

**MOST TECHNICAL LITERATURE ON HIGH-SPEED DESIGN FOCUSES ON TERMINATION, RINGING, AND CROSSTALK. DESPITE SIGNAL INTEGRITY'S IMPORTANCE, INADEQUATE TIMING MARGINS CAUSE MANY MORE ERRORS IN TODAY'S 100-MHz DIGITAL DESIGNS.**

# Practical timing analysis for 100-MHz digital designs

**A**S INCREASING CHIP COMPLEXITY, high clock rates, and analog signal-integrity issues complicate digital design, time-to-market pressures continue to shorten development schedules. These factors present increasing challenges to digital-design engineers, who must spend more time understanding software and system-level issues and have less time for details such as timing analysis. Because you can't ignore board-level propagation delays in today's more-than-100-MHz designs, detailed timing analysis is essential.

This analysis develops general timing-margin equations, which include the effects of clock skew and propagation delay. A fast and effective procedure for analyzing timing margins uses a spreadsheet instead of traditional timing diagrams. This approach quickly identifies approximate timing margins early in the design. A little extra work can improve the margins, reduce the parts cost, or reduce engineering or pc-board-layout effort.

## DIGITAL-TIMING REVIEW

A good starting point is a review of the timing requirements for a typical synchronous digital connection, such as the one in **Figure 1**. In this circuit, the driver, IC<sub>1</sub>, produces valid data after a delay of T<sub>CO</sub> from the rising clock edge. The data must arrive at IC<sub>2</sub> in time to meet the receiver's input setup-and-hold-time requirements.

**Figure 2** illustrates the timing relationships among the driver clock-to-output delay, T<sub>CO</sub>; the receiver setup time, T<sub>SU</sub>; and the receiver hold time, T<sub>H</sub>. Given these timing parameters, which the IC manufacturer normally specifies, you can establish the following relationships for maximum clock frequency, setup margin, and hold margin:

$$\begin{aligned} \text{MINIMUM CLOCK PERIOD} &= T_{CO(\text{MAX})} + T_{SU} \\ \text{SETUP MARGIN} &= \text{CLOCK PERIOD} - T_{CO(\text{MAX})} - T_{SU} \end{aligned}$$

and


$$\text{HOLD MARGIN} = T_{CO(\text{MIN})} - T_H$$

where T<sub>CO(MIN)</sub> and T<sub>CO(MAX)</sub> are the manufacturer-specified minimum and maximum clock-to-output propagation-delay values for IC<sub>1</sub>'s output, and T<sub>SU</sub> and T<sub>H</sub> are the manufacturer-specified minimum setup-and-hold times for IC<sub>2</sub>'s input. Note that although you can always increase the clock period to increase the setup margin, increasing the clock period does not affect the hold margin, leading to the following important result: Input-hold margin is independent of clock frequency.

This result indicates the importance of verifying hold-time requirements during a project's design phase, because you can't eliminate hold-time violations by simply lowering the clock frequency.

## BOARD-LEVEL PROPAGATION DELAY

The simple example of **Figure 1** neglects data-propagation delay and clock skew between the transmitting and receiving ICs. In real digital systems in which clock signals reach frequencies of 100 MHz and higher, board-level-propagation or TOF (time-



**Figure 1**

**This basic synchronous digital circuit includes only a driver and a receiver, both of which depend upon a common clock.**

of-flight) delays are significant. In fact, at these high speeds, controlling clock skew is often necessary for producing adequate timing margins. To perform a realistic timing analysis, you must therefore modify the basic timing model to include clock skew and propagation delay. To model TOF and clock skew, the diagram in **Figure 3** introduces separate signals for clock and data at the driver and receiver. In this diagram, CLK<sub>1</sub> and DATA<sub>1</sub> represent signals at the driver, and CLK<sub>2</sub> and DATA<sub>2</sub> represent signals at the receiver.

