

Interconnect design with VLSI CMOS

by R. F. Sechler

Historically, high-performance logic circuit interchip design has focused on bipolar emitter-coupled logic (ECL) circuits and signals, but VLSI CMOS has attained performance levels at which problems unique to its characteristics must be addressed for design optimization. In this paper, CMOS interchip circuit models are applied to develop packaging and wiring constraints for synchronous communication.

Introduction

For more than 30 years, the highest-performance computing systems have been designed around ECL circuits. An extensive body of analysis has centered on ECL switching in transmission-line environments. In IBM, ECL circuits were in use by 1960, but were not extensively documented until completion of the System/360™ Model 91 [1]. Since then, both transmission-line switching theory [2] and operation with ECL circuits have generated an extensive literature [3, 4].

The operation of transmission-line networks has a strong dependence upon the electrical characteristics of the interconnected circuits, but there are substantial differences between ECL and CMOS. Because of the nonsymmetry of its source impedance under positive or negative transitions, ECL is poorly suited to series (source) termination. However, its combination of small-signal-swing receiver and low-output impedance driver is a good combination for parallel (load) termination. As a result, the published body of analysis and design is

primarily oriented toward parallel-terminated transmission lines.

CMOS, with its large-signal-swing receiver and relatively high-output impedance driver, is less suitable than ECL for parallel termination. However, its symmetrical source impedance, comparable in magnitude to transmission-line characteristic impedance, is ideal for series termination. In extreme cases, series termination is slower than parallel [5], but its ease of implementation combined with its lower power and current makes series the preferred CMOS termination.


Our discussion here is therefore oriented toward series termination of symmetrical CMOS circuits. Although the analysis is generic, specific examples from IBM POWER, POWER2™, and PowerPC™ workstations [5] are applied. Also, since parallel termination is within the capability of contemporary CMOS, I will mention where it may have application. I conclude by projecting the likely direction for future CMOS high-performance synchronous interconnect circuits.

Synchronous signal control

Computing system signals can be divided into three classes—clocks, synchronous data, and asynchronous controls. Of the three, synchronous data exceed the other two by more than two orders of magnitude. For this reason a large body of application rules has been developed for their implementation, and for this reason we consider them in depth here.

A simple unidirectional network delay bound for series termination is shown in Reference [5], Figure 4, Equation (1). A more comprehensive model for a multisource

©Copyright 1995 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.



Transmission line ($L_1 + L_2 + L_3$): Z_0 (impedance), T_0 (delay)
Sources (S), Receivers (R), $A = R_d/Z_0$

Signals	Path length	Composite signal
1, 2	$L_2, L_2 + 2L_3$	$2/(2A + 1)$
5, 6 (VR_1)	$2(L_2 + L_3) + L_2,$ $3(L_2 + L_3) + L_3$	$-2/(2A + 1)^2$
3, 4	$2L_1 + L_2,$ $2L_1 + L_2 + 2L_3$	$+4A/(2A + 1)^2 = E_s$
7, 8 (VR_2)	$4L_1 + L_2,$ $4L_1 + L_2 + 2L_3$	$-2A/(2A + 1)^3 = E_t$

Figure 1

Series-terminated transmission-line switching signals.




Figure 2

Capacitive cluster noise network.

synchronous network is shown here in **Figure 1**. Delay (timing completion) is dependent on reinforcing reflections from the extremities of the transmission line. An M -drop delay equation from source to receive for the longest reflected signal path becomes

$$T_d \leq [2T_0(L_1 + L_3) + T_0(L_2)] + (MZ_0C). \quad (1)$$

The delay is composed of two terms. The first is the time for the latest switching signal to travel from source to

receive. The second is the summation of all reactive time constants arriving with or before this latest switching signal. This is the usual form for interchip delay equations. Enhanced accuracy may be realized by combining the characteristic capacitance of the transmission line with the load capacitance [1-4]. However, for our purposes Equation (1) is satisfactory.

Our goal is to develop packaging constraints and interconnection rules for which the delay bound of Equation (1) applies. A first subject to consider is driver and receiver models. Active drivers are treated as resistors with time-varying voltage sources, receivers and inactive drivers as capacitors in parallel with p-n diodes (Figure 1). These approximations are satisfactory in most cases. We comment where they are invalid.

Also, we explained [5] that signal control of synchronous data naturally separates into pre- and post-switching problems. Control of delay within the confines of Equation (1) falls into two categories. The first includes the conditions for which the network of Figure 1 correctly represents delay. These are a validation of network conditions *preceding* switching. The second includes the conditions under which the switching signal is certain to retain its level after timing completion. These are a validation of network conditions *succeeding* switching. We address both separately.

