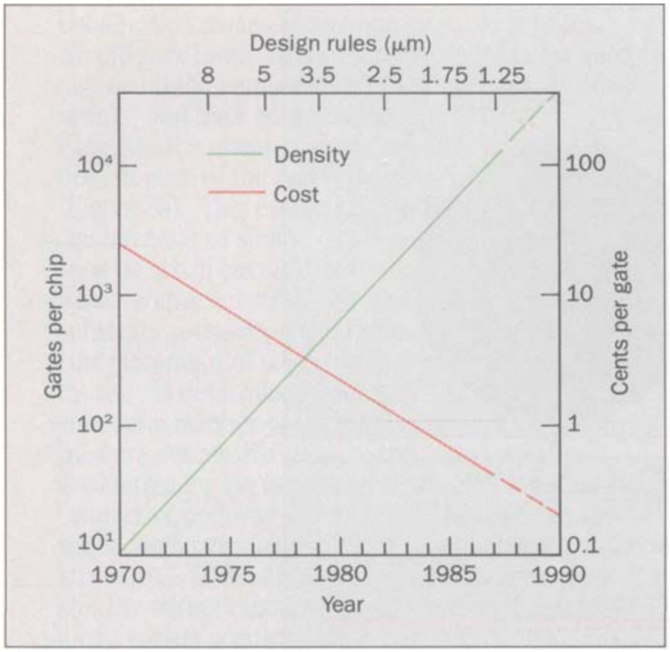

Figure 2. VLSI capability.

Table I. Impact on System Products

	VLSI	Photonics	Software control
Circuit density	↑	↑	—
Interconnect density	↑	↑	—
Memory size and density	↑	↑	↑
Thermal density	↑	↓	—
System capability	↑	↑	↑
System complexity	↑	↑	↑
Cost per function	↓	↓	↓

Interconnection Hierarchy

The interconnection hierarchy is the sequence of levels or layers of hardware that join ICs and other components to form an electronic system. Interconnection is the bridge between the silicon chip and a functional equipment frame.

Figure 3 shows the various interconnection levels that constitute this bridge.

- *Level 0* consists of the gate-to-gate interconnections within the silicon chips.
- *Level 1* is the packaging of silicon chips, typically into DIPs, small-outline ICs (SOICs), or chip carriers (CCs). This level's primary function is to provide electrical fan-out from the chip I/O (input/output) pads to the package leads on larger spacings. In addition, the package provides environmental protection; thermal paths for heat removal; and a more rugged device for easier handling, assembly, and testing. Occasionally, level 1 is eliminated by attaching the chip to the next level using wire bonding, beam-lead attachment, tape automated bonding (TAB), the "flip-chip" approach (controlled collapse bonding), or wafer scale integration.[5] This has been done most commonly on ceramic hybrids but is increasingly used at the circuit-pack level.
- *Level 2* is the circuit-pack level of interconnection. Printed conductor paths connect the device leads of components to the board and to the electrical connector for off-the-board interconnection.
- *Level 3* is the shelf level of packaging. It provides circuit pack to circuit pack interconnection, most commonly using a backplane and connector system.
- *Levels 4 and 5* are the unit and frame levels of interconnection. They provide the backplane to backplane and unit to unit connections, and most commonly use discrete or ribbon cabling.

To assess the required technological advances, we will concentrate on levels 1, 2 and 3, as shown schematically in Figure 4. Innovations provided to meet the needs at these three levels are most critical to supporting the advances in VLSI, photonics, and software control.

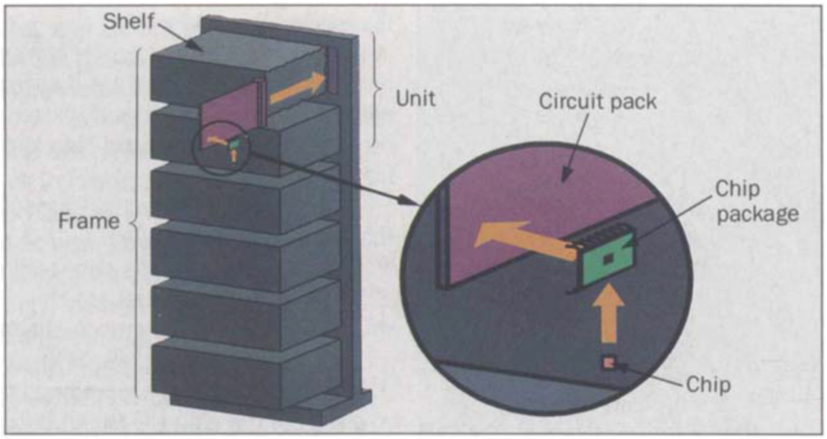

Figure 3. Packaging and interconnecting of electronics in equipment frames.

