Modeling and characterization of long on-chip interconnections for high-performance microprocessors

by A. Deutsch
G. V. Kopcsay
C. W. Surovic
B. J. Rubin
L. M. Terman
R. P. Dunne, Jr.
T. A. Gallo
R. H. Dennard

Long on-chip interconnections with dimensions larger than the minimum ground rules are rigorously analyzed and experimentally characterized for the first time. A test vehicle has been built and characterized with representative wiring found in high-performance CMOS microprocessor chips (line lengths of 0.8–1.6 cm, and line widths of 0.9–4.8 μm using a five-metal-layer structure). The need for distributed *RLC* transmission-line representation is highlighted through measured and simulated results. By showing the problems encountered when using long, nonuniform on-chip transmission lines, guidelines are developed to take advantage of the lower-resistance interconnections for use in high-speed, cycle-determining paths. Such guidelines are given both for current and optimized wiring practices and for cross-sectional structures.

Introduction

Continued advances in integrated circuit technology are making it possible to build large microprocessor chips that contain both the CPU and cache circuitry. The largest chips today are already approaching 20 mm on a side, and operate at clock rates up to 300 MHz [1]. Complex microprocessor chips consist of four to nine million transistors fabricated on a single die. To interconnect this large number of devices, typical chips include three to five conductor layers with high-density wiring having less than 1-μm line width and spacing. Critical interconnections such as clock lines, control lines, and data lines between processor and cache may traverse up to one chip edge, and thus be 1–2 cm in length. Delay on such lines is likely to limit the achievement of even higher microprocessor clock frequencies.

The above requirements result in lines with very small cross-sectional dimensions that have very high resistance, R, ranging from 35 to 500 Ω/cm. The 0.5-μm CMOS devices used to build 2–3-ns-cycle-time processors have

0018-8646/95/$3.00 © 1995 IBM

initial signal transitions of 100–200 ps. Our study has shown that even a 1-cm-long line can have a propagation delay of 130–370 ps. The propagated signal rise time varies from t_r = 100–900 ps. Thus, the propagation delay is comparable to the rise time. Stated differently, the line length is comparable to the signal wavelength, λ, which is of the order of 0.8 to 2.5 cm ($\lambda \cong vt_r$). This implies that transmission-line properties have to be taken into account [2]. To achieve short delays on such long lines, on-chip low-impedance drivers must be used. A long line no longer appears in the traditional way [3] as only a simple capacitive load. Because of the long lengths and high resistance, even a lumped *RC*-circuit representation is no longer applicable. In some cases, the complete electromagnetic wave properties for these transmission lines [4] must be taken into account, and the terminal voltages and currents must be determined using the distributed R, L, C, and G line parameters. Respectively, R, L, C, and G are the line resistance, inductance, capacitance, and dielectric conductance per unit length. The nonideal current return path through the on-chip power and ground reference buses results in increased inductance and resistance for these interconnections. It is shown that such lines have exponentially increasing propagation delay with line length. Moreover, the high wiring density and fast signal-switching speeds result in increased crosstalk [5], and large delay tolerance is caused by adjacent line coupling.

This study is focused on representative long, on-chip wiring that provides the global interconnection between functional logic units. These lines have not previously been analyzed in great depth, since they represent a small portion of the total wiring demand. However, the performance of such interconnections is crucial in critical paths. The paper is organized as follows. Modeling and experimental results are shown from specially designed test structures that address each of the relevant performance issues for lines that are either 0.8 or 1.6 cm long, and have 0.9-, 2.7-, or 4.8-μm widths. First the structures are described, and then detailed models are developed on the basis of rigorous three-dimensional line parameter extraction; lossy, coupled transmission-line representation is used for signal simulation. The models are then validated by comparing measured and simulated signal propagation delay, rise time, and crosstalk. The models are also used to evaluate the accuracy of a simplified *RC*-circuit calculation of delay and crosstalk. We explore chip wiring improvements such as the use of some less resistive wiring for long lines, careful placement of power buses, and restrictions on coupled line length. We conclude with simulated examples shown with representative driver circuit speeds and output impedance levels.

Test vehicle design

• *Long on-chip wiring characteristics*

Any two uniform parallel conductors that are used to transmit electromagnetic energy can be considered transmission lines [4]. In ideal cases, both the electric (E) and magnetic (H) field vectors are perpendicular to the direction of propagation and the solutions to Maxwell's equations for E and H are TEM (transverse electromagnetic mode) waves. A rigorous analysis of such lines must solve for the voltage and current waves. Since most digital pulse applications are concerned with signals at the two ends of the transmission lines, terminal voltages (V) and currents (I) are generally determined. Both V and I can be calculated using either distributed-circuit parameter or discrete, lumped-parameter representation. Both V and I are expressed in terms of the characteristic line parameters per unit length, R, L, C, and G. It is shown in [4] that for ideal lines the circuit representation is consistent with the rigorous analysis based on Maxwell's equations. For lossless lines $R = G = 0$, $Z_0 = \sqrt{L/C}$, and propagation delay $\tau = \sqrt{LC} = \sqrt{\varepsilon\mu}$. The TEM approximation is equivalent to a condition in which the line cross section is much smaller than the wavelength λ. In practical wiring configurations, the conductors have a finite conductivity, the dielectric around the conductors can be inhomogenous (air and SiO_2, for example), the line capacitance is increased by the orthogonal wiring layers and the presence of the silicon substrate, and the inductance and resistance are increased by the nonuniform power bus layout. Such effects give rise to E or H field components along the direction of propagation. If the line cross section and the extent of these nonuniformities remain a small fraction of the wavelength for the frequency range of interest, the solutions to Maxwell's equations are still essentially TEM waves, or so-called quasi-TEM. For the structures described in this study, the wavelength is 0.8 to 2.5 cm ($\lambda \cong vt_r$). Even for a five-layer wiring structure, the vertical height is only around 10 μm, and the largest separation between a signal line and a power bus is of the order of 170 μm; therefore, the quasi-TEM solution is still valid. Approximate representations of the lossy, quasi-TEM lines are easily handled by circuit simulator programs such as SPICE [6] or ASTAP [7].

For very short lines ($l \ll \lambda$) such as those used between logic gates, and for slow devices, a single *RC* section model has generally been used [3] to characterize on-chip interconnects. As the line lengths and circuit speed increase, i.e., as the wavelength decreases, a distributed N-section *RC* representation should be used for accurate delay predictions, with each subsection being a small fraction of the wavelength. As R, ω ($= 2\pi/t_r$), and l

increase, the inductive component of the impedance becomes relevant, and the interconnections are treated as transmission lines with distributed R, L, C, and G parameters. As speeds continue to increase, the frequency dependence of these parameters may have to be included as well, mostly as $R(f)$ and $L(f)$. **Table 1** shows an example for the simulated response of a 1.6-cm-long line having R = 174.0 Ω/cm and C = 1.83 pF/cm. The line is being driven by a 10-Ω impedance driver having 2.5-V swing and 100-ps initial rise time. The lines were designed to have 2.7-μm width and spacing and to have an orthogonal-bus reference that is described later. The predicted signal delay and rise time are compared for lossy transmission-line models, for complete RLC models without frequency dependence, or constant R and L, and for various distributed π-section RC representations. Although the rise-time discrepancy is quite large, a ten-section RC circuit gives a fairly good delay prediction, and the error becomes significant only for fewer subsections. Delay is measured at the midpoint of the signal. The crosstalk at the near end of an adjacent quiet line (NEN) has significant error for the RC-circuit simulation case. It is shown that the R, L, C representation becomes increasingly necessary as the lines are placed further from their reference power bus.

Unlike most packaging interconnections [5], on-chip wiring has line resistance Rl which ranges from being comparable to the characteristic impedance, Z_0, to much greater than Z_0. It is shown in [5] that for $Rl \leq 2Z_0$, the lines behave like an LC circuit with fast-rising response. As Rl exceeds $2Z_0$, the rise time is slowed, becoming similar to the response of a slow RC circuit. The delay on such lines increases exponentially with length. The lines in this study have Rl in the range of $0.6Z_0$–$10Z_0$. **Figure 1** shows typical response for R = 35.5 Ω/cm, Z_0 = 32.0 Ω, and lengths l = 0.5–2.0 cm. The driver circuit output has 10-Ω impedance and 100-ps rise time. The lines are open-ended, and the transition from LC-type behavior on short lines to slow RC-line behavior on long lines is evident. As the rise-time slowdown caused by the RC-like behavior of the upper portion of the signal approaches the 50% signal level, which is the receiver circuit switching threshold, the delay increases substantially. **Figure 2** highlights this effect by showing the delay for lines having R = 35.5, 167.0, and 501.5 Ω/cm.

Most FET circuits driving the short wires between logic gates have an output impedance Z_D which is much higher than the line impedance Z_0, and the interconnections appear as lumped capacitive loads. High-impedance devices initially inject only part of the available signal voltage swing into the line $[Z_0/(Z_0 + Z_D)V_{DD}]$, and steady-state levels (generally V_{DD} or 0 V for CMOS circuits) are reached only after several round-trip delays from the open end where doubling occurs. The long interconnections

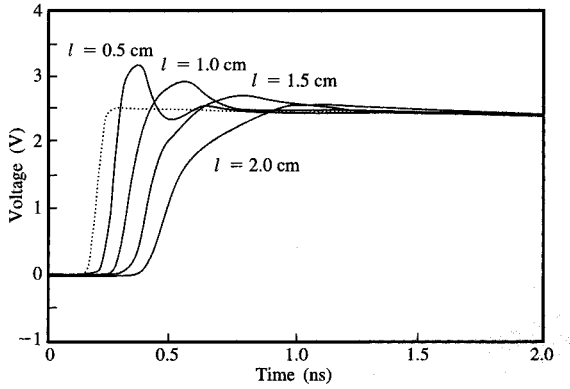

Figure 1

Signal propagation simulated for 100-ps-rise-time, 10-Ω source (dotted curve) on line lengths l = 0.5–2.0 cm having the design of site 3B with a parallel nearby power bus on M5 and R = 35.5 Ω/cm.

Table 1 Simulated delay, rise time, and *NEN* crosstalk for site 2A with orthogonal ground-bus return. A distributed RC representation with 1, 3, 5, 10, or 50 subsections is compared to transmission-line modeling with constant RLC or $R(f)L(f)C$. Lines are 2.7 μm wide and 1.6 cm long on M4 layer, with R = 174.0 Ω/cm, C = 1.83 pF/cm, driven with a 10-Ω driver, with 100-ps rise time and 2.5-V swing. The differences Δ are with respect to the $R(f)L(f)C$ case.

Circuit	*Delay* (ps)	Δ (%)	*Rise time* (ps)	Δ (%)	*NEN* (mV)	Δ (%)
RC-1	652	65	2065	152	158	−49
RC-3	484	22	1233	51	179	−42
RC-5	443	12	1096	34	182	−41
RC-10	412	4	1001	22	184	−40
RC-50	386	−3	931	14	185	−40
RLC	381	−4	829	1	211	−31
$R(f)L(f)C$	396		819		307	

used in the critical paths cannot afford such slow-rising signals, since the delay would be prohibitive for 3-ns or shorter cycle times. Low-impedance driver circuits must be used, in much the same manner that off-chip circuits are designed for chip-to-chip package wiring.

Wiring in multilayer structures has orthogonal conductors in adjacent layers which increase the line capacitance and therefore the signal delay. Unlike chip-to-chip interconnections, on-chip lines have very nonuniform transmission-line cross sections. Typical printed-circuit-board or MCM wiring [5] uses multiple stacks of X- and Y-directed lines between reference planes. Most on-chip power distribution relies on wide power buses placed sparsely on the topmost layer. This layer has thick

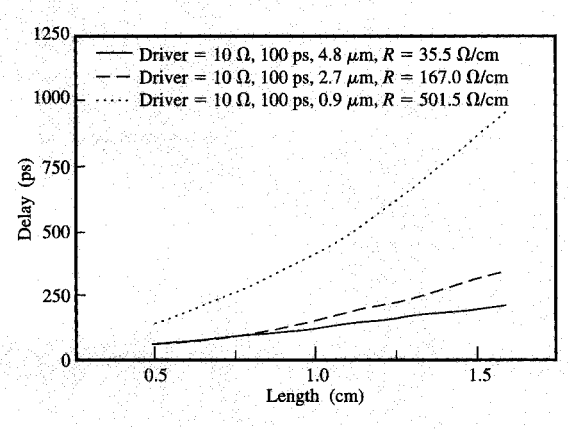

Figure 2

Simulated propagation delay using 10-Ω, 100-ps-rise-time driver circuit for 4.8-μm (solid curve), 2.7-μm (dashed curve), and 0.9-μm (dotted curve) lines with R = 35.5, 167.0, and 501.5 Ω/cm, respectively.

metallization, as shown in **Figure 3(a)**. Narrower buses on the lower levels distribute the power from the top to the individual logic cells. Such a nonuniform reference mesh results in increased inductance and resistance, which in turn increase the propagation delay and line-to-line crosstalk.