Three networks, shown in **Figures 1, 2, and 3**, define the extreme signal control environments. Readers knowledgeable with transmission-line interconnections will be familiar with these networks [1, 3, 4]. However, the characteristics of VLSI CMOS combined with our implementation produce unique design limits. Figure 1 is a network which can be used to assess the effect of the ratio of source impedance to transmission-line impedance on pre- and post-switching signals. Figure 2 is a network which can be used to assess post-switching noise from capacitive loading. Figure 3 is a network for assessing post-switching inductive and coupling noise among active lines through connectors. In each case our analysis considers solely the dominant network effect. These approximations are useful both for understanding and for analysis, because one or another of these elements tends to dominate a specific environment.

Before proceeding, it is worth considering just what signal levels are needed to complete switching and maintain switched levels. Our high-frequency design utilizes simple CMOS receiver amplifiers. The intended effect is to maximize noise tolerance (40-45% of signal) and minimize delay. We have two goals: First, we provide switching signals well above the switch point, perhaps 80% of V_{DD} . This ensures that transitions resulting from charging reactances are not excessive. Second, following switching we ensure switched levels at least 67% of V_{DD} , providing about 10% noise tolerance from other sources.

• Delay

Two conditions must be met in order to ensure that Equation (1) represents delay. One is that the network of Figure 1 can be approximated as a single transmission line to which sources (*S*) and receivers (*R*) are attached. The second is that the sum of the two transmitted signals is sufficient to ensure switching.

The first condition is violated by additional transmission-line segments connecting sources or receivers [1, 4]. Since transmission-line stubs are essential, the indicated rule is to restrict their maximum length. One rule of thumb is that the electrical length not exceed 1/6th the signal transition. Then the primary reflection from the stub arrives before switching of any receiver is initiated, and the stub may be approximated as a lumped capacitance. (In practice, stub rules must provide for segment wiring across connector boundaries, compromising an optimum.)

The second condition is satisfied by a low source impedance. Referring to Figure 1, we see four components to the signal—two incident and two primary reflections. In the network shown, it is not possible to ensure timing completion prior to receipt of signals from the left. The reason is that during switching the network forms an oscillatory system between the source and the open-ended line on the right. The transmission-line segment has an effective source impedance equal to the driver in parallel with the line segment. Therefore, lowering the driver impedance reduces intermediate signal levels.

Thus, a timing equation based upon receipt of primary signals from both ends of the line is essential to a series-terminated system. In Figure 1, the longest switching signal length is $2L_1 + L_2 + 2L_3$. If we let $L_1 > L_2 + L_3$, the last primary signal occurs after the first reflection from the source, VR_1 . Following switching, if $L_1 < L_2 + L_3$, there is a second reflection, VR_2 . Applying transmission-line algebra, we have

$$E_s/V_{DD} = 8A/(2A + 1)^2, \quad (2)$$

$$E_t/V_{DD} = 4A(4A + 1)/(2A + 1)^3, \quad (3)$$


with $A = R_d/Z_0$, for the switching and switched signal, respectively.

For a fully resistive source, switching and switched signals meet our criteria for $A \leq 1.33$. In some designs, all source resistance is intrinsic to the transistor. Output current then saturates by $V_{DD}/2$, converting the source into a current source. When a combined current source-resistive approximation is used, the upper bound for the linear portion of source resistance is a ratio not exceeding 1.

• Post-switching noise

Transmission-line matching


We have described a series-terminated transmission-line system having a delay bound based upon Equation (1).




I_i, I_r : Ground components of incident, reflected switching current

$$R_c \ll (R_d + Z_0)/N, R_c C_c \gg NL_c/(R_d + Z_0)$$

(a)



(b)



(c)

Figure 3

Coupling noise among active lines: (a) connector/source model; (b) equivalent circuit; (c) signal timing.

Following switching, transmission-line controls must maintain the signal well above receiver thresholds. Transmission-line mismatches and reactive elements contribute to this signal control problem.

Transmission-line matching is determined by the ratio between source resistance and characteristic impedance. This ratio can vary by up to a factor of three. Then, if one specifies a typical value of 1 for A , its maximum is 1.5, its minimum 0.5. The first value produces excess delay, expressed by Equations (2) and (3) above. The second produces underdamped network transient ringing. To constrain this second extreme within desired delay bounds, diode clamping has been widely used in series-terminated systems [6].

