

Aspects of the electrical design and analyses of the printed circuit boards of the IBM Enterprise System/9000 water-cooled processors

by L. E. Boone
M. R. Brinthaupt
J. A. Malack
J. Pavlik

Aspects of the electrical design and analyses of the multilayer printed circuit boards of the IBM Enterprise System/9000™ water-cooled processors are discussed. The design and analyses pertained to achieving an increase in wirability and a decrease in voltage drops, power loss, simultaneous switching noise, variation in characteristic impedance, and reflections due to the presence of stubs.

Introduction

The approach used for designing the printed circuit boards of the IBM Enterprise System/9000™ (ES/9000™) water-cooled processors was introduced in 1981 [1] with the IBM 3081 processors. The boards of the 3081 processors were very effective for interconnecting their thermal conduction modules and served as controlled electrical environments for signal transmission [2].

Since 1981, the requirements defining the electrical environment of a mainframe processor have changed

©Copyright 1992 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.




Figure 1

Cross section of the printed circuit board of the ES/9000 water-cooled processors. From portion of a figure in [5], reproduced with permission.

rapidly, driven by increases in circuit switching speeds and density, increases in associated power dissipation, and decreases in signal pulse rise times [3]. These changes have resulted in more demanding and controlled board requirements [4]. Noise margins and allocations needed to be reduced, physical tolerances on board parameters tightened, and new electrical design constraints satisfied.

In this paper we discuss aspects of the electrical design and analyses of the boards of the ES/9000 water-cooled processors. The following aspects are discussed:

1. Changes in board cross section to improve wirability, optimize voltage drops, and reduce power dissipation.
2. Reductions in the nominal line width and tolerance of their conducting lines, in order to reduce board thickness while controlling characteristic impedance; this also resulted in less coupled noise.
3. Reductions in simultaneous switching noise; of concern because of the decrease in processor cycle times and the increase in the number of circuits switching simultaneously.
4. Detection and elimination of stubs (partially connected segments of wiring); necessary because the decreased rise time of the circuits makes it more likely that the presence of stubs can result in deleterious reflections.

Although many details have been omitted because they are currently confidential, it is hoped that in reading this paper the reader will gain a useful understanding of the key electrical considerations that had to be addressed in order to achieve those advances.

Power distribution

The design of the cross section of the boards of ES/9000 water-cooled processors was based on the designs of boards of the 3081 and 3090™ processors (also water-cooled). The number of signal planes was increased from six to ten in order to achieve increased wirability, and power plane thickness was increased in order to provide

Table 1 Comparative dimensions and electrical characteristics of printed circuit boards of several IBM processors.

Printed circuit board	3081 (water-cooled)	3090 (water-cooled)	ES/9000 water-cooled models
Size (mm)	600 × 700	600 × 700	600 × 700
Thickness (mm)	4.50	5.19	7.44
Drill size (mm)	0.46	0.46	0.46
Grid (mm)	2.5 × 1.25	2.5 × 1.25	2.5 × 1.25
Lines per channel (1.25 mm)	2	2	2, 1
Line width (mm)	0.102 ± 0.025	0.102 ± 0.025	0.081 ± 0.013
Line separation (mm)	0.152	0.154	0.193
Line thickness (mm)	0.04 ± 0.009	0.04 ± 0.009	0.043 ± 0.0076
Line to power plane (mm)	0.315	0.315	0.285
Signal planes	6	6	10
Power planes	12	16	12
Power plane thickness (oz)	1	1	2
Board voltage levels	-4.25	-2.2	-2.2
	-1.5	-0.7	-0.7
	Ground	Ground	Ground
		+1.4	+1.4
Board current (A)	500	1400	~3400
Characteristic impedance (Ω)	78 + 12, -7	78 + 12, -7	80 + 10, -9
Near-end coupled noise (mV)	175 (max)	175 (max)	129 (max)
Series resistance (mΩ/cm)	46 + 22, -11	46 + 22, -11	54 + 18, -11

more current-carrying capability [5]. Comparative dimensions and electrical parameters are listed in **Table 1**.

The power planes also provide a reference for the signal transmission lines on the signal planes (designated as S1, S2, ..., S10) [6], as indicated in **Figure 1**. The cross section shown is designated as a 10S/12P (10-signal and 12-power-plane) configuration. Each signal-plane pair is separated by a pair of power planes. The four voltage levels needed for circuit operation, V_{CC} , V_i , V_r , and V_{EE} , are distributed to the power pins of the thermal conduction modules by the internal power planes of the board.

The outermost layers are designated as the module layer and the EC (engineering change) layer. The EC layer contains a pattern which, when used with twisted pair cable, permits the repair or modification of the signal transmission lines.

The commoning bus must provide power at the four circuit voltage levels. The layer that provides the V_{CC} voltage level must be capable of delivering over 3400 A, since it constitutes the return to the power supply. Two voltage connections are made at the input of each bus to guarantee uniform current flow and minimal voltage differential at the commoning bus outputs.

