

Metastability challenges for 65nm and beyond; Simulation and measurements

Salomon Beer¹, Ran Ginosar¹, Jerome Cox², Tom Chaney² and David M. Zar²

¹EE Dept, Technion—Israel Institute of Technology, Haifa, Israel,

²Blendics Inc. St. Louis, Missouri,
{sbeer@tx, ran@ee}.technion.ac.il

Abstract— Recent synchronizer metastability measurements indicate degradation of MTBF with technology scaling, calling for measurement and calibration circuits in 65nm and below. Degradation of parameters can be even worse if the system is operated at extreme supply voltages and temperature conditions. In this work we study the behavior of synchronizers in a broad range of supply voltage and temperature corners. A digital on-chip measurement system is presented that helps to characterize synchronizers in future technologies and a new calibrating system is shown that accounts for changes in delay values due to supply voltage and temperature changes. We present a detailed comparison of measurements and simulations for a fabricated 65nm bulk CMOS circuit and discuss implications of the measurements for synchronization systems in 65nm and beyond. We propose an adaptive self-calibrating synchronizer to account for supply voltage, temperature, global process variations and DVFS.

INTRODUCTION

Multiple-clock System on Chip (SoC) designs requires synchronization when transferring signals and data among clock domains and when receiving asynchronous inputs. Such synchronizations are often susceptible to metastability effects [1], which may propagate into the receiving circuit and may cause malfunctions. To mitigate the effects associated with metastability, latches and flip flops are often used to synchronize the data [2], in schemes such as pipelined flip flops. There is, however, a certain probability that the circuit will not resolve its metastable state correctly within the allowed time. To enable assessing the risk, and to enable the design of reliable synchronizers, models describing the failure mechanisms for latches and flip flops have been developed [1][2][3]. Most models express the risk of not resolving metastability in terms of the mean time between failures (MTBF) of the circuit (1)

$$MTBF = \frac{e^{S/\tau}}{T_w \times F_C \times F_D} \quad (1)$$

where F_C and F_D are the clock and data transition frequencies, respectively, S is a pre-determined time allowed for metastability resolution, τ is the resolution time constant, and T_w is a parameter describing a vulnerable time window which is experimentally determined.

Desirable values of MTBF depend on the application and range from several years upwards. The parameter τ is

predominant in synchronizer characterization since its effect on MTBF is exponential (1). Evidently, as technology scales, F_C and F_D increase and to maintain high MTBF τ must decrease as well. In the past, τ was believed to improve with technology scaling [4]. However, recent measurements [5][6] indicate that scaling trends of synchronizers should be re-considered: The need to obtain full characterization may be imperative in technologies beyond 45nm. Because of this, simulation methods that can reliably predict measurements are growing in importance. The physical testing of a system at the extremes of its operating conditions is the industry standard method for validating proper operation. For digital systems that are at risk of a metastability failure, the risk of entering metastability may be higher in extreme PVT corners. Synchronizer parameters τ and T_w in (1), can be seen as dependent on supply voltage and temperature; $\tau(V_{DD}, T), T_w(V_{DD}, T)$. As a result, careful simulation of the system design at several points throughout its operating region combined with verification is proposed as a dependable approach to the detection of potential metastability failures. This paper describes a fully digital on-chip characterization circuit to measure synchronization performance that is an improvement of the circuit shown in [7]. We show a new calibration circuit to compare measurements and simulations over a wide range of supply voltage and temperature corners for a standard library flip flop. We do a careful analysis of measurements and simulations and show good correlation between them. A 100x73 μm^2 on-chip digital measurement circuit was fabricated in a low power 65nm bulk CMOS process and was tested to prove its utility for measurements and for validating simulation results.

1 PROPOSED CHARACTERIZATION METHOD

The measurement consists of sampling the output of the flip-flop-under-test twice: first (X in Figure 1) by the clock delayed by a factor DL (by means of a delay line) and second (Y) by the negative edge of the clock. The two samples are compared by a XOR gate Figure 1(a). A metastability event that resolves during the time window between the delayed and negative edges of the clock (namely an event that did not resolve within the allotted time DL) increments the counter (the number of events expected to resolve after the negative edge of the clock is

negligible, since a low frequency clock is used). The measurement continues for time period T and thus MTBF is T divided by the counter value. The data and clock are produced by two uncorrelated sources.




Figure 1. (a) Metastability measurement circuit (b) Waveforms explaining the operation of measuring method.