**Figure 4** shows the timing diagram for this model, including the additional signals to represent clock and data timing at the receiver as well as at the transmitter. TOF<sub>D</sub> and T<sub>SKEWC</sub> represent propagation delay and clock skew in the equations for setup margin, hold margin, and minimum clock period:


$$\text{SETUP MARGIN} = \frac{\text{CLOCK PERIOD} - T_{\text{CO(MAX)}} - T_{\text{SU}} - \text{TOF}_D + T_{\text{SKEWC}}}{\text{TOF}_D + T_{\text{SKEWC}}} \quad (1)$$

$$\text{HOLD MARGIN} = \frac{T_{\text{CO(MIN)}} - T_{\text{H}} + \text{TOF}_D - T_{\text{SKEWC}}}{\text{TOF}_D - T_{\text{SKEWC}}} \quad \text{and} \quad (2)$$

$$\text{MINIMUM CLOCK PERIOD} = T_{\text{CO(MAX)}} - T_{\text{CO(MIN)}} + T_{\text{SU}} + T_{\text{H}}$$

where TOF<sub>D</sub> is the propagation delay, or TOF for the data path between IC<sub>1</sub> and IC<sub>2</sub>, T<sub>SKEWC</sub> is the clock skew from IC<sub>1</sub> and

**Figure 2**



In the circuit of **Figure 1**, the driver, IC<sub>1</sub>, produces valid data after a delay of T<sub>CO</sub> from the rising clock edge. The data must arrive at IC<sub>2</sub> in time to meet the receiver's input setup-and-hold-time requirements.

IC<sub>2</sub>, defined as positive when IC<sub>2</sub> is clocked later in time than IC<sub>1</sub>. You can establish a lower bound on the clock period by setting the setup-and-hold margins to 0 in **equations 1** and **2**, yielding:

$$\text{CLOCK PERIOD} = T_{\text{CO(MAX)}} + T_{\text{SU}} + \text{TOF}_D - T_{\text{SKEWC}} \quad (3)$$

$$\text{and} \quad T_{\text{CO(MIN)}} = T_{\text{H}} = \text{TOF}_D + T_{\text{SKEWC}} \quad (4)$$

Eliminating the term (T<sub>SKEWC</sub> - TOF<sub>D</sub>) by substitution in **equations 3** and **4** leaves you with the following relationship for minimum clock period:

$$\text{MINIMUM CLOCK PERIOD} = T_{\text{CO(MAX)}} - T_{\text{CO(MIN)}} + T_{\text{SU}} + T_{\text{H}}$$

This equation indicates that if you can control clock skew, the maximum achievable clock frequency for this circuit is independent of propagation delay,

leading to the following important result: For unidirectional signaling, the uncertainty in driver propagation delay and the receiver setup-and-hold times limit the clock frequency, but the clock frequency is independent of the overall propagation-delay magnitude.

**POSITIVE SKEW**

In this example, the clock skew is positive, which means that IC<sub>2</sub> is clocked after IC<sub>1</sub>. In general, T<sub>SKEWC</sub> can be either positive or negative. As you will see, the ability to adjust T<sub>SKEWC</sub> provides considerable flexibility in optimizing timing margins. With this timing model, you can introduce delays to the clock signals, the data signals, or both to improve timing margins, thereby increasing clock frequency.

You can make the several important observations in conjunction with this timing model. Delaying the clock signal by increasing the pc-board clock-trace