The laminar buses consist of two elements, the heavy copper bus that brings current to the board and the bus tabs that are soldered to pads on the surface of the board (**Figure 1**). There are four buses mounted on each board. Each bus consists of four copper bars ranging from 3.17 to 9.56 mm in thickness, separated by insulating material. The thickness of each bar is determined by the current-carrying requirements, the voltage differential requirements, and the amount of I^2R heating permissible within thermal limitations.

The voltage tabs (bus output tabs) are a series of 0.66-mm-thick, 12-mm-wide copper tabs formed from the laminar bus. The actual number of tabs and their placement with respect to the positions of the thermal conduction modules is critical in ensuring that requirements associated with voltage differential and current limitations of the voltage connections are met. Current limitations are dictated by the thermal limits and the current-carrying capacity of the plated through-holes (PTHs) that connect the surface pad to the internal power planes.

The power planes consist of 600 by 700-mm sheets of 0.07-mm-thick (2-oz) copper separated by an epoxy/glass dielectric. Connection to the power planes is achieved using PTHs attached to the planes and projected to the surface of the board at power pads. C-springs (**Figure 1**) are used to transfer voltage levels and signals from the surface of the board to the signal and power planes of its modules. To achieve the required voltage levels at the surface, all of the PTHs must pass through sufficiently




Figure 2

Circuit noise immunity curve (not to scale), indicating sources of signal noise.

wide clearance holes in the power planes, from which they must be electrically isolated. Accordingly, 1.22-mm-wide clearance holes are drilled into the power planes. Since the holes deplete copper from the planes delivering current, the resistance between the laminar bus inputs and the springs that connect the internal planes to the modules and other components mounted on the surface of the board is thereby increased.

Voltage differential that exists because of current flow through the solid copper and the copper-depleted array of clearance holes constitutes a portion of the circuit noise immunity curve [7]. Just as coupled noise, reflections, and connector noise degrade circuit integrity, the voltage differentials associated with the power planes modify the voltage references appearing at the inputs of the board's communicating circuits. This is illustrated in **Figure 2**, which depicts the sources of signal noise (not to scale).

Voltage differential allocations ranging from 5 to 25 mV were targeted for the different voltage levels of the board. Because the dc current levels for the ES/9000 water-cooled processors represent a 2000-A increase over those of the




Figure 3

Calculated resistivity correction factor (a) with clearance holes drilled on a regular grid, (b) with clearance holes drilled on a staggered grid.

3090 processors, there was an appreciable increase in the complexity of the analysis required to ensure that the targeted allocations would be achieved.

The design and analyses of the power distribution system of the boards of the ES/9000 water-cooled processors were performed using a three-dimensional analysis program. The input to the program was a topographical profile description of the conducting media. All values except those pertaining to via resistance and current ports were inputted as sheet resistivity values. The program converted the sheet resistivity values into resistor elements connected between nodes, as would programs

such as ASTAP [8]. It was necessary to take into account the effects on sheet resistivity of temperature increases resulting from the considerable amount of power dissipation expected (because of the high current levels in the copper power planes). The temperature coefficient of bulk copper (0.0039 per °C) was used to calculate that increase.

Additionally, the sheet resistivity had to be corrected for the presence of local holes in the power planes. Accordingly, correction factors (ratios of the resistivity of power planes containing clearance holes to that of a solid power plane) were calculated. Figure 3 shows calculated values for planes having holes drilled on a regular grid and on a staggered grid (as shown on the inserts).

Corrected sheet resistivity values were then used in the power distribution analysis, permitting the voltage drops at all ports and the current through all of the vias to be calculated. The voltage drops were used to determine the minimum and maximum voltage drops occurring at module connector positions [3]. Associated via currents were calculated to ensure that they did not exceed the 100-A limitation of the bus output tabs (see Table 2). The increased current levels required proper placement of the voltage tab connections to ensure that that limitation was not exceeded.

The final phase of the power distribution analysis involved thermal analysis calculations [3]. Expected power dissipation levels were calculated for all components of the power distribution system. A summary of the calculated levels is shown in Table 3. The 2000-A increase also implied an increase proportional to the square of the current ratio in the amount of power dissipated by the distribution system. The analysis indicated that 12 sheets of 2-oz copper planes would be needed for the board cross section (Figure 1).

Signal transmission

The cross section of the boards of the ES/9000 water-cooled processors was defined by circuit and system requirements, and by manufacturing process capability [9]. The boards of the 3090 processors contained three signal plane pairs, with signal lines orthogonally placed on either side of an epoxy/glass dielectric core, interspersed with the 16 power planes. The triplate transmission line structure of Figure 4, composed of a signal core sandwiched between two adjacent power planes, was used to calculate the electrical characteristics of the signal lines. Effects on signal propagation could then be predicted using transmission line theory [10].

Calculations of noise in such a triplate structure, containing lines having the same characteristic impedance, were used to ascertain trade-offs necessary to maintain the near-end coupled noise within an acceptable range. Coupled noise is a measure of the signal energy that is

coupled to an adjacent quiet (nonswitching) line from an active line; generally a decrease in signal rise time (as in the ES/9000 water-cooled processors) causes an increase in the amount of coupled noise generated. The trade-off calculations and associated changes made it possible to reduce the near-end coupled noise by approximately 35%, to a maximum value of 129 mV, assuming a decrease in signal rise time by approximately 35%.