The entire measurement is repeated for different DL delays, obtaining a set of (DL, MTBF) readings. This set is used to compute τ and T_W . The entire circuit of Figure 1(a) is embedded on-chip. Since this measurement procedure tests for metastability at the falling edge of the clock, it examines the recovery of only the master latch and produces τ_M (τ of master latch). If $\tau_M \neq \tau_S$, the method fails to estimate τ_S (τ of slave latch) and, consequently, may give erroneous values of MTBF.

For each voltage it is possible to calculate MTBF as

$$MTBF = \frac{T}{\#(>DL)} \quad (2)$$

where $\#(>DL)$ denotes the number of resolution events that took longer than DL to resolve and T is the time period that the counter was enabled. From, (1),(2) and $S=DL$ we can derive

$$\ln(\#(>DL)) = -\frac{1}{\tau}(DL) + \ln(F_C F_D T_W T) \quad (3)$$

Thus, performing a least square fit to a line equation of DL and $\ln(\#(>DL))$ produces a line with a slope (Sl) and an intercept (Y). From (3), $\tau = -1/Sl$ and $T_W = e^{(Y)}/F_C F_D T$.

2 DELAY LINE CALIBRATION

One important step for obtaining reliable results is the delay line calibration. Both τ and T_W rely on accurate delay line values (3). A schematic of the implemented delay line is shown in Figure 2. The calibration consists of measuring the exact delay generated by the delay line in the fabricated part. For this purpose the delay line is closed into a ring oscillator (RO) and the value of the frequency generated is measured

after clock division. The period generated by the ring oscillator is given by (4), where T_{buffer} accounts for the delay of each buffer, T_{inv} is the delay of the RO inverter, T_{muxes} the delay of the two multiplexers and $T_{parasitics}$ are wiring and other parasitic delays. N is the number of buffers in the delay line, and is selected by the output multiplexer digitally.

$$T_{RO} = 2[N \cdot T_{buffer} + T_{inv} + T_{muxes} + T_{parasitics}] \quad (4)$$




Figure 2. Delay line circuit schematics and RO inverter.

To obtain the exact delay between input and output, the delay of the inverter in (4) should be estimated. Since each buffer is composed of two inverters, a good estimation of T_{inv} is obtained by calculating the slope of the plot relating N and T_{RO} using (4). Such a plot is shown in Figure 3.




Figure 3. Delay line delay versus number of buffers for different supply voltages at T=20C. T_{inv} estimated for each curve is shown.

Since τ measurement is obtained from the slope in (3) and is a differential value, it is not affected by exact values of delay of the delay line. However, T_W is based on Y (intercept), which is very sensitive to exact values of delay, then T_W values are more error prone than τ values. The error introduced in the exact delay measurement of the delay line (wiring delays, process variations, etc.) is estimated in the following section.

3 MEASUREMENT CIRCUIT

Figure 4 shows a block diagram of the complete on-chip measurement circuit, comprising a shift register that holds the configuration data, an input and clock generation unit

(ICG), a design under test unit (DUT) that includes the flip flops used as synchronizers, the delay line with the calibration unit (CDL) a measuring unit (ME) that includes the circuit of Figure 1(a), a 16-bit counter and an output serializer. A controller writes into the shift register in order to configure the DUT and the DL value. The controller sets the measuring period T , by enabling and disabling the counter. T can vary from seconds to hours. Following each measuring period, the controller initiates a serial readout of the counter value. This procedure is repeated for multiple values of DL. The entire test is performed under software control, and readings are further processed by software.




Figure 4. Block diagram of the metastability measurement circuit and system.

The test system employs a single output pin and six input pins; in practice, such a system could be designed as fully-contained built-in test without any dedicated external pins. A die micrograph of the fabricated test chip and a picture of the measuring system are shown in Figure 5.




Figure 5. (a) Board including on-chip measurement circuit with thermo stream temperature controller. (b) Micrograph of the metastability measurement circuit fabricated in 65nm CMOS.

4 RESULTS

In this section we present measurements and simulations of a library master-slave flip-flop (Figure 6), used as a synchronizer in a 65nm LP CMOS process. The measurements have been performed in sorted typical/typical (TT) parts, and represent an average of 8 measured chips. Simulations used for comparison were carried out using the method described in [8], that sweeps clock and data signals to achieve accurate τ and T_W values. The tool used to perform the simulation is described in [9].

Figure 7 shows measured count values versus delay line settings under various supply voltages. Following (3) a least squares fit is performed to obtain τ and T_W . Following this procedure, τ and T_W are obtained for each $(Temp, V_{dd})$ combination. Figure 8 shows measurement and simulation results for τ at different supply voltages and temperature nodes. Although the measured and simulated data mostly fall on top of each other, both plots are superimposed.