The arguments for dependence upon clamps is twofold: First, p-n junctions are inherent in the substrate design

of CMOS. Second, with the tolerances encountered in series termination, transient forward bias under extreme conditions is almost inevitable. Thus, it is in our interest to use these clamps to advantage.

The action of a clamp can be understood from Figure 1. If the source is at one extreme of the line, the initial signal at the other is $V/V_{DD} = 2/(1 + A)$. A high-performance design usually allows A to fall well below 1, forward-biasing a clamp. The current drawn by the diode diminishes reflection amplitudes. When noise sources from both reactances and transmission-line mismatches are considered, the diminished reflections alleviate noise problems.

The minimum signal for which clamping is helpful can be determined from Equation (3). There, $E_t/V_{DD} = 0.66$ at $A = 0.33$, below which clamps maintain the desired switched level. The required diode characteristic can be estimated. An incident signal produces a 25% overshoot, or about 0.8 V, at $A = 0.3$, producing $E_t/V_{DD} = 0.64$ without clamping. At this overshoot, the p-n junction current must compress reflections comparably to the signal compression below the desired 0.66. A diode current of 1.25 mA through a 50- Ω source (characteristic) impedance provides the needed compression.

Reflections are not complete with VR_1 and VR_2 . If $L_1 < L_2 + L_3$ and $A < 0.5$, clamp action maintains switched levels. Under these and other extreme conditions, the diode can draw large currents. Then, reflection control can be shown to require only that diode resistance be below the characteristic impedance. Finally, if $L_1 > L_2 + L_3$, nonclamped reflections succeeding VR_2 restore proper signal levels.

The effect of clamps on reflections is highly nonlinear. An advantage is that the typical low-resistance clamp produces no negative reflection at $3T_0L$, the time at which reactive reflections peak. If storage time is minute, the maximum reflection occurs at $5T_0L$. If storage times are large, larger reflections can be produced by the diode at $7-9T_0L$. In the extreme, their magnitudes approach the peak which would occur with no clamp present. For this reason, design with clamps requires control of their storage times.

Reactive reflections

Additional problems involve the effect of network reactances on switching signals. The reactive networks of Figures 2 and 3 produce reflections which peak at $3T_0L$. Since the series-terminated design anticipates timing completion at $2T_0L$ (T_0L in the minimum case), these reflections must be within the circuit noise tolerance.

There are two techniques to control reactive reflections through the source: One is its resistance, the other its signal transition. To quantify the second effect, we apply an exponential source signal with a time constant defined

relative to the network time constant. This allows the source to be analyzed with respect to packaging parameters creating problems. Also, the technique allows closed-form mathematical solutions to be developed with simple calculus.

An exponential source signal is a coarse approximation of the natural circuit output transition. Better approximations can be obtained by applying a trigonometric function, $V(t) = V_{DD}\sin^2(\omega t)$, $\omega t \leq 90^\circ$. However, our interest here is in a comparative analysis between circuit and package. Both the mathematics and the comparative evaluation would be unnecessarily complicated for a source signal functionally different from the network response.

In all cases, the form of the mathematical development is similar. Each network is satisfied by an equation of the form

$$dy/dt + ay = G[e^{-at}, e^{-Kat}, (at), (at)^2, \dots],$$

where y is a network voltage or current, t is time, G is a mathematical function, e is the natural logarithm base, a is the inverse of the network time constant, and K is a normalizing constant (Figures 2 and 3).

Results are expressed in the terms K and A , relating source to network for response time and resistance, respectively. The network time constant $1/a$ appears in the equations as the dimensionless number at . Noise voltages are calculated at the value of at at which noise peaks. Thus, the network time constant need not be expressed explicitly.

Capacitive clusters

The capacitive cluster (C_d) in Figure 2 produces an out-of-phase reflection V_r at $t = 3T_0L$. Given a switching signal E_s at $t = T_0L$, the equation for the switched signal at receiver R becomes

$$E_t/V_{DD} = E_s(t + 2T_0L) + V_r(t).$$

We analyze two cases: In one, the source has a time constant equal to that of the network ($K = 1$); in the other, the source time constant is a fraction of the network ($K > 1$). The two equations for the signal at the load are

$$E_s = [2/(1 + A)][1 - e^{-at}(1 + at)];$$

$$V_r = 4A/(1 + A)^2 - [2Ae^{-at}/(1 + A)^2] \cdot [(at)^2 + 2at + 2] - E_s; \quad (4)$$

$$E_s = [2/(1 + A)][1 - [(e^{-Kat} - Ke^{-at})/(1 - K)]];$$

$$V_r = \frac{4A}{(1 + A)^2} + \frac{4AK}{(1 + A)^2(1 - K)} \cdot \left[\frac{e^{-at}}{(1 - K)} + ate^{-at} + e^{-at} - \frac{e^{-Kat}}{K(1 - K)} \right] - E_s. \quad (5)$$