**TABLE 1—TIMING-MARGIN-ANALYSIS SPREADSHEET**

<b>Timing equations:</b>			
Setup margin =	clock period - T <sub>CO(MAX)</sub> - TOF - T <sub>SU</sub> + T <sub>SKEWC</sub>		
Hold margin =	T <sub>CO(MIN)</sub> + TOF - T <sub>H</sub> - T <sub>SKEWC</sub>		
T <sub>CO</sub> -	clock-to-output delay	T <sub>H</sub> -	input-hold time
T <sub>TOF</sub> -	output-signal time-of-flight	T <sub>SKEWC</sub> -	output-to-input clock skew
T <sub>SU</sub> -	input-setup time		
<b>Input Parameters:</b>			
<b>Clock frequency</b>			
<b>PC-board-trace delays</b>			
<b>IC<sub>1</sub> timing parameters</b>			
<b>IC<sub>2</sub> timing parameters</b>			
<b>IC<sub>1</sub> output to IC<sub>2</sub> input:</b>		<b>IC<sub>2</sub> output to IC<sub>1</sub> input:</b>	
Clock frequency (MHz):	100	Clock frequency (MHz):	100
Clock skew (nsec):	0.5	Clock skew (nsec):	-0.5
TOF (nsec):	0.25	TOF (nsec):	0.25
T <sub>CO</sub> output (nsec, minimum):	1.5	T <sub>CO</sub> output (nsec, minimum):	1
T <sub>CO</sub> output (nsec, maximum):	7	T <sub>CO</sub> output (nsec, maximum):	5.5
T <sub>SU</sub> input (nsec):	1.5	T <sub>SU</sub> input (nsec):	1
T <sub>H</sub> input (nsec):	0.8	T <sub>H</sub> input (nsec):	1
Clock period (nsec):	10	Clock period (nsec):	10
Setup margin (nsec):	1.75	Setup margin (nsec):	2.75
Hold margin (nsec):	0.45	Hold margin (nsec):	0.75

length relative to the data-trace length increases the setup margin at the expense of the hold margin, allowing the circuit to operate at a higher frequency. Delaying the data signal by increasing the pc-board data-trace length relative to the clock-trace length increases hold margin at the expense of setup margin. Additional hold margin may be necessary to accommodate devices with unusually long hold times. When  $T_{SKEWC}$  is equal to  $TOF_D$ , the clock skew exactly compensates for TOF, and equations 1 and 2 give the same results as the previous model, which ignores TOF. For this example, which models the special case of unidi-

rectional signaling, you can usually use a very fast clock, regardless of the propagation delay, as long as the data and clock delays are equal. In this case,  $(T_{CO(MAX)} - T_{CO(MIN)})$  rather than  $|T_{CO}|$  (the magnitude of  $T_{CO}$ ) limits the clock frequency.

As you will see, you can't draw the same conclusion in the case of bidirectional signaling. In many cases, an IC with less  $T_{CO}$  uncertainty is easier to use in a high-speed design than one with a smaller overall  $T_{CO}$ . For this reason, IC manufacturers should always specify accurate  $T_{CO(MIN)}$  values. As an example, consider a design that interfaces a device


with 0 or unspecified  $T_{CO(MIN)}$  to an SDRAM with a 1-nsec input-hold-time requirement. Whereas the SDRAM's  $T_{CO(MIN)}$  is unlikely to ever be less than 1 nsec, strictly meeting the 1-nsec SDRAM hold time requires that you add 1 nsec (about 6 in.) of board-trace delay, thereby sacrificing 1 nsec of setup margin.

**GENERALIZED TIMING MODEL**

Because most digital connections involve inputs as well as outputs on each IC, the next step is to extend the timing model to include bidirectional signaling. Figure 5 shows this straightforward extension, which adds a control-signal connection from IC<sub>2</sub> to IC<sub>1</sub>. In this figure, you can think of IC<sub>1</sub> as an SDRAM and IC<sub>2</sub> as a microprocessor. The datapath represents the SDRAM data bus, and the CTRL (control) path represents the SDRAM address and control signals.

Figure 6 shows the timing diagram for the bidirectional-signaling model, including signals CTL<sub>1</sub> and CTL<sub>2</sub>, which represent the control signals from IC<sub>2</sub> to IC<sub>1</sub>. The setup-and-hold analysis is also a straightforward extension of the unidirectional analysis:

**Figure 3**



To model TOF and clock skew, this diagram introduces separate signals for the clock and the data at the driver and receiver.