Figure 6. Library flip-flop used for simulations and measurements.




Figure 7. Measured counts versus delay for various supply voltages.




Figure 8. τ Measurements and simulation results versus supply voltage for different temperatures.

A 3D plot of the measurements is shown in Figure 9. Actual measured and simulated nodes are marked with black dots in Figure 10, while other values are interpolated. The error in τ measurement can be estimated, following (3), by the error in the least square (LS) fit. Following statistical rules [10] the LS error is given mainly by the error in the delay difference (ΔDL) and the error in the counter value difference ($\Delta \ln(\#(> DL))$).




Figure 9. τ Interpolated measurements results for different temperatures and supply voltages 3D view.




Figure 10. τ Interpolated measurements values flat view. Actual measurement grid is marked with black circles.




Figure 11. τ Measurements and simulation error for different supply voltage and temperature nodes.




Figure 12. T_W Interpolated measurements and simulation results for different temperatures and supply voltages 3D view.




Figure 13. T_W Interpolated measurements values flat view. Actual measurement grid is marked with black circles.




Figure 14. T_W Measurements and simulation error for different supply voltage and temperature nodes.

Since ΔDL is a delay difference, any absolute delay errors introduced by the last 3 terms in (4) cancel out. Then, the error in the delay measurement is given by the instrument precision for measuring RO frequency which is about 5%. The error in the counter value calculation is given by the time the counter is enabled, since the counter is enabled by an external signal the time difference of the enable signal is only one clock cycle (16ns) which is negligible compared to the experiment time (about 1 minute) and hence the relative error in the counter value is negligible. The least square error calculated with the measured data was about 5.2% which corresponds to the predictions presented. Figure 11 shows the error in τ between measurements and simulations. For each node the delay line is calibrated. The difference error between measurements and simulations ranges between 1.3% to -3.2% through the entire range of supply voltage and temperature nodes, which is between the margins of the measurement error described above. Figure

12 shows T_W measured and simulated results. Figure 13 shows in black circles actual measured nodes, other points correspond to polynomial interpolated values. Measured values of T_W are highly affected by absolute errors in the delay measurements. Delay errors can be caused by insertion delay difference between clock to the DUT and clock to delay line (Figure 1), this insertion delay is reduced to minimum in the design phase, though some parasitic still exists (part of that is accounted for by $T_{parasitics}$ in (4)). This error contribution, is not taken in account in the measurement of DL but has a high impact on the calculation of Y in (3). Due to high process variations in deep sub-micron technologies, the delay of the RO inverter can differ highly from the mean inverter delay of the other inverters in the delay line. Those variations may reach 40% in 65nm, and that is why the measured value of T_W displays such a high variability.




Figure 15. Number of FF's in synchronizer for $F_c = 100 \text{ Mhz}$ and $F_d = 100 \text{ Mhz}$.




Figure 16. Number of FF's in synchronizer for $F_c = 500 \text{ Mhz}$ and $F_d = 100 \text{ Mhz}$.




Figure 17. Number of FF's in synchronizer for $F_c = 1 \text{ Ghz}$ and $F_d = 100 \text{ Mhz}$.

Figure 14 shows the difference error in T_W between measurements and simulations. Note, that even for the extreme differences between measurements and simulations of T_W shown in Fig. 14, the measurement error estimate of 40% is rarely exceeded.

5 PRACTICAL IMPLICATIONS

To understand the implications of the measurements performed, we consider three synchronization scenarios and calculate the number of FF stages to be used in a synchronizer to achieve a given reliability. This calculation is based on measurements in Figures 10 and 13, the desired MTBF and the assumption that each additional stage adds a clock period to the settling time S in (1), as proposed in [11]. It is also assumed that both master and slave latches in the prospective FF's have the same τ and T_W . We omit the setup and propagation delay of each FF for the ease of the

calculation only. Figure 15 shows the number of FF's required for an MTBF of more than 50 years, having $F_d = 100 \text{ Mhz}$, $F_c = 100 \text{ Mhz}$, τ and T_W according to the measurements previously presented in Figure 9 and Figure 12. Two FF's are enough to ensure reliable synchronization when using the nominal voltage (1.1V) for the technology. For lower supply voltages, the number of synchronization stages increases to a maximum of five. We note that using the extreme value of T_W for all the calculations will produce no difference in the number of stages of the synchronizer.

