Measuring Deep Metastability

David Kinniment, Keith Heron, and Gordon Russell
Newcastle University, UK
{David.Kinniment,Keith.Heron,G.Russell}@ncl.ac.uk

Abstract

Present measurement techniques do not allow synchronizer reliability to be measured in the region of most interest, that is, beyond the first half cycle of the synchronizer clock. We describe methods of extending the measurement range, in which the number of metastable events generated is increased by four orders of magnitude, and events with long metastable times are selected from the large number of more normal events. The relationship found between input times and the resulting output times is dependent on accurate measurement of input time distributions with deviations of less than 10ps. We show how the distribution of D to Clock times at the input can be characterised in the presence of noise, and how predictions of failure rates for long synchronizer times can be made. Anomalies such as the increased failure rates in a master slave synchronizer produced by the back edge of the clock are measured.

1. Introduction

Globally asynchronous, locally synchronous (GALS) systems on chip that use IP blocks from different vendors with unrelated clocks may use many synchronizers to resynchronize the data transmitted from one IP block to another. Stoppable clocks have been proposed as a way of avoiding the time penalty resulting from the need to synchronize data, but large clock trees may make stopping the clock impractical because of the delay involved in the clock tree before the clock actually stops [1][2]. For this reason synchronization is likely to take an increasingly large proportion of the on-chip communication time in future generations of process technology. First the reliability required per synchronizer is higher, since the number of synchronizers on chip is higher. Second, process, voltage and temperature variations affect synchronizers disproportionately, and third, clock and data rates will increase all leading to long synchronization times.

Synchronization of data sent using one clock, and received by another requires a minimum of 0-1 receiving processor cycles to synchronize, and this is likely to increase to several cycles at 65 nm and below because of the time taken by a metastable level in the synchronizer to resolve to a high or low level.

For high reliability in a system with many synchronizers, each synchronizer needs to have a reliability measured in years where 1 year ~ 3.1 x 10^7 seconds. Most designs assume a synchronizer mean time between failures, $MTBF = \frac{e^{\frac{t}{\tau}}}{T_w f_c f_d}$, where f_c is the clock frequency, f_d is the data transition frequency and T_w is sometimes called the metastability window. Though this is an approximation in the deterministic region, [3], it is possible to use it to estimate the amount of time t required by the synchronizer for long term reliability. With T_w in the order of 10pS, reliabilities of more than 10^7 seconds (4 months) can only be achieved at clock and data rates of $1GH_z$ by allowing a time of longer than 32τ for metastability to settle, where τ is the resolving time constant of the synchronizer, [2].

Because the formula presented above is a simplification, [3], the value of τ often appears to vary as a function of t, [4][5]. Its starting value is sometimes higher than its final value and sometimes lower depending upon the initial conditions of the circuit. Typically circuits in which both the true and the inverse outputs of Q start low will initially take longer to resolve to a high than to a low, [3], but because the long term resolution times are unaffected by the initial transient, the initial slope is faster than the final τ. Changes in data and clock inputs during metastability can also affect resolution times.

Particularly important effects are those for circuits with multiple time constants, whose metastability trajectories may be oscillatory rather than exponential

[6][7][8][9] and the influence of the back edge of the clock in a master-slave flip-flop, which transfers a metastable state from the master to the slave device if resolution times are longer than half a cycle. Present measurement techniques do not allow these effects to be directly measured at the point of interest, which is usually beyond 15τ, thus all reliability projections are based on an extrapolation of a simplified MTBF formula, which does not take into account important effects.

In this paper we show how measurements can be made up to seven orders of magnitude in MTBF beyond current limits, and demonstrate the effect of the back edge of the clock. In section 2 we describe the most common measurement method and present new techniques aimed at increasing the number of metastable events recorded in a given time. In section 3 this method is extended by selecting only those events lasting longer than a given time; and a correction for jitter and noise in the measurement equipment is described. In section 4 we present data for a master slave flip-flop that shows the effect of the back edge of the clock, and finally we show that these techniques are capable of incorporation in an on chip measurement scheme.

2. Measuring metastability

Normally metastability measurements make use of two asynchronous oscillators driving the D and Clock inputs of a master-slave flip-flop under test.