The reflection has its maximum magnitude within the time range $1/a < t_m \leq (A + 1)/(aA) = Z_0 C_d$, and the minimum total signal occurs at $5T_0L$ when $T_0L = t_m/2$. In this condition, the incident signal has the minimum time to charge prior to receiving the primary reflection, while reinforcing secondary reflections are not yet present. However, the long-line case is of more interest, because here the reflection has the greatest impact on delay. Therefore we accept $T_0L \geq t_m$ as the minimum line length to compute switched signals.

By using Equations (4) and (5), the minimum signal at the load can be calculated as a function of the impedance and transition ratios, as shown in Table 1.

The first case, of equal time constants, could result from source transition control or restricting the maximum capacitive cluster. The second case represents a large capacitive load wired close to the source, without source transition control.

The equations become optimistic for $A \geq 1$ for a nonlinear source resistance intrinsic to the transistor. When the resistance approaches Z_0 , current saturation in the source begins to reduce the incident signal. Considering this effect, the signal value listed for $A = 1.2$ approaches that which actually occurs at $A = 1$.

From Table 1 we observe that switched signal control requires constraints on both source resistance and transition. We might set the minimum source transition to be equal to the maximum network time constant ($K \leq 1$). Then the ratio (A) between source resistance and characteristic impedance must be less than about 1.2 for a linear and 1.0 for a nonlinear source.

Coupled lines

Whereas all networks previously considered were single-line problems, the network of Figure 3 results from many lines switching concurrently and unidirectionally. Also, two related means of signal coupling contribute to the effect. One is through a common power distribution; the second is coupling from active to noise-sensitive lines, shown in the model by a back-coupling coefficient K_b . As indicated, the two reinforce each other. They also are concurrent in time when the coupled lines are physically close to the power distribution inductance. This well represents a chip carrier or other connector model.

Before proceeding, it is of interest to point out a simple calculation which highlights the troublesome nature of this network. Consider the incident and reflected currents out of and returning into the source. It can be shown with transmission-line algebra that $I_r/I_i = 2/(1 + A)$. With $A < 1$, as is likely in a high-performance system, reflected current exceeds incident current. Of course, we might note that it is dI/dt , not absolute current, to which transmitted reflections are proportional. And bus line length variations and transition dispersion reduce dI/dt . Still, larger

Table 1 Minimum voltage from capacitive reflections (Figure 2).

E_r/V_{DD}	0.75	1.0	1.2	1.4	$= R_d/Z_0 = A$
$-1 - K = 1$	0.76	0.71	→ 0.66	0.61	Equation (4)
$-2 - K = 2$	0.68	0.64	0.57		Equation (5)

reflected than incident currents in the system are evidence that design constraints are needed.

Another surprising aspect of the model is apparent. An important link in the coupled noise through the connector is a low impedance between power contacts. For the VLSI chip carrier, this low impedance (large capacitance) is provided by the chip itself. The presence of this link has strong implications for chip design.

Also, we can see from the figure that there are two reflected signals. One is from reflected current (I_r) originating on the active lines and transmitted through the power distribution and coupled lines. The other is a secondary reflection from a signal originating on the sensitized (quiet) line (I_i). This is the smaller of the two, but its peak is of sufficient amplitude at $3T_0L$ to contribute measurably to total noise. However, coincidence of all signals is very implausible. Differences between bus line lengths distribute the total signal into a wider but smaller magnitude.

To assess sensitivities, we investigate this network using effective power distribution inductance as the parameter controlling network time constants. The term combines all mutual and self inductances, including signals and power contacts. The reader will recognize that package and circuit design depend upon analysis making use of the individual terms. In a specific design, however, they tend to track one another closely. Also, the duration of measured and predicted noise in POWER systems indicated that power distribution self inductance dominated. And reasonable correlation between measured and predicted noise was achieved with estimated power distribution inductance models only.

The two signals, originating on the active and quiet lines, have equivalent transmission-line delays. Therefore, an overall equation for the switched signal level is

$$E_r/V_{DD} = V(A) + V(Q),$$

where the switched signal noise occurs on a quiet line.

