Physical and electrical design features of the IBM Enterprise System/9000 circuit module

by E. E. Davidson
P. W. Hardin
G. A. Katopis
M. G. Nealon
L. L. Wu

The requirements of the new TCM (thermal conduction module) for the IBM Enterprise System/9000™ (ES/9000™) module generated a significant number of challenges for the physical and electrical designer. For example, the need to support more circuits meant that more signal and power conductors had to be provided. In addition, the requirement for faster performance called for materials with lower dielectric constants and the use of on-module decoupling capacitors. This paper describes these changes, the design considerations that were applied to signal transmission, and the approaches that were used to contain delta-*I* and crosstalk noise in the module. Finally, the test measurements used to qualify the module are explained. The result is a TCM that more than doubles the circuit density of the TCM used for the IBM 3090™ machines, with substantially greater speed and reliability.

Introduction

To achieve the fastest performance, nearly all contemporary supercomputers and mainframes use high-speed, high-power bipolar chips on multichip modules (MCMs) [1–5]. The use of bipolar emitter-coupled logic (ECL) for mainframes has been commonplace since the early '60s, but the nearly universal acceptance of the MCM as the first- or second-level package has occurred only within the last two years. Perhaps the most important reason for this is that the high circuit densities achievable on an MCM are a prerequisite for the shorter packaging delays that are necessary for faster cycle times [6]. The breakpoint between being able to use a single-chip module (SCM) on a printed circuit board package and having to use an MCM appears to be around a 10-ns cycle or a 100-MHz clock rate, as illustrated by the Fujitsu Corporation's change from the approximately 10-ns cycle time in the SCM-based M-780 mainframe to the 10-ns cycle time in the MCM-based M-1800 machine.

During the '80s, IBM took the MCM approach, while other manufacturers of high-end machines opted for the

Table 1 Processor semiconductor and package attributes for the 3090J and the new ES/9000 Model 820 and Model 900 machines.

	3090J	*ES/9000*
SEMICONDUCTORS		
Logic		
Circuit type	ECL	ECL, DCCS*
Maximum circuits	2360	5620
Signal I/O	96 and 180	228
Total I/O	180 and 264	648
Maximum power (W)	9	27
Chip size (mm^2)	5.2 × 5.2	6.5 × 6.5
Memory		
Capacity (Kb)	32	64
Nominal access (ns)	3.5	2.5
Chip size (mm^2)	6.2 × 6.2	7 × 8.5
MODULES		
Chip sites	110	100 and 121
Substrate size (mm^2)	110 × 117	127.5 × 127.5
Chip pitch (mm)	8.5	10.8 and 9.9
Signal pins	1200	1782
Total pins	1800	2772
Maximum power (W)	850	2000
Maximum circuit density (circuits/mm^2)	20	42

*Differential cascode current switch (see [14]).

more readily available SCM package technologies. Recently, however, all of the major mainframe manufacturers have begun to pursue the MCM approach.[1] The design points differ significantly, but the same technology elements are present in all of the new high-end machine definitions. This paper describes the first major upgrades to the original IBM MCM, the thermal conduction module (TCM). The TCM was first described in 1980 as the major technology innovation in 308X machines [7].

The original TCMs consisted of 33 layers of alumina ceramic. They were 90 mm square and contained up to 133 chip sites. Direct chip attachment (DCA) using the C4 (controlled collapse chip connection) technique was employed to achieve small chip pitches and high circuit density.

The original TCM served IBM well, and was enhanced for the 3090 machine series. During its tenure, cycle times were reduced by a factor of 1.8, from 26 ns to 14.5 ns, and significant manufacturing learning led to a greatly diminished cost per chip site. Reliability also proved to be impressive for the TCM substrate: Not a single failure of the substrate or the C4 connections has ever been documented.

[1] Although the MCM packages in the new Japanese machines might be regarded as SCMs on a ceramic board, the authors choose not to subscribe to this view. We consider these new ceramic structures, some with and some without thin-film layers, to be characteristic of an MCM without direct chip attachment. Instead, interposers are mounted on top of the MCM as chip carriers and as a means for redistribution of chip I/O to the module grid.

However, at clock rates greater than 100 MHz, the current TCM technology is not adequate for the higher frequency and shorter delay requirements of the new ES/9000 Models 820 and 900.[2] Some of the shortcomings are caused by the high dielectric constant (9.5) of alumina, which leads to loadings and propagation delays that are too high for these higher-speed machines. In addition, these machines require better control of noises such as reflection, crosstalk, and delta-*I*. Finally, larger chips with more circuits require better module wiring density and better thermal fatigue characteristics for the C4 joints. **Table 1** shows the major packaged electronics differences between the 3090J machine and the new ES/9000 modules.

TCM enhancements

The new ES/9000 chip design point and the new faster cycle time requirement called for an entire set of TCM enhancements, as enumerated in Table 1. Almost all areas of module technology had to be improved, including chip I/O density, total module wirability, cooling capacity, power distribution, delay characteristics, and noise control.

The major new change for the ES/9000 module is the introduction of a cordierite glass-ceramic dielectric material with a lower dielectric constant (5.2) and a better thermal expansivity (29×10^{-7}/°C) match with silicon (26×10^{-7}/°C) to minimize stresses in the C4 joints than the alumina substrate used in the previous TCMs. The new glass-ceramic material reduces module propagation delay by 25 percent and the capacitive loading effects by 46 percent compared to the old alumina material.

Another technology innovation in the module is the use of partial thin-film redistribution (PTFR). To accomplish this, pads and thin-film wires are sandwiched between two layers of polyimide, which are placed on top of the multilevel glass-ceramic substrate. Approximately one half of the 228-signal I/O is redistributed on the PTFR layer to engineering change (EC) pads and vias that connect to the lines in the glass-ceramic region. The rest of the signal I/O is redistributed in glass-ceramic layers. The net impact of the PTFR is to significantly reduce the number of glass-ceramic redistribution layers. Fewer layers in this region of the module result in less capacitive loading between the chip output and the module transmission lines in the glass-ceramic substrate.

