XAPP094 (v3.0) February 10, 2005

Metastable Recovery in Virtex-II Pro FPGAs

Author: Peter Alfke

Summary

This application note describes the probability of a metastable event occuring in a Xilinx Virtex™-II Pro FPGA. The test circuit measures the Mean Time Between Failure (MTBF) of these metastable events. When extrapolated and compared to those of the XC4005E, these results show overall improvements in Virtex-II Pro flip-flops in terms of metastable recovery.

Introduction

Whenever a clocked flip-flop synchronizes an asynchronous input, a small probability exists that the flip-flop output will exhibit an unpredictable delay. This happens not only when the input transition violates setup and hold time specifications, but also when the transition actually occurs within the tiny timing window during which the flip-flop accepts the new input. Under these circumstances, the flip-flop can enter a symmetrically balanced transitory, or metastable, state.

While the slightest deviation from perfect balance causes 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 flip-flop's master latch, but also on the noise level within the circuit. Therefore, the delay can be described only in probabilistic terms.

The problem for system designers is not the undocumented logic level in the balanced state (it is 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 destination might clock in the final data state, while the other clocks in the intermediate metastable state.

With the help of a self-contained circuit, Virtex-II Pro CLB and IOB flip-flops were evaluated. The results show the Virtex-II Pro flip-flop to be superior in metastable recovery to that of any previously documented standard device.

Since metastability can be measured only statistically, this data was obtained by configuring the FPGA with a detector circuit, as shown in Figure 1. The flip-flop under test receives the asynchronous ~50 MHz signal on its D input and is clocked by a much higher 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 is 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 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 asserts High for one clock period each time the asynchronous input transition causes 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 are observed at clock rates below 300 MHz, since a half clock period at those frequencies is adequate for almost any metastability resolution delay. Increasing the clock rate slightly brings a sudden burst of metastable events. Careful adjustment of the clock frequency gives repeatable, reliable measurements. The measured MTBF ranged between one millisecond and one minute, and was extrapolated to beyond a million years.

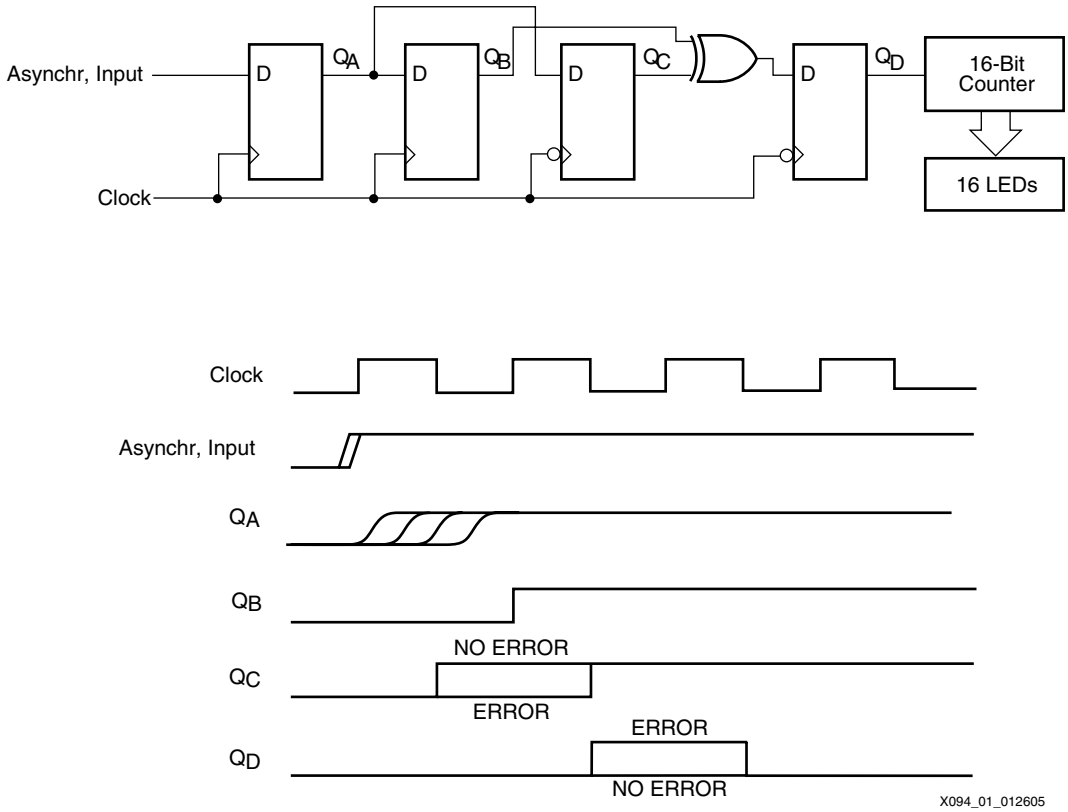

Figure 1: **Test Circuit and Timing Diagram**

Why Is There No Simpler Way to Measure the Metastability Capture Window?