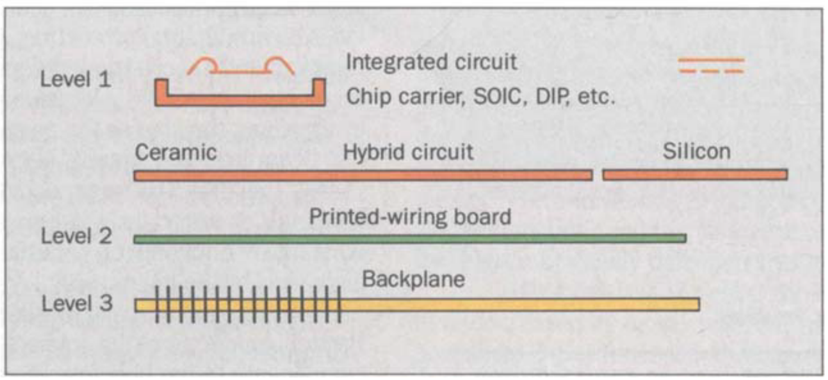

Figure 4. Interconnection levels. DIP = dual in-line package; SOIC = small outline integrated circuit.

System Performance Modeling

Effective design and development of advanced packaging architectures and hardware is simplified or, sometimes, only possible with sophisticated analytic modeling tools.

AT&T has developed an interactive systems analysis model. By applying multiple models, we can assess IC and interconnection-system design choices and architecture in terms of system performance parameters. Current models include routing, electrical, and thermal performance and are integrated into a master program to allow an interactive, coherent, and coordinated analysis. The results of these analyses can provide relative cost and performance benchmarks.

Packaging Density Analysis. The packaging density or routability analysis procedure provides guidance for matching the wiring requirements of the system's circuits to suitable interconnection structures. When this is done early in the design cycle, the results can provide valuable focus to system design activities.

The analysis procedure entails the following steps:

1. Roughly partition the circuit, whose interconnection needs are to be estimated, into manageable entities.

Usually, this means allocating components to the circuit-pack level. This step provides tentative information on the number of components, their I/O lead counts, and their general function.

2. Estimate the circuit's wiring needs (wiring congestion) in each of the two orthogonal wiring directions (Figure 5a). This prediction, which is based on average behavior of similar circuits, should be verified as soon as actual circuit data becomes available.
3. Establish the substrate's wiring capacity for a given substrate technology and design rules and for a specific placement of the circuit's components. Wiring capacity is determined geometrically. Namely, we calculate the number of usable wiring tracks that can traverse successive slices of the substrate after accounting for via usage and obstructions. (A *via* is a conductive pathway through the dielectric that connects conductors on two levels of a multilayer structure.) To illustrate wiring capacity visually, we plot the wiring capacity of each slice (in histogram form) versus position on the substrate (Figure 5b).

To assess the routability of the assumed circuit on the chosen substrate, we contrast the curve for circuit wiring demand to the histogram for substrate wiring capacity. The capacity histogram should exceed the demand curve across the substrate with margin to spare (Figure 6). The amount of margin required is a function of the efficiency of the router used, as well as other factors.

Because the circuit-wiring-demand curve assumes that the circuit in question exhibits average behavior, the actual wiring demand may be significantly different. Accordingly, this procedure is best suited to predicting relative differences between design options and assessing new technology options.

Interconnection Network Simulator. Until recently, the operating frequency of most electronic systems was low enough that we could neglect the influence of the interconnections on system performance. For example, a 10-ns (nanosecond) rise time pulse that travels through a 6-inch printed wiring path typically arrives at its endpoint in about 1 ns, long before the pulse has attained its peak value. The interconnection produces minimal perturbation on the pulse, and we can ignore its presence from a system performance viewpoint.