Additional module improvements were required to support more signal I/O on the new chips without greatly increasing the size of the chip site. The new module contains an array of C4 pads on 0.225-mm centers and a module substrate via pitch and signal wire pitch that has been reduced from the 0.5-mm value used in the 3090 TCMs to the 0.45-mm value used in the new TCM for

[2] For the purposes of this paper, further reference to the ES/9000 module will imply the particular one used in Models 820 and 900.

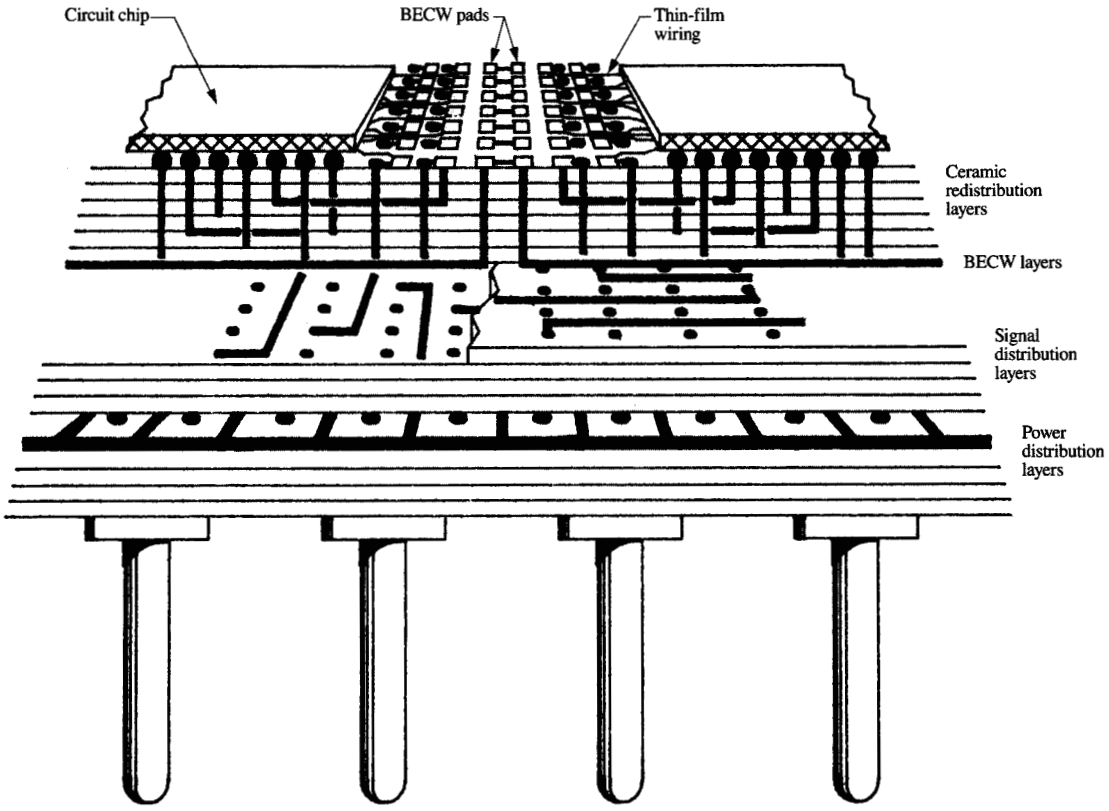

Figure 1

View of the ES/9000 module showing the top-surface features and the key regions of the cross section. From [1], reproduced with permission; © 1991 IEEE.

improved wiring capacity. To satisfy the increased wiring demand, a total of fourteen signal plane pairs are now used instead of the eight plane pairs in the 3090 module.

In addition, delta-*I* noise containment has been improved by the introduction of top-surface decoupling capacitors, which are placed in the corners of every chip site. Finally, the use of a wiring material with a lower dielectric constant has helped to lower crosstalk noise by allowing the dielectric thickness to be reduced without lowering the transmission-line characteristic impedance.

The net result of all of these improvements for the ES/9000 module is an enhanced packaged-electronics design point that has more than doubled the maximum module circuit density compared to a 3090J module. From a system perspective, this technology supports the performance requirements of the new machines.

The next section presents a more comprehensive physical and electrical description of the new module features.

Physical design

Figure 1 depicts a simplified cross-sectional view of the module substrate. The chips are C4-bonded directly to the surface through a partially depopulated 27 × 27 array of solder bumps. More than half of the signal I/O is directly connected to the top-surface engineering change (EC) pads through copper thin-film wire (TFW) on the PTFR layer. The remainder of the signals are redistributed to their EC pads in four glass-ceramic layers placed near the surface.

Because of the synergism between the thin-film and thick-film (glass-ceramic) redistribution techniques, the thick-film via punching need only occur on a staggered 0.45-mm grid instead of the straight 0.225-mm grid used for the C4 pads. As previously mentioned, the use of PTFR also reduces by a factor of two the number of thick-film redistribution layers that would have been required.

The four glass-ceramic redistribution layers are shown schematically in **Figure 2**. Below these layers are voltage planes and signal-plane pairs (P/P). Each signal P/P is a

