FUNDAMENTAL INTERCONNECTION ISSUES

Mehdi Hatamian, Lawrence A. Hornak, Trevor E. Little, Stuart K. Tewksbury, and Paul Franzon

Mehdi Hatamian *is a member of the technical staff in the VLSI Systems Research department of AT&T Bell Laboratories in Holmdel, New Jersey.* ***Lawrence A. Hornak*** *and* ***Stuart K. Tewksbury*** *are members of the technical staff in the System Architectures Research department of Bell Laboratories in Holmdel.* ***Trevor E. Little*** *is a member of the technical staff in the VLSI Design department of Bell Laboratories in Allentown, Pennsylvania.* ***Paul Franzon*** *is currently affiliated with the University of Adelaide, Australia. Formerly, he was a visiting graduate student in the System Architectures Research department of Bell Laboratories in Holmdel. Mr. Hatamian joined the company in 1982 and is working on high-speed and real-time signal processing. He has a B.S. in electrical engineering from Arya-Mehr* (continued on page 30)

This paper presents several current interconnection issues as well as future technological directions for improved interconnection/communication performance. The physical hierarchy of interconnections and the corresponding communication environment are highlighted. General issues concerning chip-to-chip and on-chip interconnections are reviewed, with particular emphasis on chip or wafer-level clock distribution. Finally, wafer-scale related system modules and optical interconnections are discussed from a "macro-integration" perspective.

Introduction

For more than twenty years, the performance and cost of digital systems have been improved mainly by focusing technology development on individual gate-level and memory bit-cell devices. For silicon metal-oxide semiconductors (MOS), physically smaller devices, with faster intrinsic speeds and lower intrinsic switching and storage energies,[1,2] have generated the current *very-large-scale-integration* (VLSI) era and promise *ultra*-large-scale integration on single, monolithic integrated circuits (ICs). The spectacular advances in semiconductor IC fabrication suggest that conventional systems can be extended indefinitely by developing ever smaller devices, and by mapping physically large systems with many ICs onto smaller realizations using only a few ICs. However, two clear limits suggest that there will be profound changes in achieving higher system performance during the next 10 to 20 years.

A Minimum Device Size for Conventional Logic. Several factors suggest a minimum channel length between 0.25 and 0.5 micrometers (μm) in silicon MOS devices for conventional, deterministic-state logic circuits.[3,4] This is a factor of about three to five smaller than present production devices. [Smaller devices may be applicable in "nondeterministic logic," where the output state of logic devices is only statistically related to the input state. Neural network-style logic circuitry, based on such nondeterministic logic, may extend device dimensions below the 0.25-μm MOSFET limit (metal oxide semiconductor field-effect transistor), a possibility not considered here.]

The approach to minimum dimensions, as reviewed by Meindl,[5] has a sobering implication. Improved system performance, in the era of fixed, minimum device size, must be achieved under conditions of fixed intrinsic device drive and speed capabilities. For higher system performance, we can focus advanced technology development on system-level, multi-chip packaging and interconnection. Improved system-level technologies will both influence and reflect innovation in system architectures. Overall, we suggest the term *macro-integration* to describe the integration of system-focused technology development and architectural innovation, in much the same way that VLSI micro-integration has integrated MOS-device-focused technologies and VLSI-optimized architectures.

Maximum Usable Intrinsic Device Speed. The total delay between devices is the sum of the intrinsic device delay (i.e., unloaded) and the intrinsic line delay (i.e., the delay added because of the interconnection line between devices). As a result, decreasing device delays below typical line delays may have little impact on system speed.[6] Speed-of-light delay is an obvious lower limit on communication delays. More importantly, interconnection lines of electronic systems must be charged and discharged, giving a dynamic switching energy

$$E_{line} \sim \tfrac{1}{2} C_{line} L V_{logic}^2 \qquad (1)$$

proportional to line capacitance per unit line length, C_{line}, line length, L, and the logic voltage swing, V_{logic}. The impact of this "interconnection switching energy," E_{line}, on system performance is similar to the impact of the intrinsic device switching energy, E_{device}, (i.e., power-delay product) on logic function performance.

Given E_{line}, decreasing the line delay requires higher power dissipation drivers, reducing the power available for logic functions. Just as device technologies have tried to reduce E_{device}, so also will system macro-integration try to lower E_{line}, primarily through shorter line lengths or, possibly, through lower voltage swings. Reducing the line energies will provide an increased line density-speed product through which communication limitations may be relaxed.[7]

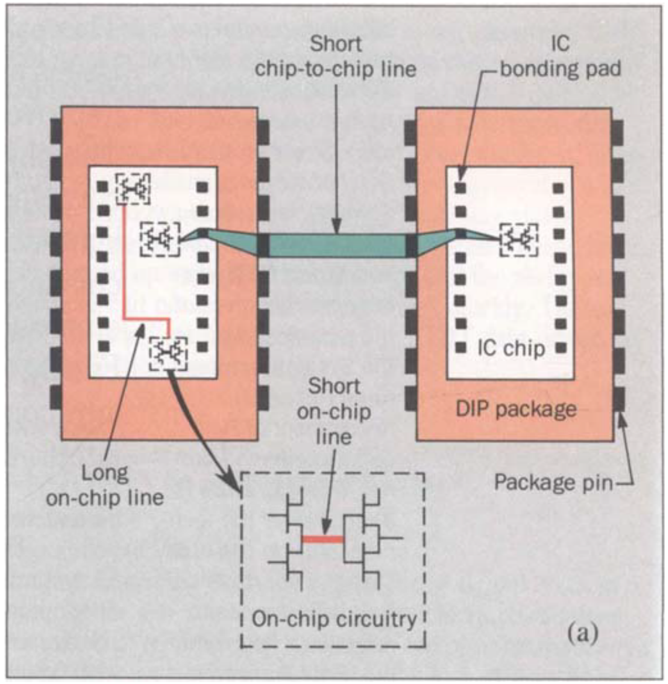

(a)

Conventional Interconnection Environments

We consider here the general physical interconnection and general functional communication environments of a multi-chip, large-scale system. Large-scale systems are emphasized because new technologies map such systems into smaller-scale systems of the future. A clear distinction between functional communication rates on communication links and electronic data rates on wires is critical in assessing the related roles of architectural and technological innovation to achieve higher system performance.

The Physical Interconnection Hierarchy. A conventional large-scale system has packaged ICs mounted on printed wiring boards (PWBs) that are plugged into card cages

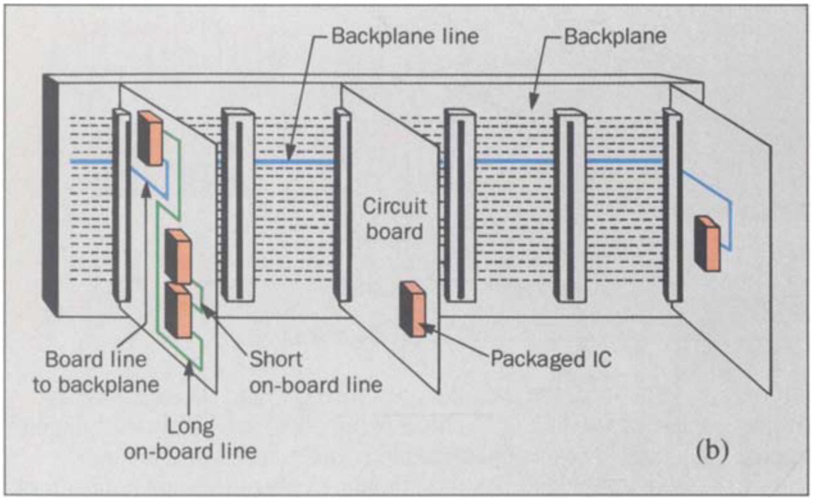

Figure 1. Link segments showing: (a) short and long on-chip interconnections as well as (b) short and long chip-to-chip interconnections, and board-to-board interconnections.

held in equipment racks, frames, etc. Figures 1a and 1b illustrate the wide range of point-to-point interconnection links in a large-scale, silicon MOS system. Each link originates on a minimum-size gate and ends on a minimum-size gate.

The source and destination of an end-to-end information transfer are I/O ports (input-output) of *computation circuitry* (circuitry that changes received information to generate new information). Depending on the number of packaging levels traversed along an end-to-end communication link, signals encounter a variety of connectors, line drivers, and perhaps logic. Here, the line drivers, multiplexing logic, and other circuitry involved in the transfer of information (without alteration) are regarded as *communication circuitry*.

Figure 1(a) shows on-chip interconnections and short on-board (chip-to-chip) interconnections; Figure 1(b) shows chip-to-chip and board-to-board interconnections. In each category, there is generally a distribution of line lengths ranging from short lines between adjacent components (between adjacent gates on ICs or between adjacent ICs on PWBs) to long lines extending across the entire packaging level (across the entire IC or the entire PWB). The numerous short lines are sensitive to density limits on interconnections. The less numerous long interconnections are sensitive to signal delays and line driver requirements.

