E. E. Davidson

Electrical Design of a High Speed Computer Package

A methodology for optimizing the design of an electrical packaging system for a high speed computer is described. The pertinent parameters are first defined and their sensitivities are derived so that the proper design trade-offs can ultimately be made. From this procedure, a set of rules is generated for driving a computer aided design system. Finally, there is a discussion of design optimization and circuit and package effects on machine performance.

Introduction

As the modern computer system places demands upon technology for higher speeds and denser integrated circuit chips, we find the engineer resorting to more systematized design techniques. For many years, the chip designer has used computer simulation to predict a circuit's performance. Until recently, however, the computer package, which is made up of the paths that interconnect the chips, has not significantly limited the machine's performance beyond the expected delay factors. As a result, the package has not gotten a great deal of design attention using simulation, with most of the package-induced problems being solved empirically on the prototype machine.

With today's level of complexity and with the long lead times involved in making circuit changes, we have found this to no longer be possible. Instead, the computer package, consisting of the chip carriers or modules, boards, connectors, and frames shown in Fig. 1, must be properly designed before the machine is built. We now investigate the electrical properties of the mechanical structure that interconnects all the LSI chips. Specifically, we examine the effects of distributing power and signal throughout the package, and we explore the effects of noise. We have developed methods for analyzing some aspects of the package, and we use simulation to determine the performance of the total package.

In this paper, we use a single chip module (SCM) planar board technology to demonstrate the procedure. The procedure, however, is general, and it can be applied to other technologies, such as card-on-board and multi-chip module on board.

The next section of this paper presents an overview of the methodology. Sections follow on power distribution, signal distribution, and noise effects. All of these effects are related through the board's characteristic impedance, which is considered next. Finally, the relationship between the package and system performance is explored.

Package design interactions

This section introduces the topics that we found must be considered by the electrical package designer to optimize his overall design. Whereas the interrelationships are described here, a more complete description of each subject is deferred until the later sections. Figure 2 shows the elements that must be considered in the electrical design of a package. The package design generally begins with such physical parameters as line widths (w), thicknesses (t), and spacings (s); the dielectric constant (ϵ_r) and its height (h); and the hole or via diameter (d) and spacings.

Given limits on these parameters, power and signal distribution structures can be designed and evaluated. The evaluation consists of converting the mechanical structure into its equivalent electrical circuit through modeling, which is usually accomplished with computer programs that yield the desired inductances, capacitances, and resistances [1–3]. Once these linear network parameters are known, any desired quantities, such as characteristic impedance (Z_0), propagation delay (T_0), and line coupling coefficients can be calculated. With this information, the electrical packaging engineer can begin to analyze the power and signal distribution systems of proposed designs.

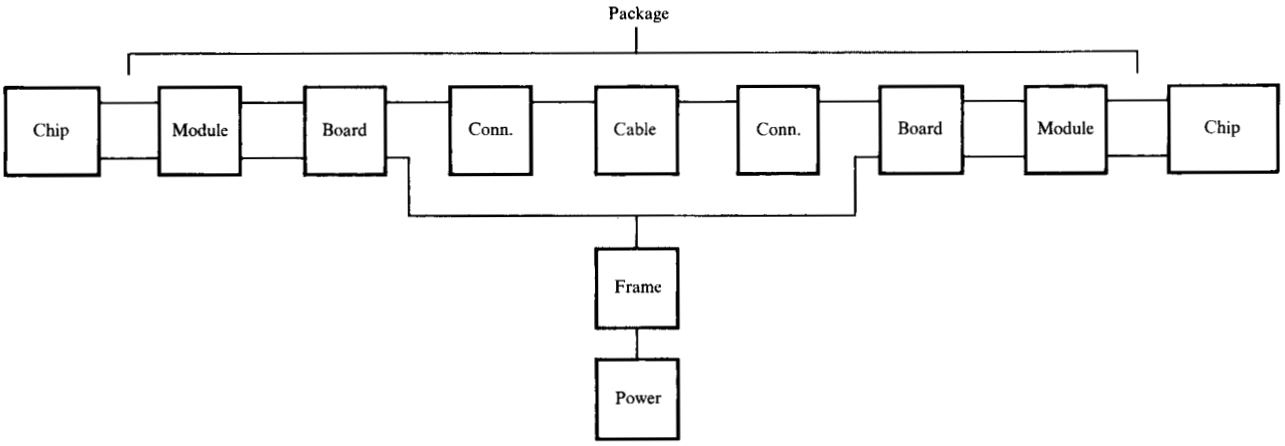

Figure 1 Computer hardware system.

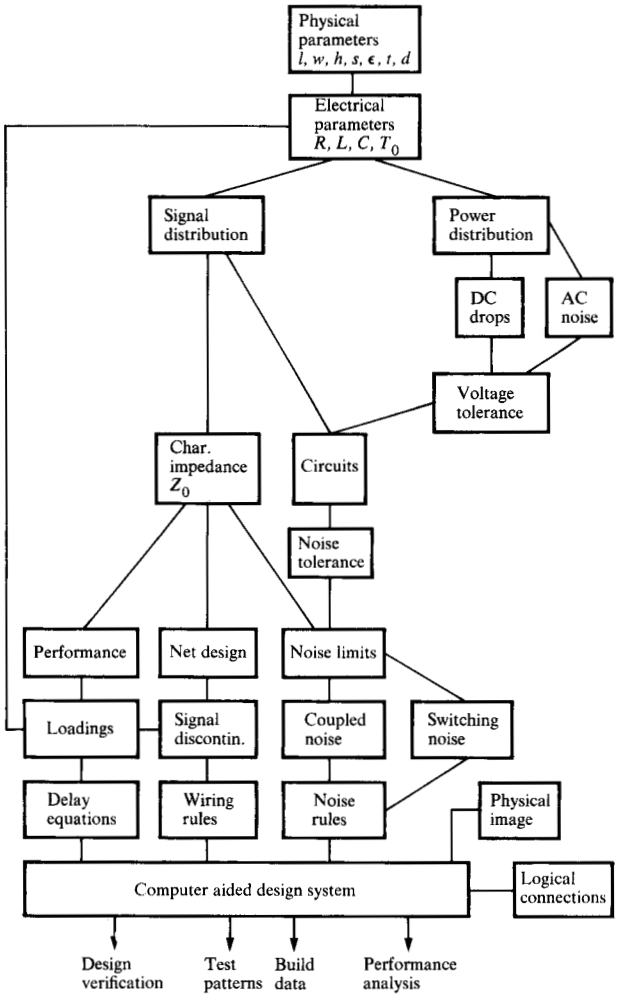

Figure 2 Procedure for the electrical design of a computer package.

We begin with the power distribution path shown in Fig. 2. Because the circuit designer needs to know the expected ranges of the supply voltages, the package designer determines the potential drops within the distribution system. Of course, a complete determination entails an evaluation of the ac variations as well as the dc drops.