If the data oscillator has a period of 99.9ns and the clock 100ns, the rising clock edge may, or may not produce a change in the output Q. Figure 1 shows the situation when the D input is low on the first clock edge, and then goes high very close to the next edge causing a change in the Q output.

We can only observe metastability if the D input is different on successive clock edges, and even then, only if it changes very close to the second clock edge, causing metastability to occur.

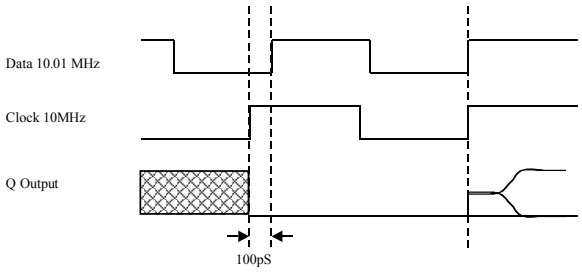

Figure 1 Two-oscillator metastability measurement

If the data and clock oscillators are not locked together, all overlap times between 0 and the 100ns cycle time for data and clock are generated with equal probability. To observe the delay due to metastability, the Q change from low to high is used to trigger the recording of each clock rising edge for a potentially metastable event. During the collection of data only those events where clock and data overlap by less than the difference between the two oscillator periods (100ns - 99.9ns) can be observed, because they are the only events which generate a change in Q. These events are then presented as a histogram of the number of events against the time from the Q output to the clock. A typical histogram is shown in Figure 2 and other examples appear in [4], and [5]. In Figure 2 the X-axis represents time from a triggering Q output back to the clock edge and therefore increasing metastability time is shown from right to left. The Y-axis is the number of events with a peak of 80,000 at 20.5ns, and the decreasing number of events recorded shows in the discrete levels of 0, 1 and 2 events at around 17 ns.

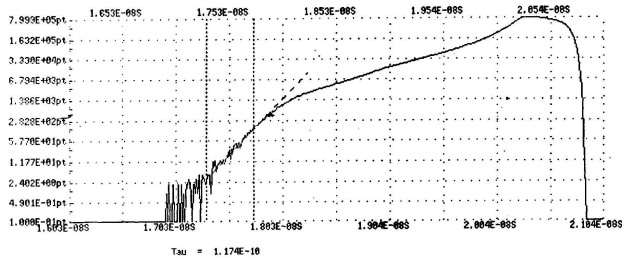

Figure 2 Event histogram for a 74F5074

Unfortunately events which result in a much longer than normal propagation delay (deep metastability) occur relatively rarely. Such an event will require much less than 100ps overlap between data and clock. Less than 100ps overlaps will occur in only 1 in 1000 of the clock cycles, so truly metastable events happen much less often than the 100 μS implied by 100ns*1000. Even then not all of these events can be collected if a general-purpose digital oscilloscope is used to collate the data. Because the oscilloscope must store, process and display the histogram, there will be a significant dead time between successive recorded events that may limit the actual events recorded to as few as 1 in 1000 of those generated. The MTBF formula shows that with $f_c = f_d = 10^7$, and $T_w = 10ps$, low probability events lasting 15τ or more will only occur approximately every $10^{-3} e^{15}$ seconds, thus to produce 1 event requires the experiment to run for about 3269 seconds (about 1 hour), and if only 1 in 1000 of the

events are recorded, the maximum metastability time that can be observed in practice is limited by the time available to record sufficient events to much less than 15τ. Increasing the data and clock frequencies can improve the number of low probability events recorded, but even so, one 15τ event only occurs in 10 hours even with frequencies of 100MHz.

In our measurements we have increased the probability of a metastable event by ensuring that the data transition is always within a small time Δt from the point of very long metastability. The assumption implicit in the MTBF formula is that on average an input transition Δt away from the balance point, that is the input time where resolution of metastability takes the longest, is given by $\Delta t = \frac{f_c . f_d}{MTBF} = \frac{T_w}{e^{\frac{t}{\tau}}}$.