This physical interconnection environment suggests the simple hierarchical model of physical interconnections shown in Figure 2. Leaf nodes are individual, monolithic digital integrated circuits, taken here to be silicon MOS, VLSI circuits. Arcs and colored ellipses of this tree-based hierarchy represent interconnection links at a given level of system packaging. Within a given level, the set of interconnections is typically highly diverse, allowing a uniform model of communication only when delays and interconnection density do not affect performance. As computation circuitry reaches higher speeds, more uniform interconnections within a given hierarchical level will help to provide a uniform communication environment.

This imposes strong constraints on architectures compatible with such uniform interconnection environments at a given level of system packaging. Black ellipses (interlevel nodes) represent connections between packaging levels. These strongly limit the interconnection lines that can extend between levels and, for very high-speed (e.g., transmission line) signal paths, introduce serious discontinuities in transmission line characteristics. Such communication limits (because of density and bandwidth

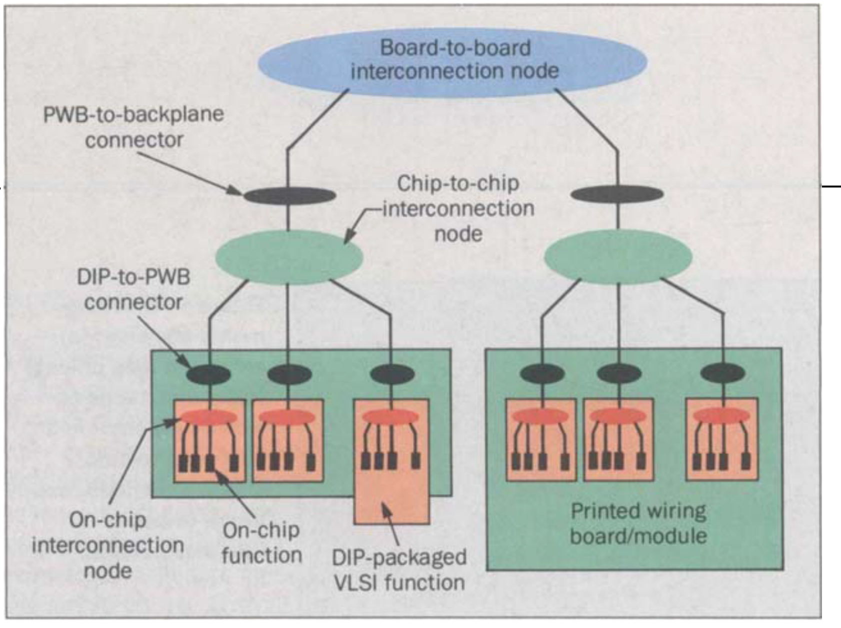

Figure 2. Tree-based, hierarchical model of physical interconnections.

limits of wires at connectors) strongly affect the partitioning of system functions among packaging levels (with the partition boundaries changing as the underlying technologies change).

In addition to passive interconnection wires, a node can include active interconnection circuitry such as drivers or multiplexers. All active electronic circuitry is necessarily at the monolithic (i.e., leaf-level) IC level of packaging. Figure 2 is consistent with this constraint for the computational circuitry but ignores this constraint for the communication circuitry. Moving drivers, multiplexers, etc. to leaf-level nodes causes the communication environment to reflect detailed implementation of the active communication circuitry.

In particular, circuitry such as line drivers and transceivers for board-to-board interconnections and often for chip-to-chip interconnections are typically power- and I/O-limited, small-scale integrated circuitry, in contrast to VLSI functional circuitry. Communication circuitry may then dominate the IC component count of a large-scale system. Communication circuitry technology that can efficiently implement high-density and high-speed communication circuitry and reduce the levels of the interconnection hierarchy will be increasingly important in avoiding the excessive sensitivity of communication environments to such details of communication circuitry implementation.

The basic physical parameters characterizing a physical interconnection link are the number of wires per link (N_l), the data rate per wire (f_l), and the signal delay over the link (τ_l). These parameters are highly interdependent. For example, changing the number of wires per logical link affects the number of wires through connectors and through interconnection circuitry. Increasing the bandwidth along an end-to-end link requires higher bandwidth wiring, higher performance drivers, and faster multiplexers. Therefore, technology advances for improved communications must generally improve the performance at each level of the interconnection hierarchy. However, it is the performance of the *communication functions*, rather than that of an individual interconnection line that is of interest in improving system performance. Higher communication bandwidth can be achieved, for example, by increasing the physical wires per communication link (holding the bandwidth per wire fixed), by increasing the bandwidth of physical wires (holding the number of wires per link fixed), or by a combination of these two approaches. This distinction between communication bandwidth and wire bandwidth is therefore essential in jointly optimizing wire bandwidth and wire density to achieve high-performance communication.

Communication Issues. The physical interconnection network just described is the medium for communication among parts of a system. To illustrate communication limits, we use the communication environment shown in Figure 3, with the physical environment of Figure 2 directly mapped into a communication environment. With the parameters given earlier, the communication bandwidth through links

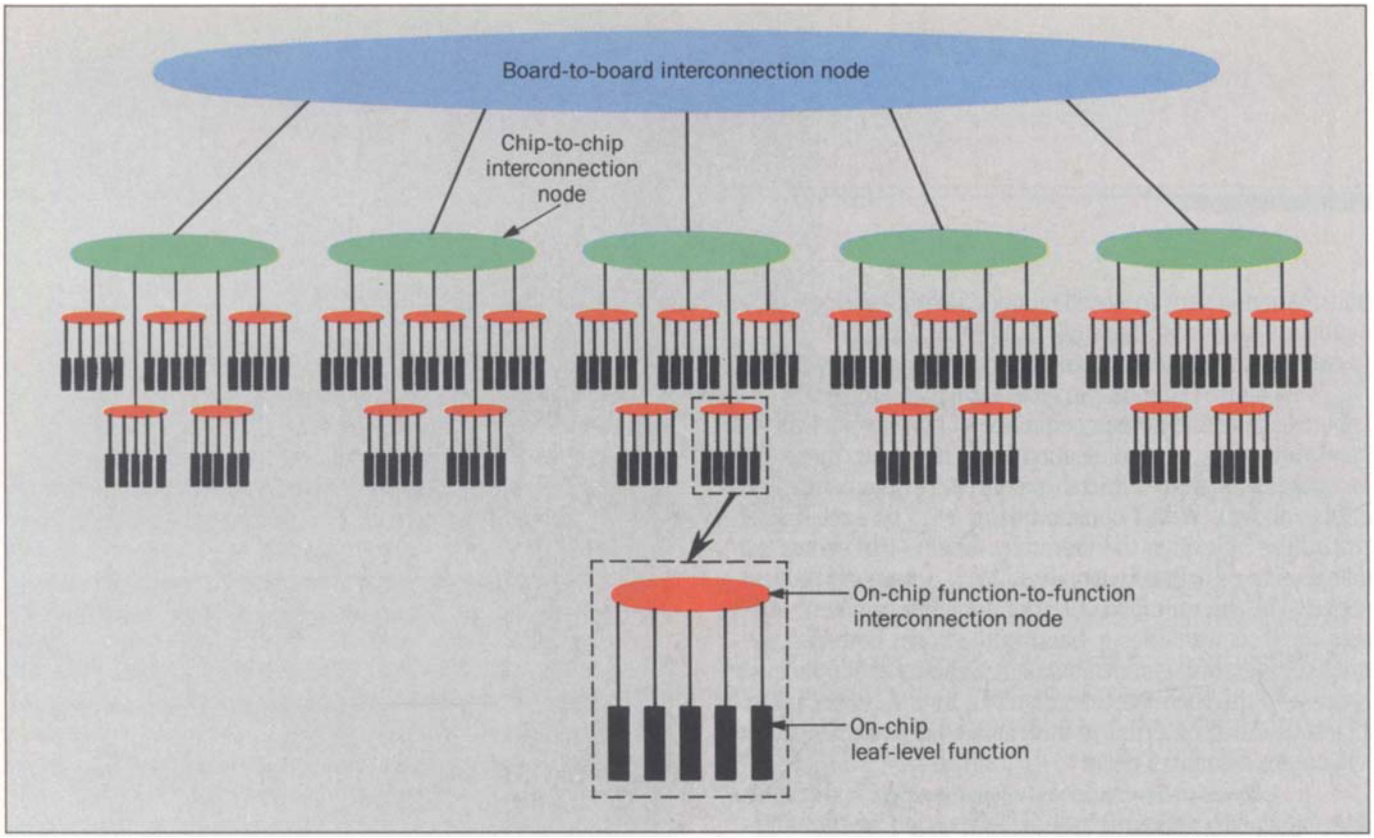

can be represented as $B_l \equiv \Gamma_l N_l f_l$ where $\Gamma_l < 1$ represents the ratio of data information to total information (control plus data) carried on the link. This assumes a uniform communication bandwidth B_{comm} throughout the hierarchy and $B_l \equiv B_{comm}$ for each communication link. Figure 3 shows five functions per chip, five chips per board, and five boards per card cage. Four general classes of communication functions are considered here.