Once the voltage ranges are known at the circuit terminals, the chip or circuit designer can begin a complete statistical design of the logic circuits using a circuit simulation program, such as ASTAP [4]. This program generates the best and worst case performance, signal swings, and noise tolerance values for the circuits. The packaging engineer is particularly interested in the characteristics of the circuits that interface with the package; *i.e.,* the off-chip drivers and receivers. As shown subsequently, these circuit values influence the requirements and trade-offs involved in optimizing the chip-package technology.

An important item in the signal distribution system, of course, is the characteristic impedance (Z_0). This parameter directly affects the performance, the net design, and the noise limits inherent in any computer packaging system.

Characteristic impedance affects the delay of a net because the loadings, which are functions of the system's electrical parameters, act in concert with the Z_0 to create a delay adder to the time-of-flight delay of the net lines and cables. An example of loading would be an SCM tapped into the middle of a transmission line net. The effect of the loading would be determined by the total capacitance of the SCM plus that of the chip at the pick-off point on the net. Since the loadings and line lengths in a net are variable, a special net delay equation is written for subsequent system timing analysis.

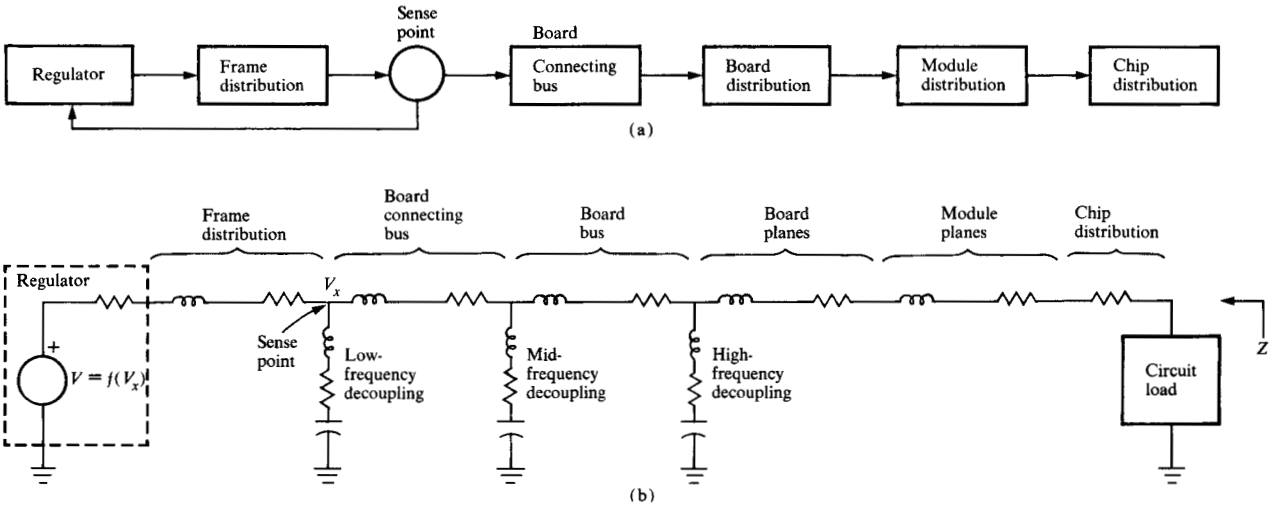

Figure 3 Block diagram (a) and equivalent circuit (b) for the power distribution system.

The loadings can also create significant reactive discontinuities on the net, which cause reflections on the transmission lines that make up the path. These reflections have to be limited to prevent false switching from occurring. The mechanism for accomplishing this is a wiring rule which controls the parameters that influence the magnitude of the reflections. The wiring rule can also be used to guarantee first incident switching; *i.e.*, the first time the signal waveform is propagated down the line, the receivers will switch and any reflections that are generated by the discontinuities will not cause the receivers to change their states.

Another effect of characteristic impedance on the signal distribution system is the generation of noise due to the switching of drivers and the coupling of transmission lines. Switching noise is caused by the inductances inherent in the power distribution system and the time rate at which the drivers switch the signal line currents. The amount of noise coupled depends on the proximity of the signal lines in the package and the time rate at which the voltage switches on these lines. Since excessive noise can cause delays and even data errors, these undesirable effects must be contained through the use of both switching rules that limit the simultaneous switching of drivers and coupled noise rules that limit the coupling lengths of parallel lines.

In the previous discussion, we alluded to delay equations, wiring rules, and noise rules. By themselves, these would serve merely as guidelines and sizing aids. However, we have integrated them into a CAD system, so they can be used to control the design.

Besides the electrical information for both the chips and package, the CAD system also needs descriptions of the physical image of each element as well as the logical interconnection information. Given all of these inputs, the design system can verify the functional performance, analyze the timings, generate test patterns for machine evaluation, and assemble data for controlling the manufacturing processes.

These concepts are explored more deeply in the ensuing sections. We next describe the design considerations for a power distribution system.

Power distribution

The power distribution system for a typical large high speed machine is diagrammed in Fig. 3(a). A regulated voltage supply feeds current into the frame distribution system, which is an assembly of laminar bus bars. At a point near the boards, a large amount of decoupling capacitance is inserted for filtering the lower frequency components of current changes caused by circuit switching throughout the system. Because voltage at this point is well filtered, it is a good place to sense the feedback voltage for controlling the regulator. Beyond the sense point, there is usually a board-connecting bus to distribute power to the board buses, and it is here that mid-frequency decoupling capacitors are located. From this point, current flows through the board bus into the board planes, where the high frequency decoupling capacitors are mounted on the boards. As power leaves the board, it enters the module through pins. Current within the module is further distributed on solid planes into the vertical vias that are connected to the chip pads. The final element for power distribution is the on-chip metal.

The electrical equivalent circuit of the power distribution system is also shown in Fig. 3(b). The mechanical structure of the system is mapped into equivalent values of inductance

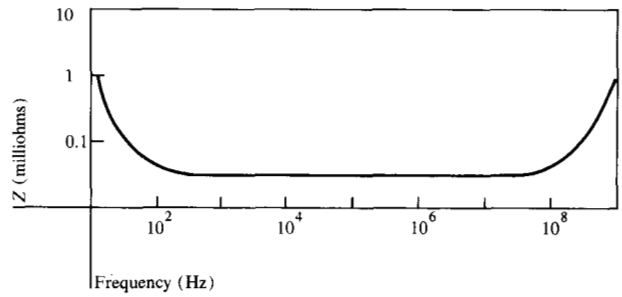

Figure 4 Output impedance *vs* frequency for the power distribution system.

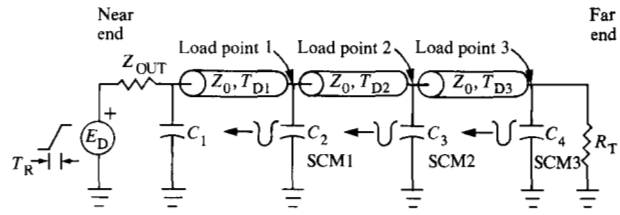

Figure 5 Discrete or distributed net configuration showing reflections from discontinuities.

and resistance by the package modeling programs. The main capacitance effect is from the low-, mid- and high-frequency discrete decoupling capacitors. Stray structural capacitances are negligible, but the inductance and resistance associated with the decoupling capacitors do tend to be significant.