If the signal lines are assumed to be loosely coupled and line time delay t_d (in units of picoseconds per mm) is assumed to be equal to or greater than half the rise time of the active signal, the saturated near-end coupled noise $V_{qnc}(\text{sat})$, the maximum amount of noise that can be coupled (without reflections) onto a quiet line, is expected to be

$$V_{qnc}(\text{sat}) = V_{\text{sig}} \frac{(K_C + K_L)}{4},$$

where V_{sig} is the signal amplitude, K_C is the capacitive coupling coefficient, and K_L is the inductive coupling coefficient.

For t_d less than half the rise time of the active signal,

$$\frac{V_{qnc}}{\ell} = \left(\frac{dV_{\text{sig}}}{dt} \right) \left(\frac{K_C + K_L}{2} \right) t_d,$$

where ℓ is the coupled length.

Another concern regarding the near-end coupled noise is the maximum allowable distance along which two parallel signal lines can share the same signal channel before it reaches its saturation level. To minimize this concern, the system used to automatically wire the signal lines on a plane is programmed to cause the lines to diverge at a specified wiring distance.

The density of the wiring on the boards of the ES/9000 water-cooled processors had to be increased by 66% because of the larger (127 × 127-mm) module size and greater number of circuits that had to be accommodated on a board [5]. The increased wiring density required an increase in the number of signal plane pairs from three (6S) to five (10S). Furthermore, the dc current level increased by approximately 150%, requiring the use of thicker power planes [5]. Some of the dielectric material used in the construction of the power core was replaced with the thicker planes.

To satisfy coupled-noise requirements, line-to-line spacing was increased by decreasing the line width (Table 1) while maintaining the center-to-center distance (pitch). The characteristic impedance Z_0 of the triplate structure can be retained at 80 Ω if corresponding changes are made in signal line width w and dielectric spacing h_1 . It can be increased by 0.21 $\Omega/\mu\text{m}$ by decreasing w , and decreased by 0.77 $\Omega/\mu\text{m}$ by decreasing h_1 . With one dielectric sheet removed, the line width needed to produce a system for




Figure 4

Triplate structure used for signal line analyses.

Table 2 Laminar bus tab currents.

Level	Tab current (A)	Delta voltage (mV)
V_{CC}	32-96	6.6-12.3
V_{EE}	54-75	6.8-17.8
V_t	46-56	8.4-11.3

Table 3 Power distribution system.

Level	V_{CC}	V_{EE}	V_t
Current (A)	1900-3400	850-2162	1200-1350
Power (W)			
Commoning bus	13.0-42.0	2.0-18.0	3.5-9.0
Laminar bus	18.0-34.8	5.0-21.0	7.5-12.0
Bus tabs	8.0-26.0	2.0-12.5	4.0-5.5
Board	15.5-51.0	7.0-38.0	8.5-12.5
Connector	20.0-63.0	7.0-41.0	12.0-18.0
Bus ear	5.0-13.0	1.0-6.5	1.5-2.0
Total	74.0-230.0	23.0-123.0	40.0-56.0

which $Z_0 = 80 \Omega$ is approximately 0.08 mm. If the center-to-center distance is unchanged, some relief is provided with respect to coupled noise. Limits had to be defined to optimize the system requirements and manufacturing processes in order to arrive at acceptable line width and tolerance levels (Table 1).

Several factors restricted the amount of latitude available for a practical design change in signal line width. The signal line width used for the boards of the 3090 processors was 0.102 mm, with a tolerance of ± 0.025 mm. The upper limit of the series line resistance precludes some combinations of signal line dimensions and dielectric spacings. If the series resistance is too high, especially for long signal lines, the signal logic level at the far-end receiver degrades because of dc losses. If the losses are excessive, the logic levels may no longer be within tolerance.

Manufacturing considerations determine the lower limit on signal line tolerance; they depend on various processing limitations, including those connected with the photographic positive used to expose the photoresist on the signal cores. The required tolerances were met after analysis and experimentation.

An impedance model based on a triplate structure was used to determine acceptable line tolerances. The model correlated well with existing capacitance programs used to calculate Z_0 . Using the model, Z_0 could be expressed as

$$Z_0 = E \frac{K}{\sqrt{\epsilon_r}} \left(\frac{w}{h_n} \right)^{K_1} \left(\frac{t}{h_n} \right)^{K_2} \left(\frac{h_1}{h_n} \right)^{K_3} \left(\frac{h_2}{h_n} \right)^{K_4} \left(\frac{h_3}{h_n} \right)^{K_5},$$

where w , t , h_1 , h_2 , h_3 , and h_n are the triplate dimensions indicated in Figure 4; the exponents are design parameters having values associated with physical parameters and their tolerances; E is an error-correction term; K is related to the impedance of the triplate structure; and ϵ_r is the assumed dielectric constant. This expression was used in conjunction with the ASTAP [8] statistical analysis program. The magnitude of Z_0 was monitored for an impedance of 67Ω or less, which was the lowest level permitted. The model was evaluated statistically using standard Monte Carlo simulations. Distributions for all parameters except signal line thickness and width were assumed to be uniform. This type of distribution provided us with a conservative set of results. The limits used in the model were the tolerance values for each of the design parameters, and the line width and thickness were each assumed to follow a normal distribution. Data for line width and thickness were obtained from an associated investigation of process parameters.