Typical silicon substrates have a resistivity of 2–3 Ω-cm and a dielectric constant of $\varepsilon_r = 11.8$. The relaxation time for charge redistribution is $\tau = \rho\varepsilon_r \cong 0.01$ ps, which is much faster than the encountered signal rise times of around 100 ps. The charges induced on the surface of the silicon substrate make its surface appear as a conductive layer that increases the signal line capacitance and thus increases delay. The very low conductivity of the silicon, however, ensures that there is negligible current flow through the substrate and that the return path is predominantly through the lower-resistance power buses. The substrate resistance is about 2000 times larger than the line resistance. If there were conduction through the substrate, the signal would be excessively degraded, giving rise to so-called "slow waves." It is shown in this study that we found no evidence of such effects.

As signal speeds increase, losses due to the skin effect could become important. It is found, however, that the skin depth δ [5] is 1.0–1.7 μm even for t_r = 50–100 ps. Skin-effect losses are significant when $\delta \ll t$ (t is the conductor thickness), which is not the case for representative multilayer on-chip wiring with t = 0.85–2.15 μm. Signal rise-time dispersion is dominated by frequency-independent series resistive losses. It is explained later, however, that the irregular power bus layout of this type of transmission line results in a frequency-dependent current distribution in the reference mesh which increases the electromagnetic energy coupling (crosstalk) between lines.

Finally, the high wiring density required for high-performance microprocessor chips results in very high electromagnetic coupling between signal lines. The energy coupled into quiet lines from adjacent active lines has typical traveling-wave properties. The noise monitored at the end near the active line excitation and that at the far end both have components of both the forward- and backward-traveling waves. The contribution magnitude is dependent on the terminations at the two ends. The far-end noise (FEN) is proportional to the difference $K_C - K_L$ in capacitive and inductive coupling, while the near-end noise (NEN) is proportional to the sum. In a homogenous structure, $K_C = K_L$ and $FEN = 0$. It is shown in [5] that $NEN \cong K_B 2\tau l V_{in}/t_r$ for $l < t_r/2\tau$ and saturates for $l > t_r/2\tau$, becoming $NEN \cong K_B V_{in}$, where $K_B = (K_C + K_L)/4$. Far-end noise is $FEN \cong K_F \tau l V_{in}/t_r$, where $K_F = [(K_C - K_L)/4]$, $K_C = C_{ij}/C_{ii}$, and $K_L = L_{ij}/L_{ii}$. It is shown in this study that in the case of short, very densely packed lines, both the capacitive (K_C) and inductive (K_L) coupling coefficients are high, but because of the short coupling length, the near-end noise never saturates [5]. The active signals do not reach full amplitude since the rise time is longer than the line delay, and *NEN* is in the range of 7–12% of the swing for two adjacent lines. Long, wider lines with wider separation have lower K_C and K_L, and although *NEN* saturates, it is still only around 4–8%. For a very nonuniform power bus layout, the long, wide lines can have very high K_L, and *NEN* can be as high as 24–30%. In all cases, *FEN* [5] is less than *NEN* but becomes significant for the last case. *FEN* increases with coupled length, but the active signals are also slowed down, thus reducing *FEN*. The very high line resistance contributes to crosstalk attenuation as well.

All of these characteristics that are specific to on-chip wiring were analyzed in detail by measurements on test sites that address each effect separately. Accurate modeling and simulation techniques have been developed that are verified by their agreement with experimental results. Subsequently, the validated models are used for performance predictions.

- *Representative on-chip structures*

The focus of this paper is on very long on-chip interconnections, which are generally placed in the upper layers of a multilayer stack. They tend to be several times wider than the minimum ground rules used for the short connections which represent the majority of all on-chip wiring. Top metal layers are also thicker because of power distribution requirements, thus affording lower-resistance

wiring channels. Critical paths, such as clock distribution nets, can be laid out with fewer neighboring orthogonal interconnections affecting their characteristics. The power distribution can comprise very wide buses, supplying V_{DD} and *GND*, for example, for CMOS devices, on the very top layer. In the layers below, alternating in direction, orthogonal narrower power buses distribute these voltages to the individual cells. A typical power mesh is shown in **Figure 3(b)** for a wiring structure having five metal layers, M1–M5, as illustrated schematically in Figure 3(a).

The routing of the wiring channels can be grouped into three broad categories. Category I comprises M3 wiring running underneath and parallel to the very wide power buses on the M5 layer labeled M5 in Figure 3(b). Since these buses can be as wide as 110 μm, they look like quasi-continuous reference planes for 1–5-μm-wide lines. Category II contains M4 wiring in channels orthogonal to these wide buses but with in-plane, narrower buses in parallel with them, completing the reference mesh. Category III has lines placed in either layer M5 or M3 and in parallel with the wide buses on M5. These lines range from being adjacent to the M5 bus, to midway between two widely spaced buses. Except for the first group, the lines generally have an irregular reference plane configuration. In a complex microprocessor chip, the long wires can be considered to be a concatenation of any of these three cases, unless custom placement is used. Each configuration can be analyzed individually.

The lines in this study were placed on M3, M4, and M5 layers, with power buses and orthogonal wiring on M2, M3, M4, and M5 layers. Most of the test sites had 1.6-cm-long lines; a few narrower lines were 0.8 cm long. Each site was 200 μm wide and isolated from any other site. There were no active devices underneath this area of the silicon chip. Each site was designed to include all the representative details of an actual microprocessor chip application. The narrow buses on the inner layers have the representative width and center-to-center separation (pitch). They are connected through the correct via pattern to the topmost layer, where the lowest-resistance reference buses are placed. Orthogonal wiring is generally included either on only one adjacent layer or both above and below the signal line layer. The signal propagation delay, rise-time degradation, and crosstalk on these sites were analyzed as a function of orthogonal layer loading, configuration of power bus layout, in-plane or intra-plane coupling arrangement, proximity to semiconductor substrate, and operating temperature. **Table 2** shows the grouping of all the sites.

Figure 4 shows a close-up of the site labeled 3C, with two coupled lines, 2.7 μm wide, with 2.7-μm separation, on the M3 layer. This is a category I design, with a solid ground plane on M5 which covers the entire site. Orthogonal to M3 signal lines are 0.9-μm-wide M4 lines on

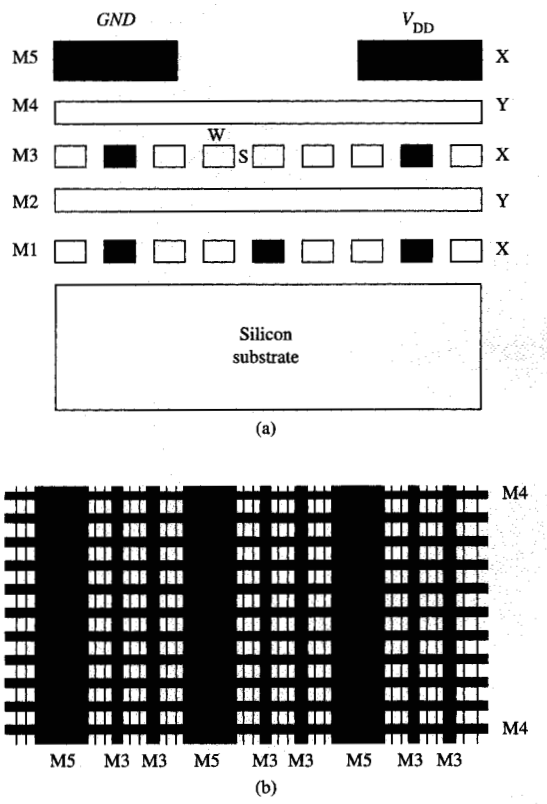

Figure 3

(a) Schematic representation of five-metal-layer on-chip wiring cross section used in the analyzed test vehicle. (b) Typical power and ground buses on layers M3, M4, and M5.

3.6-μm pitch, which represents 50% channel occupancy and is representative of average wiring usage. On the M3 layer there are also 9.9-μm wide power buses on a 63-μm pitch running alongside the signal lines to simulate the actual logic cell power bus structure. This mesh of buses connects to the top solid ground plane with the same staggered via pattern found on an actual chip where both V_{DD} and *GND* bus assignments exist. Similar designs, not shown, have 0.9-μm-wide lines on the M3 layer, with 0.9-μm spacing. In this case, the orthogonal wiring is both above and below, on M4 and M2, respectively. There are six such designs—4A, 4B, 4C, and 4A′, 4B′, 4C′—having 100%, 50%, or 0% occupancy of orthogonal wiring channels. Sites 3C and all of the 4× sites are used to investigate in-plane, line-to-line crosstalk for Category I wiring listed in Table 2. Sites 4B and 4C have duplicate cells that have an additional continuous M1 layer underneath each site. This plate isolates the lines from the effect of the silicon substrate.

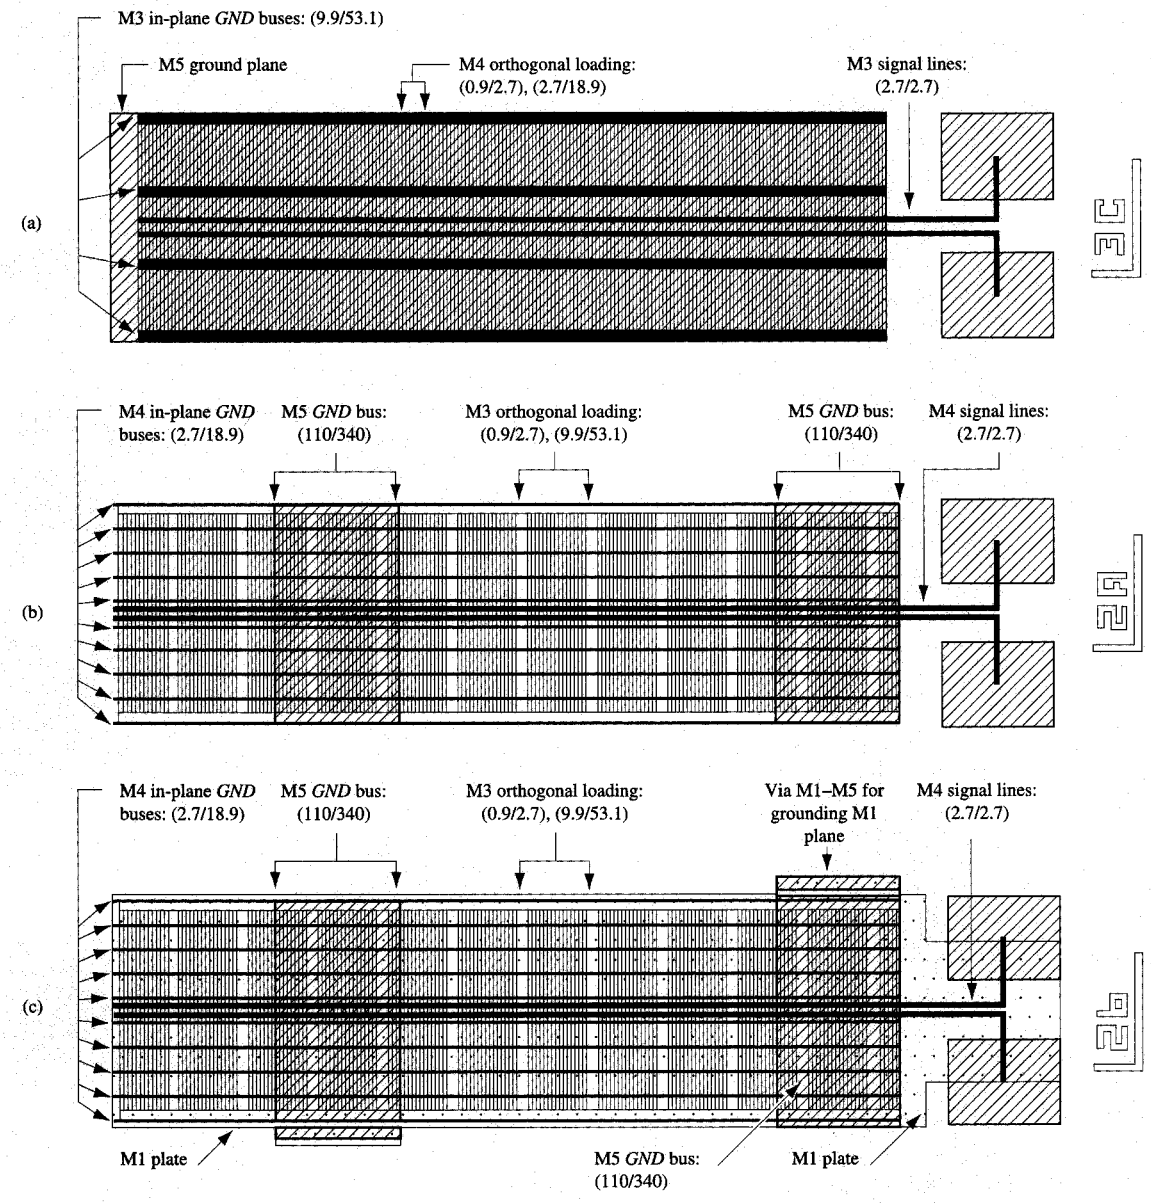

Figure 4

Layout of (a) site 3C with two 2.7-μm-wide signal lines on M3 and solid ground plane on M5; (b) site 2A with coupled 2.7-μm-wide signal lines on M4 having orthogonal 110.0-μm-wide ground buses on M5; and (c) site 2B similar to site 2A with solid plane on M1. (Numbers in parentheses are the linewidth and line spacing, respectively, both in μm.)