For a good design, the frequency-dependent driving point impedance (Z) of the power distribution system at the circuit terminals is kept very small compared to the impedance of the circuit load to avoid large potential drops in the distribution system. A typical impedance plot for a power distribution system looking back from the circuit load is depicted in Fig. 4. At very low frequencies, the network appears to be capacitive, while at very high frequencies, the network appears to be inductive. In the mid-range, the capacitance compensates the inductance, yielding a very small impedance for the power distribution system. The goal is to design the system so that the curve is flat and resistive throughout the frequency range required by the speed of the circuits.

Synthesizing the desired impedance plot in Fig. 4 requires the evaluation of the equivalent circuit shown in Fig. 3(b). This circuit model can be combined with the equivalent model of all the circuits in a simulation program such as ASTAP. At this point, a time domain analysis with a switching load, a frequency domain analysis with a quiescent load, or both, can be performed. By adding or subtracting capacitance, the best fit to the desired impedance curve can be obtained.

The model in Fig. 3 can also be analyzed for dc potential drops by studying the losses across the resistors under a condition of high circuit current. This information is used to determine the voltage tolerances required for understanding the delay and noise tolerance characteristics of the circuits.

After the design of the power distribution system has been completed and the driver and receiver characteristics have been determined, the design of the signal distribution system can commence. As depicted in Fig. 2, this leads us into the areas of net design and noise control.

Signal distribution

The signal distribution system for a high speed computer accounts for a considerable part of the total path delay. To see how these delays can be minimized, we examine the role that characteristic impedance plays in designing a signal distribution system.

As implied in Fig. 2, choosing the characteristic impedance is important to all aspects of the signal distribution system. Consequently, we must study the effects of net design, net performance, and package noise on Z_0. By superimposing these effects, we can generate a design space for selecting an appropriate value of Z_0.

First, a brief description of logic net operation and noise effects is presented. A good background discussion on this subject can be found in [5].

Consider the electrical diagram of a logic net shown in Fig. 5. This net has a driver, represented by its Thevenin equivalent circuit, that delivers a waveform to three receiving chips, each residing on a separate SCM. A capacitor represents the discontinuity that each module and chip adds to the transmission line. Resistor R_T is used to terminate the line.

When a signal is sent out by the driver, only a fraction of the unloaded voltage swing enters the line because of the voltage divider consisting of Z_0 and Z_{OUT}. The driving module capacitance, C_1, simply loads the driver, slowing it down; hence, it is not of great concern for net behavior.

As the waveform is propagated down the line, reflections are generated at each capacitive discontinuity. When the lines are long, the reflections do not affect the rise time or each other; this type of net is said to be "discretely loaded." For a line to be considered long, each segment should have transit time (T_D) that exceeds one-half of a rise time (T_R).

The reflection at each capacitive discontinuity exhibits the following properties [5]:

For

$T_R > 1.5C_DZ_0$,

$$V_r = -\frac{C_DZ_0V_i}{2T_R}, \tag{1}$$

$$T_{W50} = T_R, \tag{2}$$

$$T_{W0} = T_R + 1.5C_DZ_0, \text{ and} \tag{3}$$

$$T_{\Delta 1} = \frac{C_DZ_0}{2}, \tag{4}$$

where

C_D is the capacitance of the discontinuity,
V_r is the peak voltage of the reflection,
V_i is the incident voltage magnitude,
T_{W50} is the width of the reflection at the 50% points,
T_{W0} is the width of the reflection at the baseline, and
$T_{\Delta 1}$ is the delay added to the mainline incident signal because of the discontinuity.

Thus, as C_D and Z_0 increase, the reflections get larger and $T_{\Delta 1}$ gets longer. Limits are set on these parameters. If they were not, the reflection from load point 2 could be so large when it hits load point 1 that the receiver at load point 1 would transiently switch into its down state, causing a logical error in a downstream latch. We discuss later the writing of wiring rules to prevent this from happening.

When the line lengths get short (T_D is less than one-half of a rise time), the reflections in Fig. 5 merge together and occur during the rise time at the near end. The load capacitance is now combined with the line capacitance and treated as if it were uniformly distributed along the line. This type of net is called a "distributed" net [5].

The additional line capacitance acts to lower Z_0 and increase the propagation delay, T_0. The following formulas express this situation:

$$Z_L = \frac{Z_0}{\sqrt{1 + C_L/C_0}}, \tag{5}$$

$$T_L = T_0\sqrt{1 + C_L/C_0}, \tag{6}$$

where

Z_L is the loaded board line impedance,
Z_0 is the unloaded board line characteristic impedance,
C_L is the total load capacitance per unit length,
C_0 is the total board line capacitance per unit length,
T_L is the loaded board line propagation delay, and
T_0 is the unloaded board line propagation delay.

From Eqs. (5) and (6), an expression for the distributed delay adder due to loading can be written:

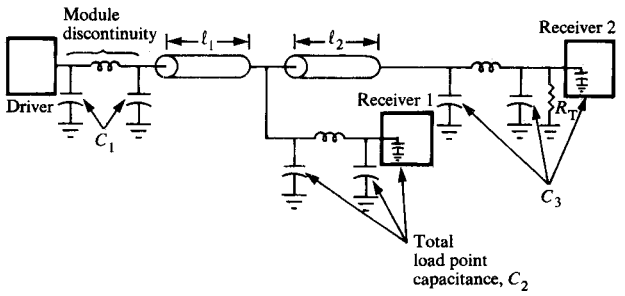

Figure 6 Equivalent circuit for a distributed net with two receiver load points.

$$T_{\Delta 2} = T_0\left(\frac{Z_0}{Z_L} - 1\right). \tag{7}$$

Treating each discontinuity as a lumped capacitance is helpful for understanding fundamental dependencies, but for actual design work more accurate modeling is required. The key missing element is the inductances of the SCMs and connectors. Once inductances are inserted into the model, the analytical approach becomes unwieldy and one must resort to circuit simulation techniques.

A complete electrical equivalent circuit of the signal path can be modeled using the previously mentioned package modeling programs. An example of the result is given in Fig. 6. These inductances and capacitances are shown lumped together into a π-section model of the module discontinuities. This yields a more complete model of a logical net with a fan-out (receiving loads) of two.

Including the module inductance reduces the loading of a purely capacitive discontinuity. This, in turn, reduces the reflection magnitude and the delay on the transmission line due to the module discontinuity. The inductance, however, makes the module signal path behave as a low pass filter with a finite cut-off frequency. The low pass filter could increase the delay of the on-module signal beyond what a pure capacitance would predict. Nevertheless, our experience has shown that, although the capacitive effects are more significant, the inductive effects are also important and that the accurate generation of wiring rules and delay equations for a nanosecond signal packaging technology requires their inclusion.