- *Internal point-to-point communications*, as represented by general interconnections among logic/memory function.
- *External-to-internal point-to-point communications*, as seen in loading program RAM (random-access memory) from a central disk storage facility.
- *Internal global communications*, as in a message broadcast by an internal function to all other functions in the system.
- *External-to-internal global communications*, as in externally supplied system clocks.

Internal point-to-point communication. For this example, all leaf-node pairs communicate at equal rates through the communication network of Figure 3 and the communication links are uniformly active. With these assumptions, Figure 4(a) shows the distribution of available bandwidths (normalized to B_{comm}/N_{leaf} with N_{leaf} the number of leaf nodes) from leaf node #0 to the other nodes. This example shows two major consequences of the physical interconnection hierarchy on the communication environment:

- Communication among neighboring leaf cells having the same parent cell is strongly favored.
- The communication bandwidth shows large discontinuities across parent cell boundaries.

Figure 3. Communication limits shown using five functions per chip, five chips per board, and five boards per card cage.

External-to-internal point-to-point communication. This type is limited by the bandwidth of the root node link. With the same external point-to-point communication rate for each leaf node, the external bandwidth to any single leaf node is B_{comm}/N_{leaf}. For large N_{leaf}, as in massively parallel computing systems, the external bandwidth to any single computing node then becomes quite small.

Internal global communication. For an internal global communication event such as "broadcasting," the same information is sent to each leaf node. With leaf node #0 sending a broadcast message to all other leaf nodes (assuming the same uniform communication bandwidth, B_{comm}, as above), the hierarchical structure efficiently distributes the single-source sequence of messages at the full bandwidth, B_{comm}, to all destinations. However, there can be appreciable delay effects (in contrast to the bandwidth effects above). With a constant delay τ_{comm} on each link, regardless of level in the hierarchy, Figure 4(b) shows the distribution of delays (normalized to τ_{comm}) among the leaf nodes. The discontinuous distribution of delays here mirrors the discontinuities in bandwidth above; however, the range of variation is much smaller. Sending the broadcast request to the root and then distributing the broadcast to all leaves (as in external-to-internal global communications) will create a uniform delay.

External-to-internal global communication. External-to-internal global communication includes such system functions as initialization commands and clocks. Broadcasting (e.g., the system clock) from the root node yields a constant delay at the full bandwidth to each of the leaf nodes under the assumptions used here. Also distributed from the root is power and ground. In both cases, there may be changes in the power handling capability of lines moving from the root to the leaves, however basic connectivity is provided by the interconnection hierarchy.

Thus, general, internal point-to-point communications can be limited by the conventional interconnection and packaging hierarchy. Technology directions are highlighted in the 4-layer model of the interconnection environment in Figure 5. Layer 2 contains the functional leaf nodes; layer 3 contains the interconnection circuitry for chip-to-chip communications. All passive chip-to-chip and higher interconnection levels of the system are in layer 4. Layer 1 contains all the interfunction, on-chip interconnections emphasizing their efficient monolithic implementation. Drivers, multiplexers, etc. for on-chip lines are considered part of the layer 2 circuitry. Including all levels of interconnection (i.e., layers 1 through 4) makes an extremely complex and

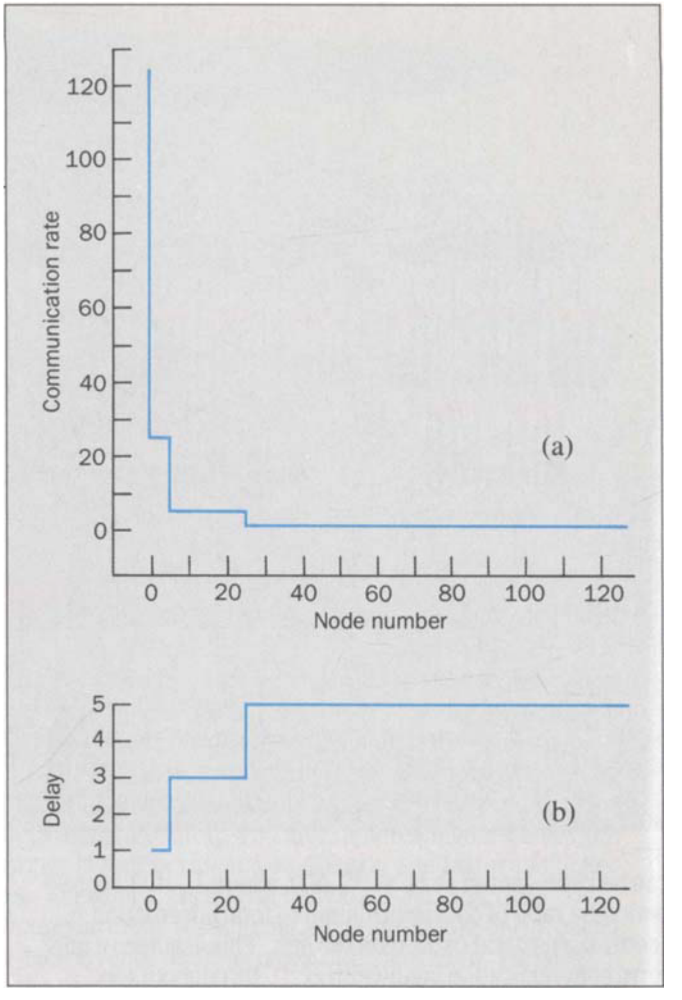

Figure 4. (a) The distribution of available bandwidths from functional leaf node 0 to leaf node J. (b) The distribution of delays (normalized to τ_{comm}) from functional leaf node #0 to node J.

highly heterogeneous environment.

The directions discussed under "New Interconnection-based System Technologies" come from two approaches. The first is to move as much as possible of the layer 4 interconnections onto either layer 1 or the lowest packaging level of layer 4, using wafer-scale, hybrid wafer-scale, or hybrid ceramic assembly techniques. The second is to exploit high-performance interconnection schemes (e.g., using optical interconnections) for the remaining, higher level interconnections in layer 4. However, such

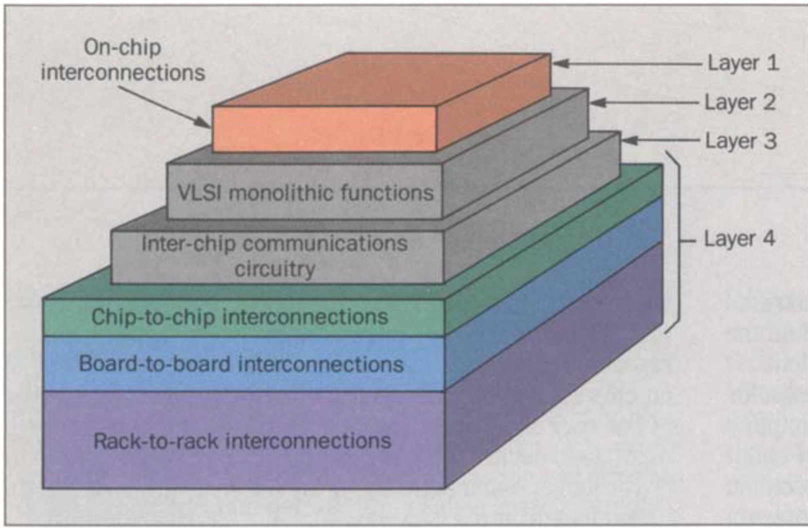

Figure 5. The structure used in Figure 1 modeled into a four-layer structure.

alternatives must be consistent with the underlying physical limits that apply to interconnections, as considered next.

Physical Performance Issues

The physical issues affecting interconnections (communication) are qualitatively similar to those affecting active devices (logic). They include:

- Power dissipation
- Throughput rate
- Delay
- Switching energy or power-delay product
- Density
- Area.

Local Physical Issues. Static power dissipation of interconnections is a major concern when low-resistance loads (termination resistance, R_{term}) are needed. Such terminations are important on transmission lines when signal rise times (τ_{rise}) are less than the signal delay (τ_{delay}) on the interconnection. The maximum signal propagation velocity is $v = c/\sqrt{\epsilon_r}$, where c is the speed of light and ϵ_r is the relative dielectric constant of the insulating material ($\epsilon_r \approx 4$ in the examples below). Lines with length L $(cm) > \tau_{rise}\,(nsec) \bullet (30/\sqrt{\epsilon_r})$ therefore are typically terminated (although resistive transmission lines can relax this constraint).[8,9] Terminated transmission lines impose relatively high-power dissipation requirements. The voltage (V_{logic}) at the termination is decreased by driver resistance (R_{driver}) and line resistance (R_{line}), relative to the supply voltage (V_{dd}), to $V_{logic} \approx V_{dd} \bullet R_{term} / (R_{term} + R_{line} + R_{driver})$. Assume R_{term} is approximately equal to $R_{term} = Z_0/\sqrt{\epsilon_r} \approx 200\Omega$ (where the impedance of free space is $Z_0 \equiv \sqrt{\mu_0/\epsilon_0} = 377\Omega$, and μ_0 and ϵ_0 are the permeability and permittivity of free space, respectively). Taking $V_{logic} = 0.5\,V_{dd}$, then the total power dissipation (driver, line, and termination) is

$$P_{total} \approx P_{driver} \bullet (1 + \frac{R_{term}}{R_{driver}} + \frac{R_{line}}{R_{driver}}) \approx \frac{0.5}{200\Omega} \bullet V_{dd}^2 \quad (2)$$

For $V_{dd} = 5V$, $P_{total} \approx 60$ mW, imposing limits, depending on thermal constraints, on the density of drivers and terminations.