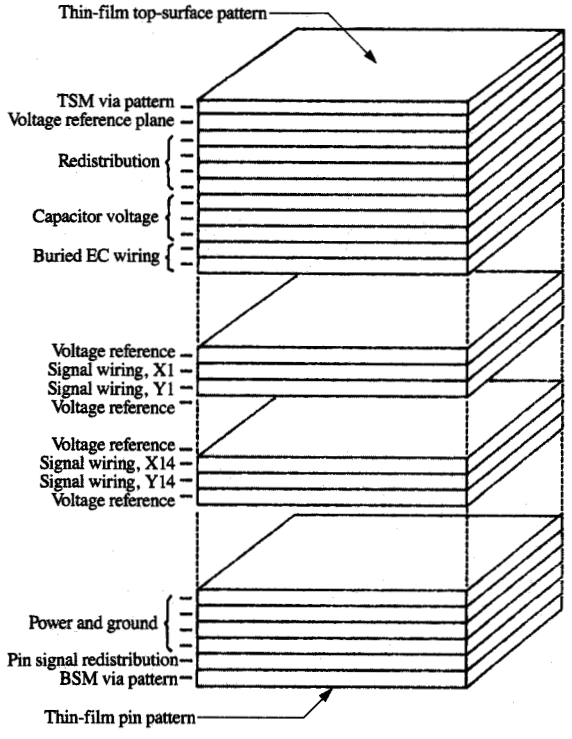

Figure 2

Cross-sectional view of the ES/9000 module substrate. From [1], reproduced with permission; © 1991 IEEE.

tri-plate structure; i.e., it consists of an orthogonal set of two thick-film signal wiring planes sandwiched between two mesh power planes. All of the conductors for the signal and power planes are made from a copper paste. The tri-plate structure provides good characteristic impedance control as well as good crosstalk control. Below the signal section are more power distribution layers. Finally, 2772 pins are attached to the bottom of the module on a staggered 2.5-mm grid. **Table 2** lists many of the physical and electrical characteristics for the substrate.

The chip site on the top surface of the module is illustrated in **Figure 3**. There are a total of 648 active signal, power, and ground C4 pads shown in the center of the diagram. Surrounding the chip pads are four rings of EC pads totaling 359 pads. Of the EC pads, 228 are signal pads, 39 are power pads, and 72 are for connections to buried engineering change wires (BECW). In addition, four EC pads are connected to module pins for board ECs, and 16 EC pads are used as tie-downs for holding surface discrete wires in place. (Board ECs are used to connect a chip I/O directly to a module pin so that a board line can be routed to a new chip I/O.) The thin-film wires for the PTFR scheme are also depicted. In each corner of the chip site is a decoupling capacitor that is shared by the three abutting chip sites.

Table 2 ES/9000 module physical and electrical parameters.

Physical	
Number of layers	63
Number of signal layers	28
TFW dimension W × H (μm)	28 × 5.5
TFW pitch (μm)	55
C4 pad pitch (μm)	225
Thick-film line dimensions W × H (μm)	89 × 25
Thick-film line pitch (μm)	450
Thick-film via pitch (μm)	450
Pin pitch (mm)	2.5 staggered
Electrical	
Relative dielectric constant	5.2
Thick-film characteristic impedance (Ω)	58
Nominal thick-film resistance (mΩ/mm)	25
Nominal thin-film resistance (mΩ/mm)	150
Nominal thin-film capacitance (fF/mm)	60
Simultaneous switching per chip	50–100
Saturated near-end crosstalk (%)*	8

*For nine coupled lines as depicted in Figure 6.

Connection between the decoupling capacitors and the chips is made by a set of voltage planes located in the redistribution section of the module, as shown in Figure 2. Each capacitor contains four sections that are used to decouple four supply voltages (including ground) in the corner of each chip site. Both the capacitor and the chip have multiple sets of connections which, in conjunction with the internal-capacitor voltage planes, create a connection between the two with a very low inductance [8]. The value of this is explained in a later section when the control of switching noise is discussed.

There are two versions of the ES/9000 module, one containing 100 chip sites and the other 121 chip sites on a smaller chip pitch. They are both 127.5 mm square. The main difference is that the 100-chip-site module has a full set of EC pads, while the 121-chip-site module has enough EC pads for only two thirds of the signal count. This design trade-off was implemented to provide the additional chip sites required for certain modules that contain large memory functions. The 100-chip module contains 121 capacitors and the 121-chip module contains 144 capacitors.

The original TCM used discrete wires on the top of the module substrate to effect EC connections. The new ES/9000 modules use a new scheme called buried engineering change wires (BECW) to implement ECs [9]. The BECW scheme employs a special signal-plane pair below the redistribution and capacitor voltage interconnection-plane pair layers, as indicated in Figure 2.