Figure 4 also shows the design of site 2A, which addresses the Category II type of referencing. Two adjacent, 2.7-μm-wide lines, with 2.7-μm separation, are placed on the M4 layer; 110.0-μm-wide power buses on a 450.0-μm pitch are on the M5 layer. The M3 layer also has orthogonal, 9.9-μm-wide buses on a 63.0-μm pitch,

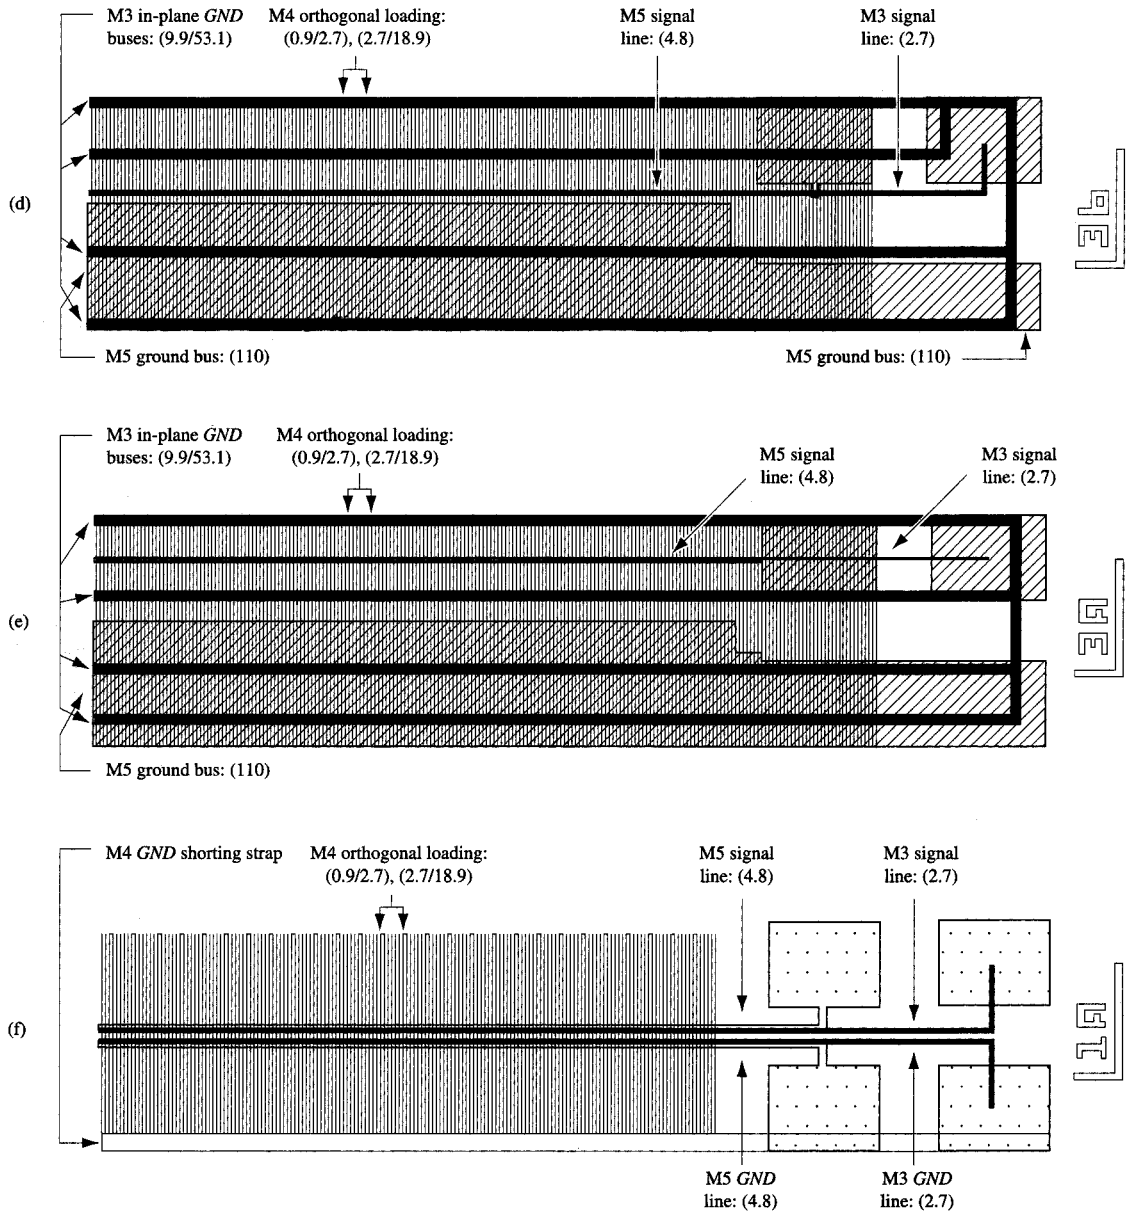

Figure 4 (continued)

Layout of (d) site 3B with 4.8- and 2.7-μm-wide parallel lines on M5 and M3, respectively, having a close wide ground bus on M5; (e) site 3A similar to 3B but with the power bus 50.0 μm away from the signal lines; and (f) site 1A with 4.8- and 2.7-μm-wide lines on M5 and M3, respectively, having a coplanar-waveguide structure, with adjacent ground lines. (Numbers in parentheses are the linewidth and line spacing, respectively, both in μm.)

and there are in-plane, 2.7-μm-wide buses on M4, on a 21.6-μm pitch alongside the signal lines. Orthogonal, 50% loading, with 0.9-μm-wide lines, is also present on the M3 layer. In this type of site, in-plane, line-to-line crosstalk in the M4 layer is being investigated in the presence of orthogonal power buses on the M5 layer. Site 2B is similar

Table 2 Description of test site designs.

Category	I	I	I	II	II	III	III	III	III	III	III
Site	3C	4B	4C	2A	2B	3A	3A	3B	3B	1A	1A
Reference	Solid *GND*	Solid *GND*	Solid *GND*	Orthogonal *GND*	Orthogonal *GND*	Parallel far *GND*	Parallel far *GND*	Parallel close *GND*	Parallel close *GND*	Coplanar wave guide	Coplanar wave guide
Signal layer	M3	M3	M3	M4	M4	M5	M3	M5	M3	M5	M3
Line width (μm)	2.7	0.9	0.9	2.7	2.7	4.8	2.7	4.8	2.7	4.8	2.7
Orthogonal wiring (%)	50	50	0	50	50	50	50	50	50	50	50

to site 2A. It includes an additional continuous M1 layer underneath the entire test site for isolation from the silicon surface.

Finally, Figure 4 shows Category III wiring. Site 3A has 4.8- and 2.7-μm-wide lines in layers M5 and M3, overlapping each other and running parallel with a 110.0-μm-wide power bus on M5. The M5 line is 50.0 μm away from the edge of the bus. Layer M4 contains 50% orthogonal wiring loading, and in-plane, 9.9-μm-wide buses present on a 63.0-μm pitch are again placed in layer M3 alongside the signal line. Orthogonal buses are present in M4, 2.7 μm wide with a 21.6-μm pitch. In this case, layer-to-layer M5-to-M3 crosstalk is being investigated. Site 3B has similar lines, but the M5 signal line is only 2.4 μm away from the wide power bus on M5. Sites 3A and 3B focus on the change in inductive current return loop for signal lines placed at varying distances from the wide reference bus on M5. Finally, site 1A has two coplanar waveguide configurations (signal–ground), overlapping in layers M5 and M3 and with orthogonal, 50% loading on the M4 layer. The lines in layers M5 and M3 are respectively 4.8 μm wide with a 2.4-μm gap and 2.7 μm wide with a 2.7-μm gap.

• *Cross-sectional analysis*

Cross-sectioning has been performed on the structure shown schematically in Figure 3; SEM photographs are shown in **Figure 5**. All of the layers are planarized up to level M4, above which the dielectric and M5 metal layer have a conformal topography, as exemplified by Figure 5(b). The thick M5 layer has substantial line width variation, as shown in Figure 5(c). All line widths w were 5.5–16.6% wider than nominal, with narrow lines exhibiting the largest variation. The metal thickness t was also 1.2–15.0% higher than nominal, while the dielectric layer was 2.4–14.0% thinner. This all contributed to increased line capacitance. Measurement of the capacitance of large parallel-plate capacitors between adjacent layers allowed the extraction of the dielectric constant for the SiO_2 insulator, which was found to be $\varepsilon_r = 4.05$ (1.25% higher than nominal). The resistivity of the AlCu signal lines, ρ, was extracted from line resistance and cross-sectional dimension measurements [$\rho = (Rwt)/l$]; ρ was found to be 7% higher than nominal, or 4.05 $\mu\,\Omega$-cm. Similar variations were observed for sites on several 8-inch-diameter wafers from two different fabrication runs that were investigated.

The line resistance was also measured over the range −160°C–100°C for line widths of 0.9, 2.7, and 4.8 μm. The 0.9- and 2.7-μm-wide lines, in sites 4C and 3C respectively, were in layer M3, and the 4.8-μm-wide line is exposed to air. All three lines exhibit very similar variations in resistivity. At −160°C, the resistivity decreased by 73%, 75%, and 76%, compared to 22°C operation for the 0.9-, 2.7-, and 4.8-μm lines, respectively. At 100°C, the increase was 28%, 28%, and 32% for the same cases. As expected, the resistivity had a nearly linear variation with temperature.

Test site modeling

• *Transmission-line modeling and simulation*

The cross-sectional dimensions were used to calculate the line parameters R, L, C, and G. The line parameters were then used in signal propagation simulations that were compared with measured waveforms. A three-dimensional, full-wave, electromagnetic analysis code [8] was used. The technique is based on the method of moments, with conductor and polarization currents represented by "rooftop" basis functions. The approach subdivides the entire structure into nonuniform, rectangular subsections, and a de-embedding procedure described in [8] allows the extraction of R, L, C, and G per unit length over a wide frequency range. Nonuniform grid subsectioning was extremely useful for the structures shown in Figure 4, where a large disparity exists between the very narrow lines (1–5 μm), and the wide power lines (2.7–110.0 μm) on large center-to-center separations (21.6–450.0 μm). **Figure 6** shows the three-dimensional configurations used in sites 2A

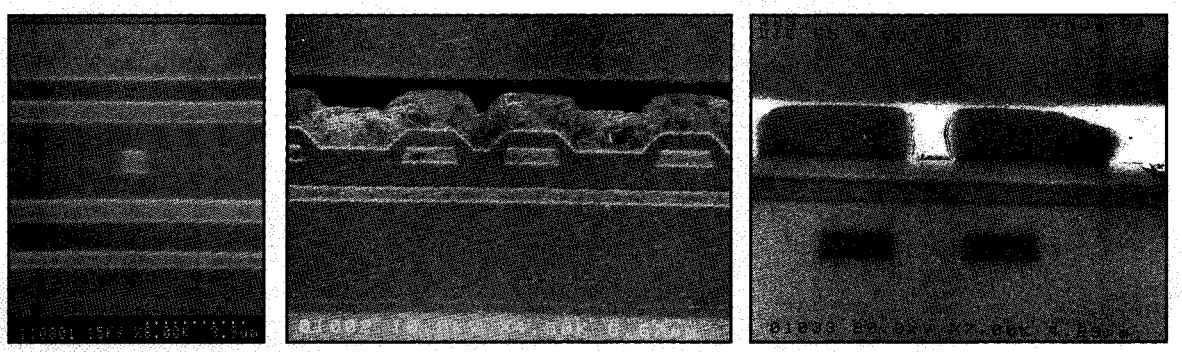

(a) (b) (c)

Figure 5

SEM photographs: (a) five metal layers showing a 1.05-μm-wide line on M3; (b) top three layers M3, M4, and M5 showing the nonplanar topography of the M5 lines; (c) site 1A cross section showing the 5.3- and 2.85-μm-wide lines on M5 and M3 and a coplanar-waveguide structure.

and 3B. The nonplanar topography on the M5 level was modeled with a rectangular staircase-like representation. All of the structures were analyzed with the code described in [8] running on an IBM RISC System/6000® Model 590 workstation with 1 GB of main storage. Most of the elements of the test configurations were calculated with only one division per conductor face or dielectric layer. The signal conductors had a maximum of three subdivisions per conductor face.

