

Introduction

The output of an edge-triggered flipflop has two valid states: high and low. To ensure reliable operation, designs must meet the flipflop's timing requirements. The input to the flipflop must be stable for a minimum time before the clock edge (register setup time or t_{SU}) and a minimum time after the clock edge (register hold time or t_{H}). Specific values for t_{SU} and t_{H} are provided in each device family data sheet in this data book, or they can be determined using the MAX+PLUS® II Timing Analyzer.


In non-synchronous systems, if the asynchronous input signals violate a flipflop's timing requirements, the output of the flipflop can become metastable. Metastable outputs oscillate or hover between high and low states for a brief period of time, which can cause system failure. Therefore, you must analyze the metastability characteristics of a device to determine the reliability of a non-synchronous design. In synchronous systems, the input signals always meet the flipflop's timing requirements; therefore, metastability does not occur.

This application note describes metastability, how it is quantified, and how to minimize its effect. It also includes metastability data for the Altera® FLEX® 10K, FLEX 8000, FLEX 6000, MAX® 9000, and MAX 7000 devices that can be used to estimate a system's mean time between failures (MTBF) when using these devices to synchronize asynchronous data.

Metastability

Violating a flipflop's setup or hold time can cause its output to become metastable. When a flipflop is in a metastable ("in between") state, the output hovers at a voltage level between high and low, causing the output transition to be delayed beyond the specified clock-to-output delay (t_{CO}). The additional time beyond t_{CO} that a metastable output takes to resolve to a stable state is called the settling time (t_{MET}). Not every transition that violates the setup or hold times results in a metastable output. The likelihood that a flipflop enters a metastable state and the time required to return to a stable state varies depending on the process technology used to manufacture the device and on the ambient conditions. Generally, flipflops will quickly return to a stable state (see Figure 1).


Figure 1. Metastability Timing Parameters



The operation of a register is analogous to a ball rolling over a frictionless hill, as shown in Figure 2. Each side of the hill represents a stable (i.e., high or low) state, and the top of the hill represents the metastable state. When a flipflop’s data input complies with minimum setup (t_{SU}) and hold (t_H) times, the output passes from one stable state to another (i.e., from high to low or low to high) without an additional delay. Analogously, the ball travels over the hill within a specified time if given enough of a push.

However, when a flipflop’s data input violates the setup or hold time, the flipflop is marginally triggered, and the output may not immediately resolve to either of the two stable states within the specified time. This marginal triggering can cause the output to glitch or to remain temporarily at a metastable state between the high and low logic levels, taking longer to return to a stable state. Either condition increases the delay from the clock transition to a stable output.

Figure 2. Effects of Violating t_{SU} & t_H Requirements



Analyzing Metastability

Metastability does not necessarily cause unpredictable system performance. If the wait time is sufficient to allow the flipflop to settle to a stable state, metastability does not affect the system; the output of the flipflop can temporarily have an undefined value, provided that it returns to a known value before the signal is evaluated. Therefore, allowing additional time (t_{MET}) for the signal to settle to a known state prevents the propagation of an undefined value to the rest of the system.

The MTBF value quantitatively shows how metastability affects your design. The MTBF provides an estimate of the mean time between the probable occurrence of two successive metastable events. The MTBF of a synchronizing flipflop can be estimated with the following formula:

$$MTBF = \frac{e^{(C_2 \times t_{MET})}}{C_1 \times f_{CLOCK} \times f_{DATA}}$$


The f_{CLOCK} parameter refers to the system clock frequency while the f_{DATA} parameter refers to the data transfer frequency. The t_{MET} parameter is the additional time allowed by the system for the flipflop to settle to a stable state. The constants C_1 and C_2 vary according to the process technology used to manufacture the device. Therefore, different devices manufactured with the same process have similar values for C_1 and C_2 .

The constants C_1 and C_2 are determined by plotting the natural log of MTBF versus t_{MET} and performing a linear regression analysis on the data. The y- intercept and slope of the resulting line determine the values of C_1 and C_2 . The formulas for the constants C_1 and C_2 are shown below:

$$C_2 = \frac{\Delta \ln(MTBF)}{\Delta t_{MET}}$$

$$C_1 = \frac{e^{(C_2 \times t_{MET})}}{MTBF \times f_{CLOCK} \times f_{DATA}}$$

Figure 3 shows the relationship between the MTBF and t_{MET} and shows how changing C_1 , C_2 , and the system frequency affects this relationship.

Figure 3. MTBF vs. t_{MET} 

The constant C_1 scales the MTBF equation linearly, shifting the entire curve up or down. The constant C_2 affects the slope of the MTBF vs. t_{MET} curve. Increasing the clock frequency shifts the entire curve to the right, lowering the MTBF value for a given settling time.

Once the values for C_1 and C_2 are determined for a particular device, you can use the MTBF equation shown on page 847 to calculate the MTBF of a system with a given settling time (t_{MET}). The t_{MET} delay is the additional time required for the flipflop to resolve to a legal state, i.e. the difference between the minimum system clock period and the actual clock period. You can also use the metastability equation to determine the t_{MET} delay required for a given MTBF value, as shown below:


$$t_{MET} = \frac{\ln(\text{MTBF} \times f_{CLOCK} \times f_{DATA} \times C_1)}{C_2}$$

Test Circuitry

Figure 4 shows the test circuit used to determine the metastability characteristics of Altera devices. In this figure, one flipflop has asynchronous clock and data inputs. The logic that generates the metastable event and the logic that detects it are both located in the device under test (DUT). The output of the synchronizing flipflop is fed directly to one of the resolving flipflops and through an inverter to the other resolving flipflop. The outputs of the resolving flipflops feed an XNOR gate that is at a high logic level when the values of the outputs (the signal and its complement) are the same. If the resolving flipflops detect that the signal and its complement have the same logical value, a metastable event has occurred and the counter is incremented.