Contrast that to today's more typical situation. A 1-ns rise time pulse in the same 6-inch printed wiring path attains its full amplitude by the time it arrives at the end of the path. Now the path's electrical performance characteristics become critical to ensuring the signal's integrity at the receiving end. It is now important to design the interconnection network concurrently and interactively with the rest of the system.

The AT&T interconnection-network electrical performance simulator functions as an analytical oscilloscope. With it, we can analyze any arbitrary linear tree-like network for its response to any arbitrary periodic-input signal. The model mathematically calculates the exact response, taking high-frequency (skin) effects into full account, and characterizes a network's performance interconnection by interconnection. The output signal from a gate can be the source signal to simulate the interconnection to the next gate's input.

Thermal Analysis Model. With increasing power densities in electronic systems, the thermal performance of these systems is of major importance. While the traditional approach of experimental modeling provides an accurate picture of thermal performance, it is often time consuming enough to compromise the delivery of a new product. However, redesign of equipment because of thermal problems that were not detected in time also is not acceptable.

The best course to avoid these problems is to have analytical tools that allow the designer to estimate the system's thermal performance early in the design process. These tools can be designed either for ease of use or for generality in handling diverse problems. The approach taken depends on a tool's intended use.

We are developing a tool that builds a thermal network model of the system in question and can successfully model a wide range of physical designs and cooling technol-

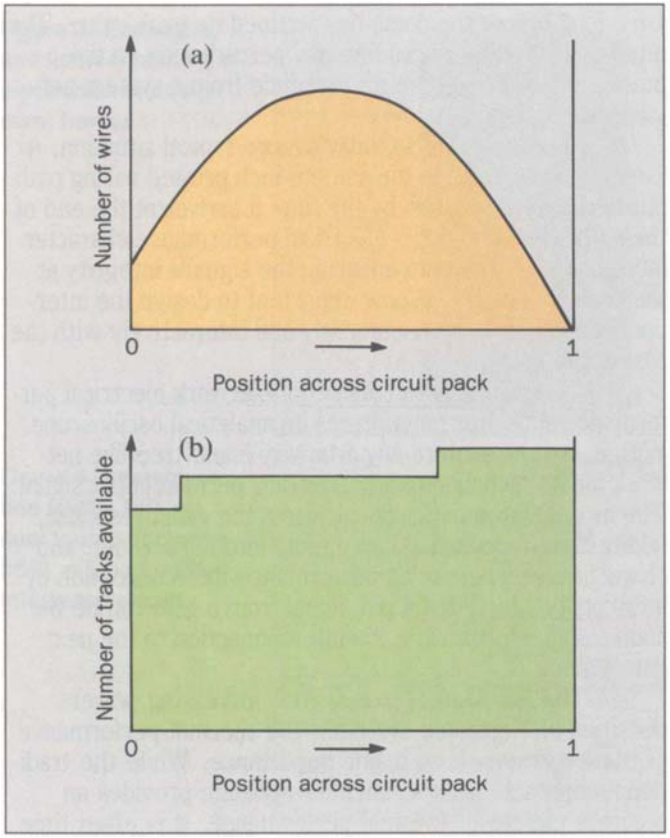

Figure 5. Packaging density. (a) Circuit wiring demand curve, done independently in *x* and *y* directions; and (b) substrate wiring capacity histogram for a particular circuit placement.