Therefore if all inputs are closer than Δt to the balance point all output events will last longer than $t = \tau . \ln \frac{T_w}{\Delta t}$. Using Δt = 100ps there are 100ns/100ps =1000 times as many metastable events as there would be if all overlaps were equally probable. This allows either the same measurements to be made in a much shorter time or more deep metastable events to be recorded in the same time.

We use a delay locked loop to hold the data input at a point that gives an 0.5 probability of a high transition in the Q output, and an 0.5 probability that it stays low. This represents a balance point that gives many very long output times. For input times within less than 0.2 ps of the balance point, the circuit outputs are non-deterministic and whether the output ends up high or low depends on the internal noise, [3],[5], as well as the external input time.

A schematic of the test set-up is shown in Figure 3. Here a 10Mz clock is passed through two closely matched paths to the data and clock inputs of the device under test.

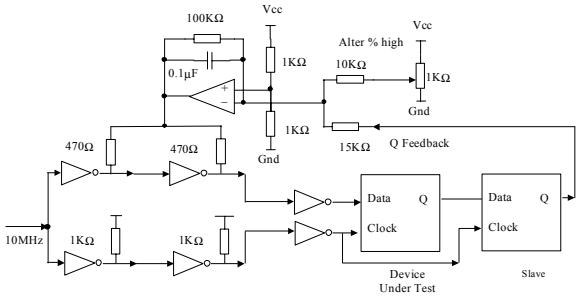

Figure 3 DLL controlling data and clock overlap

Adjusting the supply voltage to open collector inverters varies the delay in the path to the data input. If the data rising edge is slightly slow when compared with the clock edge, Q will be low on most clocks. If it is fast, Q will be mostly high. A slave flip-flop records whether the device under test has resolved high or low, and an analog integrator is used to average the proportion of high to low outputs from the slave. The arrangement forms a delay locked loop in which the voltage supply to the inverters in the data path is increased if it is too slow and reduced if it is too fast. The integrator has a second input that enables the proportion of highs and lows to be set manually. Normally we set this input voltage to $\frac{V_{high} + V_{low}}{2}$ so that the system settles to a steady state where 50% of flip-flop outputs are high and 50% low but it is possible to vary the proportions simply by increasing or reducing the voltage. A higher input forces the loop to compensate by reducing the proportion of high slave outputs, and vice versa. A separate input to the data delay path supply voltage (not shown) enables a high speed waveform to vary the delay by around + or - 100ps giving the delay distribution between clock and data shown in Figure 4

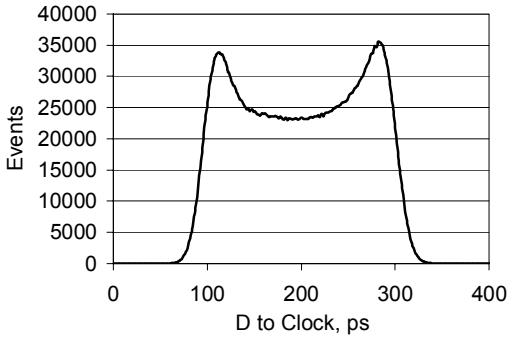

Figure 4 D to Clock distribution

This histogram was obtained using an Agilent Technologies 54855A Infiniium oscilloscope and shows a 100ps variation around an average of about 200ps, and triggering the data collection oscilloscope from the rising Q outputs gives the display shown in Figure 5.

To demonstrate feasibility of the method, we chose a 74F5074, a master-slave flip-flop specially designed to give a fast, controlled metastability response. While it is an old design, it is easily available, and comparisons can be made with published data. In Figure 5, the oscilloscope is in color grade mode, so that the density of traces at a particular point is represented by the color of the pixel

at that point on the display. A histogram of the trace density along the horizontal line is also shown in this figure, which represents the number of events passing through the pixels concerned. Because the number of input events per picosecond in the D to Clock histogram of Figure 4 varies with the D to Clock time, it is necessary to correct for the variation. In the two-oscillator measurement method, the events per pico-second is constant with respect to D to Clock time. The correction can be done, by noting that for exactly half of the input events on the D histogram Q go high, and for half, Q remains low.