It is shown in [5] that as the signal speeds increase, the skin depth δ becomes smaller than the conductor thickness, and both R and L are frequency-dependent ($R \geq \omega L$). In the case of the AlCu lines in this study with $\rho = 4.05\ \mu\Omega$-cm, $\delta = 1.85\ \mu$m at $f = 3$ GHz and $\delta = 1\ \mu$m at 10 GHz. Driver signals with 100-ps rise times contain frequency components with significant energy at 3 GHz, while for the 35-ps rise-time source used in the measurement, the frequency of interest is $f = 10$ GHz. The onset of skin effect is more noticeable for the $t = 2.1$-μm-thick M5 lines. **Table 3** shows the calculated $R(f)$ and $L(f)$ for site 2A for frequencies up to 10 GHz. A coupled line pair is analyzed, and the table shows the self R_{11}, L_{11}, and mutual terms R_{12}, L_{12} ($R_{11} = R_{22}$ and $L_{11} = L_{22}$). Site 2A has very large $R_{12}(f)$ terms due to the nonuniform current distribution in the power mesh connecting to the orthogonal wide (110.0-μm) buses on M5. The presence of the in-plane buses maintains a small variation of L_{11} and L_{12} with frequency. The resistive return path through the narrow elements of the power mesh affects R_{12} the most. It is later shown that the current crowding in the irregular mesh for sites such as 2A, 3A, 3B, and 1A contributes greatly to increased crosstalk. For example, in site 1A, R_{12} is largest because the ground return is just the coplanar conductor, adjacent to the signal line.

Table 3 Calculated $R(f)$ and $L(f)$ matrices for site 2A.

f (GHz)	R_{11} (Ω/cm)	R_{12} (Ω/cm)	L_{11} (nH/cm)	L_{12} (nH/cm)
0.001	210.60	13.60	7.50	3.94
0.010	210.70	13.60	7.50	3.86
0.100	210.90	13.80	7.40	3.78
1.000	216.60	18.20	6.50	3.02
4.000	235.50	30.50	5.60	2.39
10.00	268.30	45.40	5.40	2.23

Many techniques have been developed to incorporate the frequency-dependent losses in circuit simulator programs for time domain waveform analysis. The most common approaches compute the impulse or step response of the lines at any given time using either inverse fast Fourier transform (IFFT) [9] or inverse Laplace transform (ILT) [10] techniques to obtain the convolution integral for transition from frequency to time domain. No attempts are known, however, to analyze coupled transmission lines with the high resistive losses encountered in this study, where R ranges from 35 to 500 Ω/cm. Numerical stability problems were encountered with the code described in [10], which relies on the representation of the frequency-dependent line characteristics by a wave-shaping polynomial ratio in order to improve the efficiency of the recursive evaluation of the convolution integral in the time

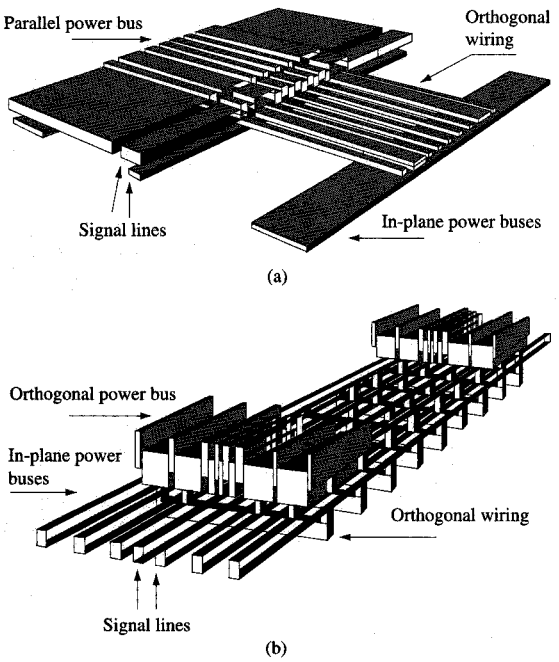

Figure 6

Three-dimensional model used to calculate $R(f)$, $L(f)$, and C for (a) site 3B and (b) site 2A. The M5 layer dimension for site 2A is shown magnified 10×.

domain. An implementation using the IFFT for obtaining the line voltages and currents in the time domain also failed to converge. These algorithms are optimized to provide a good fit to the transfer function over a wide frequency range, with special emphasis on the high-frequency region where skin-effect losses are important. In our on-chip wiring, the important region to examine is from low frequency to the point at which skin effect just becomes significant. It was proposed in [11] to represent the lines with ideal, lossless transmission lines augmented by an equivalent circuit of lumped R and L elements that synthesize the $R(f)$ and $L(f)$ frequency dependence shown in Table 3. The technique was verified for lossy waveguide structures and short, low-resistance microstrip lines. In this study, this approach was successfully extended to very high-loss transmission lines.

The symmetric coupled-line network was first reduced to two single-line circuits through a decoupled mode transformation using linear voltage-dependent current sources [9, 11]. The frequency dependence of the uncoupled lines representing the independent modes of the structure was synthesized using a distributed four-pole filter circuit, adjusted to smoothly fit the $R(f)$ and $L(f)$ extracted from the three-dimensional modeling up to 10 GHz. **Figure 7(a)** shows one subsection of the distributed network with the four filter circuits. The four circuits have a range of cutoff frequencies which cover the bandwidth of interest. The impedance $R_i + j\omega L_i$ of each section must be much less than the line impedance Z_0 in order to avoid interface reflections. The entire line representation is then broken up into (typically 10–50) subunits to achieve convergence to within engineering accuracy. The number of subunits required depends on specific line parameter values, but each unit should have electrical length much shorter than the wavelength of the propagated signal. Such an approach can easily be implemented in standard circuit simulators such as SPICE or ASTAP.

The above technique, however, is difficult to implement efficiently for the highly asymmetric cases of sites 3A, 3B, and 1A that have large differences between the self terms R_{ii}, L_{ii}, C_{ii}, and large mutual terms L_{ij}, C_{ij}, and R_{ij} because the decoupled-mode transformation becomes frequency-dependent. For these cases, the asymmetric coupled lines were modeled directly by cascaded lumped-element π-section RLC networks. Frequency dependence was represented by including additional RL elements as shown in **Figure 7(b)**. These RL element values were chosen to match the $R(f)$ and $L(f)$ of the transmission lines up to 10 GHz. More elements were used in modeling the return path, since R_{12} and L_{12} have the largest variation with frequency. The self and mutual capacitances are placed at the ends of each section. The subsection length was chosen as above to be much shorter than the relevant wavelength, and to have a series impedance much less than the line Z_0. For these same cases, the frequency-independent representation for one subsection is shown in **Figure 7(c)** as a π-network having constant resistance and inductance for both the signal line and reference return conductors.

Finally, in all cases, the accurate models of Figures 7(a)–7(c) were compared with the π-section RC representation shown in **Figure 7(d)**, where the inductive terms are omitted. The granularity requirement of the RC-circuit representation is discussed further later.

- *Modeling results*

The nonplanarity of layer M5 dramatically increases the line capacitance. Small changes in the shape of the undulation shown in Figure 5 result in large changes in calculated line capacitance. This is exacerbated by variations in line width in this layer due to the very thick metal. **Table 4** shows two examples of C_{11} and C_{12} showing the difference between measurement and calculation. Calculation was performed using both nominal and measured dimensions. Capacitance measurement was

performed at 1 MHz. The results are shown for sites 3C and 1A, with lines in the M3 or M3 and M5 layers, respectively. The dimensions extracted from the scanning electron micrographs (SEMs) shown in Figure 5 were obtained from another wafer, from the same batch as the one measured. Site 3C capacitance shows fairly similar results for the two types of calculations. The errors are consistent with the accuracy of the dimensional measurement and wafer-to-wafer variations. The capacitance measurements have an accuracy of ±1%. Site 1A initially showed a large discrepancy between calculated and measured capacitance when the nonplanar topography was not included. The agreement improved dramatically, to 6–12%, when the actual geometry was used. Similar results were obtained for all of the lines on the M5 layer in the other types of sites.

Table 5 summarizes the measured and calculated R, L, and C values for sites 1A, 2A, 2B, 3A, 3B, 3C, 4B′, and 4C′. Several observations can be made. Although the reference bus designs were different, all of the 2.7-μm-wide lines had similar self-capacitance C_{11}, because the capacitance is primarily determined by the orthogonal wiring in the adjacent layers. The self-capacitance of the 4.8-μm-wide lines is high in all cases, as expected. For this test vehicle, since the lines are not covered by the traditional overcoat layer, capacitance is slightly lower than expected on a production chip. Note the similarity between the measured and calculated capacitance values for sites 2A and 2B. Although the wiring is identical, the lines in 2B have a metal layer in M1, covering the semiconducting silicon. As was explained earlier, the silicon surface appears as a conducting layer. However, since the lines on the M4 layer are far away from the silicon surface or M1, the orthogonal wiring on M3 acts as an effective shield, and their capacitance is not affected by the inner layers or the silicon. In the case of the narrow (0.9-μm-wide) lines, there is a 35% increase in self-capacitance C_{11} between the case without orthogonal loading lines on M2 and M4 and the case with 50% occupancy. The mutual capacitance decreases by 20.6%. Although not shown in Table 5, the capacitance of the 0.9-μm lines was also not changed by the presence of a solid metal plane covering the silicon. The capacitive coupling, K_C, for the wide lines ranges from 0.08 to 0.21, while for the narrow lines with 0.9-μm width and 0.9-μm separation, K_C is very high (0.26–0.47).

Table 5 also shows the calculated high-frequency inductances L_{11} and L_{12}. The largest L_{11} value is obtained for site 3A, where the lines are far away from the in-plane wide power bus on M5. The in-plane narrow buses on the other layers, completing the reference mesh, reduce L_{11} by a small amount. The orthogonal bus referencing in site 2A generates fairly high inductance for M4 lines, while the solid ground plane in site 3C lowers L_{11} for the M3 lines.

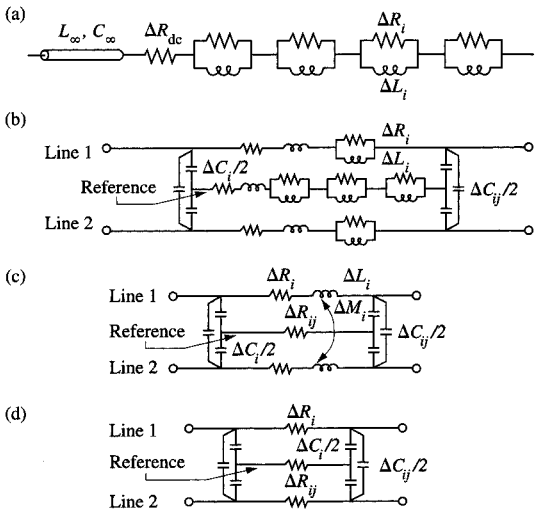

Figure 7

(a) Detailed description of the uncoupled line representation; (b) one subsection of a distributed $R(f) + L(f) + C$ synthesized network, used for simulating sites 3A, 3B, and 1A; (c) detail of a subsection for distributed $R + L + C$ network used for frequency-independent simulation; and (d) π-section $R + C$ network.