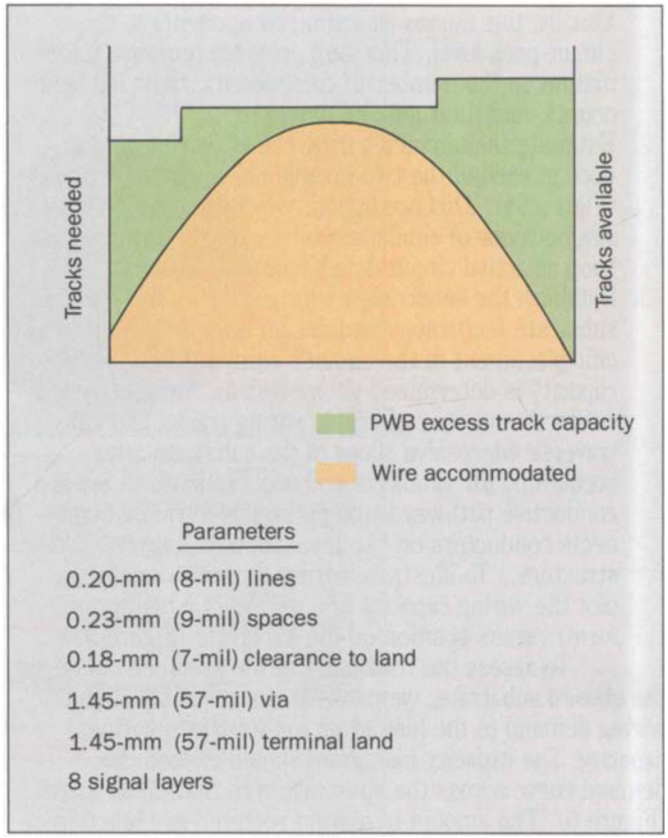

Figure 6. Curve for circuit wiring demand contrasted to histogram of substrate wiring capacity. PWB = printed-wiring board.

ogies. Individual resistances may either be entered by the user or calculated by subprograms for some specific geometries or cooling modes. This extremely flexible approach is ideally suited to the technology planning needs for which this model is designed.

This thermal analysis capability will be useful for assessing designs and concepts before experimental tests are feasible. The model will be well suited to optimizing designs through "what if" analyses. This capability is an essential support tool for developing high-performance packaging schemes, where the interconnection environment becomes a critical factor for handling high-frequency signals, dissipating high-power levels, and providing high-interconnection density in minimum space.

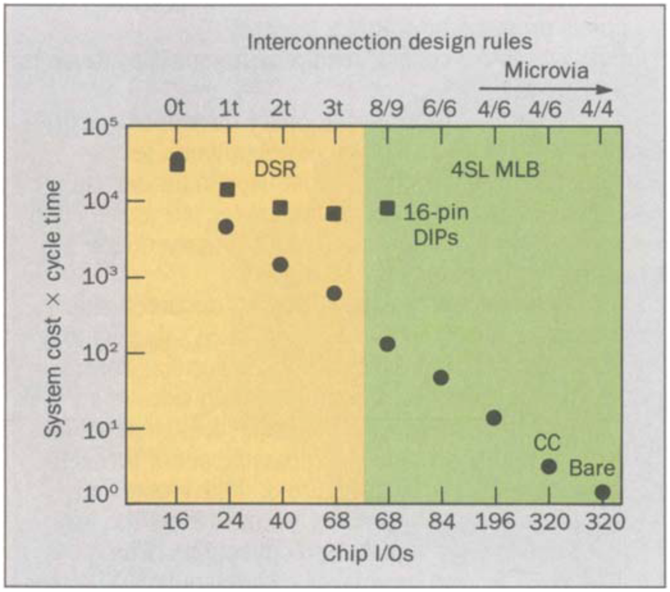

Figure 7. Effect on system cost and performance of interconnection choices for a one-million-gate system. The squares plot the cost-performance metric for a 16-pin dual in-line package (DIP), while the dots represent combinations of packaging and interconnection. CC = chip carrier; DSR = double-sided rigid board; I/O = input/output; 4SL MLB = four signal-layer multilayer board.

Cost and Performance. What is the bottom line that evolves from such forward-looking modeling activity? To obtain the maximum benefit in increased performance and reduced cost offered by VLSI's greater functionality, all levels of interconnection must be enhanced and matched to the silicon chip capability.

Figure 7 shows this for a range of IC scales of integration (with appropriate chip-package styles and I/O-lead counts) and printed-wiring-board (PWB) types to interconnect a one-million-gate random logic system.

To measure interconnection capability, we use the product of system cost and cycle time. The cycle time is a speed-performance measure that accounts for gate delays and propagation delays between gates. For each technology combination, cycle time is based on an average worst-case path length between devices and accounts for interconnection paths through level-3 interconnections.

To determine system cost, we use AT&T's Fastech® electronic packaging system, and add up all hardware costs to implement the one-million-gate system for each technology combination.