To illustrate the generation of a wiring rule, assume we want to design a distributed net and that we wish to specify the parameters of the net so as to guarantee its proper operation over the desired design range. Also assume that we want all of the nets to be first-incident. (In most cases, the latter requires that the nets be terminated near their charac-

teristic impedance.) Given this set of requirements, what sequence of steps should we follow to design this net?

To begin with, the topology shown in Fig. 6 should be coded into a circuit simulation program such as ASTAP. The blocks denoted driver, receiver 1, and receiver 2 are transient or high frequency switching models of the actual silicon circuits. The values used for the elements of the equivalent circuit model are chosen to yield a "3σ statistical worst case" condition for generating a wiring rule. This means that the limits of the wiring rules are valid for 999 out of 1000 cases as such design variables as device parameters, voltage levels, and temperature change throughout their specified ranges. A method of finding the 3σ statistical worst case circuit values is discussed in [6], and we subsequently refer to it as the "method of Chang." A true non-statistical worst case analysis, while safe and easier to perform, is too pessimistic, rarely resulting in an acceptable design.

The net output parameters that are set to their statistical worst case conditions as the input parameters are varied are the value of reflection noise and the minimum magnitude of the up or down voltage levels at a receiver input. Low receiver input voltage magnitudes cause a minimum circuit noise tolerance while the maximum reflection noise causes a noise input stress condition for the receiver. With the net input parameters set for this condition, the net line lengths and loading values are varied until the reflection noise approaches the corresponding noise tolerance. Since the length and loading values resulting from this procedure define the receiver's noise failure point, these values are incorporated into the wiring rule as the design limits.

Consider the case of the distributed net we have been using as an example. If the line segments get too long, the net will cease behaving like a distributed net. This in itself is not necessarily bad, but if the net had both distributed and discrete properties, it would be difficult to generate a simple corresponding delay equation, which is discussed later. Also, if the line segments became too short, the impedance would get too low, as predicted by Eq. (5), and the driver's output level due to its finite output impedance would become insufficient for the net. A similar situation would occur if the load point capacitance were too high. Consequently, for a properly designed distribution net, both the line length minima and maxima and the maximum load point capacitance need to be controlled. This can be expressed in the form of the following wiring rule:

$$K_1 \le \ell_i \le K_2\,,$$

$$C_j < M_1\,, \qquad (8)$$

where

K_1 and K_2 are constants with $K_1 < K_2$,
ℓ_i's are line segment lengths, as shown in Fig. 6,
C_j's are load point capacitances, as shown in Fig. 6, and
M_1 is a constant.

This procedure for generating a wiring rule can be generalized for the other net types. Once defined, the wiring rule represents a limit that the system designer must satisfy to meet the design requirements. (The procedure for incorporating these limits into the design process is subsequently discussed.)

For every net type wiring rule there is a corresponding delay equation that predicts the delay of each off-chip net as a function of the line lengths and loadings. To determine delays accurately, taking into account reflections, circuit nonlinearities, and parasitics, the circuit of Fig. 6 is again simulated by ASTAP to generate a set of data that is used for finding the coefficients in a delay equation.

The data are generated for various line lengths and loadings with all parameters at their nominal values. The outputs of these simulations are placed into a data base, which is then used by a least squares fitting routine to generate a linear delay equation that is a "best fit" for the delay values as a function of the line lengths and loadings. An example of a typical delay equation is

$$T_{D1} = a_0 + a_1\ell_1 + a_2\ell_2 + a_3C_1 + a_4C_2\,, \qquad (9)$$

where

T_{D1} is the delay to receiver 1 in Fig. 6,
ℓ_1, ℓ_2 are the lengths of the line segments,
C_1, C_2 are the total load point capacitances, and
a_i's are fitting coefficients.

A similar equation is written for the delay to receiver 2.

Equation (9) is a nominal delay equation that is placed into a special timing analysis program [7] that can time every logic/array path in the computer. The program also statistically sums the delay variations and predicts the cycle time range for the machine under design.

Because the program requires the standard deviation of the delay distribution as an input, a 3σ tolerance for each type of net delay equation is calculated. This again requires the application of Chang's method for finding the simulation input parameters that yield a slow net and a fast net. These nets are subsequently analyzed to find the positive and negative delay tolerances for each net type.

This essentially completes the description for the "net design" and "performance" paths in Fig. 2. Before we can consider the trade-offs for choosing the characteristic

impedance of an optimized electrical packaging design for a high speed computer, we study the impact that characteristic impedance has upon the noise sources in the system.

Noise limitations

Three types of noises generally concern the package designer—reflection noise, switching noise, and coupled noise. Reflection noise has already been described. Its effects are included in the wiring rules, and since it does not occur at the same time as switching noise and coupled noise, we assume it is not superimposed on these and can be treated separately. Switching noise and coupled noise, however, interact and they are treated together.

Consider Fig. 7. Two off-chip paths emanating from the same sending chip and terminating at the same receiving chip run parallel in the package. The effective inductance (L_{eff}) of the package and chip between the well-filtered board voltage planes and the circuit voltage terminal on the sending chip is also shown. In addition, electromagnetic coupling between the two transmission lines is indicated.

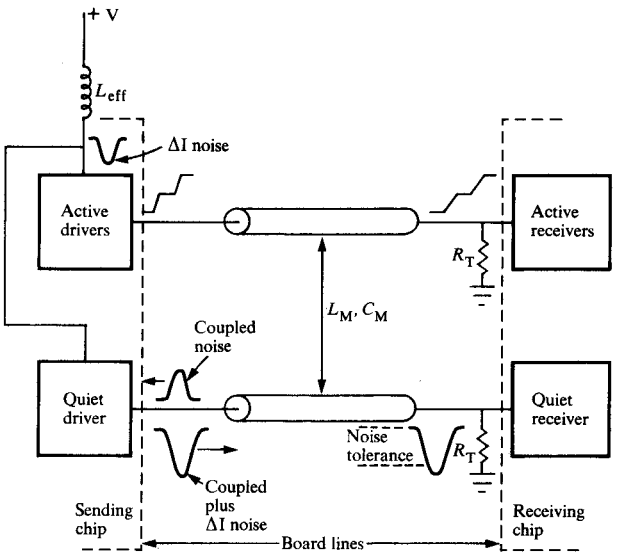

Figure 7 Causes and effects of switching and coupling noises on signal lines.

Before discussing the generation and interaction of switching noise and coupled noise, we first consider the determination of the effective package inductance.

Figure 8 shows the equivalent circuit of a SCM with a circuit load for an LSI chip. Assume that a lightly loaded logic chip has 200 quiescent circuits and 20 drivers that are simultaneously switching. (Internal circuits, even though they may be switching, are assumed to be quiescent for this analysis because their noise production is small and their net amount of up and down switching imbalance is low.) Both the quiescent load and the switching load are modeled because each one affects the amount of switching noise produced at the chip terminals.

