Microelectronics Reliability 80 (2018) 91–99

Contents lists available at ScienceDirect

Microelectronics Reliability

journal homepage: www.elsevier.com/locate/microrel

Refined metastability characterization using a time-to-digital converter

Thomas Polzer[a], Florian Huemer[b], Andreas Steininger[b,*]

[a] UAS Technikum Wien, Department of Electronics and Telecommunications, A-1200 Vienna, Austria
[b] TU Wien, Institute of Computer Engineering, Vienna A-1040, Austria

ARTICLE INFO

Keywords:
Metastability
Time-to-digital converter
TDC
Late transition detection
Carry chain

ABSTRACT

In view of the numerous clock domain crossings found in modern systems-on-chip and multicore architectures precise metastability characterization is a fundamental task. We propose a conceptually novel approach for the experimental assessment of upset rate over resolution time that is usually employed to extract the relevant characteristics. Our method is based on connecting a time-to-digital converter to the output of the flip flop under test, rather than using a phase shifted clock, as conventionally done. We present the details of an FPGA implementation of our approach and show its feasibility through an experimental evaluation, whose results favorably match those obtained by the conventional method. The benefits of the novel scheme are the ability to perform a calibration for the delay steps, a speed-up of the measurement process, and the availability of a more comprehensive and ordered measurement data set. Especially the latter can be of crucial importance when (even multiple) glitches or oscillation are suspected to result from metastability, or when the temporal distribution of upsets matters (bursts of upsets, e.g.).

1. Introduction

For various reasons, contemporary digital ASICs comprise a considerable number of clock domains, often, as a result of dynamic voltage and frequency scaling, even with variable frequency. When moving data over the clock domain boundaries, metastability issues need to be carefully considered to avoid reliability problems. This is usually achieved through the use of synchronizers, which increase the available resolution time of a flip flop in case of metastability, and thus make upsets caused by reading "undecided" values – i.e. digital signals carrying analog values between the thresholds for a well defined HI or LO logic level – less likely. In essence, these synchronizers trade performance for reliability. Therefore their careful dimensioning is crucial, which is generally based on the following equation [1]:

$$MTBU = \frac{1}{f_{clk} \cdot \lambda_{dat} \cdot T_0} \quad e^{\frac{t_{res}}{\tau}} \tag{1}$$

Here the MTBU gives the (statistical) mean time between two metastable upsets (i.e. reading undecided digital values) of a given flip flop, which corresponds to the reliability of the data transfer. It is obviously dependent on the flip flop's clock rate f_{clk} and the transition rate λ_{dat} of its (asynchronous) input data. The resolution time t_{res} is the central design parameter of a synchronizer that decides about the mentioned trade-off between reliability and performance. Typically, it is increased beyond one clock period by employing multi-stage synchronizer structures [2]. Finally, T_0 and τ are metastability characteristics of the flip flop and determined by technology and cell design. These latter parameters are essential to know for appropriate synchronizer design (i.e. appropriate choice of t_{res}). The determination of these characteristics is usually done by experimental measurement. Obtaining appropriate precision for these measurements is quite challenging in detail, as it requires time resolution in the range of picoseconds, and the process is time consuming for two reasons: Firstly, since the MTBU is a statistical quantity, a sufficient number of samples has to be collected for a valid estimation. Secondly, when "deep" metastability shall be observed (which is sometimes important to assess secondary effects like slave metastability [2]), the rate of observable events becomes arbitrarily low, which requires long experimental runs to obtain even a single sample.

In the literature numerous experimental results and setups have been reported. Virtually all of these (at least the reasonably recent ones) are based on the same principle of measurement, which will be outlined in Section 2. In this paper we present an alternative principle that circumvents some of the problems of the conventional approach and hence allows a more efficient characterization[1].

Section 3 will be dedicated to the presentation of principle and implementation of this new approach. Its experimental assessment in Section 4 will highlight its benefits as well as limitations, and these will

* Corresponding author.
E-mail addresses: thomas.polzer@technikum-wien.at (T. Polzer), fhuemer@ecs.tuwien.ac.at (F. Huemer), steininger@ecs.tuwien.ac.at (A. Steininger).
[1] This research was partially supported by the Austrian Science Fund (FWF) project SIC (grant P26436-N30).

https://doi.org/10.1016/j.microrel.2017.11.017
Received 5 August 2017; Received in revised form 26 October 2017; Accepted 17 November 2017

be discussed in Section 5. Finally Section 6 will conclude the paper.

2. Related work

Probably the most intuitive way of observing and investigating metastability is by means of an oscilloscope. Accordingly, the first systematic observations of metastable behavior have been documented by oscilloscope screenshots [3,4], and over the years several approaches involving an oscilloscope for analysis of metastable behavior have been proposed [5–8]. Their key benefit is being able to show the individual waveforms in all detail, including the actual intermediate voltage during the undecided phase. However, the extraction of metastability parameters is quite intricate, and the collection of a statistically relevant number of measurements suffers from the difficulty of automating the approach.

Circuit level simulators, such as SPICE have been successfully employed to observe the metastable behavior of an element [9–11]. While they allow convenient access without probing effects, they use idealized models that may not appropriately reflect the actual circuit parameters, and they tend to suffer from numerical problems and high sensitivity to even small changes in the chosen parameters and models [12].