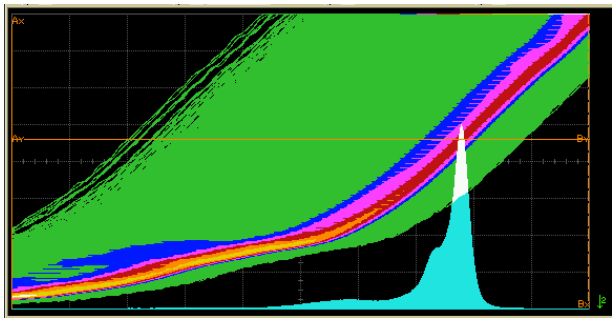

Figure 5 Display histogram

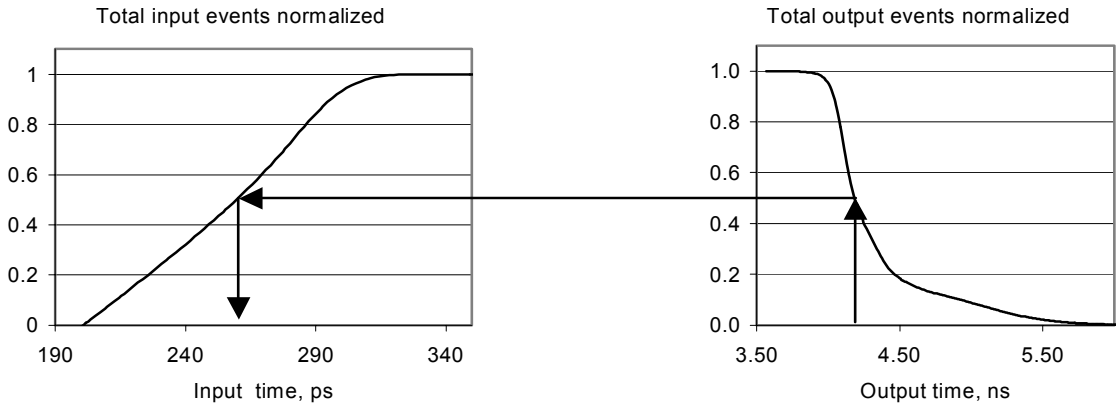

Figure 6 Output time to input time

When D is early there is a rising output on the Q histogram for every input in the D to Clock histogram. Thus if the total number of events on the D to Clock histogram are scaled from -1 to +1, those causing a high output will be in the range 0 to 1.

When the number of events on the output histogram is also scaled from 0 to 1, the correspondence between D times, and Q times can be found. This can be seen in Figure 6 where the total number of input events up to a particular input time are plotted on the left, and the total number of output events occurring after a particular output time (measured from the clock) are shown on the right. In this figure the number of events where the propagation delay is greater than 6ns is very small, but as the propagation delay falls to near the normal delay of 3.9ns the number of events rapidly increases to include all of those measured. Half of all the output events have metastability times lasting longer than 4.18 ns, and half of the input events occur with the D input between the balance point, and 260ps. Here we define the balance point to be the input time that gives an exactly equal probability of the final output being high or low, and hence in the absence of noise an input at this time would give a very long metastability time. For this device the balance point is 205ps, and in the experiment of Figure 6, 0.5 of all the events occur with a D times between 205 ps and 260 ps. The 0.5 point is therefore associated with a unique input time of $\Delta t = 260 - 205 = 55\,ps$ and a unique 4.18 ns output time. Plotting input time, Δt, against output time gives Figure 7 where the input time has been adjusted by 205 ps to give a zero at the balance point. This graph shows that any input less than 55ps away from the balance point will give an output lasting longer than 4.18 ns, so if the synchronization time is set to 4.18 ns failures will occur in a synchronizer with inputs less than 55ps. The mapping process of Figure 6 effectively linearizes the number of events against Δt so that the number of events per picosecond in Figure 7 is constant. It is important to note that this graph is statistical in nature; it is not possible to measure any single input to much better than 1 ps because of the presence of thermal noise. Each point on the graph, however, represents more than one event. It is

characterised by the number of events that occurred after the output time given by its X coordinate. The Y coordinate of the point is the input time just greater than the same number of input events. Because input event times are distributed uniformly over the input event time range, typically 100ps, random noise has no effect on the input distribution, there are still a constant number of events per picosecond. The accuracy of the graph however, depends on the number of events, the more events in the range, the greater the effective accuracy, and the smaller the input times that can be plotted.