Two line widths, 0.081 mm and 0.084 mm, with tolerances from 0.0036 to 0.0203 mm were used in the model. The 0.084-mm line width was eliminated with one failure in 500 cases, where Z_0 was less than 67Ω . The

0.081-mm line width was selected, with an optimum tolerance of 0.013 mm. Tolerances less than 0.013 mm did not produce failures, but the numbers of signal lines having characteristic impedances less than 71Ω increased. This was attributed to the unchanged tolerances of the remaining parameters.

Simultaneous switching noise

Simultaneous switching noise is generated from the large change in power supply currents demanded by simultaneously switching driver circuits. The power distribution system of a board is inherently inductive and, therefore, resists instantaneous current changes; this induces voltages in the power distribution system. The induced voltages, or simultaneous switching noises, manifest themselves as voltage spikes that appear on the power distribution planes of the board and propagate throughout the transmission line system formed by the parallel distribution planes.

Usually simultaneous switching noise is roughly proportional to the amount of switching current and inversely proportional to the time required for the current to change; these relate directly to the number of simultaneously switching drivers present and to the rise time of the driver output signals. Simultaneous switching noise can be controlled by power supply bypass capacitors that provide a localized energy source and lower the effective impedance between the voltage distribution planes [11].

Excessive switching noise can affect system performance by introducing voltage variations sufficiently large at the receiver circuit to produce logic errors. Two basic failure mechanisms are encountered. First, if the magnitude and timing of the switching noise are proper, the reference voltage for the receiver circuit can be affected, causing the relative receiver input voltage to become latched at the incorrect logic level. Second, if sufficient switching noise is "fed through" a nonswitching driver, a portion of the switching noise can travel from the driver's output to a receiver's input and latch at the incorrect logic level. To ensure that various board designs were in compliance with simultaneous-switching noise allocations, an analysis was performed on the basis of noise measurements performed with an open-bath test system. Because the open-bath system could not accommodate a fully functional ES/9000 water-cooled processor board, results obtained from a functionally limited test board were enhanced via software to more accurately represent noise on one that is fully functional.

As already indicated, the circuits used in the IBM ES/9000 water-cooled processors make use of four distinct voltage levels [5]. Measurements of simultaneous-switching noise at the board level were performed using a product-level board and two thermal conduction modules

containing a complement of specially designed ECL test chips. Because the modules of these processors are water-cooled, direct access to chip test points was not possible. Therefore, the open-bath test system was configured to permit removal of the cooling plates of the modules, thereby allowing access to the chips while in operation within temperature specifications. The board and its attached modules were immersed in a liquid dielectric material.

All test nets were driven from one module, continued through circuit traces in the board, and terminated on the second module. The test board contained a full complement of 232 tantalum decoupling capacitors, and the test modules contained a full complement of ceramic decoupling capacitors [12]. Additionally, test hardware provided for up to 284 emitter-follower (EF), 288 push-pull (PP), and 40 bi-directional (BiDi) drivers, switching simultaneously at less than 100 ps of skew.

Because of constraints of the open-bath tester, access to the back of the board could not be gained when a conventional voltage probe was used. Therefore, all simultaneous-switching noise measurements on the test board were performed using a 50- Ω miniature coaxial cable attached directly to voltage pads on its bottom side.

The results obtained provided a good indication of how the test board performed. However, a board in an ES/9000 water-cooled processor contains more modules, each having more and denser chips, and, hence, more switching drivers. Furthermore, the inclusion of more modules and chips implies more of a dc current load. As board loading increases, the effective impedance seen between board voltage plane pairs decreases, resulting in a decrease in simultaneous-switching noise. However, the presence of more drivers results in an increase in such noise because of the increase in switching activity.

After all test waveforms were processed, the worst-case noise was found to occur when all 284 EF drivers were switched. Therefore, the EF measurements were used as the basis for all analyses of simultaneous-switching noise. The noise due to a single EF driver switching on the test board was calculated by dividing measured voltage values by the number of switching EF drivers (284).

Noise values were further reduced to account for the differences in dc current loading between the test board and a board of a fully functional ES/9000 water-cooled processor.

In the ES/9000 water-cooled processors, use is made of several other types of off-module drivers in addition to the EF type. For each type, there exists a simultaneous-switching coefficient that provides an indication of the amount of simultaneous-switching noise produced on the modules when the driver switched. Those coefficients were used in an analysis to ensure that differing modules did not

exceed their simultaneous-switching-noise allocations, assuming that the ratios of switching coefficients among driver types were the same at the module and board levels. A comparison of the ratios of EF-related to BiDi-related noise measurements at the board level and EF-related to BiDi-related switching coefficients validated this assumption. Therefore, to determine the worst-case plane-to-plane noise at the system level due to one of any type of driver, the EF-related values were multiplied by the ratio of the coefficient for the driver type of interest and the coefficient for the EF driver.