Table 4 Self and mutual capacitance values measured and calculated using nominal cross-sectional dimensions, and with dimensions obtained from SEMs shown in Figure 5 for sites 3C and 1A.

Site 3C: Solid *GND*, M3, 2.7 μm		
Measured	C_{11} = 1.750 pF/cm	C_{12} = 0.220 pF/cm
Nominal dimensions (%)	+4.0	−10.5
SEM dimensions (%)	+8.0	−12.6
Site 1A: Coplanar waveguide *GND*, M5/M3, 4.8/2.7 μm		
Measured	C_{11} = 3.690/1.740 pF/cm	C_{12} = 0.242 pF/cm
Nominal dimensions (%)	−43.0/−2.5	−46.0
SEM dimensions (%)	−6.0/+8.0	+12.0

In the case of narrow lines in the M3 layer, with 0.9-μm width and 0.9-μm separation, their large distance from the solid reference plane on M5 results in high inductance. The

Table 5 Resistance, capacitance, inductance, and impedance measured and calculated for eight of the sites.

Sites	*Resistance* (Ω/cm)		*Capacitance* (pF/cm)				*Calculated inductance* (nH/cm)		*Calculated coupling*		*Calculated impedance* Z_0 (Ω)
			Calculated		*Measured*						
	Calculated	*Measured*	*C11*	*C12*	*C11*	*C12*	*L11*	*L12*	K_C	K_L	
					Length = 1.6 cm						
3C	167.0	166.0	1.89	0.194	1.75	0.220	3.19	0.70	0.10	0.22	41
2A	174.0	166.2	1.83	0.167	2.05	0.137	4.49	1.84	0.09	0.41	50
2B	174.0	166.2	1.80	0.165	2.05	0.137	3.64	1.49	0.09	0.41	45
3A-M5	35.5	34.6	4.32	0.350	4.18	0.260	6.20	4.91	0.08	0.73	38
3A-M3	167.8	163.9	1.67	0.350	1.74	0.260	6.84	4.91	0.21	0.67	64
3B-M5	35.5	35.9	3.36	0.323	3.55	0.255	3.45	2.25	0.10	0.65	32
3B-M3	167.8	165.1	1.86	0.323	1.70	0.255	5.07	2.25	0.17	0.44	52
1A-M5	38.5	35.8	3.46	0.270	3.69	0.242	3.21	1.94	0.08	0.60	31
1A-M3	167.8	165.7	1.87	0.270	1.74	0.242	4.30	1.94	0.14	0.45	48
					Length = 0.8 cm						
4B′	501.5	492	1.71	0.454	1.88	0.514	4.25	2.00	0.26	0.46	50
4C′	501.5	492	1.44	0.670	1.39	0.647	4.32	2.20	0.47	0.50	55

inductive coupling K_L also shows clear differences for the different designs. It is smallest, $K_L = 0.22$, for the solid ground plane reference case of site 3C. It increases to 0.41 for the orthogonal ground bus, and rises as high as 0.73 for the line which is 50 μm away from the wide power bus in site 3A. The inductive coupling is also high for sites 4B′ and 4C′.

Experimental characterization

• *Measurement setup*

Time-domain measurements were performed using the setup shown in **Figure 8**. A 20-GHz sampling oscilloscope, HP Model 54120A, was used to launch a 200-mV swing, 35-ps-rise-time step excitation. Such a source simulates high-speed logic signals. Custom-built, coaxial probes [12] were used to detect the transmitted signal and the crosstalk on the quiet adjacent line. **Figure 9** shows the equivalent circuit representation of the measurement with the probe connected at the far end of the quiet line for far-end-noise evaluation. In all cases, the source impedance and the termination are each 50 Ω. Signal propagation is repeated on short structures that are identical with the long lines and have either 0.06- or 0.09-cm length. The difference in the measured delay on the long and short lines yields the propagation per unit length, without the error introduced by pads and probes. Broadband measurements of attenuation, phase constant, and impedance were performed using the short-pulse propagation technique described in [13]. In this case, the step source was first differentiated using a passive impulse-forming network, Picosecond Pulse Lab Model 5210, placed in-line with the input probe. A 35-ps-wide pulse was then launched on the test line and the identical short structure. The relative comparison of the Fourier transforms for the two lengths of identical lines yields $\alpha(f)$, $\beta(f)$, and $Z_0(f)$. The technique works best if the pulses are shorter than the round-trip propagation delay on the lines in order to eliminate reflections from pads and probes. The propagation delays on the short structures are much less than needed; however, the end pad and probe discontinuities were found to be acceptably small, and the measurements were relatively accurate. Finally, the time-domain and short-pulse measurements were repeated over the −160°C to 100°C range. The lower limit is important for investigation of the effect of decreased AlCu resistivity on performance. The higher limit is close to the operating temperature of high-power microprocessor chips.

• *Experimental results*

Time-domain measurements

Time-domain reflectometry (TDR) measurements were performed on all of the structures by monitoring the reflection at the input source from the line impedance relative to the 50-Ω source. The results were not very useful. It is explained in [5] that the line impedance increases substantially along the length of the line because of the very high resistive losses. The source rise time was subsequently so degraded that the reflection from the line was buried in the slow round-trip response.

Signal propagation delay, rise-time degradation, and crosstalk were investigated by monitoring the transmitted signal at the end of either the active line or the quiet line, as shown in Figure 9. **Table 6** lists the results obtained at room temperature for sites 3C, 2A, 2B, 3A, 3B, 1A, 4B′, and 4C′, while **Table 7** highlights some of the results at −160°C and 100°C. All of our measurements were performed with a 50-Ω terminator, while in general on-chip interconnects are open-ended. These results can thus be viewed as the fastest possible response of these lines, since the 35-ps-source rise time is nearly an ideal step excitation.

The propagation delay was measured at the 50% level of the input and output signals. Sites 3C, 3A (M3), 3B (M3), and 1A (M3) had 1.59-cm-long lines; sites 2A and 2B had 1.6-cm lines; sites 3A (M5) and 3B (M5) had 1.56-cm lines; and sites 4B′ and 4C′ had 0.8-cm lines. The longest delay was obtained for the 4.8-μm-wide lines on the M5 level that were far away from the parallel wide ground reference bus. These lines have both high capacitance and inductance, which offset the benefit of the much lower resistance. The 2.7-μm line on M3 has even higher inductance and four times higher resistance, which results in slightly higher delay, even though the capacitance is 2.4 times smaller. The 2.7-μm lines with orthogonal bus reference on M4 (sites 2A and 2B) had slightly lower delay. The fact that the delays are so similar for both 2A and 2B is evidence for the absence of any slow-wave effects. Even though the reference mesh is very hollow, the orthogonal loading lines on M3 shield the M4 lines from any other effects below. This is proved by the capacitance values, which are equal for 2A and 2B. As explained earlier, the power mesh provides a much lower resistance path than the semiconducting substrate. The slight differences in performance are due to dimensional differences. Even for sites 4C and 4C′, with 0.9-μm lines on M3 with no loading lines on M2, there was no evidence of signal slowdown. These 0.8-cm-long lines had 82- and 84-ps propagation delays, and 119- and 121-ps rise times, respectively. As expected, the delay increased monotonically with increased orthogonal loading, and there was no evidence of unusual effects caused by the substrate. The lines in 3B had lower delay than the 3A-design lines because both C and L are lower. The inductance is much lower because of the close proximity to the parallel ground bus, which is 2.4 μm away. The capacitance is lower for the M5 line because of the difference in line and insulator dimensions. The lines with a solid ground plane reference (3C) and the coplanar waveguide structure (1A) had the smallest delay. However, 1A lines on M5 had the slowest rise time, followed by 1A (M3) and then 2B and 2A. (1A lines are slowed down by dispersion caused by the very lossy ground conductor.) Both the delay and rise time were much lower for the

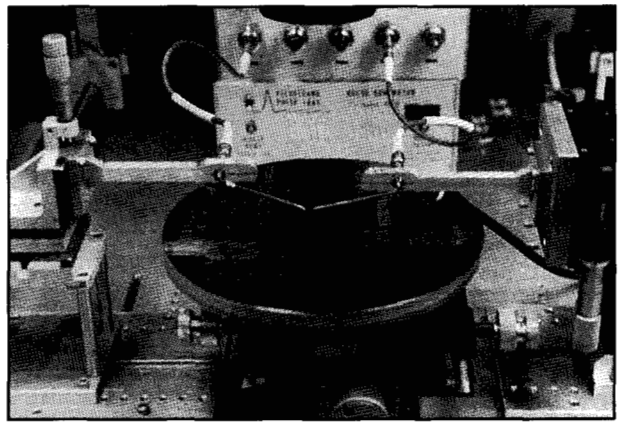

Figure 8

View of experimental setup; a 20-GHz sampling oscilloscope was used for time-domain measurements of 8-inch-diameter wafers.

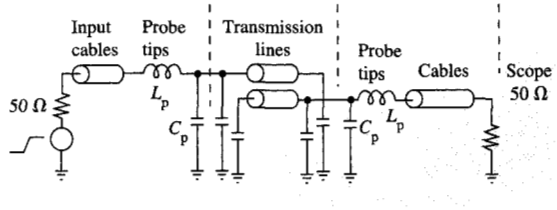

Figure 9

(a) Equivalent circuit representation of test circuit, showing input probes exciting the active signal and output probes attached at the far end of the quiet line for *FEN* crosstalk measurement. The probe discontinuities are C_p = 0.03 pF and L_p = 0.16 nH.

0.8-cm lines, even though the resistance is 501.5 Ω/cm. The line delay is always greater than half the rise time, but it is greater than t_r only for site 3A (M5).

There is a 2.9× reduction in rise time for the signal propagated on the 2.7-μm line in site 3C at −160°C, compared to room temperature (22°C). As a consequence, the delay is lowered by 1.5×. At 100°C, delay and rise time increase by 15% and 16%, respectively. This confirms the RC-line behavior explained earlier due to the very large R ($Rl \cong 6.4Z_0$). At −160°C, $Rl \cong 1.6Z_0$, which means that LC-line behavior is present and the rise time is much faster. At −160°C, all of the designs show a rise time that is faster than propagation delay, as seen in Table 7. The resistance-dominated performance is also confirmed by the short-pulse measurements explained above. At

Table 6 Propagation delay, rise time, and crosstalk measured and simulated for the sites of Table 2. Measurements were made with 35-ps rise time, 200-mV swing excitation, 50-Ω source and termination.

Sites	*Propagation delay* (ps)		*Rise time* (ps)		*Crosstalk* (mV)	
	Simulated	*Measured*	*Simulated*	*Measured*	*Simulated*	*Measured*
			Length = 1.6 cm			
3C	196	232	322	324	−0.4/6.0	0.0/6.2
2A	218	271	317	345	−3.9/5.4	−3.8/4.0
2B	225	269	346	378	−3.2/4.5	−1.8/3.9
3A-M5	284	287	324	261	−13.0/23.0	−14.0/14.0
3A-M3	250	291	186	227	−13.0/23.0	−14.0/15.0
3B-M5	214	242	319	281	−11.0/14.0	−11.0/10.0
3B-M3	219	247	303	292	−11.0/14.0	−11.0/10.0
1A-M5	218	227	316	427	−11.0/9.0	−3.3/6.0
1A-M3	219	237	317	381	−11.0/9.0	−3.3/6.0
			Length = 0.8 cm			
4B′	103	121	193	175	0.0/11.0	0.0/12.7
4C′	64	84	129	121	0.0/21.8	0.0/19.7

−160°C, the propagated pulses on 0.09- and 1.6-cm-long lines in site 3C had widths of 34.4 and 52.8 ps, respectively. The same pulses widened substantially at room temperature, to 35.6 and 86 ps. Because the signals were extremely slow, there was great loss in amplitude. **Figure 10** shows the measured attenuation for site 3C lines at room temperature and −160°C. The calculated values represent the odd and even mode components. The measurement was made on only one of the two coupled lines with the neighboring line unexcited, which is why the experimental data lie between the two calculated values. The bandwidth of the measurement is much higher at −160°C (about 14 GHz) compared to about 10 GHz for +22°C.