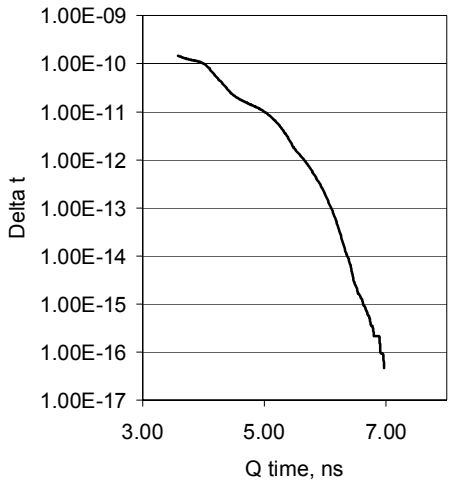

Figure 7 Standardized metastability characteristic

A significant advantage of this technique is that it allows the results to be presented in an easily understood standardized form, independent of oscillator frequency or number of events. With knowledge of the clock and data frequencies any point on the y-axis, Δ_t, can also be converted to give the mean time between failures, since a single input will give an overlap less Δ_t, in any particular clock cycle with a probability $\Delta_t f_c$, and over a time T, there will be Tf_d inputs. It follows that $MTBF = \frac{1}{\Delta_t f_c f_d}$ and given the values of f_c and f_d in a system, the scale can easily be converted from Δ_t to $MTBF$

3. Deep metastability

If the source of data delay variation is removed, the number of input events close to the balance point will be increased, and therefore the ability to measure longer metastable times correspondingly enhanced. Figure 8 shows the distribution of input times that results.

Here the deviation from the central value is about 12ps, and is very similar to a distribution that would be produced by random noise alone. This is partly because the flip-flop output value, high or low, is at least partly determined by internal thermal noise, and partly because there is a significant noise element in the oscilloscope measurements.

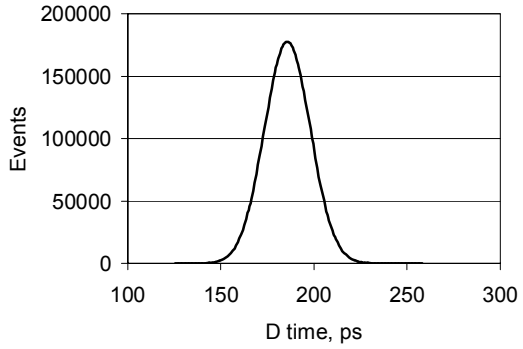

Figure 8 Event histogram with only noise variation on the input.

The measurement noise can be estimated by producing a histogram of jitter on the rise of the clock. When triggering on the same rising transition as the measurement it is possible to observe that at the trigger point the measurement deviation is very low, but away from that point it spreads out to around 9 ps. The specification of the 54855A gives a similar value for this type of measurement of 9.2ps.

Because of this relatively large measurement noise component we cannot reliably use Figure 8 to assign input times to output times.

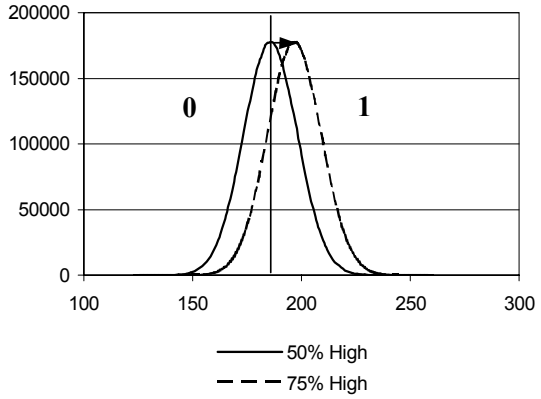

Figure 9 Shifting the input distribution.