The tests indicate that a metastable delay increment of 500 ps occurs about once per second. During that time, the 300 MHz clock makes 100 million attempts to synchronize the rising and falling data edge. Only one of these results in a 500 ps delay increment. Since data and clock are asynchronous, the capture window for this short metastable event is 3 ns ÷ 100 million = 0.03 femtoseconds. No practical way exists to zero in on such a narrow window, and random testing is the only successful methodology.

Metastability Calculations

The MTBF can be defined only 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

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

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. Some researchers list 1/K2 as a time constant, tau (τ).

Metastability Measurements

The circuit in Figure 1 was implemented in an XC2VP4 device using 0.13 micron, 9-layer-metal technology. Two different implementations put QA, the flip-flop under test, into a CLB and into an IOB.

Metastable measurement results are listed in Table 1 and are plotted in Figure 2. The time plotted on the horizontal axis includes the clock-to-out delay of QA, plus a short interconnect delay, plus the setup time at the input of the QC flip-flop. The extrapolated results are far superior to any metastable data published previously. When granted 2 ns of extra settling delay, the problems caused by metastability are almost eliminated, as their MTBF exceeds millions of years.

Table 1: **Virtex-II Pro Metastability Measurements and Calculations (1997 data for XC4005E added for comparison)**

Device	XC2VP4			XC2VP4			XC4005E –3
V_{CC} (V)	1.5	1.35	1.65	1.5	1.35	1.65	5.0
Flip-Flop Type	CLB	CLB	CLB	IOB	IOB	IOB	CLB
Low Frequency (MHz)	300	310	390	310	300	420	109
Half Period (ps)	1667	1613	1283	1613	1667	1190	4587
MTBF1 (ms)	60,000	20,000	60,000	30,000	30,000	30,000	1,000
High Frequency (MHz)	390	420	490	420	430	500	124.4
Half Period (ps)	1282	1190	1020	1190	1163	1000	4019
MTBF2 (ms)	1.69	1.046	5.16	0.987	1.84	6.96	0.016
Half Period Difference (ps)	385	423	262	423	504	190	568
Ln (MTBF1 / MTBF2)	10.478	9.86	9.36	10.322	9.70	8.37	11.09
K2 (per ns)	27.2	23.3	35.7	24.4	19.24	44.05	19.52
1 / K2 = tau (ps)	36.8	42.9	28.0	41.0	52.0	22.7	51.2
MTBF multiply / 100 (ps)	**15.2**	**10.3**	**35.6**	**11.5**	**6.85**	**81.8**	**7.04**

Table 1 lists the experimental results from which the exponential factor K2 was derived. The clock frequency was adjusted manually, while counting errors. Measurements were taken at room temperature, but testing at V_{CC} extremes gives an indication of performance at higher and lower temperature.