**TABLE 2—TIMING MARGINS FOR SELECTED TOF AND CLOCK-SKEW VALUES**

TOF	Clock skew	IC <sub>1</sub> setup	IC <sub>1</sub> hold	IC <sub>2</sub> setup	IC <sub>2</sub> hold
0.25	-0.5	3.75	-0.25	0.75	1.45
0.25	-0.4	3.65	-0.15	0.85	1.35
0.25	-0.3	3.55	-0.05	0.95	1.25
0.25	-0.2	3.45	0.05	1.05	1.15
0.25	-0.1	3.35	0.15	1.15	1.05
0.25	0	3.25	0.25	1.25	0.95
0.25	0.1	3.15	0.35	1.35	0.85
0.25	0.2	3.05	0.45	1.45	0.75
0.25	0.3	2.95	0.55	1.55	0.65
0.25	0.4	2.85	0.65	1.65	0.55
0.25	0.5	2.75	0.75	1.75	0.45
0.25	0.6	2.65	0.85	1.85	0.35
0.25	0.7	2.55	0.95	1.95	0.25
0.25	0.8	2.45	1.05	2.05	0.15
0.68	-0.5	3.32	0.18	0.32	1.88
0.68	-0.4	3.22	0.28	0.42	1.78
0.68	-0.3	3.12	0.38	0.52	1.68
0.68	-0.2	3.02	0.48	0.62	1.58
0.68	-0.1	2.92	0.58	0.72	1.48
0.68	0	2.82	0.68	0.82	1.38
0.68	0.1	2.72	0.78	0.92	1.28
0.68	0.2	2.62	0.88	1.02	1.18
0.68	0.3	2.52	0.98	1.12	1.08
0.68	0.4	2.42	1.08	1.22	0.98
0.68	0.5	2.32	1.18	1.32	0.88
0.68	0.6	2.22	1.28	1.42	0.78
0.68	0.7	2.12	1.38	1.52	0.68
0.68	0.8	2.02	1.48	1.62	0.58

$$IC_2 \text{ SETUP MARGIN} = \text{CLOCK PERIOD} - T_{COU1(MAX)} - T_{SUU2} - TOF + T_{SKEWC}$$

$$IC_2 \text{ HOLD MARGIN} = T_{COU1(MIN)} - T_{HU2} + TOF - T_{SKEWC}$$

$$IC_1 \text{ SETUP MARGIN} = \text{CLOCK PERIOD} - T_{COU2(MAX)} - T_{SUU1} - TOF - T_{SKEWC}$$

and

$$IC_1 \text{ HOLD MARGIN} = T_{COU2(MIN)} - T_{HU1} + TOF + T_{SKEWC}$$

The timing diagram in Figure 6 is too complicated for everyday use. Moreover, there is no longer a closed expression for minimum clock period because there are too many parameters. Fortunately, you can easily perform timing analysis for this model with a spreadsheet. Before beginning the spreadsheet analysis, observe the following about the bidirectional-signaling model:

- Delaying the clock signal to IC<sub>2</sub> by increasing the pc-board clock-trace

length increases data-setup margin at IC<sub>2</sub> and increases CTL-hold margin at IC<sub>1</sub> at the expense of decreased data-hold margin at IC<sub>2</sub> and decreased CTL-setup margin at IC<sub>1</sub>.

- Delaying the clock signal to IC<sub>1</sub> by increasing the pc-board clock-trace length increases CTL-setup margin at IC<sub>1</sub> and increases data-hold margin at IC<sub>2</sub> at the expense of decreased CTL-hold margin at IC<sub>1</sub> and decreased data-setup margin at IC<sub>2</sub>.

- Increasing the TOF for a signal increases its hold-time margin at the expense of setup margin.

In general, when both devices'  $T_{CO}$ ,  $T_{SU}$ , and  $T_H$  specifications are the same, setting  $T_{SKEWC}$  to 0 yields optimal timing margins. When the devices have different timing requirements, you can usually introduce clock skew to improve the timing margins. Note also that when the  $T_{CO(MIN)}$  and TOF are both small, meeting input-hold-time requirements can be difficult. This observation suggests that closer is not necessarily better from a timing standpoint. In any case, when timing margins are tight, you should always do a detailed analysis to verify timing and to establish detailed layout instructions.


### ESTIMATING TIME OF FLIGHT

To determine and optimize timing margins, you need a reliable method for estimating TOF. Initially apply this method to estimate preliminary timing margins from parts-placement data and then later to create pc-board-layout instructions. The two primary components in TOF are propagation delay in the pc-board traces and capacitive-loading delay. Trace delay varies only slightly because of variations in the pc-board dielectric and characteristic impedance.