The interconnection and packaging technologies that were evaluated included:

- Double-sided rigid (DSR) and four signal-layer multilayer boards (4SL MLBs) as printed wiring substrates
- Design rules of zero-, one-, and two-conductor tracks between plated through holes on 2.54-mm (100-mil) centers on DSR substrates
- Design rules of 0.20-mm (8-mil) lines and 0.23-mm (9-mil) spaces down to 0.10-mm (4-mil) lines and spaces, with either standard vias or microvias for signal-layer interconnection on MLB substrates
- Chip packages from DIPs (16 to 40 I/O leads) through perimeter-lead chip carriers (68 and 84 I/O leads) to area-array packages (196 and 320 I/O leads).

A calculation was also done for an unpackaged (bare) chip attached directly to the PWB.

The upper curve in Figure 7 plots the value of the cost-performance metric, for a fixed package (16-pin DIP) and varied substrate type and design rules. Clearly, only limited improvement can be achieved if one parameter is held fixed. In fact, if we limit chip packaging to 16-pin DIPs, we reach a point of diminishing return where we incur a penalty in enhancing the substrate capability from DSR to 4SL MLB.

The lower curve represents an optimum combination of packaging and interconnection capability to minimize the metric and match the capability offered by ICs available as a function of time over recent history. One can see an impressive five orders of magnitude improve-

Figure 8. **Routability of a surface-mounted chip carrier (CC) versus a dual in-line package (DIP), which requires through holes.**

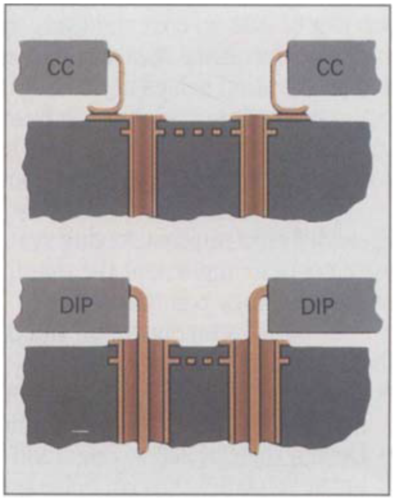

ment in performance using technologies currently in hand or soon envisioned as essential.

Packaging Requirements

Using information from analysis of trends and analytical modeling of system performance needs, one can focus on the necessary technology developments in materials and hardware. From a system-level viewpoint, the following factors are likely to control the directions in which solutions are developed.

- Interconnection levels will be selectively eliminated in specific applications for cost and performance reasons, but the traditional interconnection levels will usually be retained.
- The ever-widening gap between the cost and performance needs of consumer-grade products and high-performance processor systems precludes developing common technology for both market segments. That is, the "trickle-down" philosophy is becoming less valid and system differentiation is required.
- Within a given performance range, standardization of hardware design will become prevalent with feature-rich options provided by software control.
- More aggressive cooling requirements must continue to be met.

Given that these evolutionary directions are driving interconnection technology developments, let us examine more explicitly what must happen for certain materials and packaging requirements for interconnection media. Details on many developments surveyed here are provided in other papers in this issue.[6,7]

Surface-Mount Technology. Surface-mounted chip packaging is a major tool for meeting VLSI's density and performance needs. The package's smaller size contributes, as does component assembly to both sides of a PWB. Also, interconnecting surface-mount devices requires less PWB resources than would a comparable set of through-hole devices, as depicted in Figure 8. This greater efficiency occurs because few holes are needed all the way through the board for signal interconnections. The required via holes can be smaller and, in some PWB structures, localized to just between the two signal layers being interconnected. Thus, the footprint of the via space is available for routing on all other signal layers.[8]

These factors lead naturally to designing structures that are optimized for surface mounting (and are less suitable for through-hole components). Figure 9 shows such a structure. It differs from conventional PWB structures in that layers are formed sequentially, much as on ICs. This then leads to the capability for sequential vias—vias that connect any two or more sequential layers. Thus, PWB space savings come about in two ways!

- Holes can be made with a smaller diameter if no pins are inserted.
- Vias can be placed between just the layers where they are needed, not through the whole board.