Consequently, by far the most popular method for assessing metastability characteristics is based on the following observation:

When plotting $ln(MTBU)$ over t_{res}, Eq. (1) becomes a straight line, whose gradient and offset are determined by τ and T_0, respectively. The same holds true for the upset rate which is the inverse of the mean time between upsets: $UR = 1/MTBU$. This suggests the following approach for experimentally determining the metastability characteristics: For given fixed clock and data rates the upset rate is measured for different settings of the resolution time. From these points the graph $UR(t_{res})$ can be plotted and the characteristics extracted. In fact this is the preferred principle of the measurement approaches for τ and T_0 in the literature. Note that here metastability is measured indirectly by observing the delay in the output transition it causes, and "upset" is defined as a delay that exceeds the available resolution time.

Concerning the practical realization of the experiment, providing clock and data is fairly easy. The only important thing here is to take care that these are uncorrelated, such that from time to time metastable upsets actually occur as a consequence of the uniformly distributed phase difference [5], and hence the assumptions made by Veendrick in the derivation of Eq. (1) hold.

Alternatively, the phase shift between clock and data can be carefully controlled to produce a maximum number of metastable upsets. Due to the influence of jitter, noise, voltage and parameter variations, etc. this is a very delicate task that requires special provisions like, e.g., using a delay locked loop [6,13]. The obvious advantage of this method is to generate a relatively high yield of (even deep) metastable upsets, which allows their detailed study. The correlation between clock and data transitions established by the phase control, however, rules out a characterization based on the application of Eq. (1). So this approach is not useful for our purpose.

Measuring the upset rate is also fairly easy; it just takes a counter and a time measurement. A more difficult question is what to actually count, i.e., how to identify an upset. This is usually done by sampling the output of the flip flop under test once after t_{res} and once much later (when it can be safely assumed that metastability has resolved already), and comparing these samples: A mismatch indicates that the final output state (defined by the reference provided through the second sample) had not been reached after t_{res}, hence the flip flop under test (further abbreviated as UUT for "unit under test") must have been metastable then.

Note that this is actually a conjecture that follows from the assumption that there is only one transition leading from the initial state to the final one. Should there, however, be multiple transitions, then this conjecture may be invalid. This inability to handle multiple transitions correctly is a severe deficiency of the conventional approach.

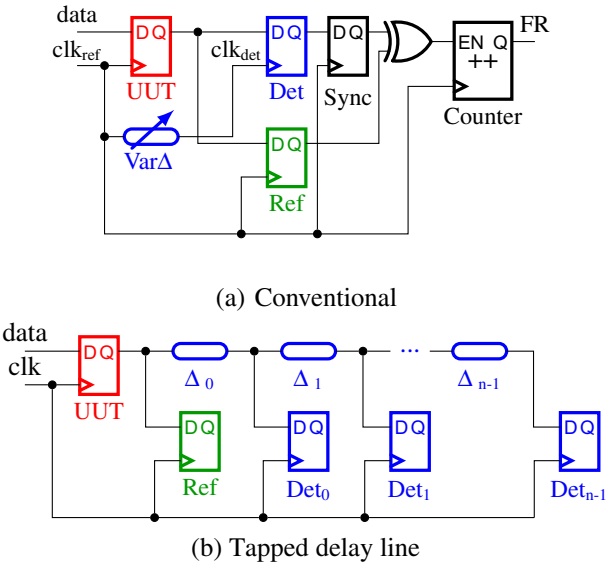

Fig. 1. Basic circuit for metastability characterization.

The most tricky thing in practice, however, is the provision of t_{res}. Note that, with contemporary technology (say 28 nm) virtually all metastable states of the UUT usually resolve within less than 1 ns, and only very few last longer (recall from Eq. (1) that with higher resolution time MTBU becomes *exponentially* longer). As a result, to produce the $UR(t_{res})$ graph with reasonable quality, t_{res} must be controlled with a resolution and known with an accuracy in the picosecond range, which is definitely non-trivial. In the approaches published so far, t_{res} has been realized as a time shift between the UUT clock and the clock of the flip flop taking the first sample of the UUT's output (*Det* in Fig. 1a). This yields the measurement circuit illustrated in Fig. 1a.

The controlled delay VarΔ required for this purpose has been practically implemented in a number of different ways: In [14] an external pulse generator is used, which, apart from requiring such external measurement equipment, is not applicable any more for modern ASIC and FPGA technologies, as the timing uncertainties introduced by the external cabling are no more acceptable. By triggering the flipflop under test with the rising clock edge, while sampling its output with the falling one, [15] utilizes the HI-pulse width of the clock for determining the resolution time. Again, this approach scales badly for recent technologies, as the appropriate control and variation of clock pulse width in the ps-range is difficult to accomplish. In a similar way, [16] employ rising and falling edges to control the resolution time on their FPGA target, but even though they carefully manipulate the routing of the flip flops to achieve the required changes, they obtain only a very limited number of measurement points.

The starved buffer proposed in [13] allows precise tuning of a single (uncalibrated) delay, which was very useful in the context of a closed control loop for driving the flip flop into deep metastability, where the exact knowledge of the delay was not required. On an FPGA-platform, however, starved buffers are not available. Moreover, we would need to use multiple of these buffers, calibrate them and coordinate their control, which is a significant effort.

In [17] the authors propose the use of an inverter chain with fixed delay, and they vary the clock frequency to obtain different resolution times. This, however, also changes f_{clk} in addition to t_{res} in Eq. (1), which makes the extraction of τ more cumbersome. The same is true for [15,18] where rising and falling edges are utilized to determine the resolution time, but with a fixed duty cycle and a variable clock frequency.

The tapped chain of logic functions from [19], the tapped inverter chain from [11], and apparently also the "delay line" mentioned in [20] are very close in nature to our proposed approach. As will be outlined in