For every wiring net of each board, information regarding relevant distances, switching times, etc. was used to quantify the simultaneous-switching activity of a board and calculate its simultaneous-switching noise levels on a module-by-module basis. For this purpose, each thermal conduction module of a board was treated as a point source of board switching noise.

Module self-noise was calculated over time for each thermal conduction module of a given board. It was defined as noise at the board that could be attributed to all drivers and receivers that switch on a module. Next, module transmitted noise was calculated over time for each module. Such noise was defined as the noise at the board that was transmitted through the board from a module site. Finally, module self-noise and transmitted noise were summed over time to calculate the total simultaneous-switching noise of a board expected at each module location.

Module self-noise was calculated for every module of each board as follows: For each net on the board of interest, the driving and receiving module designations were examined to determine the type of net present. Nets were divided into three categories: 1) those that remain on the board (involving a driver and receiver); 2) those that drive off the board (involving only a driver); 3) those that start off the board and end at a receiver on the board of interest (involving only a receiver).

For a net that remained on the board, the noise contribution at the driving module due to the driver was determined by constructing a timing diagram in the form of a step function that was assumed to be initially off, on at the early switching time, and off at the late switching time. It was assumed that the driver could switch at any time between the early and late switching times. The amplitude of the step function was determined from a knowledge of the driver type and the noise due to one switching driver, calculated using the ratio of simultaneous-switching coefficients, as discussed above.

The noise contribution at the receiving module due to the receiver was determined using the same timing diagram used for the driver. The diagram, however, was delayed by the amount of time necessary for the output of the driver to arrive at the receiving module. This delay was

determined from the sum of the net length at the board level multiplied by a typical board signal line delay, and the driving and receiving module lengths multiplied by a typical module signal line delay. For this analysis, the receiver was assumed to cause the same noise at the receiving module as the driver caused at the driving module.

For a net that drove off the board, the noise contribution at the driving module due to the driver was determined in the same manner as that for the driver of an on-board net. Similarly, the noise contribution at the receiving module of a net that received on the board was determined in the same manner as that for the receiver of an on-board net.

A timing diagram was produced for each net of each board and for each of the six plane-to-plane voltage pairs. The diagrams were then summed over time by module location to complete the self-noise calculations.

Module transmitted noise was calculated for every module of each board as follows: For each module on the board of interest, the distance between the geometric centers of the modules was determined. It was then assumed that the noise produced from a module would travel through the board to another module site via TEM-mode propagation, at 67 ps/cm (assuming a dielectric constant of 4.0). An estimate of the time necessary for the noise from one module to travel to another was then obtained by multiplying the TEM-mode delay by the center-to-center distances. Additionally, the transmitted-noise magnitude, as seen at a remote site, was reduced in order to account for transmission losses. Accordingly, the self-noise diagrams for each module site were delayed and reduced in magnitude, thus completing the transmitted-noise calculations. Finally, self-noise and transmitted-noise diagrams were combined to produce the total expected simultaneous-switching noise of the board at each module site.

The maximum simultaneous-switching noise between the voltage planes of all analyzed boards was found to be below the simultaneous-switching-noise allocation. If noise levels had exceeded the specification, additional decoupling capacitors could have been added. Finally, the presence of the on-module decoupling capacitors significantly reduced the simultaneous-switching noise at the board level.

Detection of stubs

The presence of stubs (partially connected segments of wiring) may affect the transmission properties of the signal lines in a printed circuit board at high frequencies (MHz) or when the lines are propagating fast pulses having sharp rising leading edges. The stubs in printed circuit boards can be less than 2.5 cm or greater than 30 cm in length. Their effect on signal transmission varies depending on their lengths compared to the signal lines to which they are

attached. For long signal lines, a short stub has negligible loading effects and may be tolerated. However, for short signal lines, a long stub can have significant loading effects, and thereby affect line functionality.

The potential formation of buried line stubs is inherent in the manufacturing and/or repair of printed circuit boards [8]. Stubs can form as a result of the following:

1. *Masking of a short.* Masking of an electrical short that exists between two adjacent internal signal nets by the presence of an open. If an open has been detected on one of the nets, the open condition will drive repair action after lamination to correct the open but not remove the short. This short creates a stub on the repaired net.
2. *Incomplete deep deletion.* Board repair involving the removal of the barrel of a PTH by drilling. This isolates the internal net from the endpoints of the PTH. Although external wires are bonded to the surface pads of the board to replace the internal nets, the removal of the PTH does not result in isolation of the internal net which, when connected to its surface pad, thus leads to the creation of a stub.
3. *Void formation and reconnection at early net repair deletion site.* Such repair (prior to lamination) involves deletion directly on signal planes of the defective endpoints of an internal net. Epoxy reflow during lamination is used to fill the area between the ends of the net and the nearby PTH sites. However, incomplete epoxy reflow can leave voids in the epoxy which, upon through-hole plating, can fill with conductive material, connecting the net to a PTH. The net thus becomes a stub.
4. *Missed deletion.* Failure to delete all of the above-cited endpoints, possibly as a result of data or equipment error, will result in connection of a defective internal net to a PTH after through-hole plating. Such an undeleted net then becomes a stub.
5. *Short between functional net and early net repair net* [13]. A short that exists between a site of early net repair net and an adjacent internal net or PTH. Such a short can be due to the presence of extraneous foreign material connecting the nets, plating epoxy voids, or crevices along planes of epoxy delamination.