Although surface-mount technology is a step forward to meeting density and performance needs, it generates a new set of problems in the areas of assembly technology and reliability. Device placement and solder reflow technologies must be refined to maintain high

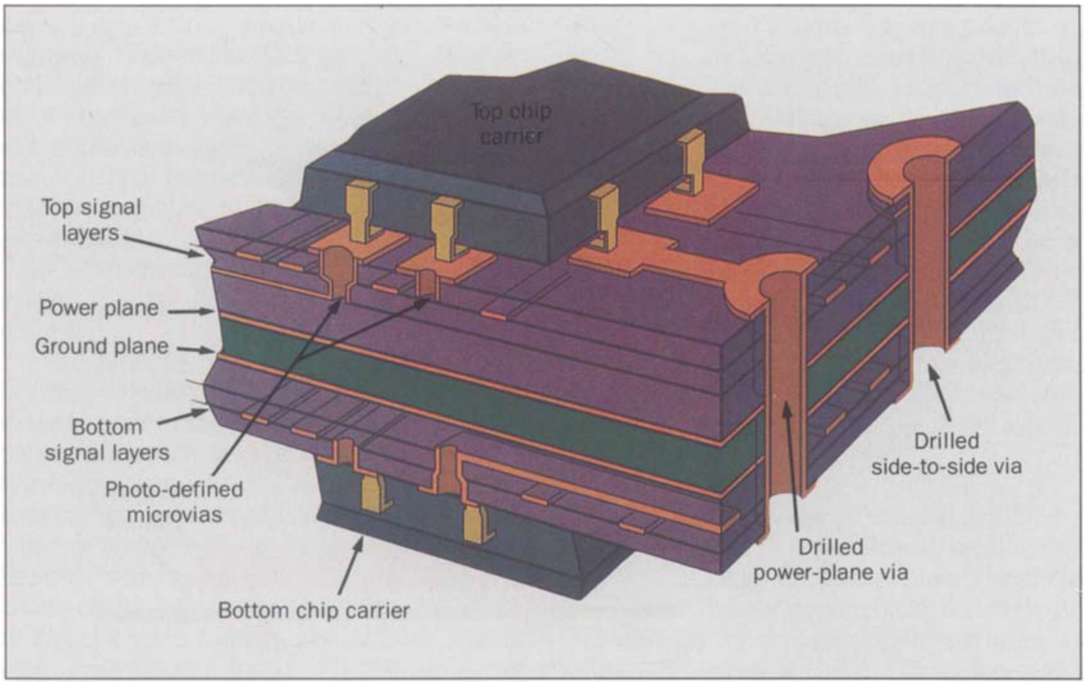

Figure 9. A nonreinforced, photo-defined dielectric, high-density multilayer board. This device is optimized for surface mounting.

assembly yield.[9] Surface-mount passive discretes and some packaged ICs must be improved to survive the shock of a molten solder wave.

Finally, solder is being forced to do mechanical attachment, in addition to its more historical electrical interconnection function. The poor fatigue and creep properties of solder limit its usefulness in surviving large thermal swings, especially when the devices and the substrate on which they are attached have significantly different thermal-expansion coefficients.[10]
This has been a strong motivator for using leaded, rather than leadless, packages to take advantage of the compliance in the copper lead and minimize the strain in the solder joint.

An alternate material system of interest to accomplish this function is conductive polymer films.[11] These composite materials are being developed in both elastomeric and adhesive forms to provide both separable and permanent interconnections at the second and third levels, as depicted in Figure 10. The interconnections consist of a metallic phase embedded in a dielectric polymer matrix and are conducting in one or two directions and insulating in the other orthogonal directions. Thus, they have the potential to duplicate solder's function.

Optical Interconnection. Optics will become part of the interconnection infrastructure for some classes of high-performance systems. This will happen because optical systems have some inherent advantages over electronic systems in very-high-speed and large-bandwidth applications.

Some of the characteristics that give optical systems this advantage are:

Figure 10. Conductive polymer interconnection between a chip carrier (CC) and printed-wiring board (PWB).