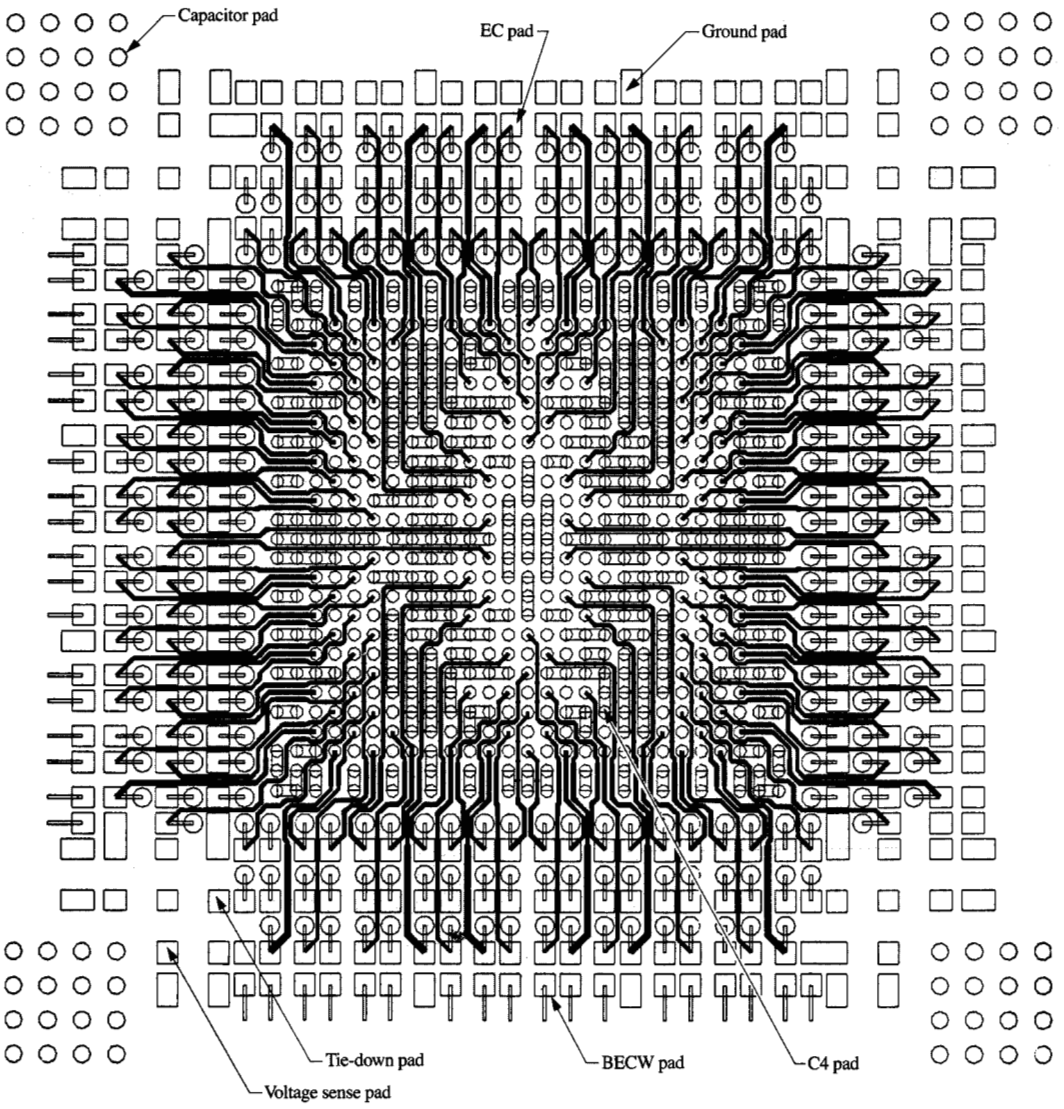

Figure 3

Top-surface view showing the thin-film redistribution wire pattern and the engineering change pad pattern. Capacitor pads are shown in the corners. From [1], reproduced with permission; © 1991 IEEE.

The older, discrete-wire scheme for ECs was abandoned because the faster signals in the new machine require better control of characteristic impedance and crosstalk than could be realized with that method. Other techniques such as top-surface twisted-pair and coaxial wires were also investigated, but BECW was judged to be the best trade-off between electrical and manufacturability requirements. The problem with using a controlled-impedance wire on top of the module was that there was insufficient room for all of the ground pads required to properly terminate the transmission lines on the top surface.

Access to the BECW signal lines is through special top-surface pads and vias, as shown in Figures 1, 3, 4, and 5. After connection has been made to a BECW circuit, the line is routed as shown in **Figure 4**. Whenever the direction

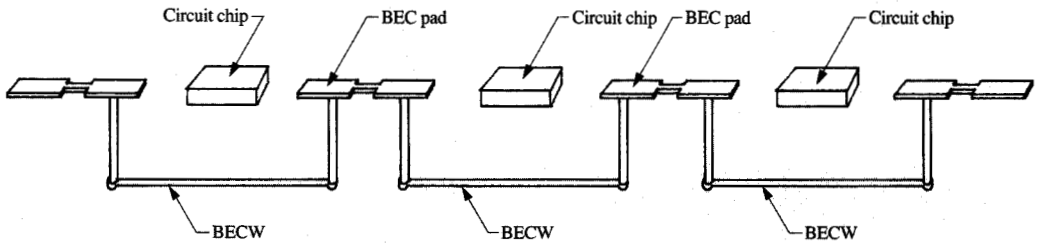

Figure 4

Buried EC wire signal path. From [1], reproduced with permission; © 1991 IEEE.

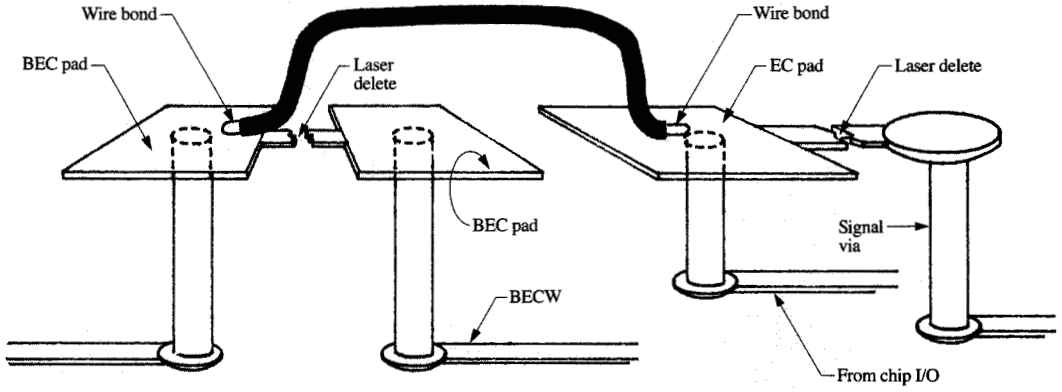

Figure 5

Structure for deleting a connection to an EC pad and adding a wire connection to a buried EC pad. From [1], reproduced with permission; © 1991 IEEE.

of the line is changed, the delete strap between BECW pads is broken, and a short surface discrete wire is connected to a pad which connects to an orthogonal line on the other BECW layer.

The procedure to effect an EC is illustrated in **Figure 5** and is described as follows:

1. After choosing the relevant EC and BECW pads, laser-delete the straps to both pads. (The strap connected to the EC pad goes to a signal via that is connected to a regular signal line, as shown in Figure 1. Similarly, the strap connected to the BECW pad is connected to a pad that goes to a BECW line, as shown in Figure 4. Both have to be broken, as shown in Figure 5.)
2. Bond a short discrete wire between the two pads.
3. Route the EC wire in the BECW layers using top-surface discrete wires to change direction as necessary.

This completes the physical description of the module substrate. The next section describes the electrical design specifics, including noise control and net design.