The SCM model is shown to be only inductive for power distribution because the capacitive effects are too small to have a significant impact for the current levels involved. The module pins, planes, and interconnection vias are all modeled using the package analysis programs discussed earlier. All self and mutual (M) inductance effects between the various power paths are included. The signal paths for the 20 drivers are also effectively modeled, including all inductive and capacitive self and mutual effects. As depicted, the signal lines feed terminated board transmission lines when they leave the SCM.

Once this SCM model is available, the chip can be analyzed in its true environment, instead of the idealized voltage source environment in which circuit designers normally operate. By including the quiescent load, the on-chip impedance between the voltage nodes is represented.

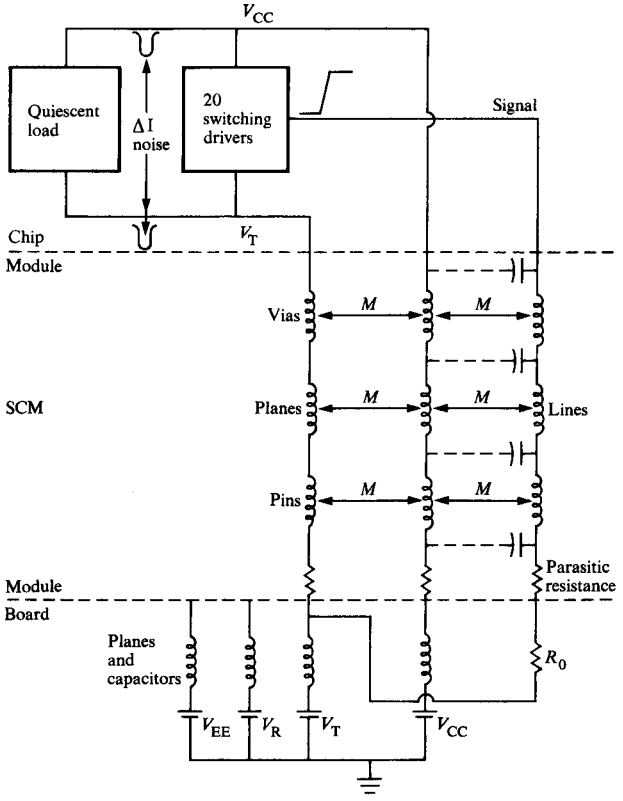

Figure 8 Equivalent circuit of the package used for determining the L_{eff}.

Since this impedance is small, it has a great shunting effect that significantly affects the noise levels on the voltage lines.

As an illustration, consider what happens when the switching load changes state. Assume that current is drawn from the V_{CC} supply through the driving circuits into the signal lines. This current through the module inductances induces a voltage across the module that is proportional to the rate of change of the current, so circuit speed is of primary importance. The induced noise is shown on V_{CC}, where one assumes it would occur. The noise, however, appears on the V_T supply at the chip as well as at all the other supplies because of the on-chip shunts and the mutual coupling within the SCM. All of these affect the resultant noise, compelling the package analyzer to include all of them if he is to accurately predict noise levels.

The effective package inductance cannot be directly determined by manipulating and simplifying the equivalent circuit of the module because of the model complexity caused by the parallel paths, the mutual effects, and the effect of the chip load. To overcome these problems, the circuit in Fig. 8 can be simulated in ASTAP with the resultant noise waveform generated. From the simulation, the noise magnitude and the current slew rate can be determined. Substituting these values into the following formula yields the effective inductance

$$L_{\text{eff}} = \frac{V_n}{N \dfrac{di}{dt}}, \tag{10}$$

where

V_n is the noise magnitude,
N is the number of simultaneously switching drivers, and
di/dt is the current slew rate of one driver.

Using this back door approach, the value of L_{eff} that produces the correct noise can be placed into the equivalent circuit for explaining the noise effects in the system. This has been done, as indicated in Fig. 7.

Figure 7 shows a set of active, simultaneously switching drivers and one quiet driver with its associated network. For illustrative purposes, we again assume that there are 20 simultaneously switching drivers. When these drivers switch, a negative going inductive voltage appears across the effective package inductance. This voltage is commonly referred to as switching noise or ΔI noise.

Switching noise has two primary effects upon system operation. First, it can cause larger off-chip delays than are normally accounted for. In addition, switching noise can cause the loss of data integrity within the system.

Consider that the ΔI noise is on the order of 300 mV high and 0.5 ns wide at the 50% points. Most driver circuits allow the positive supply noise to feed through to the signal output. When this noise is superimposed upon a normal up-going output transition, a flat spot develops in the waveform. If this distortion occurs below the receiver threshold level, increased delay, on the order of a few hundred picoseconds, will occur on every off-chip net that emanates from the sending chip. Given that a machine critical path could have three or four off-chip nets, the overall cycle time could be affected by more than a nanosecond or about 3% for a 30-ns machine. Because of the magnitude of this effect, Eq. (9) is rewritten with an appended term as follows:

$$T_{DI} = a_0 + a_1\ell_1 + a_2\ell_2 + a_3C_1 + a_4C_2 + \sum_{k=1}^{m} D_k N_K, \tag{11}$$

where

m is the number of different driver types,
D_k is the noise penalty coefficient in ps per switch for each different driver type on a given chip, and
N_K is the number of each driver type that is simultaneously switching.

Besides increased net delay, ΔI noise can also cause improper data to be stored in a system register. Consider the following case while referring to Fig. 7. The ΔI noise caused by the active drivers switching into the up state is a negative going voltage pulse that appears at the quiet receiver input as a spurious signal. When this noise exceeds the receiver's noise tolerance, the receiver output will be greater than its input. As a result, the noise can be propagated through the logic on the chip until it reaches a latch. If this latch is improperly set by the noise, a datum error has occurred.

Also, any transmission line terminators on the same chip that switch at or near the same time that drivers are switching can add to the total amount of ΔI noise. Consequently, they should be considered along with the drivers as ΔI noise contributors.

Compounding the problems with ΔI noise is coupled noise or cross talk, caused by the electromagnetic interactions between signal lines in close proximity. In certain situations, it is likely that coupled noise will add to ΔI noise.

Assume that the nets in Fig. 7 are part of a wide data bus system; *i.e.*, the single quiet line is in a sea of actively switching lines. Each line begins at the same sending chip and terminates at the same receiving chip. It is also reasonable to assume that all of the lines are laid out in a large parallel path on the board. This physical situation describes

the worst case noise superposition problem in a computer package because coupled noise is coincident with switching noise.

When the active drivers associated with a bus simultaneously switch, the ΔI noise is created at the output of the quiet driver as previously discussed. At the same time, coupling between the active lines and the quiet line creates coupled noise on the quiet line.

Since the board is a homogeneous structure, all of the coupled noise propagates towards the near end of the quiet line. With positive going active transitions, the coupled noise is also positive. For a given Z_0, the magnitude of the coupled noise is proportional to the mutual inductances and capacitances between the lines and the voltage slew rate of the active lines, while the width is proportional to the coupled line lengths and the rise time [5]. Depending upon the values of these parameters, the coupled noise can be large in amplitude and width.