For example, a thermal limit of 1 W / cm^2 would limit the density of driver/receiver pairs to about 16 / cm^2. Decreasing the voltage levels dramatically relaxes this limit (e.g., @ $V_{dd} = 0.5V$, 1600 / cm^2 driver/receiver pairs are allowed by the 1 W / cm^2 thermal limit).

If terminations are not needed (e.g., rise times are longer than signal delays), a power-speed tradeoff is imposed by the the "interconnection switching energy" [E_{line}, defined in Equation (1)]. The energy-per-unit line length used below is $\epsilon_{line} \equiv E_{line} / L$. Again, reducing voltage swings has advantages, dramatically decreasing the interconnection switching energy and allowing higher speeds at the same level of power dissipation. Scaling interconnections to smaller widths does not provide the continual reduction in C_{line} (and thus in ϵ_{line}) that might be expected. As the ratio of line width (w_{line}) to dielectric thickness decreases, fringing fields eventually dominate the interconnection capacitance of an isolated line.[10] As the line spacing decreases, line-to-line capacitance also becomes important. These limit line capacitances, even for $w_{line} \rightarrow 0$, to approximately $C_{line} \geqslant 1$ picofarads/cm (pF/cm). With $V_{logic} = 2.5V$, $\epsilon_{line} \geqslant 3$ picojoules/cm (pJ/cm). Given the small intrinsic switching energy (power-delay product) of the active devices used for logic, the interconnection dynamic energy is a major concern. For example, with device energy $E_{device} \approx 50$ femtojoules (f J), $L \bullet \epsilon_{line} > E_{device}$ for line length $L > 160$ μm @ 2.5V logic swings.

The transmission line delay, τ_{delay}, of signals with $\tau_{rise} \ll \tau_{delay}$ is

$$\tau_{delay} = \left[\frac{\sqrt{\epsilon_r}}{c}\right] \bullet L \qquad (3)$$

for ideal transmission lines.

Maximum bandwidth is generally limited by drivers and receivers (although line dispersion, not considered here, imposes an upper limit on usable bandwidth; that limit decreases with increasing line length). Circuit board interconnection wiring is generally low resistance; transmission line propagation delays provide a simple but reasonable model of interconnection delays. The ideal transmission line has a delay proportional to line length, favoring short lines on circuit boards for small delays and avoiding the need for terminations.

On-chip interconnection lines generally have appreciable resistance per unit length. The signal delay on such lines has the general form

$$\tau_{delay} \approx R_{driver}C_{load} + (R_{driver}C_{line} + R_{line}C_{load})L + \tfrac{1}{2}R_{line}C_{line}L^2 \qquad (4)$$

for lumped RC (resistance capacitance), distributed RC, and capacitive interconnection lines. Actual coefficients of the various terms depend on the model; however, the general dependence on driver resistance, line resistance and capacitance, and load capacitance is about as given.[12] The on-chip lines display three characteristic length regions.

- For very short lines, the delay is constant (equal to the intrinsic device delay, $R_{driver}C_{load}$).
- For longer line length, the delay becomes proportional to line length. If $R_{line}C_{load} \ll R_{driver}C_{line}$ (typical for most on-chip metal interconnections), then the delay term is linear in L and can be considered as resulting from a signal velocity v^*, with delay given by L / v^*. This "effective velocity," assuming $C_{line} = 1\, pF / cm$, is

$$v^* = \frac{Z_0}{R_{driver}} 3 \times 10^6\, cm/sec$$

  a factor of about 10^3 to 10^4 below the speed of light.
- For long lines, the delay increases quadratically with line length, a serious limitation for critical timing paths.

To avoid high static-power dissipation on longer low-resistance lines, signal rise times can be degraded to avoid the need for terminated lines. As the signal rise time is degraded, the effect of propagation delay becomes relatively less severe. Reduced rise times, and the lower data rate, do not necessarily reduce communication bandwidth because several lower bandwidth interconnection lines could be used in place of a single high bandwidth interconnection line to obtain the same net data rate. However, the increase in interconnections must be consistent with the limits on interconnection wires.[13]

Aside from the short interconnections appearing among logic gates within a logic function, interconnections generally connect one planar circuit area with another planar circuit area. Whereas the complexity of a logic function increases with that area, the number of external interconnections to that area typically is proportional to the edge dimension of the circuit area. This implies a decreasing connectivity (I/O pins per logic gate) as the functional logic complexity of a function module (IC chip, PWB, etc.)

increases.

If ρ_{line} is the pitch between adjacent lines, the interconnection density N^*_{int}(lines per cm) is $N^*_{int} = (1/\rho_{line})$. For on-chip interconnections using 1-μm design rules, $\rho_{line} \approx 2.5\mu m$ giving $N^*_{int} \approx 4 \times 10^3$ lines/cm. For circuit board interconnections with 10-mil pitch, $N^*_{int} \approx 40$ lines/cm. This great disparity between on-chip and circuit board interconnection density (for a single layer, although multiple layers could be used in both cases) greatly limits the use of higher density interconnections as an alternative to high data rates on circuit board interconnections.

A further limit is the pitch (ρ_{pad}) of bonding pads of an IC. With $\rho_{pad} \approx 6 mils$, (1 mil = 0.0254 millimeters) $N^*_{int} \approx 70$ lines/cm through an IC edge. The silicon wafer-scale and thin-film ceramic hybrid circuits discussed later would provide much higher interconnection densities (e.g., $N^*_{int} \approx 200$ lines/cm if $\rho_{line} = \rho_{pad} = 2$ mils). Combining the higher interconnection densities with shorter interconnections (to avoid terminations and unnecessary static power dissipation) should support significantly higher communication rates.

Finally, we look at the planar area required for interconnections. As the area needed for interconnection increases, the separation between active components generally becomes larger, increasing interconnection delays and switching energies. Multi-layer metalization for circuit boards, ceramic substrates, and WSI-related modules (wafer-scale integration modules) can greatly reduce the area for interconnections, an important factor in scaling large systems to physically smaller systems.

Selected On-Chip Issues. The performance issues of on-chip interconnections center upon the interconnection function (data, clock/control and power/ground). Power and ground connections extend to each active logical gate in the circuit, imposing particularly severe layout constraints. The power distribution network must be designed to eliminate voltage drops, which would degrade circuit operating margins. Aluminum wires are sensitive to electromigration deterioration at high current levels and must be sized (limiting the maximum current density) to provide the desired long-term reliability. Placing the power distribution network on upper levels of a multi-layer metalization technology allows wide power distribution lines (for small voltage drops and for avoiding electromigration damage) without imposing an excessive area overhead. It also frees the lower metalization levels for the dense, and generally narrower, signal lines. VLSI multi-layer metalization[14] also greatly reduces the routing constraints encountered during circuit layout.

On-chip data interconnections have the advantage of high densities and relatively short length. For silicon MOS, higher power drivers (typically using a cascade arrangement)[15] are necessary to maintain signal rise times on lines of length $L \gg E^{(min)}_{device}/\epsilon_{line}$. Here, $E^{(min)}_{device}$ is the minimum intrinsic switching energy of the logic devices. A long on-chip line also can be divided into a series of shorter segments, each driven by a relatively low power driver,[16] rather than using a single cascade driver. Such distributed drivers are the equivalent (except for the small scale of ICs) of the conventional "repeatered" lines of telephone networks. For very high-speed circuitry, driver power dissipation becomes a major limitation, restricting the density of active circuitry relative to lower speed circuits. In addition, the crosstalk noise on the closely spaced lines increases rapidly with increasing data rates (particularly on busses extending over substantial parts of the IC). This typically requires larger spacings between signal lines at very high signal rates. Such general on-chip interconnection issues have been extensively discussed and are not considered here.

However, a special performance issue impacts very high speed, synchronous systems. In particular, clock skew generally has been viewed as a severe limitation on maximum throughput rate and has been used to advocate optical clocking and other clock distribution schemes. The development of clock skew limits by Hatamian and Cash[17] is highlighted here as an example in which layout and architecture can relieve what is widely viewed as a fundamental limit.