Electrical design

The goal of the electrical design of the module is to meet the performance and functionality requirements for the machine. This design work affects both the physical design of the module and the hardware design of the machine. Many design conflicts arise from the need to make both

a fast machine and a manufacturable module, and the electrical designer often acts as a mediator to resolve such conflicts. In the end, design constraints resulting from these trade-offs and compromises are handled with a set of design rules that are used with a computer-aided design (CAD) system to control the final hardware design.

The electrical designers begin by providing a set of specifications to the module physical designers, to ensure that the substrate has the right electrical characteristics for all the application goals. This interaction between electrical and physical designers affects the way the power and signal distribution systems are implemented. Almost all aspects of the substrate physical design are affected. Categories involved include the basic module cross section, line dimensions and spacings, layer thicknesses, decoupling capacitors, pin and via assignments, EC scheme, power-plane geometries, and materials used.

All performance-related parameters, such as signal speed, simultaneous switching, and module delay, had to be enhanced in the ES/9000 module in order to meet the performance requirements. This affected the rules that control noise and signal distribution on the module, as discussed below.

Reliable operation of a high-performance mainframe requires the control of electrical noise. Noise, in general, is a composite of unwanted spurious signals that can jeopardize the functional integrity of the machine. Its operational impact is insidious, and excessive noise must be avoided at all costs [10].

Noise is generated by three basic sources. Reflection noise is created by transmission-line discontinuities on the signal interconnection nets, and is controlled by wiring rules. "Delta-*I*" noise is created by current changes that inevitably occur as circuits switch simultaneously [11]. These current changes induce a voltage transient across the predominantly inductive power-distribution system. The final source of noise, crosstalk, is caused by the electromagnetic interaction between nearby lines and vias.

Each of these sources separately can cause data failures. When they are superimposed, they can seriously impair the operational integrity of the system. Limiting total noise to an acceptable level is a mandatory part of designing a high-performance digital system.

- *Reflection noise*

Wiring rules resulting from circuit simulations are used to design the module interconnection nets so that first-incident switching occurs [12]. To accomplish first-incident switching, reflections are never allowed to penetrate into the receiver threshold band. Even though receiver noise tolerance can be low for other noises while larger reflections are present, it is assumed that reflection noises do not algebraically add to other noises because they generally are not time-coherent with delta-*I* noise and crosstalk. The main impact of reflections is to add net delay.

In general, because large reflections cannot be tolerated, they limit the quality and quantity of allowed net topologies. To minimize reflections, a controlled characteristic impedance is provided by the tri-plate structures in the signal P/P section of the module. These lines are terminated in either resistive or diode terminators that reside on the chips [13]. Even with terminators, reflections do occur because of the capacitive-stub discontinuities that occur on the transmission system between the chip loads. These are caused by the paths between the driver and receiver circuits and the module transmission lines. The contributing components of these capacitances are on-chip lines, PTFR and substrate redistribution lines, and substrate vias.

In the ES/9000 module, signal rise and fall times are approximately half of the corresponding 3090™ values. To maintain similar wiring rules in terms of line lengths, fan-out, and loadings, all stub capacitances had to be cut in half. The use of PTFR and glass-ceramic allowed the designers to achieve this goal.

Many types of net topologies are allowed, including single-load point-to-point, multiple-load near-end cluster, far-end cluster, distributed, and discretely distributed nets [14]. Driver types include emitter followers, push-pull, and bidirectional push-pull. All terminators are contained on the chips. They include 58 Ω to V_T for emitter followers and 58 Ω, 80 Ω, and diodes to ground for push-pulls. (Three power-supply voltages are provided: V_{CC} = 1.4 V, V_T = −0.7 V, and V_{EE} = −2.2 V. The system reference voltage is ground.) The 80-Ω terminator is used for the higher-impedance board and cable nets. Diode terminators are used frequently on the module surface to reduce overall power dissipation.

With this wide variety of net, driver, and terminator types, the technology designers have given the system designers sufficient flexibility to wire between chips without running into major I/O and wiring capacity limits. The wiring rules also have a companion set of delay equations, which are used by the system designers to do timing analysis. Like the wiring rules, the delay equations are derived from circuit simulations. The result for a delay equation is a first-order polynomial where each coefficient is determined by a multivariable linear regression technique from the simulation data. The delay equation represents nominal manufacturing data. Tolerance data are also supplied such that statistical timing analysis can be performed by the system designers.

- *Delta-I (switching) noise*

The most difficult part of doing module electrical design is to control the switching noise environment. Advances in mainframe technology have increased module stress with

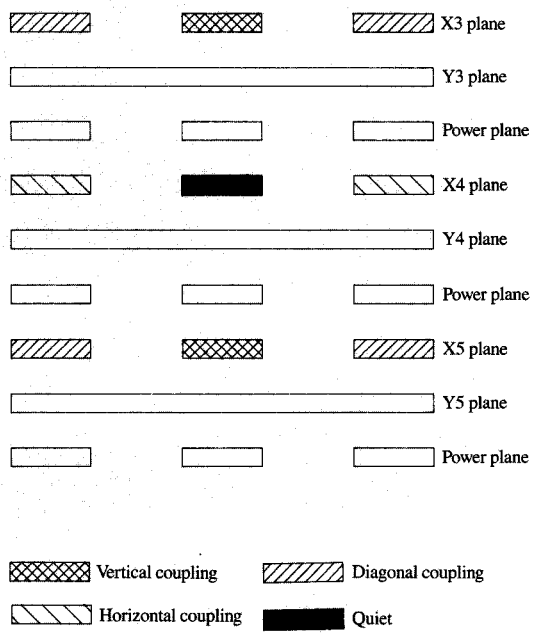

Figure 6

Cross-sectional view of three plane-pairs in the glass-ceramic substrate showing the signal lines that couple energy into a quiet line. From [1], reproduced with permission; © 1991 IEEE.

respect to noise control. Signals keep getting faster, and the larger chip circuit capacities create a requirement for more simultaneous switching. To support the higher circuit densities, higher package densities for I/O and wiring are provided. In addition, the substrate plane count is increased for more wiring capacity. The thicker and denser substrate can worsen the noise situation by driving up effective inductances and via-to-via coupling coefficients arising from power distribution.