When the coupled noise traveling towards the near end of the quiet line hits the driver, it is reflected back towards the far end. Since a good driver has a low output impedance, this reflection will be negative and large. After the occurrence of the near-end, quiet-line reflection, the coupled noise has the same polarity as the ΔI noise. If the coupling begins close to the drivers, the coupled noise will be superimposed on the ΔI noise, and both of them will be propagated towards the far-end quiet receiver.

Since parallel data paths in a bus configuration are common, the coincidence of switching noise and coupling noise has to be accounted for during the design of a high speed computer. The maximum coupled noise can be modeled using the aforementioned package analysis programs and circuit simulation programs. A maximum value can be determined because coupled noise reaches an amplitude saturation when the line delay of the coupled length equals half of the active waveform's rise time [5]. If coupled noise is excessive, a rule that limits the coupling lengths can be written to yield a lower value. Once this value is determined, it can be subtracted from the value for the receiver's noise tolerance. The residual noise tolerance is allocated for the ΔI noise.

To limit ΔI noise, the amount of simultaneous switching of drivers on a given chip must be limited. These limits are expressed by an equation of the following form:

$$\sum_{i=1}^{p} K_i N_i \leq (NT - CN), \tag{12}$$

where

- p is the number of different types of drivers and terminators on a given chip,
- K_i is the noise coefficient in volts/switch for a given type of driver or terminator,
- N_i is the number of simultaneously switching drivers or terminators of a given type on a given chip,
- NT is the noise tolerance of the receiver in volts, and
- CN is the coupled noise allowance in volts.

Equation (12) should be calculated under 3σ statistical conditions; *i.e.*, the simulation model uses the method of Chang to approximate an overall 99.9 percentile design, meaning that only one in a thousand cases over the entire design space should fail. Again, it should be emphasized that calculating Eq. (12) under true worst case conditions would yield an overly pessimistic and restrictive design.

The simultaneous switching limit equation is applied to every chip in the system to guarantee that errors due to noise do not occur within the computer. To accomplish this, the CAD system used tests every chip in the system against Eq. (12).

As discussed, the simultaneous switching equation exactly describes the wide bus situation. It assumes that the switching noise combines with the coupling noise and that they both arrive at a latch that is timed to receive a datum at the same instant. For a bus in a critical machine path, this is realistic. Since large high speed computers tend to have many parallel data buses, there is justification for following this type of noise design procedure.

The design space

Sufficient information is now available for analyzing the various factors that contribute to designing a computer package. The key parameter that ties everything together is the board's characteristic impedance. We now consider the various dependencies upon Z_0 and generate a design space that allows the selection of an optimized value for Z_0.

The dependence of the transmission line delay adder due to loading effects upon board impedance is shown in Fig. 9. Notice that there is an upturn in the delay adder for both low and high values of Z_0. The upturn for high values of Z_0 is explained by Eqs. (4) and (7) for discrete and distributed lines, respectively. The upturn for low values of Z_0 is a little more difficult to explain.

To begin with, every driver has a finite output impedance (Z_D). As a result, the voltage input delivered to the transmission line is a function of the input voltage divider consisting of Z_D and Z_0. For a given Z_D, the input voltage to the line is reduced as Z_0 gets smaller. Since the rise time (T_R) remains the same, the amount of time required for the line voltage to reach the receiver threshold level increases as the line input voltage decreases.

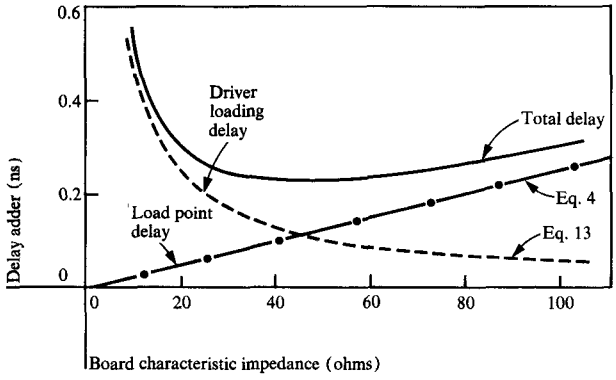

Figure 9 Delay adder *vs* characteristic impedance for a discrete net, with $C_D = 5$ pF, $T_R = 1$ ns, and $Z_D = 10$ ohms.

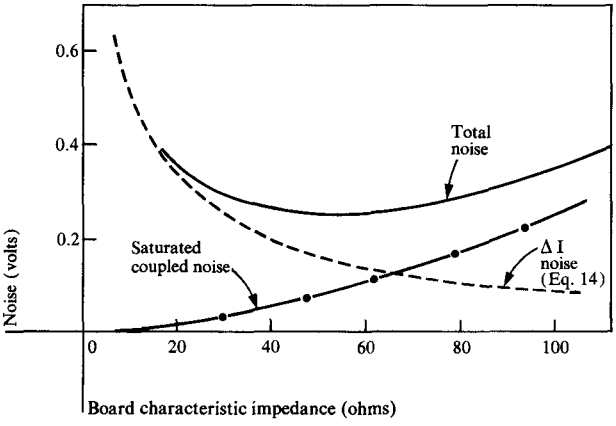

Figure 10 Noise amplitude *vs* characteristic impedance with two active lines and one quiet line. Parameter values are $L_{eff} = 0.5$ nH, $N = 20$, $V_D = 1$ V, $T_R = 1$ ns, $Z_D = 10$ ohms, $W = 100$ μm, $s = 400$ μm, $t = 25$ μm, and $\epsilon_r = 4.0$.

If we assume that the receiver threshold level is half of the driver's open circuit swing, then the following expression can be written for the delay adder, $T_{\Delta 3}$, due to the driver's output impedance:

$$T_{\Delta 3} = \frac{T_R Z_D}{2Z_0}. \tag{13}$$

By summing this equation with Eq. (4) for a discretely loaded net for the conditions shown in Fig. 9, the total net delay adder dependence upon the board characteristic impedance can be derived. Note that at low values for Z_0 the inability to drive the line with sufficient voltage causes the delay adder to go up drastically. At high values of Z_0, the delay adder increases linearly with the slope, being determined by the value of the loading capacitance for the SCM plus the chip input. In between the low and high values of Z_0, there is a region where the delay adder is minimized.

Noise can be drawn as a function of board characteristic impedance. For a given value of package L_{eff}, the ΔI noise ($V_{\Delta I}$) tracks with Z_0 as follows:

$$V_{\Delta I} = N \frac{L_{eff} V_D}{(Z_0 + Z_D) T_R}, \tag{14}$$

where

N is the number of simultaneously switching drivers, and V_D is the driver's open circuit voltage amplitude.

The coupled noise is not easily expressed mathematically because it is a complex function of the line geometries. Using the previously mentioned modeling and simulation techniques, the curve for near-end coupled noise shown in Fig. 10 can be generated for the conditions shown.