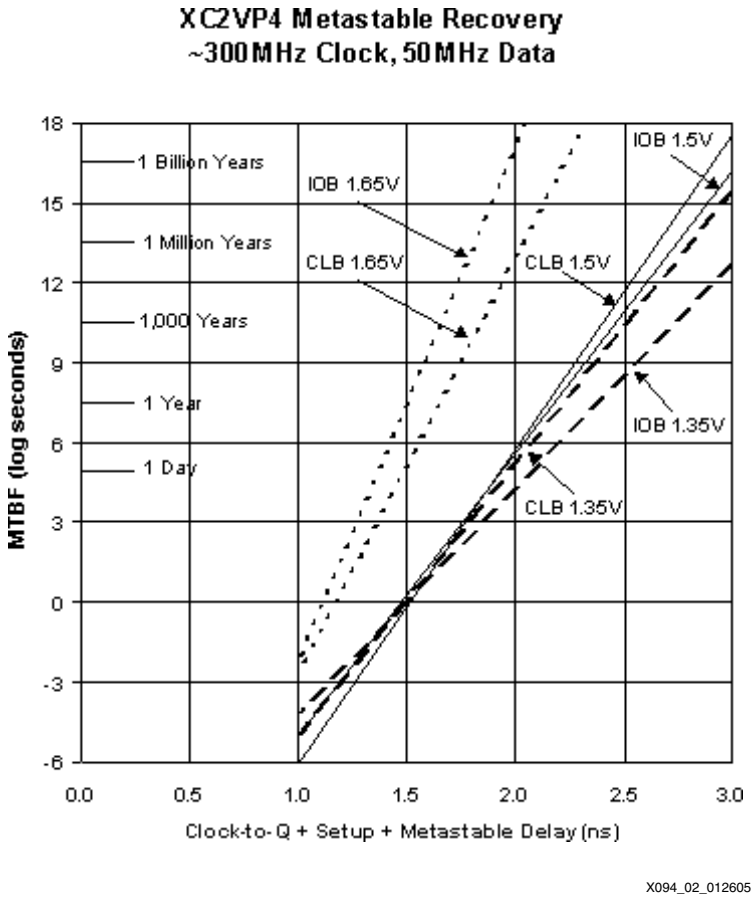

Figure 2: **XC2VP4 Metastable Recovery**

Xilinx began measuring and publishing metastable results as early as 1988. The earliest data was published in 1989, and was expanded in 1997 to cover newer devices. To illustrate the improvement in the Virtex-II Pro flip-flop, Figure 3 combines the new results with the 1997 results, all normalized to a 1 MHz asynchronous data rate and a 100 MHz clock rate. (Recalculating MTBF for lower data and clock frequencies increases MTBF but does not affect the slope on the logarithmic scale. It effectively moves the lines to the left to compensate for the 150 times smaller product of the two frequencies, causing an MTBF that is 150 times longer.)

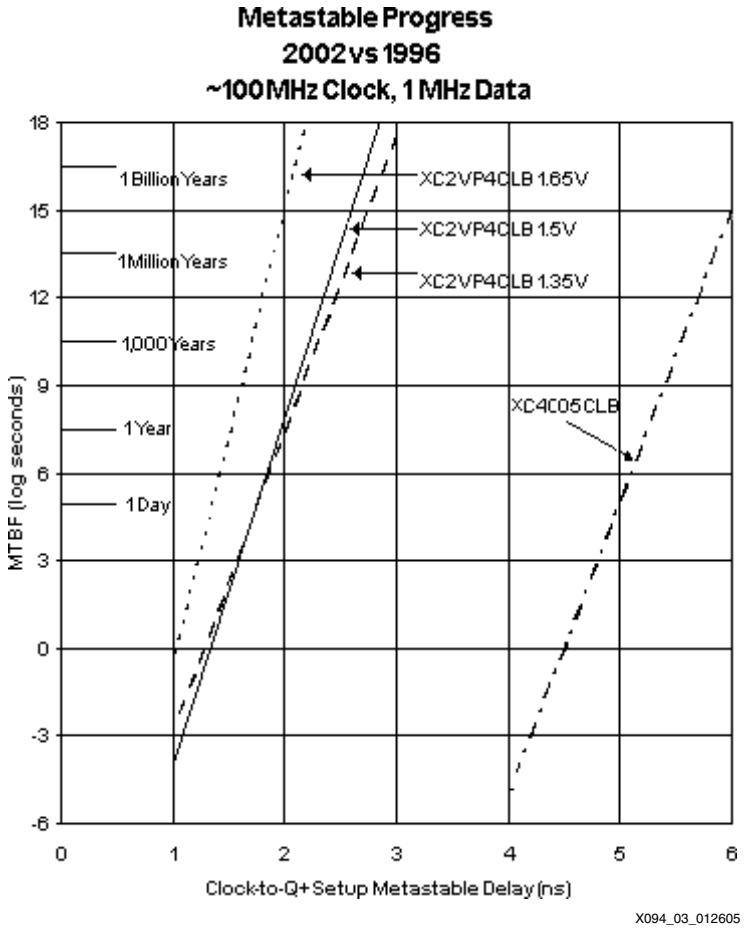

Figure 3: **Metastable Progress**

For other operating conditions, divide MTBF by the product of the two frequencies. For an ~10 MHz asynchronous input synchronized by a 200 MHz clock, the MTBF is 20 times shorter than plotted; for a ~50 kHz signal synchronized by a 1 MHz clock, the MTBF is 2000 times longer than plotted in Figure 3.

Conclusion

Asynchronous inputs can result in unpredictable delays at the synchronizer flip-flop output. Modern CMOS circuits are so fast that this metastable delay can safely be ignored for clock rates below 200 MHz. The user must, however, minimize the routing delay between the synchronizer and the next-level flip-flop, so as not to squander the advantage gained by the fast settling time.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
11/24/97	2.1	Initial Xilinx release.
02/10/05	3.0	Extensive rewrite.