To compensate for the ever-greater demands placed upon them, the module electrical designers have had to insist upon physical improvements. They have required thinner dielectrics, more voltage planes, and enough vias for shielding. Partial thin-film redistribution (PTFR) was also required to help to reduce the substrate thickness. In addition to these requirements, the inclusion of on-module decoupling capacitors was specified for the ES/9000 module to reduce overall effective inductance and support larger amounts of simultaneous switching.

As previously described, the decoupling capacitors are placed in the corner of every chip site. Each capacitor comprises four sections in order to provide for complete noise filtering of all the voltages [8]. Each chip quadrant is in close proximity to a decoupling capacitor to keep the connecting inductances very low. Simultaneous switching or delta-*I* noise was more than cut in half by the use of capacitors.

- *Crosstalk*

In a complex structure such as the ES/9000 module substrate, the evaluation and control of crosstalk is a difficult task. The structure in **Figure 6** shows three P/P of signal wiring. The middle conductor in the "X4" level is designated as the "quiet" line. Coupling to the two adjacent X4 conductors is relatively strong; we label this horizontal coupling. There is also signal coupling to the other two "X" direction P/Ps, even though they are separated from the X4 P/P by power planes. The latter, because of all the vias that pass through them, are constructed as orthogonal mesh planes. Electromagnetic field lines pass through power-plane apertures, allowing coupling to the adjacent parallel P/Ps. Coupling to the lines that run in the same direction and are directly above and below the quiet conductor is called vertical coupling. Similarly, the four corner conductors depicted are considered to be diagonally coupled. All of these couplings mean that the designer must pay attention to the eight codirectional lines that surround the quiet line when assessing total crosstalk.

Coupled noise or crosstalk comes in two varieties, near-end (NE), or backward, and far-end (FE), or forward. Near-end crosstalk is caused by the sum of the inductive and capacitive coupling, while far-end crosstalk is caused by the difference between the two [10]. In a homogenous medium (a single dielectric), FE crosstalk is supposed to be zero. Unfortunately, in a structure such as a TCM substrate, the presence of orthogonal vias and lines increases the capacitance of a given line without increasing its inductance. This modifies the normal crosstalk results from homogeneity, and one finds a substantial amount of FE crosstalk [15]. The presence of these other lines and vias (OLV) also lowers the characteristic impedance and the propagation velocity of the signal transmission lines.

Near-end crosstalk has the characteristic that it grows as the coupled line length grows until the coupled length corresponds to the time it takes the signal to propagate half of a rise time. At this point, the crosstalk amplitude saturates, but its pulse width grows as the coupled length continues to increase [10]. **Figure 7** shows the difference in total NE crosstalk for the eight surrounding lines between the 3090 module and the new ES/9000 module. The NE noise is plotted as a function of the coupled length. Note that for the ES/9000 case, the saturation amplitude and length are both less than for the old module. This result is a function of many parameters, including dielectric constant, rise time, and coupled-line geometries. The total NE

crosstalk for the ES/9000 module is approximately 8% of the total signal swing.

Although this percentage seems low enough to avoid a crosstalk problem, it is not the only consideration. In other areas of the module, lines can also couple into the same quiet line. Superposition occurs, and the resultant noise can exceed the NE saturation noise of a single set of coupled lines. Add FE noise to the NE noise and the situation becomes even more untenable. To further complicate the noise problem, delta-*I* noise can add to the crosstalk [10]. The logical consequence of all of this noise superposition is that the designer must monitor it to ensure that data errors due to noise do not occur.

To accomplish this, a CAD tool has been developed which works by keeping track of all crosstalk and delta-*I* noises that can couple onto a quiet line. It calculates the amplitudes of the noises, the timing relationships, and the effect of transmission-line reflections on the coupled and delta-*I* noises. The program keeps a running total as a function of time of the noises at the input of any receiving circuit in the system. It can even take into account statistical amplitude and timing variations caused by manufacturing tolerances.

Whenever the program determines that the receiver noise tolerance has been exceeded, a noise failure is indicated. The designer is then obligated to reduce the delta-*I* noise by cutting back on the number of simultaneously switching drivers and terminators and/or to reduce the crosstalk by rerouting lines.

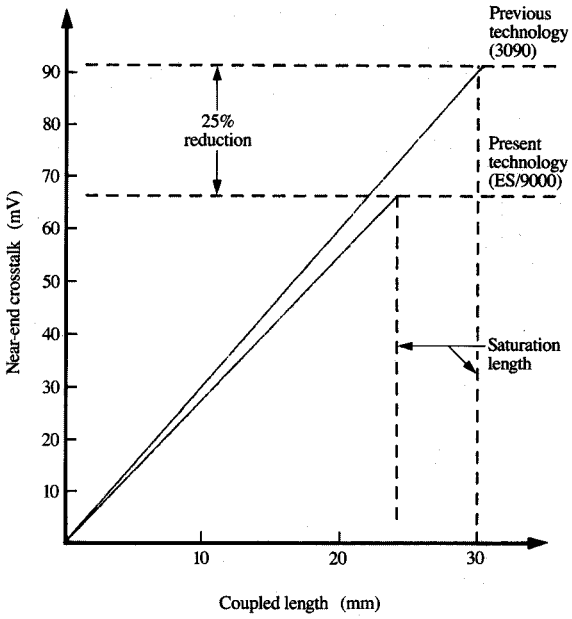

Figure 7

Diagram illustrating the differences in near-end crosstalk as a function of coupled length for the previous 3090 module and the new ES/9000 module. From [1], reproduced with permission; © 1991 IEEE.