To overcome this problem we biased the DLL to produce a range of different proportions of high and low outputs from the device under test. If instead of having an integrator bias of $\frac{V_{high} + V_{low}}{2}$ giving 50% high values we set the bias voltage to $\frac{V_{high} + 3V_{low}}{4}$ we get 75% high values, and the center point of the D distribution moves earlier by exactly the time required to shift 25% of the events from the 0 side of the balance point to the 1 side as shown in Figure 9.

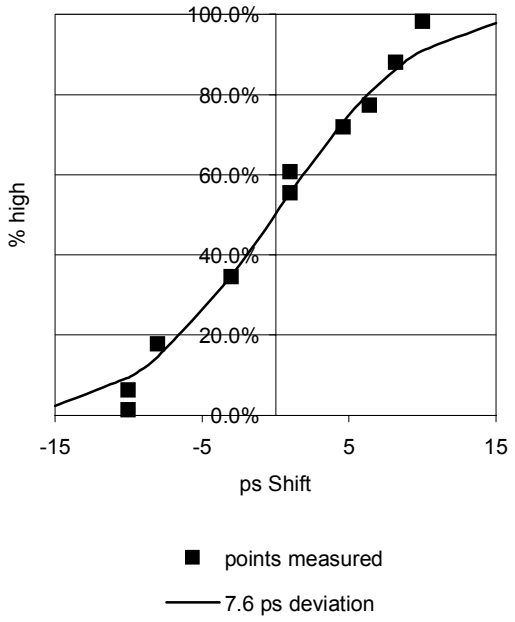

Figure 10 Measurement of actual D distribution

The shifts required to give different probabilities of high outputs can now be plotted on a graph showing % high points against resulting shift, Figure 10. If we assume that the time of input events follows a normal distribution, we can compare these shifts with distributions having different values of standard deviation, σ. The line with the closest fit to the points on Figure 10 represents the cumulative probability of a high output for a random input time deviation of 7.6 ps, so we can conclude that the actual distribution has a deviation of this value.

The corrected input time distribution corresponding to this is shown in Figure 11. For comparison the original histogram measured on the oscilloscope is also shown, demonstrating that measurement noise makes a significant contribution to the raw data. Combining the oscilloscope deviation of 9.2 ps with 7.6 ps as the root sum of squares gives a result of 11.9 ps, very close to that observed.

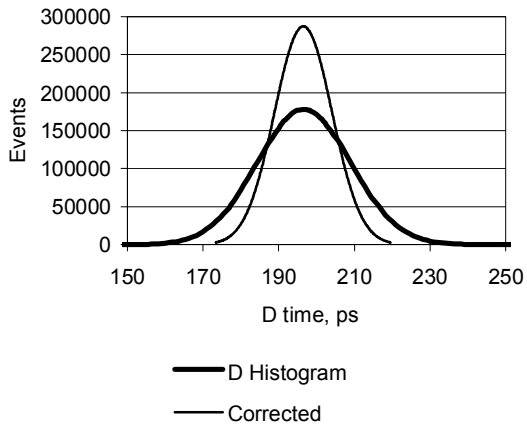

Figure 11 Corrected input time distribution

This method of generating and measuring time distributions at the pico-second level in the presence of noise allows us to take our measurements another order of magnitude further, as shown in Figure 12, by the 7.6ps line.

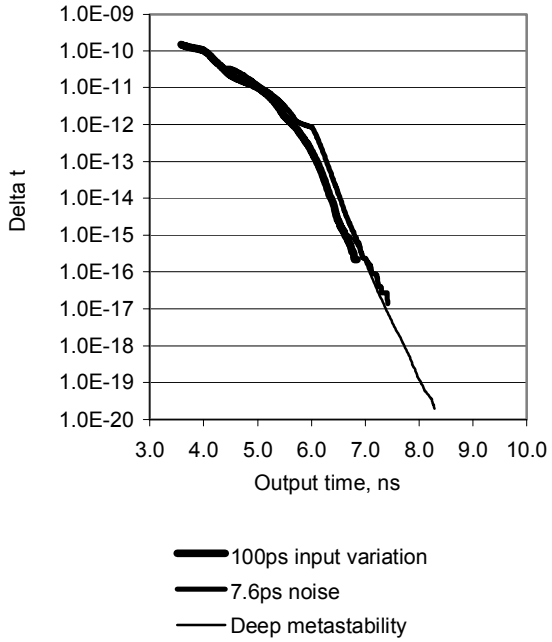

Figure 12 100ps, 7.6ps deviation, and deep metastability plots

This figure also shows the results of triggering the oscilloscope only on Q outputs that appear after 6.5 ns. We do this by clocking the Q output into two flip-flops, one at 6.5ns, and the other at 50ns. The