All of these sources of stub formation generally affect the electrical characteristics of the signal transmission lines: their characteristic impedance, propagation delay, and coupled noise.

Each signal line must be terminated by a resistance that is equal to its characteristic impedance Z_0 in order to maximize energy transfer and minimize reflection. However, if a stub is connected in parallel with the matched load, the impedance match is compromised. In this case, the transmission line is now terminated by an




Figure 5

Receiver input requirements, depicted as required signal input voltage levels.




Figure 6

Representation of a stub connected to a transmission line (a) at the near end of the transmission line, (b) at its midpoint.

equivalent impedance equal to the parallel combination of Z_0 and the impedance of the stub.

For proper circuit operation at the board level, all board parameters must satisfy circuit design requirements. This is typified by input requirements such as the *least positive up level* (LPUL) and the *most negative down level* (MNDL) requirements, as shown in **Figure 5**. To meet required circuit levels, the impedances of the transmission lines must be within the range of impedances called for in the design of the driver circuits. This can be offset if stubs are present at driver outputs or receiver inputs.

If a stub is introduced into the signal line of a network, a reflection is produced on the transmission line that reduces the active signal voltage level. The stub can be located anywhere along the transmission line; its position determines the time and duration of the reflection seen at the receiver input.

To calculate the amplitude of the output signal with a stub present, the reflection coefficient ρ (the ratio of the input voltage and reflected voltage at the point of mismatch) must be determined. Calculated output waveforms arising because of mismatches in characteristic impedance resulting from the presence of a stub at the near end and midsection of the transmission lines of **Figure 6** are shown in **Figure 7** for stub lengths of 0, 2.5, 5.0, and 25.0 cm. Assuming the presence of a 1.0-V input pulse in an 80- Ω transmission line, $\rho = -0.333$, and $Z_{\text{stub}} = Z_0$, we obtain

$$V_{\text{reflected}} = \rho V_{\text{incident}} = -0.333 \text{ V},$$

$$V_{\text{resultant}} = (1 + \rho)V_{\text{incident}} = 0.666 \text{ V}.$$

The position and time duration of the reflection are dependent on the length of the section of transmission line preceding the stub and the length of the stub. The waveforms are outputs of a network analysis program. If the stub is assumed to be located at the near end of the transmission line, a reflection occurs at the rise time and lasts for twice the total time delay of the stub. For example, in the first configuration with a stub of 25 cm in length, the output is expected to rise to 666 mV and remain there for 3.4 ns, then continue to rise until reaching the incident level of 1 V. If the stub length equals the mainline transmission line, the resultant output pulse will never reach the incident voltage level. Such a condition of insufficient pulse amplitude could prevent the receiving circuit from switching.

With processor cycle times decreasing and pulse rise times decreasing, it is imperative that signals arrive at a receiver circuit within defined times and contain the proper information at the time the circuit is clocked. Stubs affect not only the information carried, but also the time at which an intended signal arrives at the receiver. **Figure 8** shows that a logic circuit designed to switch at 6.5 ns may not switch until 3.0 ns later. In the ES/9000 processor, such a delay may cause improper voltages to appear during the




Figure 7

Calculated reflections on a signal resulting from the presence of stubs of different lengths, connected as indicated in Figure 6.

clocking interval. If the stub is on a clock line, the clocking pulse may be delayed by 3.0 ns (or more, depending on the length of the stub).

The impact of a stub on propagation delay becomes apparent when twice the delay of the stub approaches the rise time of the driving circuit. At this point, perturbations can be seen on the rise time. The shorter the rise time, the more concern over the length of the stub. With shorter rise times, even short stubs can have detrimental effects. This is shown in **Figure 9**, which shows calculated effects on a signal at a far-end receiver of the presence of a stub at a

near-end driver. As can be seen, small discontinuities begin to appear at rise times below 600 ps.

Coupled noise may also be affected by the presence of stubs. Such noise results from capacitive and inductive coupling from a line that is switching between logic levels (active line) to another quiescent (quiet line) at a logic level. Assuming that the signal line structures in a board are triplate structures in a homogeneous material, their capacitive coupling coefficients K_C should be approximately equal to their inductive coupling coefficients K_L . The coefficients may be calculated using many of the capacitance analysis tools. Once the capacitive components are known, the capacitive coupling coefficients can be calculated from the relation

$$K_C = \frac{C_{12}}{C_{12} + C_{1g}},$$

where C_{12} is the line-to-line capacitance and C_{1g} is the capacitance of that line with respect to a reference conductor.