On-chip clock signals must generally be distributed from an I/O pad connection to all flip-flops and latches of the circuit. This "global" and widely distributed signal locally manages the movement of data between registers of the circuit function. While clocking all registers at precisely the same instant is often desirable, it is often neither practical nor necessary. In Figure 6, a data transfer takes place between two modules of a synchronous system. Each module is composed of a computation logic section, CL, and a register, R, clocked by the global clock signal. The following time parameters are defined. t_{pl} and t_{sl} are the propagation delay (from input change to earliest change in output) and the settling time (from input change to stable output) of the computation logic. t_{pr} and t_{sr} are the propagation delay and the settling time of the register. t_{ck} is the clock arrival time at the local module with respect to a global reference, and t_{pi} is the propagation delay of the interconnection between the two modules. The propagation delays are typically much smaller than the settling times. We consider an edge-triggered clocking scheme, although similar results are obtained for other conventional CMOS clocking schemes.

Consider the transfer of data between CL_1 and CL_2 in Figure 6. At time $t = t_0$, the output of CL_1 begins to be loaded into R_1, and at time $t = t_0 + \delta$, the output of CL_2 begins to be loaded into R_2, where δ is the clock skew between the two modules defined as :

$$\delta = t_{ck2} - t_{ck1} \tag{5}$$

Note that clock skew is defined here on the basis of local relative timing, rather than relative to some global timing event. With this local relative definition, δ can be either positive or negative.

At time $t = t_0 + t_{pr} + t_{pi} + t_{pl}$, the response to the change in R_1 has propagated all the way to the input of R_2. For the data transfer to take place properly, this event at $t = t_0 + t_{pr} + t_{pi} + t_{pl}$ must occur after the loading time of R_2 at $t_0 + \delta$. The following condition must then be satisfied:

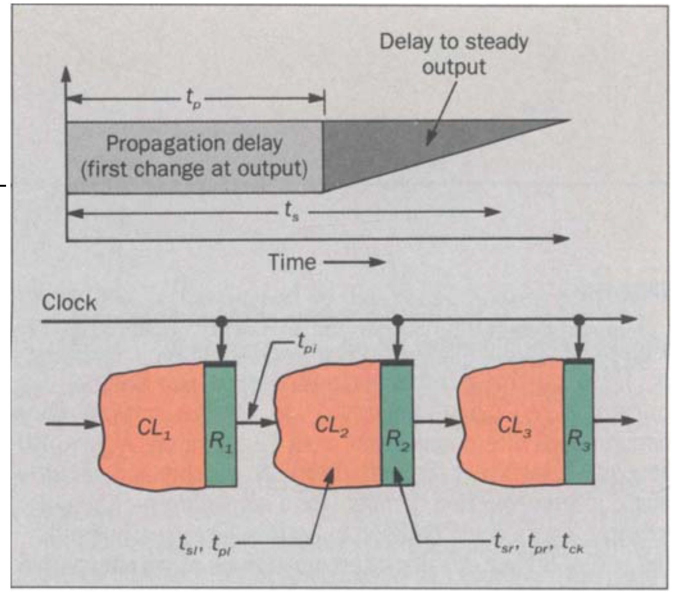

Figure 6. Data transfer between two modules of a synchronous system.

$$\delta < t_{pr} + t_{pi} + t_{pl} \tag{6}$$

In other words, the clock skew must be less than the sum of the propagation delay times of the register, interconnection, and logic. In the case of negative δ, because the propagation delays are always positive, condition (6) is always satisfied.

However, negative clock skew degrades the maximum throughput rate (i.e., the minimum clock period). The clock period T must be large enough to allow for the computation to take place. The total computation (or settling) time is $t_c = t_{sr} + t_{pi} + t_{sl}$. For both positive and negative clock skew, we must then have

$$T > t_c - \delta = t_{sr} + t_{pi} + t_{sl} - \delta \tag{7}$$

These relations (6) and (7) give the following results.

- If clock skew is everywhere negative, then condition (6) is always satisfied and the propagation delays of register, logic, and interconnection cannot cause errors; the system will always operate properly provided the period is large enough. The penalty is loss of throughput because the clock period must be increased by $|\delta|$.
- If clock skew is everywhere positive, then condition (6) must be satisfied to prevent system failure. However,

throughput is improved because the period can be shortened by δ (a case where skew helps the throughput!). Obviously the improvement in throughput as δ is increased is very limited because soon δ violates condition (6).

- If both positive and negative clock skew appear in a synchronous system with feedback, then condition (6) must be satisfied to prevent failure. The throughput rate is also degraded because of the negative skews present. The clock distribution network can still be designed to achieve maxium throughput rates while assuring correct operation.

For a properly structured digital circuit function and a suitable layout, both of data transfers and clock distribution, maximum clock rates can then be achieved, even though clock transition times differ at various points in the circuit. It is the combination of an architecture and layout with an understanding of the presumed limit that relaxes an otherwise severe physical constraint.

Selected Chip-to-Chip Issues. Chip-to-chip interconnections are subject to the interconnection density limit, noted earlier, at the chip edges. In addition, those chip edge boundaries, by imposing pad and package capacitance, introduce an abrupt discontinuity in the interconnection energy that must be overcome to extend the signal to the circuit board. The combination of (1) limited interconnection density from the chip, (2) the energy discontinuity at the chip-to-circuit board interface, (3) the generally longer lines encountered once on the circuit board, (4) the large area penalty imposed by small-scale integrated line drivers on circuit boards and (5) the larger "feature size" of on-board interconnection lines make the chip-to-circuit board interface particularly troublesome.

Perhaps ideally, the circuit board line drivers would be integrated on the IC chip. However, this has two serious limits. First, the high driver power dissipation restricts the number of output pins (and reduces the remaining power dissipation allotment for the functional circuitry). Second, the inductance on power and ground connections from the circuit board to the IC creates the so-called ΔI noise when the output signals are switched.[18] This ΔI noise induces local power and ground voltage glitches that increase quadratically with the output data rate.

Despite serious performance limits, use of conventional IC packaging, such as dual-in-line packages (DIPs), and printed wiring boards has several practical advantages. It allows repair of the circuit board easily at the single IC level, repairing either circuit board wires or replacing defective ICs. Design changes can be achieved by cutting and adding wires (i.e., "white wires"). Wire-wrapped circuit boards can be custom populated with components and interconnections for convenient prototyping before completing the design of a printed wiring board. Finally, the relatively large dimensions of lines and components makes automated assembly (using surface-mount components) relatively straightforward.

The advanced hybrid circuit and WSI-based system technologies forfeit these advantages, at least to some extent. In particular, a dramatic reduction of the scale of off-chip features would make prototyping, assembly, and repair more difficult than in conventional system-level technologies. In defense of reducing the scale of circuit board and higher system-level features, it can be argued that the development of VLSI has had a similar qualitative effect. The loss of some of the customization/repair capability seen by moving lower-scale integrated components into VLSI is compensated by (1) a higher degree of programmability and versatility of the "discrete" VLSI functions as seen in contemporary digital signal processors (DSPs),[19] (2) the development of powerful CAD tools to reduce significantly the design effort[20] (arguably making VLSI design less complex than conventional circuit-board based system design), and (3) a relatively uniform physical environment allowing accurate and relatively straightforward simulation of circuit behavior. Similar developments for future systems are likely to make advanced system-level technologies such as wafer-scale and hybrid ceramic modules more practical and cost effective in the future.

New Interconnection-based Technologies

Emerging interconnection-based technologies are considered briefly here. Hybrid circuit techniques (using either ceramic[21,22] or silicon[23-25] substrates) can reduce the size of multi-chip system modules. *Hybrid wafer-scale integration* (H-WSI)[6] monolithically integrates drivers, multiplexers, and other "glue" circuitry on the substrate, providing an active device layer separate from the VLSI ICs (as in Figure 5 to implement layer 3 of the interconnection model). Monolithic *wafer-scale integration* (WSI[26,27]) provides a potentially greater reduction in interconnection limits, although that potential must be weighed against several practical limitations. The hybrid silicon technologies, as well as high-speed monolithic gallium arsenide (GaAs) technologies[28,29] may be combined with optical interconnections to achieve higher bandwidth communications.

As in conventional telephony networks, such optical interconnections undoubtedly will appear first in the area of "long-distance" interconnections (higher levels of the interconnection hierarchy) but will extend to lower levels as the technologies progress. Also interesting is the possible development of dielectric-separated, multi-level silicon layers for three-dimensional circuits.[30-32] Growth of GaAs on silicon substrates, combined with the lower temperature processing of GaAs circuitry, suggests the possibility of integrating both high-speed GaAs logic and optoelectronics on VLSI silicon wafers.[33-35]

Advanced Ceramic and Silicon Substrate Hybrid Circuits. Hybrid circuit technologies here refer to the mounting of unpackaged ICs on an interconnection substrate. Ceramic substrates and thin-film ceramic technologies are well established. Silicon substrates, which have important advantages, are developing rapidly for hybrid circuit substrate applications. Advanced hybrid circuits could reduce line pitches to less than 1 to 2 mils, a factor of about 10 better than printed wiring boards. There would then be more parallel data communications and less multiplexing, which would, in turn, relax bit rates on I/O lines for many system functions or provide much higher net communication rates for other system functions.

