



Metastable Recovery

XAPP 094 November 24, 1997 (Version 2.1)

Application Note By Peter Alfke and Brian Philofsky

Introduction

Whenever a clocked flip-flop synchronizes an asynchronous input, there is a small probability that the flip-flop output will exhibit an unpredictable delay. This happens when the input transition not only violates the setup and hold-time specifications, but actually occurs within the tiny timing window where the flip-flop accepts the new input. Under these circumstances, the flip-flop can enter a symmetrically balanced transitory state, called metastable (meta = between).

While the slightest deviation from perfect balance will cause the output to revert to one of its two stable states, the delay in doing so depends not only on the gain-bandwidth product of the circuit, but also on how perfect the balance is, and on the noise level within the circuit; the delay can, therefore, only be described in statistical terms.

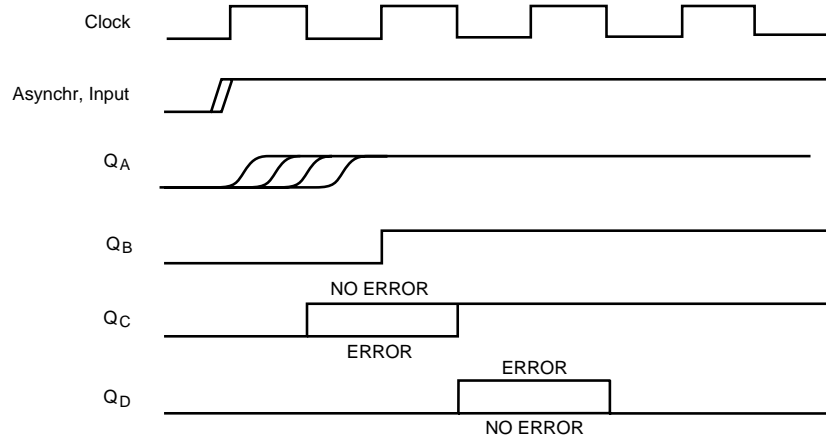

The problem for the system designer is not the illegal logic level in the balanced state (it's easy enough to translate that to either a 0 or a 1), but the unpredictable timing of the final change to a valid logic state. If the metastable flip-flop drives two destinations with differing path delays, one des-

tinuation might clock in the final data state while the other does not.

With the help of a self-contained circuit, Xilinx evaluated the XC4000 and XC3000-series flip-flops. The result of this evaluation shows the Xilinx flip-flop to be superior in metastable performance to many popular MSI and PLD devices.

Since metastability can only be measured statistically, this data was obtained by configuring several different Xilinx FPGAs with a detector circuit shown in [Figure 1](#). The flip-flop under test receives the asynchronous ~1-MHz signal on its D input, and is clocked by a much higher manually adjustable frequency. The output QA feeds two flip-flops in parallel, one (QB) being clocked by the same clock edge, the other (QC) being clocked by the opposite clock edge. When clocked at a low frequency, each input change gets captured by the rising clock edge and appears first on QA, then, after the falling clock edge, on QC, and finally, after the subsequent rising clock edge, on QB.

If a metastable event in the first flip-flop increases the settling time on QA so much that QC misses the change, but QB still captures it on the next rising clock edge, this error



X5985

Figure 1: Test Circuit and Timing Diagram

can be detected by feeding the XOR of QB and QC into a falling-edge triggered flip-flop. Its output (QD) is normally Low, but goes High for one clock period each time the asynchronous input transition caused such a metastable delay in QA. The frequency of metastable events can be observed with a 16-bit counter driven by QD.

By changing the clock frequency, and thus the clock half-period, the amount of acceptable metastable delay on the QA output can be varied, and the resulting frequency of metastable events can be observed on the counter outputs.

As expected, no metastable events were observed at clock rates below 70 MHz for the XC4005-6, or below 100 MHz for the XC4005E-3, since a half clock period at those frequencies is adequate for almost any metastability-resolution delay. Increasing the clock rate slightly brought a sudden burst of metastable events. Careful adjustment of the clock frequency gave repeatable, reliable measurements.

Metastability Measurements

The circuit of [Figure 1](#) was implemented in five different Xilinx devices: two cutting-edge devices using 0.5 micron, 3-layer-metal technology, the XC4005E-3 and the XC3142A-09, one device, the XC5206 using 0.6 micron, 3-layer-metal, and, for comparison purposes, also in two older-technology devices, the XC4005-6 and the XC3042-70.