By combining the ΔI noise, Eq. (14), with the curve for coupled noise, the total noise curve shown in Fig. 10 can be drawn. At low values of Z_0, the ΔI noise dominates; at high values the coupled noise dominates. Again, there is a minimum region for the noise as there was for the delay adder.

The one missing element for choosing a good range for Z_0 is the dependence that noise tolerance has upon board impedance. This dependence is not easily determined analytically. It can, however, be determined through circuit simulation. The procedure involves analyzing the receiver's transfer characteristic at a particular set of up and down levels [8]. The resultant noise tolerance is valid for these input levels. However, knowing that the noise tolerance tracks directly with the input levels, we can draw the desired dependence upon Z_0 once we know one point on the curve.

Assume that the output from the simulation predicts a noise tolerance of 0.3 volt for the input corresponding to a 50-ohm board impedance system. Now, as the Z_0 varies, the line current varies and the voltage drop across the driver's output impedance (Z_D) varies. This causes the receiver input levels to change, which alters the noise tolerance. Using standard voltage divider techniques, a dependence for noise tolerance as a function of Z_0 can be stated as

$$NT = NT' + \left(\frac{Z_0 V_D}{Z_0 + Z_D} - \frac{Z_0' V_D}{Z_0' + Z_D} \right), \tag{15}$$

where NT' is the value of noise tolerance at Z_0'. For our example and under the conditions shown in Fig. 11, a curve of Eq. (15) is drawn for *NT versus* Z_0.

An interesting design space exists when the total delay adder curve from Fig. 9, the total noise curve from Fig. 10,

and the noise tolerance curve from Fig. 11 are superimposed, as in Fig. 12. Now we can study the interaction among delay adder, noise level, and noise tolerance for choosing an optimized Z_0 for the system.

To begin with, only Z_0's in the range where noise tolerance is greater than total noise are considered. This region exists between 40 and 105 ohms for our example. The other goal is to pick a Z_0 that minimizes the delay adder. For this particular example, the delay adder is a slowly increasing function of Z_0 in the region that satisfies the noise requirements, so picking a higher than minimum Z_0 would not be a great penalty. Also, other factors such as circuit power dissipation and improved noise tolerance encourage one to pick a Z_0 that is higher than 40 ohms. The larger delay adder discourages choosing high values for Z_0. Taking everything into consideration, it seems that a choice between 60 and 80 ohms would result in some additional noise margin at a slight increase in delay adder and that this range would be preferred. However, a choice within the 50 to 100 ohm range would be acceptable.

An interesting conclusion from this analysis is that, even though characteristic impedance is an important parameter for the packaging system, it appears that choices over a large range are acceptable in practice. The main limitation seems to be below 40 ohms, where all the interesting characteristics are degraded very rapidly.

System effects

It is difficult to accurately assess the effects that the package has upon overall system performance because of the large number of variables involved in determining path delays. Besides the silicon delays, key delay factors in the package are influenced by the way the system designer partitions the logic. Within the partitions, the placement of the chips can affect the overall delays. The physical distances involved, such as the module-to-module and the board-to-board spacings, also strongly affect delays.

The electrical packaging designer has little control over delay effects. Basically the packaging engineer is charged only with guaranteeing that the nets function properly in the 3σ statistical design domain. After he designs these nets and publishes delay equations, it is the system designer who uses this information in determining path delays.

One area where the package designer does have some effect upon delay is in making each logic net a first incident path. If receivers do not settle down in a known state until after a far-end net reflection has occurred, more path delay will be introduced. Another place that the package design affects delay is in the amount of simultaneous switching

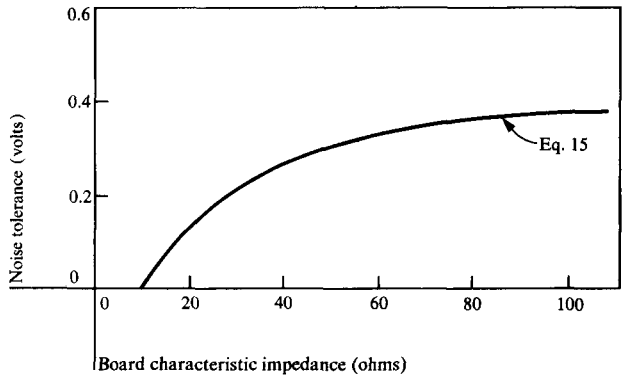

Figure 11 Noise tolerance *vs* characteristic impedance with V_D = 1 volt, Z_D = 10 ohms, Z_0' = 50 ohms, and NT' = 0.3 volt.

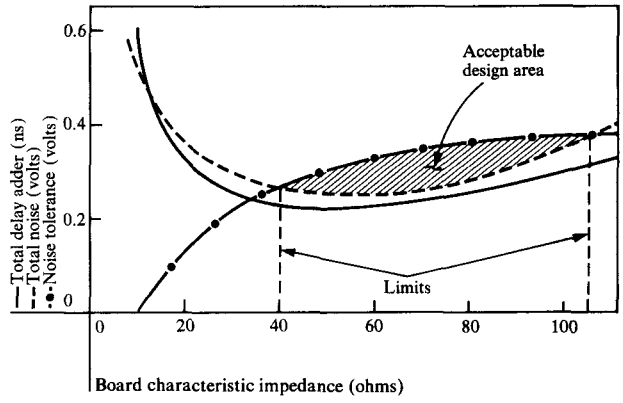

Figure 12 Design space for characteristic impedance.

allowed by the switching rules. This is because of the waveform distortion introduced by ΔI noise as described in Eq. (11).

Some insight can be given into the constituents of computer cycle time from a study done for the critical paths in a large high speed machine with a single-chip-module-on-a-planar-board packaging technology which has two types of paths: on-board and off-board. The former path is between SCM's on the same board; the latter is between boards and involves cables.

The results from analyzing output from a path delay calculator are shown in Fig. 13. Typical inter-module and inter-board critical paths are depicted with their delay components given as a fraction of the normalized machine cycle time. Notice that critical paths completed on one board

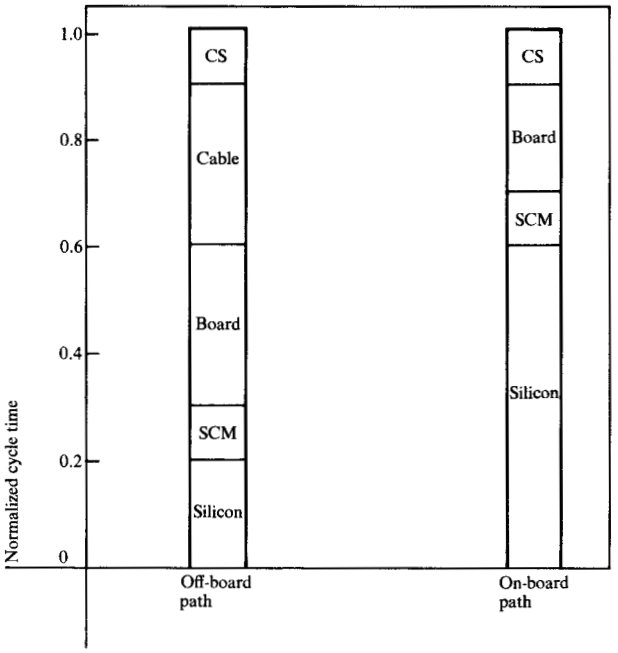

Figure 13 Machine delay components for characteristic paths.

