

IN A DESIGN IN WHICH YOU MUST REDUCE TIGHT

TIMING, ROUTING ALL TIMING-SENSITIVE LINES

IN BURIED STRIPLINE LAYERS MINIMIZES ONE

SOURCE OF JITTER AND LOWERS THE OVERALL

REQUIRED TIMING BUDGET.

Analysis of board layout helps cure jitter problems

ALL HIGH-SPEED-digital-product designers face the critical electrical-performance challenges of meeting a timing budget, meeting a noise budget, and passing an EMC (electromagnetic-compliance)-certification test. Designers need to account for a number of factors when calculating a timing budget. Most high-speed, digital products are synchronous, clocked systems, and they require that a series of operations happen within one clock cycle. These operations include all the gate-switching delays within one logic depth, the intrachip propagation delays, the interchip propagation delays, the rise time or charging delays from the interconnections, the setup-and-hold times, and the skews between the clock and the data lines. The timing budget allocates how much time is assigned for each source of delay.

"Jitter" is the variation in arrival time of a clock or data edge from cycle to cycle. The amount of jitter for a given signal may vary from cycle to cycle. Signal-propagation-delay sources can be either random, when you cannot predict the jitter, or deterministic, when you can predict the amount of jitter. Finally, you must add a "margin" to the timing budget; margin accounts for the inability to accurately predict all the other terms. The larger the uncertainty, the larger the margin you need to be confident of the product's correct operation.

As the clock frequency increases, the clock period becomes shorter. Designers try to meet the goal of ever-higher clock frequencies by minimizing each term in the timing budget. The better you can accurately predict each element of the timing budget, the smaller the margin you need. A smaller margin allows you to shorten the period and still meet the timing budget.

One type of deterministic jitter arises from crosstalk between adjacent lines in a bus even in low-loss interconnects. With models of the interconnect

that allow prediction of the crosstalk and a good simulator that includes these crosstalk effects in the timing, you can accurately predict the deterministic jitter from crosstalk and minimize the jitter budget and reduce the margin. By understanding the source of the crosstalk-induced jitter, you can minimize its impact and make the cost-versus-performance trade-offs to optimize the final product.

CROSSTALK-INDUCED JITTER

Crosstalk-induced jitter occurs when the simultaneous signals on adjacent bit lines affect the arrival time of the signal on one bit line. If designers omit this form of deterministic jitter from their timing budgets, the product may fail. You can use a tightly coupled 3-bit bus as an example (**Figure 1**). In this bus, each line is a 50Ω, 9-in.-long microstrip, and the line and space are both 5 mils. You need only a 3-bit bus to understand the jitter problem because you need be concerned only with the center line and the track immediately adjacent. The effect of the tracks beyond is negligible compared with the effect from the coupling with the adjacent lines.

The center bit line is the "victim" line. You can simulate the arrival time of its signal at the receiver under three relevant conditions. You can simulate both "aggressor" lines—the adjacent bit lines on either side of the victim—to be




Figure 1 Two adjacent aggressor lines surround the victim line. The three microstrip traces have 5-mil lines and spaces.

off, switching with the same bit pattern as the victim line, or switching with the complementary bit pattern to that of the victim line.

The simulation shows that the arrival time of the victim line varies depending on the condition. The total delay time for the victim bit line is about 1.5 nsec when the other lines are off. The other two cases show a variation of the arrival time of ± 0.1 nsec, or about 7% the total delay time, depending on the bit pattern on the bus (Figure 2).

When the aggressor lines drive the same bit values as those of the victim line, the delay on the victim line is longer, whereas when the aggressors have the complementary bit values, the delay of the signal on the victim line is shorter. The coupled capacitive or inductive currents add to the quiet line, affecting the arrival time and the crosstalk-induced deterministic jitter. Although the far-end crosstalk and crosstalk-induced deterministic jitter share a common origin, the magnitude of the jitter is independent of the total crosstalk and independent of the rise time. Variations in the propagation speed of the signal create crosstalk-deterministic jitter, and you can eliminate it by careful design. It is possible to have a system with a high degree of crosstalk but no deterministic jitter.

SIGNAL SPEED AND DETERMINISTIC JITTER

More than 95% of pc boards use FR4 material. In surface traces, or microstrip topology, the effective dielectric constant the signal sees is a combination of the contribution of the bulk dielectric constant of the laminate, typically on the order of 4.5 for FR4, and the air, with a dielectric constant of 1. In general, the exact value of the effective dielectric constant is a complicated function of the

cross-section geometry and the bulk dielectric constants.

The propagation delay of a signal down a uniform transmission line relates solely to the effective dielectric constant the signal sees as it propagates. If the effective dielectric constant, k_{EFF} is about 3.5, as might be the case in an FR4 microstrip, the propagation speed of the signal is about $12 \text{ in./nsec}/\sqrt{(3.5)} = 6.4 \text{ in./nsec}$. A 9-in.-long line, for example, would have a time delay of about $\text{length}/6.4 \text{ in./nsec} = 9/6.4 = 1.4 \text{ nsec}$. If

you add to this delay the roughly 0.1-nsec delay from the rise-time degradation from the 2-pF input-gate capacitance, you arrive at a total propagation delay of about 1.5 nsec.

Any variation in the effective dielectric constant, due to the bit pattern on the bus, affects the speed of the signal on any data line and affects the propagation delay, which contributes to jitter.