Again we consider two cases. One has a source time constant equal to the power distribution, the second a time constant equal to a multiple of the power distribution. The calculation is for a long-line environment where preceding signals have unlimited time to settle:

$$\begin{aligned} V(A) &= 2e^{-at} [at/(1+A)]^2, \\ V(Q) &= 2ate^{-at} (1-A)/(1+A)^2, \end{aligned} \quad (6)$$

Table 2 Maximum coupled noise voltage (Figure 3).

E_t/V_{DD}	0.6	0.667	0.75	1.0	$= R_d/Z_0 = A$
-1 - $K = 1$	0.50	0.45	0.39	0.3	Equation (6)
-2 - $K = 0.5$	0.37	0.34	0.3		Equation (7)

Table 3 Noise reduction from parallel bias resistance.

	$B = 2$	3	4
(1) $M = B(B - 1)/(B + 1)^2$	= 0.22	0.38	0.48
(2) $M = [2B/(2B + 1)]^3$	= 0.51	0.63	0.70
(3) $M = 4B^2(2B - 1)/(2B + 1)^3$	= 0.38	0.52	0.61

$$V(A) = \frac{4K}{(1+A)^2(1-K)^2} \{e^{-Kat} - e^{-at}[1 + at(1-K)]\},$$

$$V(Q) = \frac{2K(1-A)}{(1-K)(1+A)^2} (e^{-Kat} - e^{-at}). \quad (7)$$

From Equations (6) and (7) we can calculate maximum noise. For Equation (6), the maximum occurs at $t = 2/a$; for Equation (7), slightly above $2/a$. **Table 2** shows the reflection signal peaks.

Our first observation from the table is that source resistance has an opposite effect from that resulting from the capacitive cluster. Here, decreasing source resistance increases the reflection amplitude. Although clamps reduce reflections for $A < 0.67$, by this level reflections are excessive. Therefore, adequate switched signal levels require a power distribution time constant shorter than source transition.

Noise suppression can be achieved through design control on circuits or package. The most effective approach for the former is restrictions on simultaneously accessed outputs. But this degrades performance, and our desired approach is control within the package. Another design technique, discussed below, is available to control noise without degrading performance.

Resistive control of reactive reflections

Both source resistance and transition control have obvious limitations for managing reflected noise. For either, to avoid a major impact on system performance, we accept a limited increase in delay from the source. The problem is complicated by the difficulty of designing and modeling the networks of Figures 2 and 3.

In fact, in POWER systems design, difficulties with the approach became clear. Original designs to 41 MHz were tuned to 62.5 MHz by chip design mapping into a higher-performance semiconductor, with little package or circuit redesign. With this tuning strategy, long-line problems from reduced source transition eventually result. For our

design, source transition was halved, increasing reactive noise substantially.

The approach is understood by reconsidering the signals in Figures 2 and 3. Incident signals traversing the transmission line once are succeeded by reflections resulting from three line excursions. A bias resistor placed on the line will degrade the incident signal once and the reflection two to three times. This resistance need not act like a matched parallel terminator to reduce noise. Resistances exceeding $3Z_0$ are effective in increasing the ratio of switching to reflected signals.

For example, we can consider bias resistors R_0 placed in any of three positions in Figure 3. For the first, resistors are placed at the receivers. For the second, they are placed in the middle of the line. In the third, the resistors are split, and $2R_0$ is placed at each end of the line. For bias resistance $R_0 = BZ_0$, the noise multiplier is shown in **Table 3**.

The first approach is practical in a net for which the problem is unidirectional. (The net itself may not be unidirectional.) The second is a convenient solution for a (marginal) bidirectional problem. The third is a general-purpose bidirectional design solution.

The noise level may be increased by nonzero quiescent currents in the line. If the resistor is biased to $3V_{DD}/5$, the quiet signal level is $V_q/V_{DD} = (3A/5)/(A+B)$. For example, if $A = 2/3$, $B = 3$, the quiescent voltage is 0.11. When the double-resistor design solution (3) is applied, the peak level for $K = 1$ in Table 2 is 0.34, 76% of the original 0.45.

Bias resistors can be added to the completed system without source redesign. Since they affect delay minimally, they are a useful system design aid. However, given the large number of lines potentially involved, they are a poor substitute for reactance control. They are effective in marginal situations at the extremes of product capability.