It is likely that the pad pitch can also be reduced significantly, allowing more edge I/O pads per unit chip length. With finer pitch lines, area-based I/O pads used with flip-chip mounting schemes become attractive for high-density I/O lines. Such techniques provide a general scaling down of the size of circuit boards, providing not only narrower lines but also closer spacing of components with a corresponding decrease in line lengths.

Finally, capacitive and inductive loading effects imposed by IC packaging can be considerably reduced by eliminating the first level of packaging. Although the full potential of advanced hybrid packaging will have to be developed, the advantages are already seen in the ceramic substrate packaging approach[21] and the silicon substrate, multi-chip packaging modules.[25]

Passive substrates retain one limitation of conventional packaging, namely inefficient integration of drivers, multiplexers, transceivers, and other communication circuitry. Such small-scale circuit functions might be integrated on VLSI circuits, although power limits and I/O limits are serious constraints. Alternatively, small-scale integrated ICs (such as conventional bipolar "glue circuitry") could be mounted on the hybrid substrate, although the large number of small ICs would waste substrate area and potentially dominate substrate area for very wide data paths. Integration of small-scale circuit functions (for communication circuitry) directly on the silicon substrate provides an attractive evolutionary direction not provided by ceramic substrates. Distributed drivers, pipelined high-speed interconnections, bus control circuitry, multiplexer-demultiplexers, etc. could all potentially be regarded as substrate functions[6] rather than integrated VLSI or separate, discretely mounted, small-scale integrated functions.

Such substrate-integrated communication circuitry, combined with mounted VLSI functional circuitry, is called hybrid wafer-scale integration here, emphasizing that the limitations of WSI circuitry apply to the communications circuitry but not to the VLSI computation circuitry. Beyond simple integration of conventional glue circuitry and mounting of silicon VLSI circuits, the H-WSI approach allows a mixture of different technologies (whereas monolithic WSI, considered below, is limited generally to a single technology). In particular, silicon VLSI memory and logic, silicon bipolar drivers, GaAs high-speed electronic and optoelectronic devices and optical waveguides can be combined in a compact realization.

Beyond the advantages provided by planar hybrid circuits, further performance advantages would be obtained by three-dimensional stacks of hybrid circuits.[6,23,36] The large-aspect ratio (diameter to thickness) of silicon and ceramic substrates may lead to much shorter interconnection lines vertically among stacked H-WSI boards than among similar boards mounted on a planar large-area circuit card. This is illustrated in Figure 7(a), (b), and (c), which suggests a future, high-performance computing node within a mesh-connected array of computing sites.[37] In these examples, the processing unit [and related CPU (central processing unit) circuitry] is shown mounted on a silicon hybrid substrate. The large local memory [including RAM for user memory and EEPROM (electrically erasable programmable read-only memory) for solid-state disk memory] is shown using wafer-scale integration. Connection of the processing node to the network is shown with an intelligent distributed networking switch mounted on a silicon substrate. Special microconnectors and microcables, perhaps using micromachined silicon features for connectors and TAB-like cables (tape automated bonding), would provide very high-density interconnections between adjacent wafer stacks. Such three-dimensional interconnected modules and high-density interconnections are likely to emerge over the next 10 years, providing the very compact systems sought using system macro-integration. Optical vertical connections in such stacked modules are discussed later.

Monolithic Wafer-Scale Integrated Circuits. There are several array-based system functions, including conventional memory ICs, systolic arrays,[38] image processing arrays,[39] and cellular logic arrays. When fabricated on a silicon wafer, the ICs for these functions are organized in a regular spatial array, often mimicking the final array structure of packaged ICs. Wafer-scale integration (WSI) of such arrays merely adds the cell interconnections on the wafer, avoiding the need to separate the individual die. By avoiding the packaging of individual array cells, WSI allows individual cell sizes to range from very small cells to cells comparable to or larger than conventional ICs. WSI implementation of an entire system array function, combined with efficient I/O lines into that monolithic function, therefore achieves a potentially attractive approach to high-performance system integration.

However, fabrication defects and various circuit faults cause some of the WSI circuitry to be defective. Efficient schemes to detect and localize the defective circuitry and either repair or replace defective cells are therefore needed.[40] Such repair and replacement schemes almost invariably degrade the potential advantages of WSI relative to discrete ICs. However, independent of the issues confronting WSI, large-scale systems in general will increasingly require some level of fault tolerance,[41,42] given: (a) the vast number of active devices and interconnections used, (b) the difficulty of achieving full production-level testing of complex system components, and (c) the need for graceful system degradation during in-service failures. There are two general classes of functions from the perspective of fault tolerance: functions requiring replacement of defective circuitry (class 1), and functions in which defective circuitry can be ignored (class 2). High-performance systolic arrays designed for a specific $N \times N$ array function typically require replacement of defective cells and represent class 1 functions. Fabricating an $M \times M$ physical array (with $M > N$), the $N \times N$ array is obtained by connecting a set of N^2 functional cells in the $M \times M$ array. Connections among functional cells can be provided either with physical addition or deletion of links (structuring/restructuring) or with closing or opening of logical switches (configuration/reconfiguration). With low-cell yields, individual functional cells can be selectively connected to a separate interconnection network using either physical structuring or logical configuring of cell-to-

Figure 7. A future mesh-connected multicomputer with (a) a compact, stacked wafer-scale computing module, (b) high-density module-to-module interconnection microcables, and (c) silicon micromachined cable connectors.

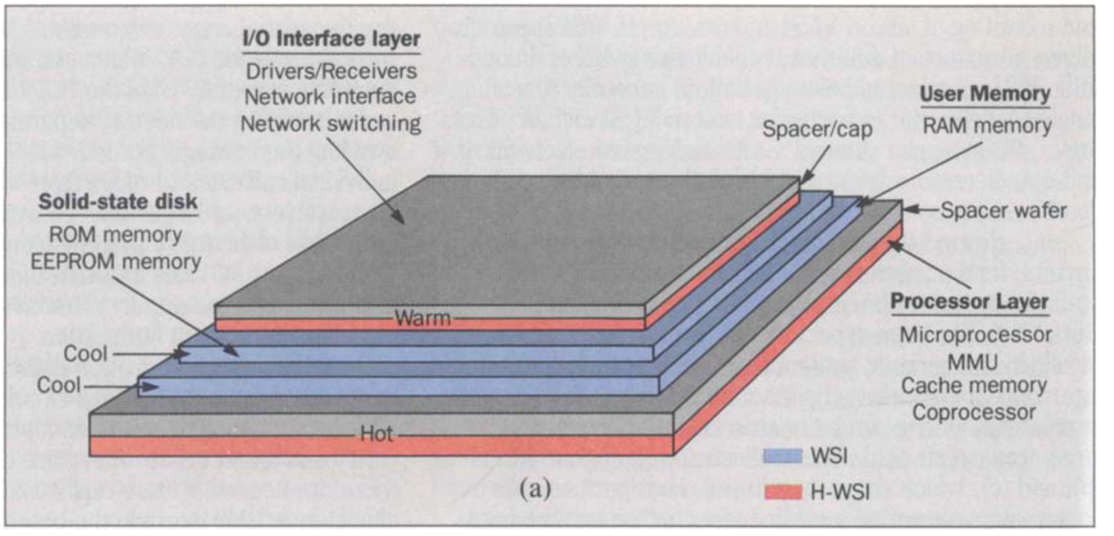

network links. With higher cell yields, fully connected arrays can be fabricated initially and the network modified to bypass faulty cells using either physical restructuring or logical reconfiguration of the network. Other system functions (class 2) can simply ignore defective cells, so long as that information is available to the system. For example, a WSI "disk" memory with defective memory cells can be used by declaring memory blocks bad in the file system.[26] A mesh network connecting a set of general-purpose parallel processors can simply use the network to bypass faulty network switching/computing nodes. Both classes of functions are subject to critical features that must be functional for the wafer to be functional. Global clocks and power/ground nets are examples of such critical features. By providing high-yield critical features and a mechanism to accommodate faulty circuitry, WSI provides efficient integration of compatible system functions.

However, system architectures compatible with WSI must avoid conditions that significantly lower yields. Monolithic circuits mixing high-performance system functions [e.g., high-performance CPU logic, high-density dynamic RAM, or ROM] requiring different device structures and process schedules typically have more complex processing and correspondingly lower yields. Furthermore, the most aggressive IC technologies have low yields at the beginning of the learning curve and those low yields restrict WSI applications. WSI and H-WSI are therefore complementary approaches, with H-WSI implementing system functions not compatible with the yield limits of WSI.

Optical Interconnections. Figure 1 showed various segments (links) of an electrical interconnection path terminating on minimum geometry drivers and receivers (and contained within a locally implemented system). Optical interconnections, extensively used for long-distance telephony lightwave communications, have been extended down to shorter distance, local area networks connecting several locally implemented systems. Further advances in photonic and electro-optical components will extend optical interconnections into even shorter, intrasystem interconnections. Optical data links will undoubtedly appear first at the higher levels of the layer 4 interconnections as shown in Figure 5.