Figure 4. Metastability Test Circuit

All logic is part of the device under test.



Because the resolving flipflops are clocked by the falling clock edge, the required settling time can be controlled by changing the clock high time (Δt). The settling time t_{MET} can be determined with the equation below. The t_{ACNT} delay is the minimum clock period, which is equal to the minimum delay from the clock edge to the output of the synchronizing flipflop, plus the delay from the output of the synchronizing flipflop to the input of the resolving flipflop, plus the setup time of the resolving flipflop. The t_{MET} parameter is the minimum time allowed under normal operation of the circuit:

$$t_{MET} = \Delta t - t_{ACNT}$$

Metastability Characteristics of Altera Devices

Figure 5 shows the metastability characteristics of FLEX 10K, FLEX 8000, FLEX 6000, MAX 9000, and MAX 7000 devices. For all devices, f_{DATA} is 1 MHz and f_{CLOCK} is 10 MHz.

Figure 5. Metastability Characteristics of Altera Devices




Table 1 summarizes the values of C_1 and C_2 for Altera's FLEX 10K, FLEX 8000, FLEX 6000, MAX 9000, and MAX 7000 devices.

Applying the Metastability Equation

Device	C ₁	C ₂
FLEX 10K	1.01×10^{-13}	1.268×10^{10}
FLEX 8000	1.01×10^{-13}	1.268×10^{10}
FLEX 6000	1.01×10^{-13}	1.268×10^{10}
MAX 9000	2.98×10^{-17}	5.023×10^9
MAX 7000	2.98×10^{-17}	5.023×10^9

You can use the C₁ and C₂ values listed in Table 1 to calculate the MTBF for a specific settling time, or you can calculate the minimum settling time for a specific MTBF. For example, the equation below calculates the t_{MET} needed to ensure a MTBF of one year (approximately 3×10^7 seconds) for an EPF8452A device with a data frequency of 2 MHz and a clock frequency of 10 MHz.

$$t_{MET} = \frac{\ln(3 \times 10^7) + \ln[(10 \times 10^6)(2 \times 10^6)(1.01 \times 10^{-13})]}{1.268 \times 10^{10}} = 1.41 \text{ ns}$$

When a MTBF of one year is required, 1.41 ns should be added to the clock-to-output delay (t_{CO}) of the synchronizing flipflop when performing timing analysis. Due to the logarithmic relationship between the MTBF and t_{MET} , small changes in t_{MET} dramatically affect the MTBF. If the required MTBF increases from one year to 10 years in the example shown above, the t_{MET} delay increases to only 1.59 ns.

Figures 6 and 7 show the t_{MET} delay required for FLEX 10K, FLEX 8000, FLEX 6000, MAX 9000, and MAX 7000 devices when f_{DATA} is one half of f_{CLOCK} . Because the MTBF is inversely proportional to the value ($f_{CLOCK} \times f_{DATA}$), these figures can be used to find the MTBF for many designs. Metastability is probabilistic, and MTBF values are mean values calculated for a limited sample of devices and should only be used to estimate t_{MET} delays.

Figure 6. FLEX 10K, FLEX 8000 & FLEX 6000 MTBF Values





Figure 7. MAX 9000 & MAX 7000 MTBF Values



12

Device
Operation


Avoiding Metastability

Several techniques can be used to reduce metastability in a system. If an asynchronous signal is fed to several flipflops, the probability that a metastable event will occur greatly increases because there are more flipflops that could become metastable. In this case, you can avoid metastability by using the output of the synchronizing flipflop throughout the system rather than the asynchronous signal.

You can also avoid the negative effects of metastability by adding the t_{MET} calculated for a specific MTBF to the worst-case timing delay calculations, giving the output of the synchronizing flipflops time to settle. Faster devices provide faster t_{CO} and t_{SU} times, which provide additional time for the t_{MET} delay without sacrificing overall system speed.


To reduce the effects of metastability, designers most commonly use a multiple-stage synchronizer in which two or more flipflops are cascaded to form a synchronization circuit (see Figure 8). If the synchronizing flipflop produces a metastable output, the metastable signal may resolve before it is clocked by the second flipflop. This method does not guarantee that the second flipflop will not clock an undefined value, but it dramatically increases the probability that the data will go to a valid state before it reaches the rest of the circuit.

Figure 8. Multiple-Stage Synchronizer



One drawback of the multiple-stage synchronizer is that it takes longer for the system to respond to an asynchronous input. A solution to this problem is to use the output of a ClockBoost™ clock doubler to clock the two synchronizing flipflops. See Figure 9. This approach allows the system to respond to an asynchronous input within one system clock cycle, while still improving MTBF. Although the ClockBoost clock could decrease the MTBF, this effect is more than offset by the two synchronizing flipflops.

Figure 9. Multiple-Stage Synchronizer with ClockBoost Circuitry



Conclusion

Metastability is a phenomenon that only affects flipflops used to synchronize data from asynchronous systems. The metastability characteristics for a particular device depend on the process technology used to manufacture the device and on ambient conditions. Altera devices have very good metastability characteristics; you only need to add a small t_{MET} delay to the t_{CO} delay to achieve a high MTBF value.