When a microstrip trace is far away from its neighbors, a signal propagating down the line has some field lines going only through the laminate material and some field lines going partly through the air (Figure 3). These field lines that extend into the air are also called fringe field lines. In a 50 Ω microstrip in FR4, the fringe fields are responsible for about half the capacitance between the signal and the return path. The only accurate way of calculating the effective dielectric constant is to use a 2-D field solver. A field solver takes the fringe-field distributions into account and accurately predicts the effective dielectric constant and propagation speed.

Suppose that two adjacent signal lines lie close enough to the victim line that their fringe fields significantly interact. Significant fringe-field interaction generates crosstalk, and the acceptable crosstalk determines how close the traces can be to each

other. The presence of adjacent, aggressor, signal lines affects the fringe-field distribution of the victim line, depending on the bit pattern on the two aggressor lines.

When the aggressors carry the same bit, all three signal lines have the same voltage, and few fringe-field lines exist in the air above the traces, because they all have the same voltage, and no field lines exist between conductors with the same voltage. Most of the fringe-field lines from the victim line to the return plane are through the bulk material.




Figure 2 The Mentor Graphics HyperLynx simulator displays the received signal on the victim line for the three cases of the aggressors off, the aggressors switching with the same bit pattern, and the aggressors switching with the complementary bit pattern to that of the victim line.




Figure 3 Electric-field distribution around a 50 Ω microstrip in FR4 shows the fringe fields outside the width of the trace.




Figure 4 Electric-field distributions around three, closely coupled, 50 Ω microstrip traces in FR4 show the fringe fields changing with the bit patterns. Aggressor lines have the same bit pattern as the victim (a) and complementary bit patterns (b).

When the adjacent aggressor bit lines have the opposite bit from that of the victim line, a large voltage difference occurs between the victim line and the aggressor lines, and a large fringe field exists between the lines. The higher the fraction of field lines from the victim line to the return path in the air, compared with through the bulk material, the lower the effective dielectric constant (Figure 4). A change in the effective dielectric constants changes the delay of the signal on the victim line, depending on the bit pattern on the aggressors. A change in the bit pattern on the aggressor lines, compared with the victim line, changes the time delay on the victim line. When the aggressors have the same bit pattern as the victim, the effective dielectric constant is higher, the propagation speed is lower, and the delay is longer. When the aggressors have the opposite bit pattern from that of the victim line, the effective dielectric constant is lower, the propagation speed is higher, and the time delay on the victim line is shorter.

One way of reducing deterministic jitter from crosstalk is to reduce the overlap of the fringe fields by spreading the traces farther apart. For example, by increasing the spacing between all the lines to twice the line width, a common design rule for acceptable near-end crosstalk, you can reduce the deterministic jitter to less than 60 psec out of 1.5 nsec, or about 4% (Figure 5). Too close a spacing between victim and aggressor lines may cause a crosstalk problem in addition to deterministic jitter. To decrease both crosstalk and jitter, the spacing should be as large as possible, or at least twice the line width. However, this approach means using a lower density—and, perhaps, more expensive—board. You need an EDA tool that accurately predicts crosstalk and deterministic jitter to evaluate cost-ver-




Figure 5 Doubling the spacing between victim and aggressor lines reduces jitter.




Figure 6 In a tightly coupled stripline, all field lines see exactly the same dielectric constant, independent of their distribution.




Figure 7 The signal at the receiver of a victim line in a tightly coupled stripline shows that jitter has disappeared. Though significant crosstalk exists, no deterministic jitter occurs.

sus-performance trade-offs. A simulator that integrates a 2-D field solver should automatically take into account not only the noise effects, but also the deterministic jitter associated with crosstalk.

If variation in fringe fields in the air causes deterministic jitter from crosstalk, no deterministic jitter occurs in a structure with homogenous dielectric distributions, such as fully embedded microstrip or stripline. In a stripline geometry, the fringe fields between the

victim and the aggressor lines might closely interact, but whatever the field distributions, the field lines from the victim trace always see the same dielectric constant, that of the bulk laminate (Figure 6). For the extreme case in which spacing is equal to the line width in a 50Ω stripline, the effective dielectric constant is always the bulk dielectric constant, independent of the bit pattern on the aggressor lines. As the bit pattern changes, no jitter occurs on the victim line (Figure 7). Though a lot of crosstalk voltage occurs on the victim line, no jitter occurs in the arrival time of the signal on the victim line. □

AUTHORS' BIOGRAPHIES
Eric Bogatin is chief technology officer at Synergetix Inc (Kansas City, KS). He has a bachelor's degree in physics from the Massachusetts Institute of Technology (Cambridge) and a master's and a doctorate in physics from the University of Arizona (Tucson). His interests include writing popular-science and technology articles and writing science fiction. See his Web site at www.Bog.Eng.com.

Gene Garat is technical marketing engineer for Mentor Graphics' HyperLynx product line. He works with customers and signal-integrity professionals to solve board-level signal-integrity and EMI problems through the use of Mentor Graphics' simulation tools, such as HyperLynx. He has a bachelor's degree in mechanical engineering from California State Polytechnic University—Pomona. His spare-time interests include playing racquetball, skiing, motorcycle riding, and golfing.

TALK TO US
Post comments via TalkBack at the online version of this article at www.edn.com.