You can usually get results accurate to within about 10% by multiplying 170 psec/in. (about 6 in./nsec) by the trace length in inches. Accuracy is important, because you use trace length to control clock skew and TOF delay when you optimize timing margins.


Capacitive loading, however, wreaks havoc in designs by causing waveform distortion and introducing delays that depend on bus topology and driver characteristics. Loading delays are less predictable than propagation delays because of variations in capacitance, but these de-

**Figure 4**



The timing diagram for Figure 3 includes additional signals to represent clock and data timing at the receiver as well as at the transmitter.

**Figure 5**



This straightforward extension of Figure 4 connects a control signal from IC<sub>2</sub> to IC<sub>1</sub> to enable bidirectional signaling.

lays always increase TOF, leading to another important point: TOF delays caused by capacitive loading always decrease setup margin and increase hold margin.

In TOF calculations, don't combine loading delays with propagation delays; estimate them separately and then provide sufficient setup margin in the final design to compensate for the delays. For complicated bus structures that have multiple loads and tight timing margins or when driving SDRAMs and DIMMs, a transmission-line simulator is the best way to ensure accurate TOF calculations. However, for simple bus structures with light loads, you can simply allow 50 psec of additional setup margin for each picofarad of load capacitance. The following analysis uses this rule of thumb and defers detailed TOF calculations.

### SPREADSHEET ANALYSIS

You now have a model that you can use to perform detailed timing analyses on most high-speed digital designs, but the associated timing diagrams are confusing, and creating them is time-consuming. The equations for timing margins can save a lot of time because you can use them in a spreadsheet analysis, such as the one in **Table 1**. You can download the Excel spreadsheet in the **table** from the

web version of this article at [www.edn.com](http://www.edn.com).

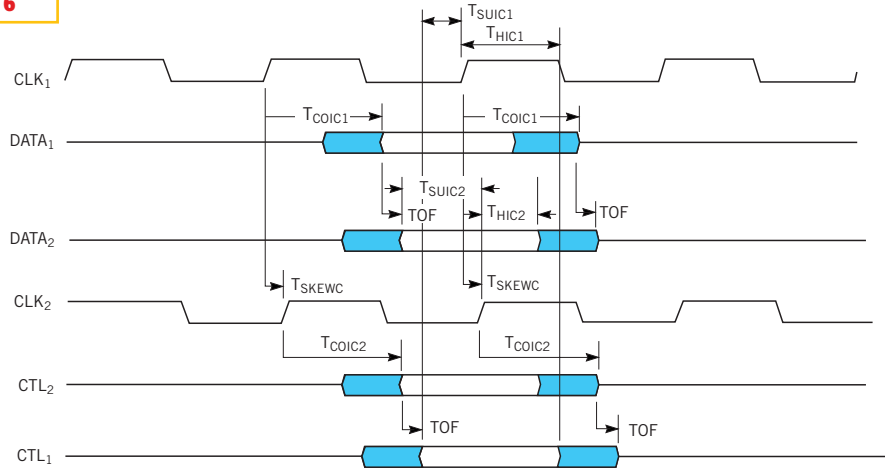
In this spreadsheet, the equations for setup-and-hold margin are located in the cells at the bottom of the page. These equations reference  $T_{SU}$ ,  $T_H$ ,  $T_{SKEWC}$ , TOF, and clock period from the user input area (colored cells). You can easily perform a what-if analysis by entering different values of TOF and clock skew and observing the resulting setup-and-hold margins.

Using this spreadsheet, enter the manufacturer's timing parameters, estimate TOF using preliminary placement data, and then adjust the clock skew to optimize the timing margins. You can summarize the spreadsheet design procedure as follows:

1. Enter the clock frequency and each device's manufacturer-specified values of  $T_{CO(MIN)}$ ,  $T_{CO(MAX)}$ ,  $T_{SU}$ , and  $T_H$ .
2. Using the preliminary parts-placement and pc-board-constraint data, estimate the trace length and TOF for the signals of interest. Enter the TOF into the spreadsheet.
3. Adjust the clock skew (and possibly the TOF) to optimize the timing margins for both devices.
4. Create pc-board-layout instructions to implement the required clock skew and TOF delays.