Several wide data buses in POWER and POWER2 systems have bias resistors to suppress noise. In most cases, the problem attacked is connector noise. In one case a capacitive cluster dominated. A four-word bus with a large source capacitance (paralleled memory cards) transfers data from memory into cache. To ensure stable operation, initial system tests included stress data patterns on all major buses. When capacitive clusters dominate, the sensitized line switches, reducing but not eliminating inductive noise. When inductive coupling dominates, the sensitized line is quiet, eliminating capacitive noise.

• Synchronous signal control summary

We have defined a series-terminated transmission-line system having extreme delay limited by Equation (1), for interconnect nets of Figures 1-3. Signal control is divided into pre- and post-switching problems. From our analysis, a set of application rules can be proposed.

Two constraints apply for pre-switching control (Figure 1). The first is a maximum stub length, limited by rule of thumb to one-sixth the typical signal transition. (In symmetrical networks, considerable leeway can be applied.) The second is a maximum ratio of source resistance to transmission-line impedance. For a linear source, the limit is 1.33. For nonlinear sources having most resistance within the transistor (current saturation by $V_{DD}/2$), a maximum about 1.0 is a better guide.

Additional constraints apply for post-switching control:

- Figure 1: $A \geq 0.5$, or clamp control.
- Figure 2: $A \leq 1.25$ (1.0 for nonlinear source);
 $T_r \geq R_d Z_0 C_d / (R_d + Z_0)$.
- Figure 3: $A \geq 0.67$, or clamp control;
 $NL_e / (R_d + Z_0) \leq T_r / 2$.

The clamp diode must draw approximately 1.25 mA minimum at 750 mV forward bias ($V_{DD}/4$), and have resistance not exceeding Z_0 ; T_r is the natural (unloaded) transition (time constant) of the source signal.

A series-terminated system meeting the constraints above can provide the desired delay and noise management. There is a signal control technique available to compensate for inaccurate modeling or system performance tuning beyond original designs. The resistor biasing described previously can reduce post-switching noise by 15–25% in marginal environments.

In any specific design, the application rules postulated here could be developed with the aid of circuit analysis programs. Difficulties with interpretation would be readily resolved. These include stub lengths, the relationship between source and transmission-line impedance, and the relationship between source and network response times. The values presented here, however, are representative of those encountered in POWER and POWER2 interconnect designs.

Discussion

• Delay minimization

In a synchronous system, noise is synchronized from the clock and is therefore time-determinant. As a result, there is some freedom to choose interchip timing completion. For series-terminated nets, we chose two transmission-line delays, but could have selected three or more. The greater the number, the simpler become the constraints on noise, but the lower the maximum frequency supported. However, two line delays represent the *minimum* for our series-terminated system.

It is, of course, possible to reduce interchip timing below two transmission-line delays. To determine the performance limits of CMOS, it is necessary to investigate this potential. To do so, we first classify synchronous interconnections according to two properties.

The first property is the transmission-line switching mode, which can be either reflected or incident signal. We repeat the delay bound from Reference [5], Equation (1) and Figure 4.

Delay $\leq$ transmission time + reactive time constants

$$T_d \leq 2T_0 L + MZ_0 C \quad (\text{reflected signal delay}) \quad (8a)$$

$$T_d \leq T_0 L + MZ_0 C/2 \quad (\text{incident signal delay}) \quad (8b)$$

Second, the networks can be identified according to three categories: 1) Bidirectional, with source and receive at opposite extremities of the line. Examples are processor data buses [5]. 2) Single-source (unidirectional), with multiple receivers. Examples are address buses [5]. 3) Multiple-source (bidirectional) and multiple-receiver. Examples include the POWER system memory bus and the PowerPC bus [5].

If parallel-terminated, the first and third networks must have terminators at both transmission-line extremities. For the second, optimal wiring would place the source at one transmission-line extremity and the terminator at the other.

The simplest interpretation of the equation is that (8a) represents series and (8b) parallel termination. Actually, Equation (8a) represents series termination having adjacent sources and receivers, or parallel termination during transfer of nonadjacent source control. Equation (8b) represents single-source parallel termination or series termination with source and receive at opposite extremes of the transmission line.

We have, then, two important observations. First, to support multiple-source parallel-terminated delay, an unused cycle is needed during transfer of source control. Second, for series-terminated nets incident switching is supported, at least for Category 1.