Crosstalk

Table 6 and Table 7 list the measured crosstalk at the far end of the quiet line. The negative values represent the *FEN* noise caused by the forward-traveling wave [5]. The positive values are the reflected signals from the near end of the line, which is open-ended. The test end is terminated in 50 Ω. The highest *FEN* is obtained for the highly inhomogenous case of site 3A, and especially on the M5 layer, for which K_C, K_L, and $K_C - K_L$ are extremely high; 3B results are slightly lower. Most of the other cases exhibited no measurable far-end noise due to very slow rise times on the active line. The reflection from the near end, which is generated by the backward-traveling wave [5], is highest for the very narrow (0.9-μm-wide) lines with 0.9-μm spacing. As expected, K_C decreases with increased orthogonal loading (from 0 to 50% in sites 4B′ and 4C′), resulting in lower *NEN*. Table 7 shows the disadvantage caused by lower signal attenuation and dispersion. The crosstalk increases substantially with lowered resistance. The faster rise times increase both *FEN* and *NEN*.

Figure 11(a) shows measured and simulated waveforms on the 1.6-cm-long, 2.7-μm-wide lines on the M3 layer in site 3C. The actual cross section was 2.85 × 0.85 μm. The simulated waveforms are shown both with and without taking frequency-dependent losses into account. In this case, the improvement in agreement with measurement goes from −19.0% to −14.0% for delay, and from −5.7% to −3.6% for rise time, for constant R and L versus $R(f)$ and $L(f)$. **Figure 11(b)** shows the comparison for the TDR trace and the *FEN*.

The agreement between measurement and simulation shown in Table 6 is consistent with the experimental uncertainty in line dimensions and the difficulty in modeling the nonplanar topography of M5 with the limited granularity that could be accommodated in reasonable CPU time and storage.

Application to practical on-chip interconnect simulation

The accurate transmission-line models we developed were then used to simulate representative configurations of long, on-chip interconnections driven by fast, low-impedance driver circuits. **Figures 12(a)** and **12(b)** show the circuits

Table 7 Measured propagation delay, rise time, and crosstalk variation as a function of operating temperature for five of the sites. Sites 4A, 4B, and 4C are of category I, with 0.9-μm single lines on the M3 layer, with 100, 50, and 0% orthogonal wiring in the M2 and M4 layers.

Sites	*Propagation delay* (ps)			*Rise time* (ps)			*Crosstalk* (mV)		
	−160°C	*22°C*	*100°C*	*−160°C*	*22°C*	*100°C*	*−160°C*	*22°C*	*100°C*
				Length = 1.6 cm					
3C	151	232	269	110	324	367	−4.8/16.0	0.0/6.2	0.0/5.8
2A	184	271	303	154	345	383	−12.4/7.8	−3.8/4.0	−2.4/3.0
				Length = 0.8 cm					
4B′	79	121	144	66	175	212	0.0/29.4	0.0/12.7	0.0/10.0
4A	82	127	144	71	171	207	—	—	—
4B	81	117	141	73	164	200	—	—	—
4C	62	82	93	53	119	142	—	—	—

used for signal propagation and *FEN* or *NEN* crosstalk monitoring, respectively. The coupled lines are either open-ended, assuming an active and a quiet driver at the near end, for worst-case *FEN*, or with the quiet line having the reverse arrangement for worst-case *NEN*. **Figures 13(a)** and **13(b)** show simulated waveforms for the two types of circuits using a 10-Ω impedance driver circuit with 100-ps rise time, driving the M5 line of site 3A, which has R = 35.5 Ω/cm and length of 1.6 cm. The crosstalk is monitored on the M3 line, which has R = 168.0 Ω/cm. Simulated waveforms are shown for three cases: a transmission-line model with $R(f)$ and $L(f)$ dependence represented by the circuit shown in Figure 7(b); a model with constant R and L of the type shown in Figure 7(c); and finally a distributed RC circuit with 50 subdivisions, as in Figure 7(d). It was shown in Table 4 that both K_C and K_L are high. This explains the large underestimation of *NEN* and the absence of *FEN* for the *RC* circuit, which does not include the K_L contribution. Moreover, the delay is underpredicted because in this case $Rl = 1.4Z_0$ (Rl = 56.5 Ω, Z_0 = 38.0 Ω), and the line has an LC-type characteristic. The rise time is fast (t_r = 254 ps) and shorter than the line delay of 305 ps. The *RC* representation is no longer valid, and the full transmission-line characteristics have to be taken into account. The propagated signal even has a slight overshoot due to the doubling at the open end of the line, which is not shown for the *RC* circuit.

Table 8 shows the results using ideal drivers with 10-Ω impedance, 100-ps transition, and 2.5-V input swing for sites 3C, 2A, and 3A. The K_B [$K_B = (K_C + K_L)/4$] in these cases is 0.08, 0.13, and 0.2, respectively, which would imply a saturated NEN noise ($NEN \cong K_B V_{in}$, where V_{in} is the signal on the active line) of 199, 313, and 500 mV, respectively. In the case of the solid ground plane reference of site 3C (Table 8), the delay error is very small (4–6%) for either constant parameter representation or RC-type circuit. The constant RLC representation underestimates crosstalk by 22% compared to the case in which frequency-dependent losses are taken into account. The RC-circuit crosstalk is in error by 31% compared to the same case. Since $\tau l < t_r/2$ (τl is the propagation delay), *NEN* is not saturated and is low. In the case of site 2A with R = 174.0 Ω/cm (Table 8), the lines have the

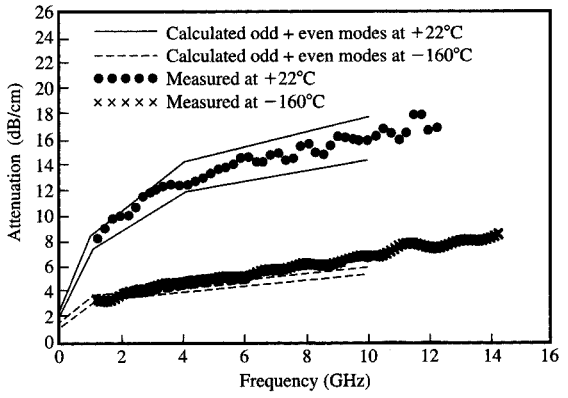

Figure 10

Calculated and measured attenuation using the short-pulse propagation technique [13] at −160°C and +22°C. Calculated values show the odd and even mode components for the two coupled lines of site 3C. Measurement was made on only one line.

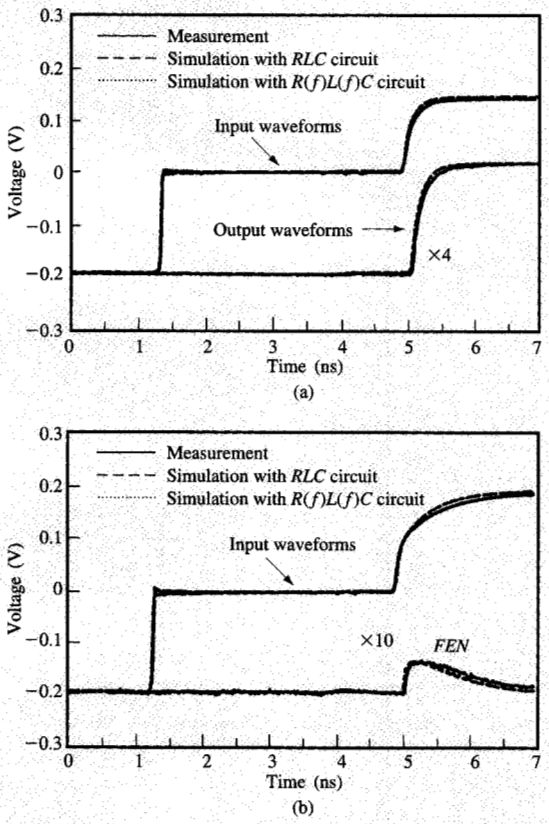

Figure 11

(a) Measured (solid curves) and simulated waveforms using frequency-independent R and L parameters (dashed curves), and frequency-dependent $R(f)$ and $L(f)$ (dotted curves), for 1.6-cm-long, 2.7-μm-wide M3 lines with solid ground on M5 (site 3C). Input monitored at sending oscilloscope channel; output measured at far end of the line and terminated with 50 Ω. (b) Measured (solid curves) and simulated TDR and *FEN* waveforms using frequency-independent R and L parameters (dashed curves), and frequency-dependent $R(f)$ and $L(f)$ (dotted curves), for 1.6-cm-long, 2.7-μm-wide M3 lines with solid ground on M5 (site 3C). Input monitored at sending oscilloscope channel on active line; output measured at far end of adjacent quiet line and terminated with 50 Ω as shown in Figure 9.

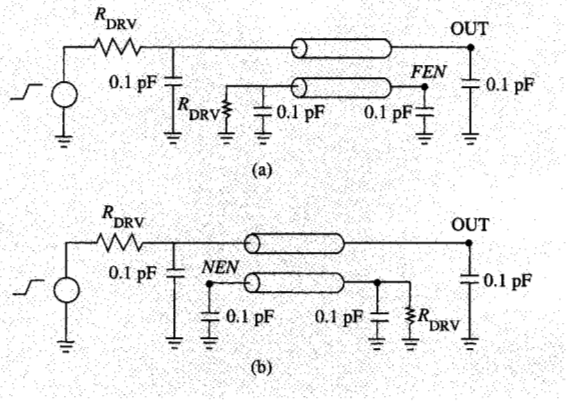

Figure 12

Circuit representation for analyzing worst-case crosstalk and signal propagation delay and rise-time degradation with various driver circuit impedance R_{DRV}: (a) far-end coupled noise, *FEN*; (b) near-end coupled noise, *NEN*.

much slower response ($Rl = 5.3Z_0$) of a typical *RC* line. The large resistive drop in the mesh that carries the return current to the orthogonal power buses contributes additional crosstalk when $R(f)$ and $L(f)$ dependencies are included [using the model of Figure 7(a)], as shown in **Figure 14.** The underestimation of *NEN* is 31% to 40% for constant parameter and *RC* representations, respectively. Finally, site 3A (Table 8) with $R = 35.5$ Ω/cm has $K_B = 0.2$. Here the overshoots shown for the 100-ps rise time in Figure 13 are generating additional crosstalk, which is higher than the saturated value of 500 mV. Both *FEN* and *NEN* are very large. Since we are considering two coupled lines, this result clearly suggests the need for coupled length restrictions for long parallel lines. The *RC* representation in this case underestimates delay by up to 31% and crosstalk up to 62%. The line delay is larger than the signal rise time, and transmission-line representation is essential. While the less resistive lines have the advantage of propagating the fastest signals, they generate the most noise. It should also be noted that for commonly used wiring configurations, such as data buses, the crosstalk could approach twice the values reported here. For example, three long coupled lines in the same layer would allow two active lines to couple noise into a quiet line. Coupling from a layer above or below could also increase the noise, and in this case coupled length restrictions would also have to be imposed to avoid logic failures.

The distributed *RC* representation is adequate for most delay prediction, particularly for the shorter, highly resistive interconnections. It was explained earlier and in [3] that the number of subdivisions depends on the line R and C values and the driver and load conditions. It was further recommended in [3] to use a concatenation of T-sections rather than uniform *RC* ladder networks. For the 2A case shown in Table 1, for example, it was found that a ten-section π-network representation gave only a 4% overestimation in delay compared to an accurate model, which is acceptable. The one-section *RC* circuit, however, had a 65% overestimation. In the case of design 3C with solid ground plane on M5, a ten-section *RC* ladder

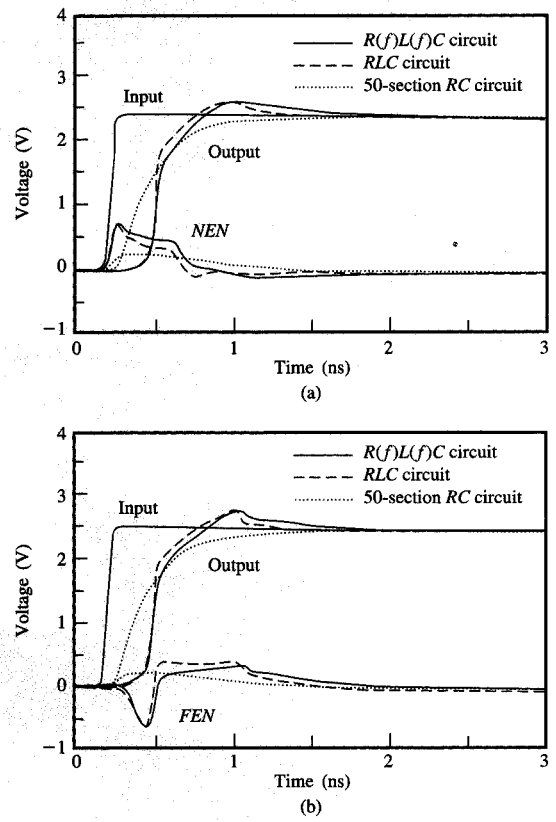

Figure 13

Simulated propagated signal and (a) *NEN* and (b) *FEN* crosstalk using the circuit of Figure 12 for 4.8-μm-wide, 1.6-cm-long lines on M5 having parallel, distant, wide ground bus on M5 (site 3A). The input circuit has 10-Ω impedance and 100-ps rise time. Simulation was done using frequency-dependent $R(f)$ and $L(f)$ (solid curves), constant R and L parameters (dashed curves), and 50-section distributed *RC* representation (dotted curves).