Figure 16 presents a similar case but with $F_c = 500 \text{ Mhz}$. For this case the number of FF's needed using nominal voltage varies from three to four depending on the temperature. It is possible to assume that when a circuit is in functional operation, the self-heat generated, will make junction temperature of internal nodes in the circuits higher than zero. The black line in Figure 16 indicates a 0-degree

junction temperature limit below which actual practical operation is unlikely to be encountered. Under this constraint, the maximum number of FF's for the worst case is around 15. Figure 17 is a similar case with $F_C = 1\text{GHz}$. In this case the number of FF's needed increases significantly even for temperatures above zero degrees. Designing for worst case in this scenario will result in extremely high latency which may be deterrent in most applications.

As is evident from previous examples, synchronization stages are a function of supply voltage and temperature. Those parameters tend to change dynamically within functional operation. Since a large number of FF's is needed in corner cases only, it is possible to dynamically adapt the number of synchronization stages based on τ measurements. With this adaptation, a tradeoff between latency and reliability can be achieved for typical scenarios without compromising reliability in corner cases. A scheme of an adaptive synchronizer is shown in Figure 18. A τ measurement circuit (TMC, similar to Figure 4) is used to measure τ dynamically. Measured values are then passed to the stage calculation unit (SC), which is in charge of calculating the MTBF for the selected number of stages from the measured τ , T_w and F_C (note that F_D can be approximated by F_C to produce a lower bound in the MTBF calculation). The number of stages is digitally selected by controlling the output multiplexor. Also note, that a single adaptive synchronizer is suitable to account for global-process variations, operating-condition changes and for dynamic voltage and frequency scaling (DVFS). For the case of in-die process variations the scheme will not produce reliable synchronization, since the measured FF's are not the same as used in the synchronization chain. A complete description of the adaptation circuit is beyond the scope of this publication and will be addressed in future works.




Figure 18. Proposed adaptive synchronizer

6 SUMMARY

We showed a circuit to extract synchronization parameters (τ and T_w) in a 65nm process, with high efficiency, for a wide range of supply voltages and temperatures. These

measurements are designed to validate simulations of the parameters under varying process, supply voltage and temperature (PVT) conditions. We have made measurements and simulations for a wide range of supply voltage and temperature corners and confirmed that the worst case scenario is achieved for low supply voltages and low temperatures. We compared measurements to simulations and showed that τ can be predicted with an error of less than 5%. T_w was predicted with higher variability, but its impact on MTBF is significantly smaller than that of τ .

We proposed a self-calibrating synchronizer design based on our measurement circuit, which can be adjusted automatically and dynamically to varying circumstances such as global process variations and environmental conditions (voltage and temperature), as well as to dynamic voltage and frequency scaling (DVFS) systems where voltage and clock frequencies may vary from time to time. Such a self-calibration circuit may change the number of serial FFs included in synchronizers to assure the desired level of reliability under changing conditions.

7 ACKNOWLEDGEMENT

The work of Salomon Beer was supported in part by HPI institute for scalable computing. This work was supported in part by the National Science Foundation under Grant No. 0924010

8 REFERENCES

- [1] L. Kleeman and A. Cantoni, "Metastable behavior in Digital Systems", IEEE Design & Test of Computers, 4, 6, 4-19, 1987.
- [2] R. Ginosar, "Metastability and Synchronizers: A Tutorial," IEEE D&T, 28(5):23-35, 2011.
- [3] C. Dike and E. Burton, "Miller and noise effects in synchronizing flip-flop" IEEE JSSC, Vol. 34 No. 6, pp. 849-855, June 1999.
- [4] M.S. Baghini, M.P. Desai, "Impact of technology scaling on metastability performance of CMOS synchronizing latches", Proceedings of ASP-DAC/VLSI Design 2002. pp.317-22, 2002.
- [5] S. Beer, R. Ginosar, *et al.*, "The Devolution of synchronizers", Proceedings of the ASYNC 2010, in press.
- [6] D. Chen, D. Singh *et al.*, "A comprehensive approach to modelling, characterizing and optimizing for metastability in FPGAs," FPGA 2010
- [7] S. Beer, R. Ginosar, *et al.*, "An on-chip metastability measurement circuit to characterize synchronization behavior in 65nm", (ISCAS) – May 2011
- [8] S. Yang and M. Greenstreet, "Computing synchronizer failure probabilities," DATE 2007.
- [9] MetaACE, a Blendics product. See: <http://blendics.com/index.php/blendics-products/>
- [10] B Bevington, P. R. Data Reduction and Error Analysis for the Physical Sciences. New York: McGraw-Hill, 1969.
- [11] D. Kinniment, Synchronization and Arbitration in Digital Systems, Wiley 2007.