The values in this spreadsheet example come from a design that uses a Texas Instruments TMS320C6211 DSP with PC100 SDRAM operating at 100 MHz. For this design, the microprocessor is close to the SDRAM, so you can use a minimum value of 0.25 nsec for TOF, which corresponds to a minimum trace length of about 1.5 in. Because capacitive loading increases the hold margin and decreases the setup margin, the strategy is to maximize the setup margin and then estimate the additional capacitive loading and trace delay that the design supports. Ideally, you want enough margin to use an automatic router on the SDRAM-data and -control signals.

**Figure 6**



The timing diagram for the bidirectional-signaling model of Figure 5 includes signals CTL<sub>1</sub> and CTL<sub>2</sub>, which represent the control signals from IC<sub>2</sub> to IC<sub>1</sub>.

Suppose that instead of requiring the pc-board designer to match data- and control-signal lengths (a time-consuming task), you constrain these lengths to fall between 1.5 and 4 in. (You might expect this type of variation in trace length from the output of an automatic routing program.) To analyze the timing for these constraints, you can examine the worst-case margins for the longest and shortest traces with various values of clock skew. You can then choose a clock skew that gives acceptable overall margins.

Table 2 lists the setup-and-hold margins for several values of clock skew for TOF values of 0.25 nsec (1.5 in.) and 0.68

nsec (4 in.). Recall that you chose these two TOF values as minimum and maximum trace lengths for the SDRAM-data and -control signals. The table shows that IC<sub>2</sub>'s setup-and-hold margins are the limiting parameters. The highlighted rows show the setup-and-hold margins for the 1.5- and 4-in. traces with 0.5 nsec of positive clock skew. If you introduce 0.5 nsec of positive clock skew between IC<sub>1</sub> and IC<sub>2</sub>, you have worst-case margins of 1.32 and 0.45 nsec for setup-and-hold time, respectively.

**CLOCK SKEW AND TOF**

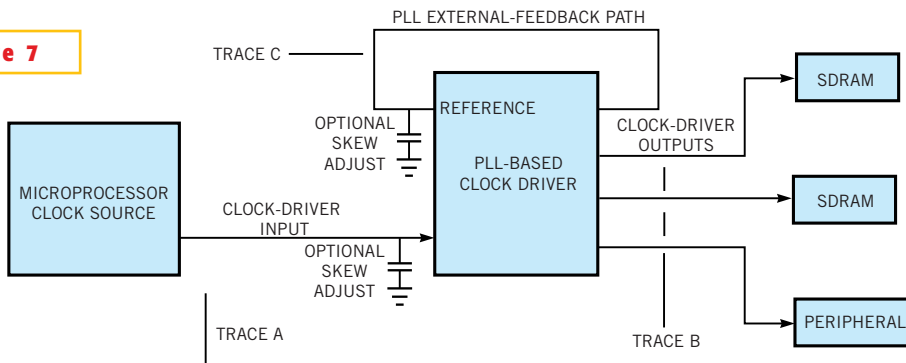
After determining timing margins and the required implementation delays,

your next task is to design the clock architecture and create pc-board-layout instructions. In general, the best way to make clock-skew and TOF adjustments is to control the length of the associated signal traces during pc-board layout. Controlling trace lengths is a painful task for pc-board designers. In addition, one of the primary design goals is to minimize pc-board-layout constraints. Careful planning during parts placement is particularly important to achieving this goal.

Figure 7 shows the clock architecture for a typical 100-MHz-microprocessor SDRAM interface. The microprocessor supplies a bus clock, which serves as the timing reference for all data and control signals. A PLL clock driver distributes this clock to the SDRAMs and other peripherals. In this architecture, the PLL contains an external feedback path, which is the key to controlling clock skew between the microprocessor and its peripherals.