As is the case for reflections and delay, the critical concern regarding noise is when 2τ of the coupled section equals the rise time of the circuit. At this point, the coupled noise of the circuit reaches saturation. Prediction of the coupled noise becomes very difficult in systems where the cycle times are less than 2τ of the system data




Figure 8

Calculated amplitudes of signals at a far-end receiver, depicting the effects of stubs of different lengths connected at the near end of a driver.

and clock nets, as for the ES/9000 water-cooled processors. Because of varying net lengths and reflections from different-level quiet receivers, noises from previous system cycles may be present at the time of new switching noise. In this case, noise sources adding algebraically can result in complex fluctuations that add to other switching noise sources. To perform correctly, circuits must adhere to stringent noise tolerance curves.

Figure 10 shows two possible situations that may occur for a stub connected to a printed circuit line for which the stub lies in a wiring channel adjacent to an active line.





Figure 9

Calculated amplitudes of signals at a far-end receiver. The signals are assumed to have different rise times. In (a) no stubs are assumed to be present. In (b) a 2.5-cm-long stub is assumed to be connected at the near-end driver.




Figure 10

Representation of a stub connected to the transmission line of a quiet near-end driver (a) at the near-end driver, (b) at the midpoint of the line. An active line is nearby.

(The active line that couples noises to the stub and the quiet mainline are usually different lines.) The coupling coefficients in each case, calculated for typical circuit line geometries, are 0.332 for both K_L and K_C . These values were determined by assuming appropriate board cross-section design parameters. The calculated saturated noise for a 1.0-V transition was found to be

$$V_{\text{qnc}}(\text{sat}) = 1000 \text{ mV} \frac{(0.332 + 0.332)}{4} = 166 \text{ mV}.$$

In the absence of the stub, the near-end noise remains at the saturated level for $2\tau + t_r$ (rise time) of the main coupled section. In the presence of the stub, the amount of noise increases, provided the stub/active line combination switching is in the same direction. If the stub/active noise combination is of sufficient length to create a saturated condition, the two noise sources are additive. The presence of the stub leads to a modified reflection coefficient at the near-end receiver. For illustrative purposes, if the impedance of the stub is assumed to be 80Ω , equal to that of the termination on the near end of the quiet net, $\rho = -0.333$. This implies that one third of




Figure 11

Calculated amplitudes of coupled noise at near-end driver of Figure 10, in the presence of stubs of different lengths, connected as indicated in that figure.

the incident noise should be reflected at the point of the stub connection, keeping in mind that the quiet line termination is also at the near end of the coupled noise present on the stub.

To calculate the amplitude of the near-end noise as a result of the presence of the stub, the noises due to the two sources must be added algebraically. The results of a simulation of the nets are shown in **Figure 11**. The increased noise is expected to be present for a time equal to t_r + twice the time delay associated with the stub. After this time, the resultant noise should return to a voltage level slightly less than the incident voltage level (arising from multiple reflections).

As a result of the above analyses, the following methods for stub detection were developed:

1. A capacitance measurement method that is applicable for use in connection with early net repair sites [14] and single-pin plated through-holes.
2. A frequency response method that treats the stub as an open, unterminated transmission line, utilizing the frequency response behavior of a circuit line under test to indicate the presence or absence of an attached stub.
3. A time-domain reflectometry method for measuring the characteristic impedance of a line; deviations below a design value indicate the existence of an attached stub.

By using the above testing methods, stubs can be removed once they are detected. Furthermore, the analysis of stub data thus obtained has facilitated the implementation of process changes to reduce their occurrence.

Summary

Increases in circuit switching speeds and density, increases in associated power dissipation, and decreases in signal pulse rise times have placed demands on the electrical performance of the printed circuit boards in mainframe processors. Aspects of the electrical design and analysis of the printed circuit boards used in the IBM Enterprise System/9000 water-cooled processors have been discussed. The design and analysis have led to improvements in power distribution, signal transmission, simultaneous switching noise, and stub detection and elimination.

Acknowledgments

The authors would like to acknowledge the contributions of the following persons: David Hawken and Jeffrey McKeveny for their printed circuit line data, Bruce Chamberlin for his mechanical support, Robert Ruane for his stub formation mechanisms and descriptions, Jagdish Gupta for the model used to determine signal line tolerances, and Paul Gebler and his department for use of their open-bath tester and their associated technical support.

Enterprise System/9000, ES/9000, and 3090 are trademarks of International Business Machines Corporation.

References

1. Roy F. Bonner, John A. Asselta, and Frank W. Haining, "High Performance Printed Circuit Board for the IBM 3081 Processor," presented at the special session of the IEEE 31st Electronic Components Conference, Atlanta, GA, May 11, 1981.
2. Roy F. Bonner, John A. Asselta, and Frank W. Haining, "Advanced Printed-Circuit Board Design for High-Performance Computer Applications," *IBM J. Res. Develop.* **26**, 297-305 (May 1982).