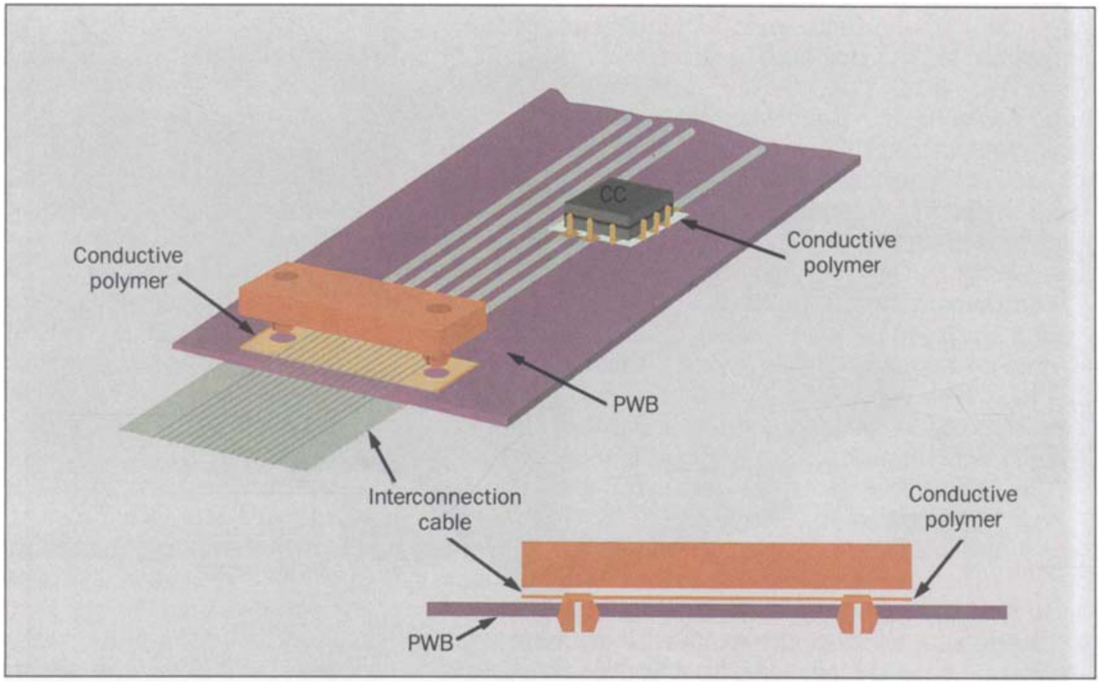

- *Low crosstalk*—Well-designed fiber cables have negligible light leakage from fiber to fiber, even at very high data rates. By contrast, electrical interactions between leads in planar structures become ever greater as data rates increase, to the point that each lead must have a coaxial sheath. This results in a bulky and costly structure.
- *Low attenuation*—A signal propagated through a medium becomes attenuated over distance. In long-haul communications, we counteract this by spacing repeaters or relay stations periodically along the transmission line. For electrical and radio transmission media, repeater spacings are about tens of kilometers. In lightguide systems, repeater spacings can be about 100 km, which translates to a large savings.
- *High bandwidth and low dispersion*—Properly designed optical lightguides maintain the spectral purity of transmitted light very well. Therefore, they can support the transmission of very-high-bandwidth signals (at least ten times better than comparably sized electrical cables). This, coupled with the low attenuation, makes for a cost-effective transmission medium.
- *EMI immunity*—Electronics are placed in many environments that have a high level of electromagnetic interference (EMI). Thus, spurious signals can falsely trigger the circuits. Optical systems can function in such environments with total immunity.

However, to date, optical systems have not penetrated the vast electronics market in a big way. A few examples of penetration into level-3 interconnections to support fiber-optic communications can be found. An optical array connector that effectively brings optical fibers directly onto a circuit pack in the same way as electronic signals is discussed in other papers.[6,12] However, large scale use of optics at circuit-pack levels awaits the development of optical devices, board-level waveguides, and efficient technology for coupling waveguides to devices. If true optical-device developments prove successful and provide the equivalent of "optical silicon" for switching and processing functions, the need for optical to electrical conversion will be circumvented. This will greatly improve system speed and potentially provide a more cost-effective system.

If optical devices are realized, we will need new board-level lightguide materials and photo-definition processes to pattern the lightguides. Electro-optic nonlinear polymer materials are envisioned as a potential solution for board-level waveguides. However, discrete fibers are expected to be the most likely near-term solution to provide this board-level, device-to-device interconnect path. In either case, an effective means to connect the waveguides to the devices will also be needed. This is the analog of a solder joint in the electronic world, and significant development is foreseen to provide as cost-effective a connection process.