oscilloscope is triggered only if the first flip-flop gives a low, and the second a high. Since digital oscilloscopes collect data continuously, it is possible to measure events occurring well before the trigger - in this case 60ns before. Because only a small number of events last longer than 6.5ns, the trigger rate is 1000 time slower on average, and almost all meaningful events are captured. There are now far more useful events which therefore give greater time accuracy, and enable us to go a further three decades down in input time. Unfortunately we do not know exactly how many of the input events lead to output events over 6.5ns. To plot these results, the actual number of output events occurring over a period of time is recorded by counting the number of oscilloscope triggers over a specific time in a separate counter, not part of the oscilloscope functionality. Input events causing rising output times longer than 6.5ns have a very low probability, and the probability of any rising clock edge leading to a trigger during the measurement period is given by $P = \frac{OutputTriggerRate}{ClockRate}$.

Around input time = 0, the normalised curve of total input events against input time on the left hand side of Figure 6 is very linear, that is $\frac{d(TotalInputEvents)}{dInputTime} = K$ where both P and K are constants. By a similar argument to that presented in section 2, any change in number of output events due to a change in output time is reflected in a corresponding change in input time. In this case the input curve is linear, so it is only necessary to multiply the total output events at any output time on the right hand side of Figure 6 by the constant $\frac{P}{K}$ to get the equivalent input time. The result is shown in the deep metastability curve of Figure 12 that reaches down to almost 10^{-20}s

The graph shows all three measurement techniques applied to a 74F5074 device run at 10MHz for approximately 1000 seconds in each mode. The resulting range of Δt from 10^{-10}s to 10^{-20}s is 5 decades (11τ) more than existing methods as represented by Figure 2. It is interesting to note that τ for this rather complex bipolar device has at least 3 different values, 350ps between 4n and 6 ns, 120ps between 6.2ns and 6.8ns, then 140ps beyond 7ns. This last value cannot be seen using conventional measurement methods. The fairly distinct breakpoint at 6 ns corresponds to at input of about 1ps, and is therefore in the deterministic region [5].

4. Back edge

To confirm that our results are not artefacts of the measurement method we also measured a more conventional device.

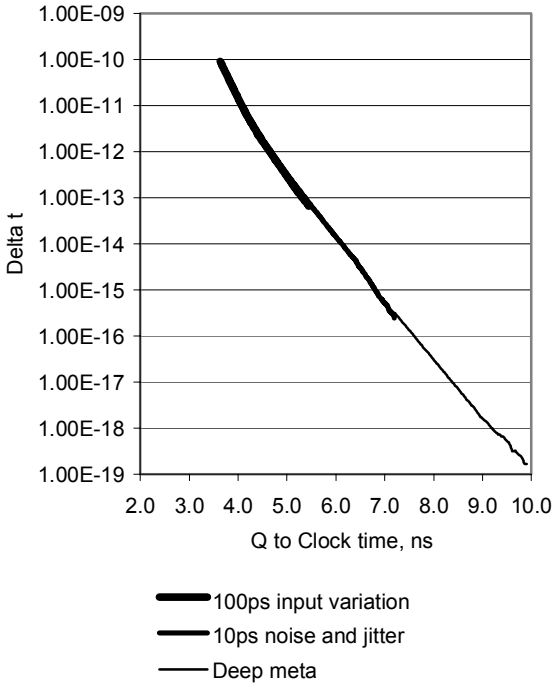

Figure 13 CMOS master slave device

Figure 13 shows the plots of output time against input time overlap for a 74ACT74 from Philips. This is a CMOS master-slave circuit, with a τ of about 350ps.