Table 8 Simulated delay, rise time, and crosstalk, using the circuits shown in Figure 12, for designs with solid *GND* (3C), orthogonal *GND* (2A), and parallel far *GND* bus (3A). All of the lines are 1.6 cm long and are driven by 10-Ω driver with 100-ps rise time and 2.5-V swing. The calculated propagation delays are τ = 77.6, 90.6, and 164 ps/cm, respectively.

Solid GND (3C)

Simulation method	*Delay* (ps)	Δ (%)	*Rise time* (ps)	Δ (%)	*NE noise* (mV)	Δ (%)	*FE noise* (mV)
$R(f)L(f)C$	366	—	778	—	192	—	−13
RLC	352	−4	783	1	150	−22	−6
50 *RC* sections	346	−6	832	7	133	−31	<1

Orthogonal GND (2A)

Simulation method	*Delay* (ps)	Δ (%)	*Rise time* (ps)	Δ (%)	*NE noise* (mV)	Δ (%)	*FE noise* (mV)
$R(f)L(f)C$	390	—	804	—	307	—	−109
RLC	375	−4	814	1	211	−31	−58
50 *RC* sections	380	−3	916	14	185	−40	<1

Parallel far GND (3A)

Simulation method	*Delay* (ps)	Δ (%)	*Rise time* (ps)	Δ (%)	*NE noise* (mV)	Δ (%)	*FE noise* (mV)
$R(f)L(f)C$	305	—	254	—	754	—	−630
RLC	290	−5	201	−21	749	−1	−613
50 *RC* sections	210	−31	515	103	287	−62	<1

network overestimates delay by 8%, while a π-network underestimates by 8% only at a one-section representation. Even for the three-section case, the π-network error is only −1.7%. Comparison was made for the delay and rise-time overestimation for three lines, namely from sites 4B′, 3C, and 3B with R = 501.5 Ω/cm, R = 167.0 Ω/cm, and R = 35.5 Ω/cm, when modeled with ten- or 50-section *RC* circuit. The largest error was only 7.6% and was negligible for the most resistive lines ($Rl = 10Z_0$).

It was pointed out earlier (Figure 2) that narrow lines, with $R \geq 500$ Ω/cm, should not be used for lengths greater than about 1 cm. If such lines are replaced by less resistive lines on a thicker top layer, such as M5 (R = 35.5 Ω/cm) at 1.6 cm, there is a delay reduction of 660 ps, which represents a significant portion of the processor cycle time.

Figure 15 shows the *NEN* prediction for lines in sites 4B′, 3C, and 3A, with R = 501.5, 167.0, and 35.5 Ω/cm. The very resistive, 0.9-μm lines have rise times of 109–2420 ps for lengths of 0.2 to 1.6 cm; as a consequence, although K_B is very high (0.18), crosstalk is only 7–12%. The lines of 3C have lower K_B of 0.08, the rise times are 53 to 778 ps, and the delay barely exceeds $t_r/2$. The noise is saturated for length greater than 0.5 cm and is only 4–8%. Finally, the least resistive lines, with very high K_B of 0.20 and R = 35.5 Ω/cm, have rise times of 48–254 ps. Since the delay is much greater than $t_r/2$, the noise saturates easily and has the range of 24–30%.

The strong capacitive coupling between the very narrow and closely spaced lines of site 4B′ (0.9-μm width and separation) affects the delay of adjacent lines as well. **Figure 16** illustrates this effect for three lines being driven by one positive transition only, by three positive transitions, and by two positive and one negative transition on the center line. When the three lines are driven in

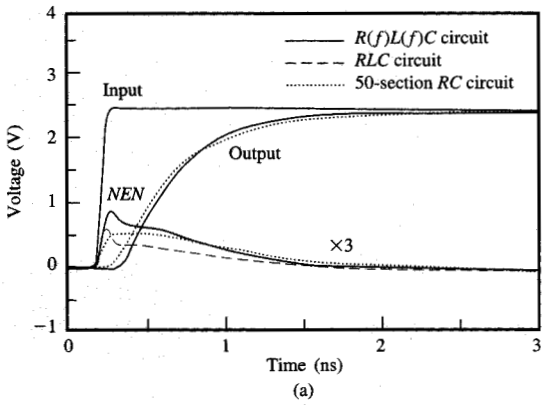

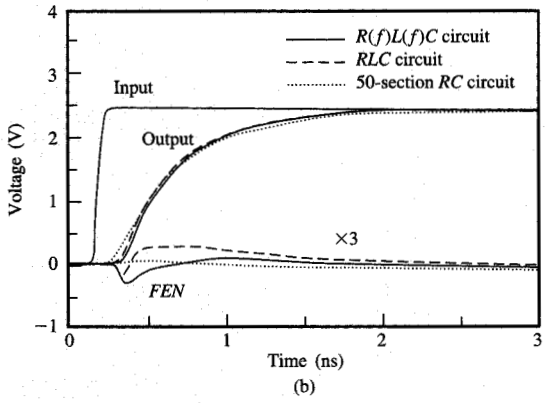

Figure 14

Simulated propagated signal and (a) *NEN* and (b) *FEN* crosstalk using the circuit of Figure 12 for 2.7-μm-wide, 1.6-cm-long lines on M4 having orthogonal wide ground buses on M5 (site 2A). The input circuit has 10-Ω impedance and 100-ps rise time. Simulation was done using frequency-dependent $R(f)$ and $L(f)$ (solid curves), constant R and L parameters (dashed curves), and 50-section distributed RC representation (dotted curves).

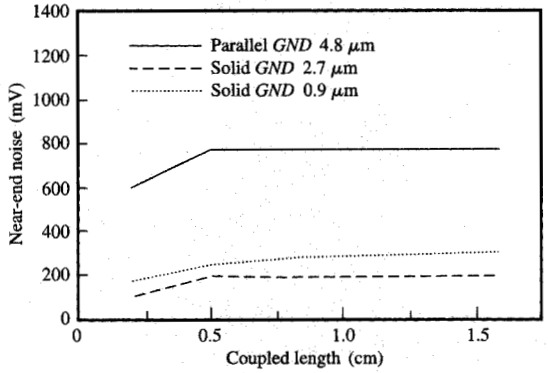

Figure 15

Simulated *NEN* crosstalk using the circuit of Figure 12 for lengths l = 0.2 to 1.6 cm for 4.8-μm (solid curve, site 3A), 2.7-μm (dashed curve, site 3C), and 0.9-μm (dotted curve, site 4B′)-wide lines. The input source has 10-Ω impedance and 100-ps rise time.

phase, the mutual capacitance is not charged, and the effective capacitance of the center line is $C_{22} - 2C_{12}$. When the center line is driven out of phase with the outer lines, the effective capacitance is $C_{22} + 2C_{12}$. For opposite-polarity signals, delay is increased by 45% compared to the single-transition case. Similarly, delay is decreased by 29% when equal polarities are used ($K_C = 0.24$).

Summary and recommendations

A rigorous study has been made of the characteristics of representative long, on-chip interconnections. The performance of such lines is becoming more important as most of the high-performance processor and cache circuits are integrated on very large chips. Signal propagation measurements were compared with results from three-dimensional modeling. Simulation of these interconnections represented as transmission lines with frequency-dependent losses gave excellent agreement with test results performed with an initial pulse excitation of 35-ps rise time. The experimental verification provided validation of the modeling and simulation techniques. These detailed models were then used to predict the performance of such lines when driven as practical on-chip interconnections. The results of these simulations were used to formulate guidelines for wiring practices and cross-section design. Here we list some of the recommendations.

It was confirmed experimentally that shrinking the line cross section in order to provide high density results in greatly deteriorated performance as line length is increased. Simulations showed that 0.25- and 0.50-cm-long, 0.9-μm-wide lines, driven by 1-kΩ gates, exhibit delays of 457 and 814 ps, which is unacceptable. Repeaters can be used, at the cost of additional chip area and complexity. In order to improve performance, low-impedance on-chip driver circuits should be used. For a 10-Ω driver circuit, the 0.5-cm line would then have a delay of 176 ps. The next step is to sacrifice some density and use wider lines, especially for cycle-time determining connections. The low-resistance lines can propagate fast signals over longer lengths. A 1.6-cm-long line with R = 35.5 Ω/cm has a 305-ps delay compared with a line with R = 501.5 Ω/cm,

for which the delay is 975 ps. For these lines, the rise times are faster than the propagation delay, and $Rl \cong Z_0$. In this case, the full transmission-line properties should be taken into account, as such lines can start to exhibit an LC-type behavior. To assess the system performance impacts using a specific semiconductor technology, these models must be extended to include actual driver and receiver circuits.

Very lossy lines can be represented by an RC circuit of the order of ten subsections. A single-section representation always overestimates delay. The distributed ten-section RC representation gives adequate delay prediction for lines that are narrower than 3 μm or $R > 60$ Ω/cm. Tightly spaced narrow lines have large variations in delay caused by strong capacitive coupling to neighboring lines, and crosstalk is significant. The less resistive transmission lines can be accurately modeled as distributed $R + L + C$ networks. The more nonuniform the current return path through the reference mesh, the more important is the need to include the inductive matrix. Although skin effect is not fully developed in lines of such small cross section, frequency-dependent modeling improves accuracy. The frequency-dependent R and L losses are more significant for the current in the lossy reference lines, and contribute to increased delay and crosstalk. A simplified distributed RC circuit underestimates delay by as much as 30% in some cases, and crosstalk by 10–70%.

To decrease crosstalk, the line-to-line separation should be larger than the line width, thus enabling the full advantage of the speed-up generated by using lower-R lines to be obtained. This is commonly the case for transmission lines used off-chip, on either the chip carrier or the printed-circuit board. Alternatively, coupled length restrictions can be imposed in the layout tool, similar to packaging practice. Delay can also be minimized by custom routing of critical lines, either under very wide power buses or close to such wide buses, or with adjacent individual wide ground conductors, as in site 1A. A generous supply of narrower power lines in all the other layers will help reduce the inductance. Alternatively, special wide paths can purposely be dedicated for routing long lines where a good reference structure is provided. Delay can be decreased by reducing not only the line resistance but also the resistance of the power mesh.

These practices have long been followed in the design of packaging interconnections with 10–100-cm lengths. Package interconnections have very well controlled transmission-line structures over their entire lengths. The observations of this study also point in the direction of lower R, either from dimensional increase or material change, and a uniform power reference such as that provided by solid planes or regular meshes. Wide lines placed on top of thicker layers, as was proposed in [14], must have planar topographies to take advantage of lower

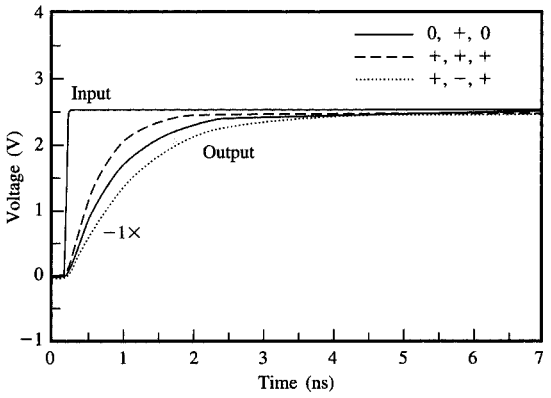

Figure 16

Simulated propagated signals on 0.25-cm-long, 0.9-μm-wide coupled lines on M3 with 0.9-μm separation, and 50% occupancy of orthogonal lines on M2 and M4 layers. Input circuit has 1000-Ω impedance and 100-ps rise time. The excitation had positive-going transition on the center line only (solid curve), positive transition on all lines (dashed curve), and positive transition on the outer two lines and negative transition on the center line (dotted curve); a 0.1-pF load is assumed at all line ends.

R and the proximity to an air interface. Simple two-dimensional modeling and RC-type representations must be replaced by accurate transmission-line analysis for longer, less resistive lines. The lack of regularity in the on-chip wiring leads to the need for layout customization, which might be less efficient or more expensive than the addition of power layers, or the use of lower-resistivity, planarized metallization such as copper, or lower-in-plane-dielectric-constant anisotropic insulators.