The PLL clock driver's output clocks advance in time by an amount equal to the feedback-path trace delay. You can compensate for the propagation delay from

**Figure 7**



In this clock architecture for a typical 100-MHz-microprocessor SDRAM interface, the microprocessor supplies a bus clock, which serves as the timing reference for all data and control signals.

the microprocessor to the peripherals to achieve zero clock skew, or you can introduce positive or negative clock skew by controlling the length of traces A, B, and C as follows:

**Case 1:**  $C = A + B$ ; skew = 0. In Case 1, the feedback path exactly compensates for the clock delay from the microprocessor to the peripherals. The clock driver's outputs advance so that the clock signals arrive at the SDRAM at the same time they leave the microprocessor.

**Case 2:**  $C < A + B$ ; Skew =  $(A + B - C) \cdot \text{Delay\_per\_unit\_distance}$ . In Case 2, the feedback path undercompensates for the clock delay from the microprocessor to the peripherals, and the clock signals arrive at the SDRAM after they leave the microprocessor.

**Case 3:**  $C > A + B$ ; Skew =  $(A + B - C) \cdot \text{Delay\_per\_unit\_distance}$ . In Case 3, the feedback path overcompensates for the clock delay from the microprocessor to the peripherals, and the clock signals arrive at the SDRAM before they leave the microprocessor. Assuming that you use this clock architecture to implement the timing shown in **Table 1**, you can use a propagation delay of 170 psec/in. to produce the following layout instructions:

1. Match trace B to within 0.5 in. and route A, B, and C so that  $(A + B - C) = 2.75$  to 3.25 in.
2. Route the bus's data and control signals so that each is 1.5 to 4 in. long.

This implementation results in clock skew of 0.47 to 0.55 nsec and TOF of 0.25 to 0.68 nsec.

Assuming clock jitter of 200 psec, in the worst case, you still have an approximately  $1.75 - (0.68 - 0.25) - 0.2 = 1.1$ -nsec setup margin and  $0.45 - 0.2 = 0.25$ -nsec hold margin. Using a value of 50 psec/pF to estimate the maximum delay from capacitive loading allows enough setup margin for bus loads of 22 pF or less. If you expect bus loading to exceed this value, you must perform a more detailed analysis, increase layout constraints, or do both. For bus loads of less than 22 pF, manually routing the clock signals should allow you to automatically route the data and control signals.

#### HOW MUCH MARGIN?

To determine the amount of setup-and-hold margin required for reliable operation, you need to consider three pri-

mary sources of uncertainty: cycle-to-cycle clock jitter, trace-length delay, and capacitive-loading delay.

All clock sources introduce cycle-to-cycle clock jitter, which causes clock-period variation. You must allow sufficient margin to prevent this variation from causing setup-and-hold-time violations. Cycle-to-cycle clock jitter affects both setup and hold margins.

Signal-trace lengths are an important component of TOF delay. You can control this uncertainty when you create pc-board-layout instructions. You generally want to allow for as much layout uncertainty as possible to minimize layout constraints, particularly if you plan to use automatic routing.

The second component of TOF delay is capacitive-loading delay. You should provide sufficient setup margin to account for the worst-case capacitive-loading delay. For complicated or heavily loaded bus structures, you should use a simulator to estimate capacitive-loading delay.

If you identify and optimize margins early in the design, you can allocate them to minimize cost, time to market, or both. As development progresses, this optimization enables you to make intelligent trade-offs between parts cost, design time, and pc-board-layout time. □

#### AUTHOR'S BIOGRAPHY

*Bob Kirstein is president of Stratus Engineering (www.stratusengineering.com), San Diego, an electronic design-consulting company, where he has worked for the past*



*eight years. He holds a BS in computer engineering from the University of California—San Diego. His 18 years' experience includes designing analog and digital boards and*

*firmware and implementing logic in ASICs and FPGAs. Among the products he has developed are several high-performance, microprocessor-based networking switches, a robotics motion controller, an automotive collision-warning system, and an FFT-based vibration analyzer. He enjoys biking, swimming, tennis, and spending time with his wife, Janice, and eight-year-old daughter, Jacquelyn. You can reach him at rkirstein@stratusengineering.com.*