To understand the first, consider the signal sequence when sources transfer control. Sources becoming inactive switch to a state (tri-state) for which they do not disturb the line. In the sequence below, sources at opposite extremes of a transmission line drive a receiver adjacent to source 1:

Cycle	Source level	Receiver level
	1 2	(adjacent to source 1)
N	0 —	0
$N + 1$	— 0	1 $\rightarrow$ 0

In cycle $N + 1$, control is transferred from source 1 to source 2. Initially, parallel terminator current switched out of source 1 creates a 1-level at the receiver. Also, source 2 initially switches little or not at all because its output level was previously 0. The positive signal from source 1 is transmitted to source 2, at which point an out-of-phase reflection is transmitted back to the receiver, providing the 0-level. However, this excursion is two transmission-line delays, more in accordance with Equation (8a) than with

(8b). Therefore, to enable incident switch delays during single-source operation, an unused cycle must be inserted during transfer of source control.

For series termination, no additional delay occurs during transfer of source control. With no quiescent current in the line, there is no extraneous signal to absorb when a source relinquishes control.

The second observation is the incident switch delay for Category 1 nets. There are two bases: First, all incident and primary reflection signals occur simultaneously at one transmission-line delay. Second, effective time constants in this net are reduced. Time constants close to the source are halved by the source impedance, and the sole receiver is at the open-ended line, maximizing its switching signal.

In considering the three bus structures, we see the advantage of CMOS symmetry. For Category 1, which dominates VLSI interconnects, series termination is always superior. Category 2, which with its single source requires no transfer cycles, can benefit from parallel termination. Its single terminator can be accommodated by contemporary CMOS. However, by customized layout of these relatively few address nets, delay minimization is usually possible with none. The layout goal would be wiring that is topologically similar to that of Category 1.

Category 3 *appears* to benefit from circuits readily parallel-terminated, because terminators would be required at both transmission-line extremities. But contemporary CMOS can provide parallel-terminated delays even for Category 3. Referring to Figure 1, we see that the "incident switching" of Equation (8b) actually allows two signals (one reflection) prior to timing completion. [As can be seen from Equation (2), reflected signal switching is contingent upon all *four* signals.] We can take advantage of this property and CMOS symmetry to provide parallel-terminated delays without matching the system.

In fact, parallel resistance well in excess of Z_0 can be combined with clamp diodes, satisfying Equation (8b) while suppressing trailing reflections resulting from a low-impedance source. Reflection control includes both source and load resistance, somewhat like the noise-suppression networks proposed to aid series termination. But to support a single transmission-line delay, the source resistance is centered below $Z_0/2$ (rather than Z_0 , as in series termination).

For unidirectional nets, a good example is the POWER system clock distribution [5]. If source resistance $Z_0/2$ drives a line terminated in $2Z_0$ to $V_{DD}/2$, an incident signal exceeding 75% of V_{DD} is provided. Succeeding reflections are well within circuit noise tolerance.

From Figure 1 we can assess the requirements to support a single transmission-line delay for Category 3 nets. The source impedance must provide an adequate switching signal, E_s , composed of any two of the four primary signals. (The minimum occurs with source and

receive at the same transmission-line extremity.) Parallel resistance must then suppress reflections sufficiently to maintain the switched signal, E_r . (This minimum occurs with receive at one transmission-line extremity and source within half a line length distant.) These criteria can be satisfied with parallel resistance of at least $3Z_0$ at the transmission-line extremities. The advantage of symmetry is that the source must drive $0.75Z_0$ transiently and $1.5Z_0$ quiescently. The source drives $Z_0/2$ overall in a matched parallel-terminated network.

Design with "incident" switching and nonmatched lines introduces the potential for transients from cycle N interfering with cycle $N + 1$. Interference is minimal as long as interchip delay does not exceed half the clock cycle. Beyond that, some excess delay may occur.

Finally, several issues weigh against parallel termination of Category 3 nets. The first is the unused cycle for transfer of source control. At high control-transfer rates, no bandwidth advantage is realized by termination. Second, bandwidth is defined by the slowest of cascaded links. In many applications, these limit bandwidth to that achievable with series termination.

We conclude that contemporary CMOS supports parallel termination, but that few interconnections benefit significantly from its use. In a circuit technology readily amenable to series termination, parallel will always be the rare exception.

• Design limits

Our interchip signal analysis illuminated packaging problems resulting from series termination. Of course, parallel termination has no shortage of design problems. The source cluster of Figure 2 has an obvious analog to a far-end cluster with parallel termination [1]. And the bus coupling of Figure 3 has a parallel termination analog [3]. However, parallel-terminated coupled noise peaks at one transmission-line delay, series-terminated at three. So, in one respect, to support the delay bounds of Equation (8), series termination *appears* to produce stricter demands upon circuits and packaging. On the other hand, series termination eases constraints on capacitive cluster wiring. With parallel termination, capacitive cluster noise results from *primary* reflections, while with series termination it results from *secondary* reflections.