3. George G. Werbizky, Steven Boyko, and Frank W. Haining, "Circuit Packaging for High Power in High Performance Printed Circuit Boards," *Proceedings of the 36th Electronic Components Conference*, IEEE, May 1986, pp. 578-594.
4. C. S. Chang, "Electrical Design of Signal Lines for Multilayer Printed Circuit Boards," *IBM J. Res. Develop.* **32**, 647-657 (1988).
5. Lloyd E. Thomas, Lawrence E. Brearley, Joseph S. Macek, and Frank W. Haining, "Second Level Packaging for the IBM ES/9000 Processors," *Proceedings of the IEPS International Electronics Packaging Conference*, San Diego, September 1991, pp. 1049-1064.
6. Donald P. Seraphim, "A New Set of Printed-Circuit Technologies for the IBM 3081 Processor Unit," *IBM J. Res. Develop.* **26**, 37-44 (1982).
7. Alonzo Martinez, Jr., "Managing dc Power Distribution in Printed Wiring Backplanes," *Proceedings of the IEPS International Electronics Packaging Conference*, San Diego, September 1991, pp. 1087-1093.
8. W. T. Weeks, A. L. Jimenez, G. W. Mahoney, D. Mehta, H. Qassemzadeh, and T. R. Scott, "Algorithm for ASTAP—A Network Analysis Program," *IEEE Trans. Circuit Theory* **CT-20**, 628-634 (November 1973).
9. J. R. Bupp, L. N. Chellis, R. E. Ruane, and J. P. Wiley, "High-Density Board Fabrication Techniques," *IBM J. Res. Develop.* **26**, 306-317 (1982).
10. R. Brown, R. Sharpe, W. Hughes, and R. Post, *Lines, Waves and Antennas, The Transmission of Electrical Energy*, John Wiley & Sons, Inc., New York, 1973, Ch. 2.
11. Oystein Pedersen, Bertel-Eivind Flaten, and Torstein Gleditsch, "Modelling of Power and Ground Planes," *Proceedings of the IEPS International Electronics Packaging Conference*, San Diego, September 1991, pp. 652-663.
12. J. N. Humenik, J. M. Oberschmidt, L. L. Wu, and S. G. Paull, "Low-Inductance Decoupling Capacitor for the Thermal Conduction Modules of the IBM Enterprise System/9000 Processors," *IBM J. Res. Develop.* **36**, 935-942 (1992, this issue).
13. John Malack and Lawrence Boone, "Stubs in Printed Circuit Boards," *Proceedings of the IEPS 1990 International Electronics Packaging Conference*, September 1990, pp. 248-262.
14. J. G. Cutillo, R. C. Frisbie, R. F. Langley, and R. E. Ruane, "Printed Circuit Early Deletion Method," *IBM Tech. Disclosure Bull.* **26**, No. 4, 22 (1983).

Received January 28, 1991; accepted for publication January 14, 1992

Lawrence E. Boone *IBM Technology Products, P.O. Box 6, Endicott, New York 13760 (LBOONE at ENDVMTKL).* Mr. Boone is an Advisory Engineer in the Endicott Technology Laboratory. He joined IBM in 1966 at the Glendale Laboratory, Glendale, New York, working in areas of hardware design and software analysis. His current work is in the electrical design and analysis of printed circuit cards and boards for large system mainframes. In 1990, Mr. Boone received a Division Award for Technical Leadership in Clark Board Electrical Development. Mr. Boone received an A.D. degree from the Pennsylvania Technical Institute.

Mark R. Brinthaup *IBM Technology Products, P.O. Box 6, Endicott, New York 13760 (BRINTHAU at ENDVMTKL, brinthaup.endicott.ibm.com).* Mr. Brinthaup is a Staff Engineer at the Endicott Technology Laboratory. His work has pertained to the analysis and design of electronic packaging systems. Mr. Brinthaup joined IBM in 1984 after receiving his B.S. in electrical and computer engineering from Clarkson University.

John A. Malack *IBM Technology Products, P.O. Box 6, Endicott, New York 13760 (MALACKJA at ENDVMTKL).* Mr. Malack is a Development Engineer and manager of an electrical packaging development group in the Endicott Technology Laboratory. He joined IBM in 1967 and has worked in areas of electromagnetic compatibility, magnetic bubble memory packaging, and printed circuit board test development. His current work is on the electrical design and development of printed circuit card and board packaging. He earned a B.S. degree in electrical engineering from Hofstra University in 1967 and an M.S. degree in electrical engineering from Syracuse University in 1973. In 1983, Mr. Malack received an IBM Outstanding Technical Achievement Award for his work on printed circuit board testing methods. He is a senior member of the Institute of Electrical and Electronics Engineers.

John Pavlik *IBM Systems Technology Division, P.O. Box 6, Endicott, New York 13760 (resigned).* Mr. Pavlik was a Staff Engineer in the electrical packaging development group in the Endicott Technology Laboratory working on the electrical design and development of printed circuit cards and boards. He received a B.T. in electrical technology in 1976 from the State University of New York at Binghamton. Mr. Pavlik joined IBM in 1968 to work on electrical connector development.