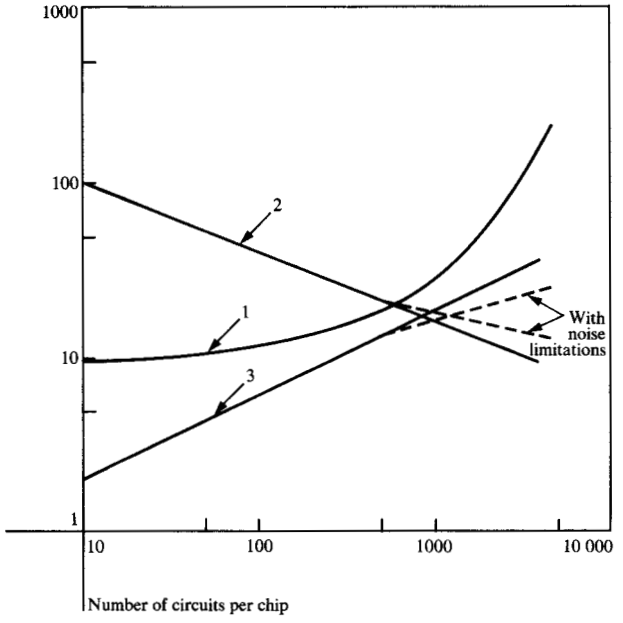

Figure 14 Impact of noise on projected system performance, where curve 1 is number of simultaneously switching drivers per chip, curve 2 is machine cycle time in ns, and curve 3 is millions of instructions per second.

exhibit a large amount of silicon delay, while critical paths completed on multiple boards exhibit a large amount of package delay. This, of course, is a result of the logic partitioning, which attempts to keep the high silicon delay paths on one board in order to minimize the machine's cycle time.

For the inter-module path, the silicon delay is 60%, the package delay is 30%, and the clock skew (uncertainty) is 10%. Package delay is 70% for the board-to-board path, silicon delay is 20%, and clock skew is 10%. From this we can see that whenever a critical path spills off a board, the package becomes the major delay element.

Another study was performed to ascertain the effect that noise has upon machine performance. In this case, a system constrained due to noise was compared to a system without noise constraints. The constraints considered were the simultaneous switching penalty due to waveform distortion and the deliberate slowing down of driver transitions to limit the noise produced. A switching penalty of 10 ps per switch and a maximum driver slew rate of 1 V/ns were used.

The results are plotted in Fig. 14. Three curves are shown: curve 1 shows the number of simultaneous switches allowed per chip; curve 2 shows the resultant cycle time prediction; and curve 3 shows the number of millions of instructions per second (MIPS). Each curve is for an IBM 3033-like machine organization and is drawn as a function of the number of circuits per chip. It was assumed that circuit delay halved for every decade jump in density.

Figure 14 indicates that, for up to 500 circuits per chip, noise was not a key problem because the circuits at this point are not stressing the package. Beyond this point, the increase in required simultaneous switching and circuit speed affects system performance due to increased package noise. As a result, the rate of performance improvement decreases. System performance continues to improve, however, because circuit delays improve (even though the driver transition times do not improve). Also, the greater circuit density means more use of the faster internal circuits and fewer module and board crossings.

Conclusions

In order to ensure the design integrity of a high speed computer, a computer aided design system has been created, as shown in Fig. 2. This software system converts the logical data flow desired by the system designer into a set of computer tapes that drive manufacturing machines. Out of these machines comes the hardware needed for assembling the new computer. Within this CAD system, all of the designs are audited to ensure that the rules are adhered to. Consequently, the rules are not mere guidelines; they must be satisfied before the design can proceed.

This type of design system comes about because the combination of LSI and high speed with large numbers of

distinct parts significantly increases the design complexity of the modern computer. With these trends, long design times become prevalent and the need to properly design the machine the first time is great. To make this happen, nothing can be left to chance; hence, the extensive use of simulation as described in this paper.

The electrical packaging engineer's role in the computer's design is to provide the rules that control the physical design of the boards and the cables and that contain the noise within the system. Without these rules, much could go wrong. As the computer industry advances into the realm of LSI, electrical package design will continue to be perfected.

In summary, a more rigorous approach to electrical computer package design than the cut-and-try method that is routinely followed by the industry has been presented. The described approach attempts to consider all the package parameters in the areas of power distribution, signal distribution, and noise design that could affect the performance and data integrity of the computer being designed.

Acknowledgments

Many of my colleagues have contributed to the knowledge used in writing this paper. During the past few years, I have benefitted from the opportunity to work with and learn from P. S. Aujla, M. Cases, R. A. Downing, M. J. Jesrani, G. A. Katopis, C. J. Parsons, and B. J. Rubin.

References

1. Pierce A. Brennan, Norman Raver, and Albert E. Ruehli, "Three-Dimensional Inductance Computations with Partial Element Equivalent Circuits," *IBM J. Res. Develop.* **23,** 661–668 (1979).
2. Albert E. Ruehli, "Survey of Computer-Aided Electrical Analysis of Integrated Circuit Interconnections," *IBM J. Res. Develop.* **23,** 626–639 (1979).
3. C. M. Sakkas, "Potential Distribution and Multi-Terminal DC Resistance Computations for LSI Technology," *IBM J. Res. Develop.* **23,** 640–651 (1979).
4. W. T. Weeks, A. J. Jimenez, G. W. Mahoney, D. A. Mehta, H. Qassemzadeh, and T. R. Scott, "Algorithms for ASTAP—A Network Analysis Program," *IEEE Trans. Circuit Theory* **CT-20,** 628–634 (1973).
5. A. Feller, H. R. Kaupp, and J. J. Digiacomo, "Crosstalk and Reflections in High-Speed Digital Systems," *AFIPS Conf. Proc., Fall Jt. Computer Conf.* **27,** 511–525 (1965).
6. F. Y. Chang, "Generation of 3-Sigma Circuit Models and Their Application to Statistical Worst-Case Analysis of Integrated Circuit Designs," 11th Asilomar Conference, IEEE Cat. No. 77CH1315-1, 29–34 (1978).
7. Robert B. Hitchcock, Sr., Gordon L. Smith, and David D. Cheng, "Timing Analysis of Computer Hardware," *IBM J. Res. Develop.* **26,** 100–105 (1982).
8. G. A. Katopis, "Methodology and Computer Aids for the Noise Tolerance of ECL Logic Gates," 12th Asilomar Conference, IEEE Cat. No. 78CH1369-8, 500–505 (1979).

Received October 20, 1980; revised December 1, 1981

The author is located at the IBM General Technology Division laboratory, East Fishkill Facility, Hopewell Junction, New York 12533.