• *Power distribution*

Another electrical design concern for the ES/9000 module is power distribution. Each chip site can contain a chip that dissipates up to 27 W of power. Corresponding in and out currents per chip site are of the order of 10 A. Such high currents could easily lead to unacceptably high substrate voltage drops unless an adequate amount of conductive resources in the form of power planes (Figure 2), module pins, vias, and chip connections are provided for power distribution. To ensure that this is indeed the case, an extensive statistical analysis of the potential drops in the substrate is performed. During this analysis, the current drawn by each chip is allowed to vary statistically in an extensive resistive-mesh model of all the power paths in a quadrant of the substrate. The resistances of the paths are also allowed to vary statistically according to the expected manufacturing tolerance. A Monte Carlo simulation is performed to determine the minimum absolute drop and the maximum absolute drop between any two chip sites in the quadrant. The minimum absolute drop is accounted for by increasing the power-regulator voltage setting, while the maximum differential drop is accounted for by including its value in the circuit power supply design tolerances.

This completes the discussion of electrical design. The next section describes the procedure that was used to verify the validity of the design practices.

Product qualification

As previously discussed, a methodology has been used that converts technology information into a format that can be used in the design of the system. This methodology has been adopted to convert electrical design constraints into a set of rules which are used to control a set of DA tools that implement the physical design of the TCM. To ensure a high-quality design, the entire design methodology was tested by evaluating a product qualification vehicle (PQV). The PQV (**Figure 8**) contained a sophisticated set of experiments that were designed to stress all of the design rules.

The PQV consisted of experiments to verify that the wiring rules, delay equations, and noise rules all described the actual hardware accurately. Other test vehicles were also used to characterize the substrate electrical parameters—to measure redistribution and via capacitances, line characteristic impedances, and propagation velocities with and without OLV effects.

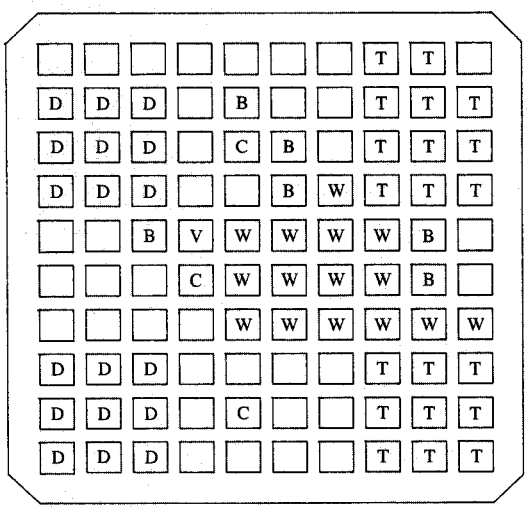

B: Chip for board wiring rule, delay equation, and delta-*I* experiments
C: Control chips
D: Driver chips for module delta-*I* experiments
T: Terminator chips for module delta-*I* experiments
V: Chip sites for module via capacitance experiments
W: Driver and terminator chips for module wiring rule experiments

Figure 8

Plan view of the experiments contained on the product qualification vehicle.

Crosstalk coefficients were also verified on one of these test vehicles.

A special automated tester called a board-level open-bath computer-aided tester (BOBCAT) was developed to test two PQVs on a system-like board. The BOBCAT was used to test modules separately and in pairs to test the operation of board nets and cable nets. It also supported board noise testing.

The BOBCAT consisted of three major elements. The first was a large tank containing a 60-cm × 70-cm board immersed in a fluorocarbon coolant. Since the EC pads shown in Figure 3 could also be used as probe points, the modules plugged into the board contained no covering cooling assembly that would block access to the pads. Cooling was done by bathing the chips in the fluorocarbon coolant which flowed continuously through the tank into the second major BOBCAT element, a large refrigeration system. The last key element of the tester was an IBM Series-1 computer-driven pattern generator and data-acquisition system. The data-acquisition system consisted of an oscilloscope and two computer-controlled robotic high-frequency probe positioners. Once the data had been gathered, the system sent the measurements to a host computer where they were compared with simulation model results. Demonstrating that measured data compared favorably with the simulation data used to generate the design rules was the essence of the product qualification process.

When product qualification was successful, the design process was essentially completed. (The auditing function implied by product qualification was performed by an independent technical organization that was distinct from the design group.)

Conclusions

The IBM Enterprise System/9000 TCM contains many new technology enhancements. The module is larger, in order to hold more circuits, and it contains a lower-dielectric material, a cordierite glass-ceramic, for better performance. Polyimide thin-film lines have been added for improved signal redistribution. To improve module noise characteristics, topside decoupling capacitors have been added.

The new module supports higher-power chips with more signal and power connections. On the bottom surface, the module connector also provides more signal and ground connections.

To ensure operational integrity, much attention was paid to the electrical design of the module. The design rules were carefully generated, taking into account the electrical properties of the substrate. A sophisticated DA program was generated to keep track of delta-*I* and crosstalk noise to prevent them from causing any data errors. Finally, all of the modeling and circuit simulation analysis was confirmed by constructing and measuring a comprehensive set of qualification vehicles. A special automated tester was used for the qualification data gathering.

The result of all these new features is a TCM that can support more function and faster cycle times than earlier versions of the TCM technology. All necessary steps were taken in the electrical design of the ES/9000 TCM to ensure that it would function as expected.

References

1. E. Davidson, P. Hardin, G. Katopis, M. Nealon, and L. Wu, "The Design of the ES/9000 Module," *IEEE Trans. Components, Hybrids, Manuf. Technol.* **14**, 744–748 (1991).
2. N. Oki, H. Murase, and M. Nishihara, "Packaging Technology for VP-2000 Series," *Nikkei Microdevices*, June 1989, pp. 50–55.
3. A. Dohya, T. Watari, and H. Nishimori, "Packaging Technology for the NEC SX-3/SX-X Supercomputer," *Proceedings of the 40th Electronic Components and Technology Conference*, May 1990, pp. 525–533.