The conclusions of this study also indicate the need for further in-depth analysis of the many effects that have been discussed here. A larger range of line lengths, cross sections, line-to-line separations, and loading and referencing arrangements must be investigated. The study must be performed with various driver and receiver circuits having a range of driving impedance and initial transitions. On the basis of such analyses, better guidelines can be formulated for the ranges where RC- or RLC-circuit representations are adequate, where different layouts are beneficial, or where improvements in technology are needed. This type of understanding can then give direction to new chip layout, parameter extraction, and timing prediction software tools, which are extremely deficient at present for use in designing long, wide on-chip interconnections.

Acknowledgments

The authors are extremely grateful to V. D. Klimanis and D. R. Knebel from IBM Poughkeepsie for their great help with the smooth transition of the digitized design layout to the mask generation and fabrication team. This work was supported by ARPA HSCD Contract No. C-556003.

References

1. W. J. Bowhill, R. L. Allmon, S. L. Bell, E. M. Cooper, D. R. Donchin, J. H. Edmondson, T. C. Fischer, P. E. Gronowski, A. K. Jain, P. L. Kroesen, B. J. Loughlin, R. P. Preston, P. I. Rubinfeld, M. J. Smith, S. C. Thierauf, and G. M. Wolrich, "A 300MHz 64b Quad-Issue CMOS RISC Microprocessor," *Digest of 1995 IEEE International Solid-State Circuits Conference, ISSCC'95*, February 1995, pp. 182–183.
2. R. E. Matick, *Transmission Lines for Digital and Communication Networks*, McGraw-Hill Book Co., Inc., New York, 1969.
3. D. D. Ling and A. E. Ruehli, "Interconnection Modelling," *Advances in CAD for VLSI, 3, Part II, Circuit Analysis, Simulation and Design*, A. E. Ruehli, Ed., North-Holland Publishing Co., Amsterdam, 1987, pp. 211–291.
4. R. E. Collin, *Field Theory of Guided Waves*, McGraw-Hill Book Co., Inc., New York, 1960, pp. 343–373.
5. A. Deutsch, G. V. Kopcsay, V. A. Ranieri, J. K. Cataldo, E. A. Galligan, W. S. Graham, R. P. McGouey, S. L. Nunes, J. R. Paraszczak, J. J. Ritsko, R. J. Serino, D. Y. Shih, and J. S. Wilczynski, "High-Speed Signal Propagation on Lossy Transmission Lines," *IBM J. Res. Develop.* **34,** 601–615 (1990).
6. T. L. Quarles, "Spice3 Version 3c1 Users' Guide," *Memorandum No. UCB/ERL M89/46*, University of California at Berkeley, April 1989.
7. *Advanced Statistical Analysis Program (ASTAP), Program Reference Manual*, Order No. SH20-1118-0; available through IBM branch offices.
8. B. J. Rubin and S. Daijavad, "Calculation of Multi-Port Parameters of Electronic Packages Using a General Purpose Electromagnetics Code," *Proceedings of the IEEE 2nd Topical Meeting on Electrical Performance of Electronic Packaging, EPEP'93, 1993*, pp. 37–39.
9. J. E. Schutt-Aine and R. Mittra, "Scattering Parameter Transient Analysis of Transmission Lines Loaded with Nonlinear Terminations," *IEEE Trans. Microwave Theory Tech.* **36,** 529–536 (1988).
10. A. J. Gruodis and C. S. Chang, "Coupled Lossy Transmission Line Characterization and Simulation," *IBM J. Res. Develop.* **25,** 25–41 (1981).
11. V. K. Tripathi and A. Hil, "Equivalent Circuit Modelling of Losses and Dispersion in Single and Coupled Lines for Microwave and Millimeter-Wave Integrated Circuits," *IEEE Trans. Microwave Theory Tech.* **36,** 256–262 (1988).
12. V. A. Ranieri, A. Deutsch, G. V. Kopcsay, and G. Arjavalingam, "A Novel 24-GHz Bandwidth Coaxial Probe," *IEEE Trans. Instr. & Meas.* **39,** 504–507 (1990).
13. A. Deutsch, G. Arjavalingam, and G. V. Kopcsay, "Characterization of Resistive Transmission Lines by Short-Pulse Propagation," *IEEE Microwave & Guided Wave Lett.* **2,** 25–27 (1992).
14. G. A. Sai-Halasz, "Performance Trends in High-End Processors," *Proc. IEEE* **83,** 20–36 (1995).

Received May 24, 1995; accepted for publication October 11, 1995

Alina Deutsch *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (DEUTSCH at YKTVMV, deutsch@watson.ibm.com).* Mrs. Deutsch received the B.S. and M.S. degrees in electrical engineering from Columbia University and from Syracuse University, respectively. She has been at IBM since 1971 and has worked in several areas, including testing of semiconductor and magnetic bubble memory devices. She is a senior engineer currently working on the design, analysis, and measurement of packaging and VLSI chip interconnections for future digital processor and communication applications. Mrs. Deutsch has written 29 technical papers, holds seven patents, and has received Outstanding Technical Achievement and Research Division Awards from IBM in 1990 and 1993. She is co-chair of the IEEE Topical Meeting on Electrical Performance of Electronic Packaging, a member of Tau Beta Pi and Eta Kappa Nu, and a senior member of the IEEE.

Gerard V. Kopcsay *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (KOPCSAY at YKTVMV, kopcsay@watson.ibm.com).* Mr. Kopcsay is a research staff member in the Packaging Technology Department. His current research interests include measurement and analysis of high-speed pulse phenomena, and technology for high-performance computers. Since joining IBM Research in 1978, he has worked on the design and analysis of computer packages. He was manager of the Package Engineering Department from 1984 to 1988. Prior to joining IBM, he participated in research and development of low-noise microwave receivers at the AIL Division of Eaton Corporation. Mr. Kopcsay received the B.E. degree from Manhattan College in 1969, and the M.S. degree from the Polytechnic Institute of Brooklyn in 1974, both in electrical engineering. He is a member of the IEEE and the American Physical Society.

Christopher W. Surovic *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (CWS at YKTVMV, cws@watson.ibm.com).* Mr. Surovic is pursuing his associate degree in mechanical science and engineering at Westchester Community College. He has been at IBM since 1989 and has worked in the area of tool and model making for about four years, making significant contributions to the development of laser ablation tools and the Smithsonian X-Ray Telescope. Mr. Surovic is a senior laboratory technician, currently working on electrical characterization and design of package structures for future computer applications. He has co-authored several technical papers in package measurements.

Barry J. Rubin *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (BRUBIN at YKTVMV, brubin@watson.ibm.com).* Dr. Rubin received a bachelor's degree in electrical engineering from the City College of New York in 1974. He then joined IBM at its East Fishkill facility in Hopewell Junction, New York. In 1978, he received an M.S. from Syracuse University, and, in 1982, a Ph.D. from the Polytechnic Institute of New York, both in electrical engineering. Dr. Rubin has worked on power transistor design, CCD technology, circuit design, and, since 1976, on electrical package analysis. He has been closely involved in the following aspects of TCM packages: signal propagation, coupled and delta-I noise, and verification of modeling. In 1986, he joined the Thomas J. Watson Research Center, where he is currently involved in the electromagnetic modeling of computer packages. Dr. Rubin has received two IBM Invention Achievement Awards.

Lewis M. Terman *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (TERMAN at YKTVMV).* Dr. Terman received the B.S., M.S., and Ph.D degrees from Stanford University in 1956, 1958, and 1961, respectively. In 1961 he joined the IBM Research Division, working first in system logic design and then on magnetic film memory development. In 1964 he began working on semiconductor memory, especially with MOS technology, and became one of its very earliest advocates, participating in the program which resulted in the first IBM MOS memory products. In the 1970s Dr. Terman worked on MOS and CCD memory, MOS logic, and MOS analog circuits. In 1979–1980 he was on the Technical Planning Staff of the Director of Research. From 1981 to 1991 he directed groups doing research on circuits, devices, and technology for advanced MOS memory and logic, and from 1988 to 1991 he was manager of VLSI Logic and Memory, directing four groups in CMOS logic and memory design and technology. From 1991 to 1993 he was a senior member of the Technical Plans and Controls staff of the IBM Research Division, and he is currently a program manager in the CMOS high-speed circuit design program. Dr. Terman is a past president of the IEEE Electron Devices Society. He is vice president of the IEEE Solid-State Circuits Council, was treasurer of the Council in 1988–1989, and was a member of the IEEE Circuits and Systems Society AdCom from 1981 to 1983. He was editor of the IEEE *Journal of Solid-State Circuits* from 1974 to 1977. He was chairman of the 1983 International Solid-State Circuits Conference, chairman/co-chairman of the 1985 and 1986 Symposium on VLSI Technology and the 1988 and 1989 Symposium on VLSI Circuits, and co-chairman of the Symposium on VLSI Technology, Systems, and Applications in Taiwan, ROC, from 1985 to 1995. Dr. Terman was chairman of the 1994 Symposium on Low Power Electronics. He was a member of the IEEE Technical Activities Board Administration Council from 1991 to 1994, and Chair of the TAB Technical Meetings Council in 1993–1994; he is currently the IEEE TAB Treasurer. Dr. Terman has received six IBM Outstanding Contribution and Outstanding Invention Awards and three IBM Corporate Awards. He was elected to the IBM Academy of Technology in 1991 and to the Academy Technology Council in 1995. He is co-chair of the Academy Technical Program Committee. Dr. Terman is the author or co-author of 35 technical articles and 24 U.S. patents. He is a Fellow of the IEEE and of the American Association for the Advancement of Science. He is the 1995 recipient of the IEEE Solid-State Circuits Technical Field Award for "leadership in the field of MOS devices and circuits for semiconductor memories."

Richard P. Dunne, Jr. *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (DUNNE at YKTVMV, dunne@watson.ibm.com).* Mr. Dunne is a senior laboratory specialist who joined the Maintenance Department at the Thomas J. Watson Research Center in 1967. After earning a B.S. degree in electrical engineering from the University of Bridgeport in 1976, he joined the Josephson Superconducting Devices group. He used the IBM IGS Design System to create mask data for the fabrication of Josephson-junction memory and logic devices. Mr. Dunne fabricated the test vehicles, performed the device and circuit tests, and maintained the cryogenic equipment. In 1983, he joined the GaAs Logic, Memory, and Power Circuit Design group, where he tested high-speed GaAs devices and circuits. Since 1993 Mr. Dunne has worked in the VLSI Design Group, using the CADENCE program for circuit design and layout.

Thomas A. Gallo *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (TAG at YKTVMV, tag@watson.ibm.com).* Mr. Gallo began his career with IBM as a co-op student in 1983. In 1985 he became a regular employee working in the area of silicide formation through the deposition of perfectly clean Si/metal interfaces. He has built his own process chambers and learned RBS, X-ray diffraction, and SEM for materials analysis. For the last six years he has performed failure analysis on most IBM BEOL work for logic and memory using optical, SEM, and electrical measurements. Mr. Gallo received an A.A.S. degree in electrical technology in 1985 from Westchester Community College and a B.S. degree in physics/mathematics in 1991 from Pace University. He is an associate member of the IEEE and a member of Tau Alpha Pi.

Robert H. Dennard *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (DENNARD at YKTVMV).* Dr. Dennard received the B.S. and M.S. degrees in electrical engineering from Southern Methodist University, Dallas, Texas, in 1954 and 1956, respectively, and the Ph.D. degree from the Carnegie Institute of Technology in 1958. Subsequently he joined the IBM Research Division, where his early experience included the study of new devices and circuits for logic and memory applications, and the development of advanced data communication techniques. Since 1963, he has been located at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, where he has been involved in microelectronics research and development since its beginning. His primary work has been in field-effect transistors (FETs) and the integrated digital circuits that use them. In 1967 he invented the dynamic RAM memory cell used in most computers today. With others, he developed the concept of FET scaling in 1972. In 1979 he was appointed an IBM Fellow. Dr. Dennard is a Fellow of the IEEE and received the IEEE Cledo Brunetti Award in 1982. He was elected to the National Academy of Engineering in 1984. Dr. Dennard received the National Medal of Technology from President Reagan in 1988 for his invention of the one-transistor dynamic memory cell. He also received the I.R.I. Achievement Award from the Industrial Research Institute in 1989, and the Harvey Prize from Technion in Haifa, Israel, in 1990.