Historically, parallel termination with incident switch delay has placed greater demands on packaging. The workstation packaging for which signal control is difficult costs less than ECL packaging, even at comparable cycle times. Our willingness to customize wiring to resolve exposures may partially explain the difference. And CMOS symmetry, along with series termination, simplifies this customization. Nevertheless, attention to packaging limitations appears essential to support continued CMOS performance growth.

Additionally, from Equations (2) and (3) we conclude that 3.3-V CMOS is the lowest supply level at which p-n junction clamps are effective. Clamps must draw appreciable current at $V_{DD}/4$ bias, which would require Schottky diodes by 2.5-V CMOS. Schottky diodes have advantages in low capacitance and storage time. However, at the highest performance, a combination of series and parallel resistive biasing may ultimately prove superior for signal control.

Concluding comments

One conclusion can be drawn from the analysis here as well as from the companion paper [5]. VLSI CMOS will demand investments in packaging to support its rapidly advancing performance. This is clear in the reactive effects on interchip signal networks, and from our prior clock distribution and power distribution analysis [5].

No circuit can circumvent the packaging problems which must be surmounted in high-performance design. However, CMOS drivers appear to me to offer advantages from their switching symmetry in two directions. This symmetry facilitates a continuum of series to parallel biasing of interchip networks. No other circuit property can compensate so well for packaging reactances and transmission-line mismatches. I include in this judgment the fabled ECL with its sub-100-mV sensing receiver.

There has been discussion of the deficiencies of CMOS receivers. Differential amplifier receivers could be employed if they enhance noise-tolerance-to-signal-swing ratios. But CMOS receivers provide an excellent ratio, with no additional power distribution investments.

Finally, concern has been expressed for the high power dissipation of CMOS drivers due to the large signal swing required by the receiver. However, the majority of CMOS interconnections will always be series-terminated, the lowest-power circuit of all.

CMOS drivers and receivers, both series- and parallel-terminated, are the right complement to CMOS logic chips.

System/360, POWER2, and PowerPC are trademarks of International Business Machines Corporation.

References

1. R. F. Sechler, A. R. Strube, and J. R. Turnbull, "ASLT Circuit Design," *IBM J. Res. Develop.* **11**, 74-85 (1967).
2. H. B. Bakoglu, *Circuit, Interconnections, and Packaging for VLSI*, VLSI System Series, Addison-Wesley Publishing Co., Reading, MA, 1990, Ch. 6.
3. E. E. Davidson, "Electrical Design of a High Speed Computer Package," *IBM J. Res. Develop.* **26**, 349-361 (1982).
4. Evan E. Davidson and George A. Katopis, *Microelectronics Packaging Handbook*, A. G. Klopfenstein, Ed., Van Nostrand Reinhold, New York, 1989, Ch. 3, pp. 111-165.
5. R. F. Sechler and G. F. Grohoski, "Design at the System Level with VLSI CMOS," *IBM J. Res. Develop.* **39**, 5-22 (1995, this issue).
6. E. E. Davidson and R. D. Lane, "Designing with Diode Transmission Line Terminators," *Electronics Magazine*, February 19, 1976, pp. 123-127.

Received August 15, 1994; accepted for publication October 14, 1994

Robert F. Sechler IBM Systems Technology and Architecture Division, 11400 Burnet Road, Austin, Texas 78758 (RSECHLER at AUSVM6). Mr. Sechler received his B.S. degree in engineering physics from Lehigh University in 1961, and his M.S. degree in engineering from the University of Vermont in 1983. In 1961 he joined IBM in the newly formed Components Division at East Fishkill, New York. During the succeeding 16 years he worked on a variety of assignments in bipolar logic circuit development. From 1964 to 1966 he was assigned to circuit design for the high-speed ECL circuits used in the IBM System/360 Model 91. From 1966 to 1971 he was a manager responsible for design of the ECL circuits used in System/370TM and advanced scientific machines; he later managed design for bipolar LSI circuits used in the IBM 4341 and 3081 processors. In 1977 he transferred to Austin to work on the small-machine semiconductor applications. Since 1983 he has been a senior technical staff member working on the implementation of high-performance technologies in advanced workstations. Mr. Sechler has several filed or issued patents and publications on computer logic circuits. He co-authored a paper on System/360 Model 91 logic circuit design, and has published recent work as part of the IBM RISC System/6000 workstation announcements in Austin.

System/370 is a trademark of International Business Machines Corporation.