Source and detector arrays, combined with micromachined silicon alignment of fibers to array devices, provide optical fiber "ribbon cables" easily interfaced to the electronics of the system.[43-45] The present high cost of

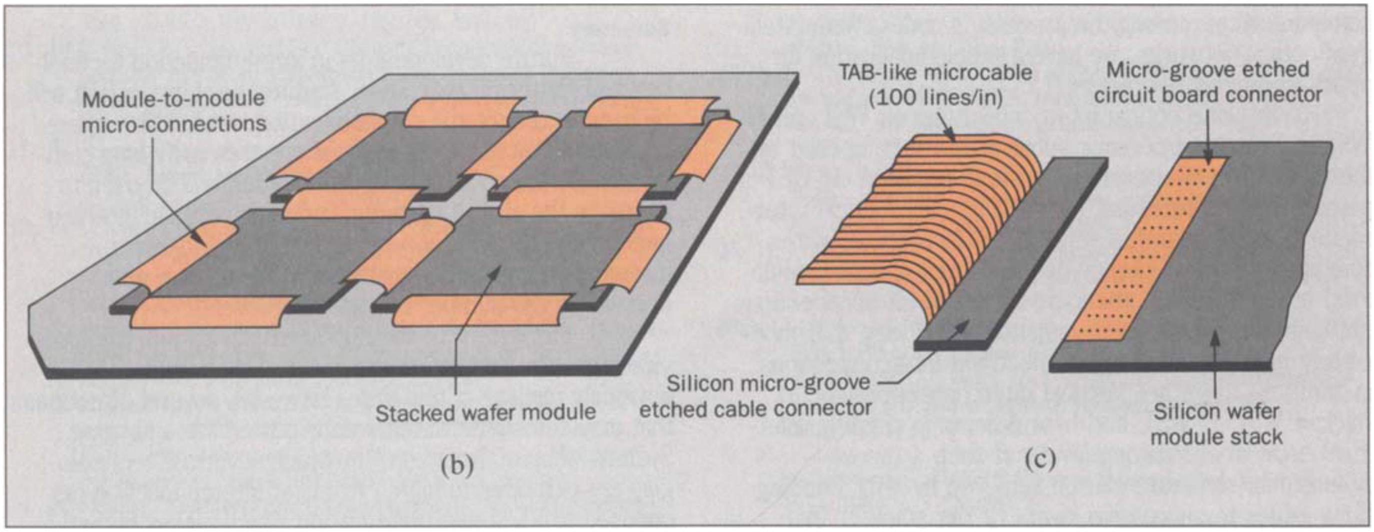

optical data links and their relatively low data rates [≈ 100 to 200 megabits per second (Mb/s)] limit present applications. However, lower cost and higher data rates will evolve over the next decade. In the near term, such higher level optical links are likely to find initial application in widely distributed control signal applications (e.g., to distribute system clocks). By providing electrical isolation, electromagnetic interference (EMI) immunity,[46] good cross-talk performance, and low dispersion (preserving sharp signal transitions), many of the severe limits presently encountered for such signals are greatly relieved. Such global control signals have the advantage of generally requiring only a single (external) optical source whereas general data interconnections would require either distributed sources or widespread distribution of external optical power to modulator "line drivers."

Beyond these applications of "long distance" optical interconnections at higher levels of layer 4 in Figure 5, optical interconnections may be used within macro-integrated systems of the future at lower levels of layer 4.[6,47] Such macro-integration is intended to move much of the longer length layer 4 electrical interconnections to lower levels of the interconnection hierarchy. However, distribution of global control signals such as clocks within those lower levels will remain a problem.

Optical clocks to control clock skew have been widely proposed. The H-WSI and WSI (with integrated or surface-mounted optical detectors) provide attractive environments for such on-wafer optical clocks. Typically, such global optical signals would be externally generated and focused on the detectors (using holographic or other techniques). Optical interconnections also may be useful for interconnection of packaged WSI-based planar and three-dimensional circuit modules. Here, the limits imposed on electrical connections by the package favor extension of the optical link into the package, rather than using separate optical and computational packaged components. A particularly interesting case arises in communication networks for highly parallel computations. If we separate the communication network and its switching nodes from the computing modules, not only optical network links but also optical logic to route data through the network's switching nodes become attractive. The electro-optical conversions are then limited to the relatively slow data rates from each computing node to its associated network switching node and the optical network can provide very high bandwidth communications among the processing nodes. With intelligent optical networks, we have a simple application for optical computing elements.[48]

In-plane optical interconnections on WSI and H-WSI wafers or on ceramic substrates are hampered by the many electro-optical conversions required. H-WSI is perhaps most interesting for such in-plane optical interconnections. Integrated waveguides for optical "wires," combined with low-energy quantum well optical modulators, are particularly attractive if optical interconnections can compete with the convenience, simplicity, and moderately high performance of electrical interconnections. In addition, there are vertical interconnections for stacked WSI, H-WSI, and hybrid ceramic circuits. Electrical connections among layers of such a three-dimensional structure can be achieved by wire bonding along edges to successive layers of the stack or by "through-substrate" electrical connections (using thermally migrated aluminum connections through silicon wafers). However, such electrical connections impose a mechanical connectivity, making assembly, disassembly, or repair difficult. Free-space vertical optical interconnections have the significant advantage of providing vertical communications without permanent bonding of the layers.[6] For H-WSI or WSI (with mounted opto-electronics), infrared, free-space optical links can pass directly through the silicon wafers (i.e., at infrared wavelengths beyond the absorption edge).[49] Integrated planar lenses (i.e., microfabricated zone plates),[50] integrated beam-stops, and other integrated optics structures can be included on the stacked wafers to reduce losses and avoid crosstalk on beam arrays.

All such low-level optical interconnections would suffer from the power dissipation and energy limits associated with integrated optical sources (LEDs or semiconductor lasers). As observed by Miller,[51] locally generating optical power in each optical line driver contrasts with the external supply of power in electronic line drivers. This suggests that optical modulators are probably preferred over integrated LED/laser sources for such low-level optical interconnections.

Summary

Future developments in interconnection technology will center on two areas. Architectural innovation will be used to reduce the communication bandwidth required by system functions and to avoid unnecessarily long communication links. Technology innovation will be used to decrease the size of systems, reducing physical interconnection lengths and thereby achieving high-density interconnections at higher levels of the system and decreasing the number of packaging levels.

The extent to which "macrofabrication" can provide higher performance and lower cost systems is obviously unclear at this time. However, several directions that may contribute to macrofabrication, are emerging. System "macrofabrication," in which advanced technologies are extended to higher levels of interconnection (as opposed to VLSI microintegration which moves higher interconnection levels onto monolithic substrates), will become increasingly important as device scaling limits are approached.

References

1. W. Fichtner et al., "0.15 μm channel length MOSFETs fabricated using e-beam lithography," *International Electron Device Meeting*, 1982, pp. 722-724.
2. H. J. Boll, H. J. Levinstein and G. E. Smith, "Silicon MOSFETs achieve high speed via submicron dimensions," *Microwave System News*, 1982, pp. 75-89.
3. J. R. Brews et al., "Generalized guide for MOSFET miniaturization," *IEEE Electron Device Letters*, Vol. EDL-1(1), 1980, pp. 2-4.
4. A. Reisman, "Device, circuit and technology scaling to micron and submicron dimension," *Proceedings of the IEEE*, Vol. 71(5), 1983, pp. 555-565.
5. J. D. Meindl, "Ultra-large scale integration," *IEEE Transactions on Electron Devices*, Vol. ED-31(1), 1984, pp. 1555-1561.
6. M. Hatamian et al., "FIR filters for high sample rate applications," *IEEE Communications*, Vol. 25(7), 1987, pp. 62–71.
7. H. Kanai, "Low energy LSI and packaging for system performance," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, Vol. CHMT-4(2), 1981, pp. 173-180.
8. R. C. Frye, private communications.
9. R. J. Antinone and G. W. Brown, "The modeling of resistive interconnects for integrated circuits," *IEEE Journal of Solid-*