High-Performance Electronic Systems. While we await advances in optical capability, near-term advances in interconnection media for very-high-speed electronic systems are essential. The most relevant parameters for substrate and packaging materials are greater dimensional stability, reduced dielectric constant, and suitability for higher operating temperatures.

The dimensional stability is required to maintain feature to feature registration with reduced interconnection dimensions and clearances. Low dielectric constant is critical to minimize propagation delays in gigahertz clock-rate systems. Control of dielectric thickness is also increasingly important to permit enhanced impedance control, a factor that requires both material and processing improvements.

Another approach to high-speed systems has been what is termed wafer-scale integration.[13,14] The technology involves interconnecting silicon ICs directly on a silicon wafer. Although a high-dielectric-constant substrate is generally being used, the system speed is maintained through extremely short interconnection paths.

Two architectures have been evaluated. In one approach, interconnection paths are patterned directly on the silicon wafer on which the transistors are formed. This has not been done successfully to date, primarily because of the extreme thermal density. The second approach has been a hybrid structure in which individual unpackaged IC chips are attached to a silicon substrate using solder-attach technology. AT&T Bell Laboratories' advanced VSLI packaging approach, described elsewhere,[15] is an example of this technology.

The Challenge of the Future

The primary message we have conveyed should be clear. Coordinated evolutionary and revolutionary advancements in interconnection technology are needed on many fronts to support the performance potential of future electronic systems. Otherwise, interconnections will become a bottleneck to maximizing system performance and minimizing system cost. Advances in materials, design, and processing are being made for devices, packaging, interconnection media, and assembly technology to sustain the performance and cost trends.

References

1. R. Bowlby, "The DIP may take its final bows," *IEEE Spectrum*, June 1985, pp. 37-42.
2. M. Anstey, "A 3-D System for Interconnection Meets the Challenge of High Density Packaging," *Electronic Packaging and Production*, Vol. 25, No. 4, April 1985, pp. 84-87.
3. E. Bloch and D. Galage, "Component Progress: Its Effect on High-Speed Computer Architecture and Machine Organization," *IEEE Computer*, April 1978, pp. 64-76.
4. J. D. Meindl, "Ultra-Large Scale Integration," *IEEE Transactions on Electron Devices*, Vol. ED-31, No. 11, November 1984, pp. 1555-1561.
5. E. R. Winkler, "Future Packages' Heat Transfer Will Affect

PCB Design," *Electronic Packaging and Production*, Vol. 25, No. 4, April 1985, pp. 104-109.
6. C. Giffels et al., "Interconnection Media," *AT&T Technical Journal*, Vol. 66, No. 4, July/August 1987, pp. 31-44.
7. S. M. Ambekar, W. E. Hamilton, and T. E. Cole, "Systems Packaging," *AT&T Technical Journal*, Vol. 66, No. 4, July/August 1987, pp. 81-95.
8. W. H. Knausenberger and N. A. Teneketges, "High Pinout IC Packaging and the Density Advantage of Surface Mounting," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, Vol. CHMT-6, No. 3, September 1983, pp. 298-304.
9. T. Dixon, "SMT Calls for a Reexamination of Soldering Techniques," *Electronic Packaging and Production*, Vol. 25, No. 2, February 1985, pp. 96-103.
10. W. Engelmaier, "Functional Cycles and Surface Mounting Attachment Reliability," *Circuit World*, Vol. 25, No. 2, Spring 1985, pp. 61-90.
11. R. Pound, "Conductive Epoxy is Tested for SMT Solder Replacement," *Electronic Packaging and Production*, Vol. 25, No. 2, February 1985, pp. 86-90.
12. W. L. Harrod and W. E. Hamilton, "The FASTECH™ Integrated Packaging System," *Solid State Technology*, Vol. 29, No. 6, June 1986, pp. 107-114.
13. B. J. Donlan et al., "Computer-Aided Design and Fabrication for Wafer Scale Integration," *VLSI Systems Design*, Vol. VI, No. 4, April 1985, pp. 34-42.
14. B. C. Cole, "Wafer-Scale Faces Pessimism," *Electronics Week*, Vol. 58, No. 13, April 1, 1985, pp. 49-53.
15. C. J. Bartlett, "Advanced Packaging for VLSI," *Solid State Technology*, Vol. 29, No. 6, June 1986, pp. 119-123.

(Manuscript received April 1, 1987)