4. F. Kobayashi, Y. Watanabe, M. Yamamoto, A. Anzai, A. Takahashi, and T. Daikoku, "Hardware Technology for the HITAC M-880 Processor Group," *Proceedings of the 41st Electronic Components and Technology Conference*, May 1991, pp. 693–703.
5. E. Canty, "VAX 9000 Systems," *DEC Product Announcement* (October 1989) and U.S. Patent 4,692,839, 1989.
6. E. Davidson, "The Performance Choice: A Study of Chip and Package Densities," *Proceedings of the 40th Electronic Components and Technology Conference*, May 1990, pp. 147–151.
7. A. Blodgett, "A Multilayer Ceramic Multichip Module," *IEEE Trans. Components, Hybrids, Manuf. Technol.* **CHMT-3**, 634–637 (May 1980).
8. J. Oberschmidt and J. Humenik, "A Low Inductance Capacitor Technology," *Proceedings of the 40th Electronic Components and Technology Conference*, May 1990, pp. 284–288.
9. P. J. Brofman, S. K. Ray, and K. F. Beckham, "Electrical Connections to the Thermal Conduction Modules of the IBM Enterprise System/9000 Water-Cooled Processors," *IBM J. Res. Develop.* **36**, 921–933 (1992, this issue).
10. *Microelectronics Packaging Handbook*, R. Tummala and E. Rymaszewski, Eds., Van Nostrand Reinhold, New York, 1989, Ch. 3, pp. 111–165.
11. G. Katopis, "Delta-I Noise Specification for a High Performance Computing Machine," *Proc. IEEE* **73**, 1405–1415 (September 1985).
12. G. Katopis, "Electrical Interconnection Design for High Performance Machines," *Proceedings of the National Electronic Packaging and Production Conference (NEPCON West 90)*, February 1990, Vol. 1, pp. 169–177.
13. E. Davidson and R. Lane, "Diodes Damp Line Reflections Without Overloading Logic," *Electronics* **49**, 123–127 (February 1976).
14. A. E. Barish, J. P. Eckhardt, M. D. Mayo, W. A. Svarczkopf, and S. P. Gaur, "Improved Performance of IBM Enterprise System/9000 Bipolar Logic Chips," *IBM J. Res. Develop.* **36**, 829–834 (1992, this issue).
15. B. Rubin, "An Electromagnetic Approach for Modeling High Performance Computer Packages," *IBM J. Res. Develop.* **34**, 585–599 (July 1990).

Received May 20, 1991; accepted for publication May 27, 1992

Evan E. Davidson *IBM Technology Products, East Fishkill facility, Hopewell Junction, New York 12533 (DAVIDSOE at FSHVMFK1, davidsoe@fshvmfk1.vnet.ibm.com).* Mr. Davidson is the manager of the High Performance Technology Department and specializes in the areas of technology applications and electrical package design for large mainframe computers. Prior to joining IBM in 1968, he was a circuit designer with Bell Telephone Laboratories. He received his B.E.E. from Rensselaer Polytechnic Institute in 1962 and his M.S.E.E. from New York University in 1964. Mr. Davidson is a Senior Member of the IEEE and a member of the Eta Kappa Nu and Tau Beta Pi engineering honor societies. He has written numerous papers and presented many papers in electrical package design. He holds various patents in the areas of memory and logic circuit design.

Peter W. Hardin *IBM Technology Products, East Fishkill facility, Hopewell Junction, New York 12533 (HARDIN at FSHVM1).* Mr. Hardin is an advisory engineer responsible for the mechanical and physical design of high-performance multilayer ceramic (MLC) electronic packages. Since joining IBM in 1970, he has worked on first- and second-level packaging designs and applications, including the thermal conduction module (TCM) for the IBM 3081 and 3090 mainframes. Mr. Hardin holds a B.S. degree in mechanical engineering from Bucknell University and is a registered Professional Engineer in the state of New York.

George A. Katopis *IBM Technology Products, East Fishkill facility, Hopewell Junction, New York 12533 (KATOPIS at FSHVMFKl, katopis@fshvmfk1.vnet.ibm.com).* Mr. Katopis is the manager of the Electrical Packaging Integration Department. He works in the areas of electrical modeling of packaging structures, noise containment, and the design of interconnection networks. He received his B.S. degree in electrical and mechanical engineering from the Technical University of Athens, Greece, in 1967, and received M.S. and M.Ph. degrees from Columbia University. Mr. Katopis is a member of the IEEE; he is also a Semiconductor Research Corporation mentor at the University of Arizona. He is the author of three papers and conference presentations on the subject of coupled and delta-*I* noise determination. Mr. Katopis is also the author of one paper on a simulation program required for the determination of noise tolerance for logic circuits, and a book on resonant circuit theory. He holds a patent on delta-*I* noise-reduction techniques for VLSI chips.

Michael G. Nealon *IBM Technology Products, East Fishkill facility, Hopewell Junction, New York 12533 (NEALON at FSHVMFK1).* Mr. Nealon is currently manager of the MLC Electrical Design and Analysis Department. Since joining IBM in 1982, he has worked on electrical design and analysis of multilayer ceramic (MLC) electronic packages. Mr. Nealon holds a B.S. degree in electrical engineering from Pratt Institute of New York.

Leon L. Wu *IBM Enterprise Systems, Neighborhood Road, Kingston, New York 12401 (LLWU at PK705VMG).* Dr. Wu is a senior technical staff member currently involved in the formulation of future high-performance computer packaging

1968, and worked from 1968 to 1981 on diffusion profile modeling, silicon defect studies, device modeling, thermal printer design, and high-performance computer packaging structure modeling. In 1981, he became a senior engineering manager of the Exploratory Packaging Department in East Fishkill. From 1987 until the present, Dr. Wu has been involved in high-performance computer technology competitive analysis and technology strategy. In 1962, Dr. Wu received his B.S. in electrical engineering from the National Taiwan University, Taipei, Taiwan. He received his M.S. in 1965 and his Ph.D. in 1968, both in electrical engineering, from Carnegie Mellon University, Pittsburgh, Pennsylvania.