State Circuits, Vol. SC-18(2), 1983, pp. 200-203.
10. H-T. Yuan, Y-T. Lin, and S-Y. Chiang, "Properties of interconnections on silicon, sapphire and semi-insulating gallium arsenide substrates," *IEEE Journal of Solid-State Circuits*, Vol. SC-17(2), 1982, pp. 269-274.
11. T. Sakurai, "Approximation of wiring delay in MOS LSI," *IEEE Journal of Solid-State Circuits*, Vol. SC-18(4), 1983, pp. 418-425.
12. D. L. Carter and D. F. Guise, "Effect of interconnect on submicron chip performance," *VLSI Design*, Vol. 4(1), 1984, pp. 63-68. See also: D. L. Carter and D. F. Guise, "Analysis of signal propagation delays and chip-level performance due to onchip interconnections," *Proceedings of the 1983 IEEE International Conference on Computer Design: VLSI in Computers*, IEEE Press, New York, 1983, pp. 218-221.
13. W. H. Knausenberger and N. A. Teneketges, "High pinout IC packaging and the density advantages of surface mounting," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, Vol. CHMT-6(3), 1983, pp. 298-304.
14. A. Sabnis, "Multilevel metalization schemes," in *VLSI Electronics: Microstructure Science*, Vol. 15, edited by N. G. Einsbruch, Academic Press, New York, 1987, pp. 293-343.
15. C. Mead and M. Rem, "Minimum propagation delays in VLSI," *IEEE Journal of Solid-State Circuits*, Vol. SC-17(4), 1982, pp. 773-775.
16. H. T. Bokoglu and J. D. Meindl, "Optimum interconnections for VLSI," *IEEE Transactions on Electron Devices*, Vol. ED-32(5), 1985, pp. 903-909.
17. M. Hatamian and G. L. Cash, "High speed signal processing, pipelining and VLSI," *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 1986, pp. 1173-1176.
18. L. W. Schaper and D. I. Amey, "Improved electrical performance required for future packaging," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, Vol CHMT-6(3), 1983, pp. 283-289.
19. W. P. Hayes et al., "A 32-bit VLSI digital signal processor," *IEEE Journal of Solid-State Circuits*, Vol. SC-20(5), 1985, pp. 998-1004.
20. B. D. Ackland et al., "Cadre - a system of cooperating VLSI design experts," *Proceedings of the International Conference on Computer Design: VLSI in Computers*, October 1985, pp. 99-104.
21. C. C. Shiflett et al., "High-density multilayer hybrid circuits made with polymer insulating layers (POLYHIC)," *Proceedings of the 1986 International Symposium on Microelectronics*, International Society for Hybrid Microelectronics, October 6-8, 1986, Reston, Virginia, pp. 481-486.
22. B. Swartz, "Review of multilayer ceramics for microelectronics packaging," *Journal of Physics and Chemistry of Solids*, Vol. 45(10), 1984, pp. 1051-1068.
23. S. K. Tewksbury et al., "Chip alignment templates for multichip module assembly," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, Vol. CHMT-10(1), 1987, pp. 111-121.
24. R. K. Spielberger et al., "Silicon-on-silicon packaging," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, Vol. CHMT-7(2), 1984, pp. 193-196.
25. H. J. Levinstein et al., "Multichip packaging technique for VLSI based systems," *Digest of Papers: IEEE 1987 International Solid-State Circuits Conference*, 1987, pp. 224-225.
26. C. Jesshope and Will Moore, eds., *Wafer Scale Integration*, Adam Hilger, Bristol, England, 1986.
27. J.I. Raffel et al., "A wafer-scale integrator using restructurable VLSI," *IEEE Transactions on Electron Devices*, Vol. ED-32(2), 1985, pp. 479-486.
28. P. O'Connor et al., "A monolithic multigigabit/sec DCFL GaAs decision circuit," *IEEE Electron Device Letters*, Vol. ELD-5(7), 1984, pp. 226-228.
29. L. E. Larsen, J. F. Jensen, and P. T. Greiling, "GaAs high speed digital IC technology: an overview," *IEEE Computer*, Vol. 19 (10), 1986, pp. 21-28.
30. H. W. Lam, "Beam recrystallized silicon-on-insulator devices," in *Electron-Beam Solid Interactions and Materials Processing*, edited by Appleton and Celler, Elsevier, New York, 1982, pp. 471-482.
31. M. W. Geis et al., "Zone-melting recrystalization of Si films with a movable-strip-heater oven," *Journal of the Electrochemical Society*, Vol. 129, 1982, pp. 2812-2818.
32. M. A. Bosch, D. Herbst, and S. K. Tewksbury, "Influence of light on NMOS transistors in laser micro-zone crystallized layers," *Microelectronics Journal*, Vol. 14(6), 1983, pp. 82-87.
33. M. A. Awal and E. H. Lee, *Materials Research Society Symposium A: Heteroepitaxy in Silicon*, Palo Alto California, April 15-18, 1986.
34. J. C. Bean, "Silicon based semiconductor heterostructures," in *Silicon Molecular Beam Epitaxy*, edited by E. Kasper and J. C. Bean, CRC Press, Boca Raton, Florida, 1987.
35. H. K. Choi et al., "Monolithic integration of GaAs/AlGaAs double-heterostructure LED's and Si MOSFET's," *IEEE Electron Device Letters*, Vol. EDL-7 (9), 1986, pp 500-502.
36. J. Grinberg, R. G. R. Mudd, and R. D. Etchells, "A cellular VLSI architecture," *IEEE Computer*, January 1984, pp. 69-81.
37. S. K. Tewksbury, L.A. Hornak, and P. Franzon, "Future physical environments and concurrent computation," *1987 Princeton Workshop on Algorithm, Architecture and Technology Issues for Models of Concurrent Computation*, September 30 - October

1, 1987, Princeton University (to be published).
38. S. K. Rao, T. Citron, and T. Kailath, "Mesh-connected processor arrays for the transitive closure problem," *Proceedings of the 24th Conference on Decision and Control*, December 1985, pp. 1565-1570.
39. M. Hatamian and G. L. Cash, "A 70-MHz 8-bit x 8-bit parallel pipelined multiplier in 2.5 micron CMOS," *IEEE Journal of Solid-State Circuits*, Vol. SC-21, 1986, pp. 505-513.
40. P. Franzon, "Interconnect strategies for fault-tolerant 2D VLSI arrays," *Proceedings of the International Conference on Computer Design: VLSI in Computers*, 1986, pp. 230-234. See also, P. Franzon, "Yield modeling for fault tolerant VLSI arrays," in *Systolic Arrays*, edited by W. Moore, A. McCabe, and R. Urquhart, Adam Hilger Press, Bristol, England, 1987, pp. 207-216.
41. V. P. Kumar and S. M. Reddy, "A fault-tolerant technique for shuffle-exchange multistage networks," *IEEE Computer*, Vol.20(6), 1987, pp. 30-40.
42. A. T. Dahbura, K. K. Sabnani, and L. L. King, "The comparison approach to multiprocessor fault diagnosis," *IEEE Transactions on Computers*, Vol. C-36(3), 1987, pp. 373-378.
43. P. P. Deimel et al., "Individually addressable monolithic 1 × 12 light-emitting diode array," *IEEE Journal of Lightwave Technology*, Vol. LT-3(5), 1985, pp. 988-991.
44. M. G. Brown et al., "Fully optically and electrically interfaced, monolithic 1 × 12 array of $In_{0.53}Ga_{0.47}As/InP$ p-i-n photodiodes," *International Electron Device Meeting*, 1984, pp. 727-728.
45. C. M. Miller, "Fiber-optic array splicing with etched silicon chips," *The Bell System Technical Journal*, Vol. 57, Part 1, 1978, pp. 75-90.
46. C. M. Lin and D. L. Carter, "Photonic I/O's at the PWB and chip level?" *Proceedings of the 6th Annual International Packaging Conference*, San Diego, California, November 17-19, 1986, pp. 374-381.
47. W. J. Goodman et al., "Optical interconnection for VLSI systems," *Proceedings of the IEEE*, Vol. 72(7), 1984, pp. 850-866.
48. A. Huang, "Architectural considerations involved in the design of an optical computer," *Proceedings of the IEEE*, Vol. 72(2), 1984, pp. 780-786.
49. L. A. Hornak et al., "Through-wafer optical interconnects for multi-wafer wafer-scale-integrated architectures," *SPIE 30th Symposium on Optical and Optoelectronic Applied Sciences and Engineering*, San Diego, California, August 17-22, 1986.
50. L. A. Hornak and S. K. Tewksbury, "On the feasibility of through-wafer optical interconnects for hybrid wafer-scale-integrated architectures," *IEEE Transactions on Electron Devices*, Vol. ED-34(7), 1987, pp. 1557-1563.
51. D. A. B. Miller et al., "Novel hybrid optically bistable switch: the quantum well self-electro-optic effect device," *Applied Physics Letters*, Vol. 45(1), July 1984, pp. 13-15.

Biographies (continued)

University of Technology, Tehran, Iran, and an M.S. and Ph.D. in electrical engineering from the University of Michigan at Ann Arbor. Mr. Hornak joined the company in 1982 and is working on optimized optical and electrical interconnection structures for hybrid and full wafer-scale-integrated multi-wafer computer architectures. He has a B.S. in physics from the State University of New York at Binghamton and an M.E. in electrical engineering from the Stevens Institute of Technology. Mr. Little joined the company in 1983. He has a B.Sc. in engineering science from Durham University, the United Kingdom, and an M.Sc. in digital techniques from Heriot-Watt University, the United Kingdom. Mr. Tewksbury joined the company in 1969 and is working on system integration and interconnection issues. He has a B.S. and a Ph.D. in physics from the University of Rochester. Mr. Franzon is completing his doctorate at the University of Adelaide, where he received his B.S. in physics and applied mathematics and his B.E. in electrical engineering.

(Manuscript received May 7, 1987)