In each device two different implementations put QA, the flip-flop under test, into an IOB and a CLB (Except for the XC5200 family which has no flip-flops in the IOB). The XC4000-series devices showed little difference between IOB and CLB behavior, but in the XC3000-series devices, the IOB flip-flops showed dramatically better metastable performance than the CLB flip-flops. This difference can be traced to subtle differences in circuit design and layout, and will guide us to further improvements in metastable performance in future designs.

Metastable measurement results are listed in [Table 1](#), and are plotted in [Figure 2](#). The results for XC4000E-3 (IOB and CLB) and for XC3100A-09 IOB flip-flops are outstanding, far superior to most metastable data published anywhere else. When granted 2 or 3 ns of extra settling delay, these devices come close to eliminating the problems caused by metastability, since their MTBF exceeds millions of years.

The older-technology devices are obviously less impressive, but they still show acceptable performance, especially in the IOB input flip-flops that are normally used to synchronize asynchronous input signals.

[Table 1](#) lists the experimental results from which the exponential factor K2 was derived. The clock frequency was adjusted manually, while observing the LSB and the MSB of the 16-bit error counter. F_L is the clock frequency that generated a ~1 Hz error rate, F_H generated a ~64,000 Hz error rate.

K2 is derived by dividing $\ln 64,000$ by the half-period difference.

Table 1: Metastable Measurement Results

Device	F_L (MHz)	F_H (MHz)	Half-period Difference (ns)	K2 (1/ ns)
XC4005E-3 IOB	111.5	131.6	0.685	16.1
XC4005E-3 CLB	109.0	124.4	0.568	19.4
XC4005-6 IOB	73.0	90.0	1.294	8.5
XC4005-6 CLB	71.2	88.8	1.392	7.9
XC5206-5 CLB	70.8	79.8	0.80	13.7
XC3142A-09 IOB	152.2	206.6	0.87	12.7
XC3142A-09 CLB	107.4	211.3	2.29	4.8
XC3042-70 IOB	46.6	61.5	2.60	4.2
XC3042-70 CLB	41.9	64.8	4.22	2.6

Metastability Calculations

The Mean Time Between Failures (MTBF) can only be defined statistically. It is inversely proportional to the product of the two frequencies involved, the clock frequency and the average frequency of the asynchronous data changes, provided that these two frequencies are independent and have no correlation.

The generally accepted equation for MTBF is

$$MTBF = \frac{e^{K2 * t}}{F1 * F2 * K1}$$

K1 represents the metastability-catching set-up time window, which describes the likelihood of going metastable.

K2 is an exponent that describes the speed with which the metastable condition is being resolved. K2 is an indication of the gain-bandwidth product in the feedback path of the master latch of the master-slave flip-flop. A small increase in K2 results in an enormous improvement in MTBF.

With $F1 = 1$ MHz, $F2 = 10$ MHz and $K1 = 0.1$ ns = 10^{-10} s:

$$MTBF \text{ (in seconds)} = 10^{-3} * e^{K2 * t}$$

Experimentally derived (see [Table 1](#)):

K2 = 16.1 per ns, for the XC4005E-3 IOB flip-flops

K2 = 19.4 per ns, for the XC4005E-3 CLB flip-flops

K2 = 8.5 per ns, for the XC4005-6 IOB flip-flops

K2 = 7.9 per ns, for the XC4005-6 CLB flip-flops

K2 = 13.7 per ns, for the XC5206-5 CLB flip-flops

K2 = 12.7 per ns, for the XC3142A-09 IOB flip-flops

K2 = 4.8 per ns, for the XC3142A-09 CLB flip-flops

K2 = 4.2 per ns, for the XC3042-70 IOB flip-flops

K2 = 2.6 per ns, for the XC3042-70 CLB flip-flops




Figure 2: Mean Time Between Failure for various IOB and CLB flip-flop outputs when synchronizing a ~1 MHz asynchronous input with a 10 MHz clock.

For other operating conditions, divide MTBF by the product of the two frequencies. For a ~10 MHz asynchronous input synchronized by a 40 MHz clock, the MTBF is 40 times

shorter than plotted; for a ~50 kHz signal synchronized by a 1 MHz clock, the MTBF is 200 times longer than plotted here.