The results are fairly typical of a CMOS device having normal Clock - Q delay of around 3.3ns and an initial fast slope, which is probably the result of asymmetrical initial conditions, [3]. The slope of the curve is constant in this case over a fairly wide range.

More interesting is the effect of the back edge of the clock. At 10MHz, on the rising edge of the clock, metastability occurs in the master, but the slave is transparent for the first half cycle (50ns). When the clock high pulse width was reduced to 5.5ns and then 4.5ns (minimum allowed in the data sheet is 5ns), we got the graphs of Figure 14: In this figure, there is a higher probability of long metastable times which is linked to the low going edge of the clock

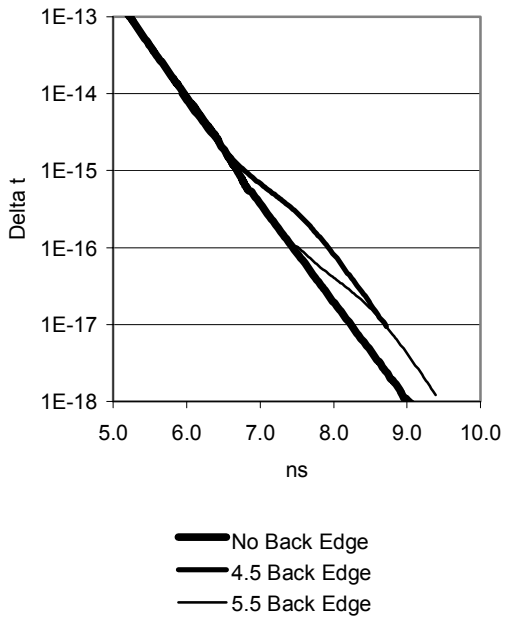

Figure 14 Effect of the back edge of the clock

The effect is also observable on the output waveforms shown in Figure 15 where the clock pulse is about 5ns in width, and the scale is 2ns per division. This shows that the circuit conditions change when the clock goes low. When the slave is transparent (between 0 and 2.8ns after the back edge) there is a bend in the Q waveform that is not present in trajectories after 2.8ns following the low going transition of the clock edge. It could be argued that this is a passive effect, such as ground bounce, but if so, it should be present both before and after the clock transition. The starting point of the rise of the Q output trajectories after the clock are also delayed by the 0.4ns seen in Figure 14.

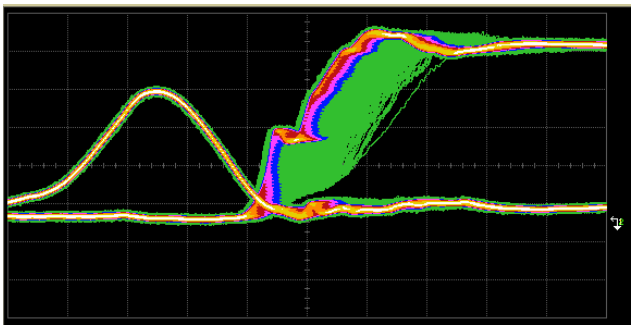

Figure 15 Clock pulse width 5nsClock

In the 74F5074, the effect is much greater, the low going clock transition being associated with a 1.5 - 2ns increase in output time for events after the back edge even though there are no anomalous output trajectories, either before or after the transition.

An analysis of why the increase in delay occurs is beyond the scope of this paper, but we can correlate it with the change in the slave from transparent to opaque. We will assume that both master and slave circuits are the same, and that in the input vs output graph of Figure 12 the master goes from transparent to opaque when the clock occurs, and the slave is transparent when the master output change appears at its input. If the input time to the master is large, the input transition passes through a master in normal transparent mode operation, but if it is very small the master will go metastable. The delay when the master is transparent is much shorter than it would be if the metastable characteristics had been projected back to the same input time, so there is an apparent extra delay due to switching in the master. The same delay can be expected when the slave switches from transparent to metastable on the falling edge of the clock. The delay difference can be found by projecting back the non-deterministic part of the curve of Figure 12 the point where $\Delta t = \tau$ as shown in Figure 16. The reason why the 74F5074 shows such a large increase, is that the difference between the transparent slave propagation time at 150ps input (less than 3.6ns) and what it would have been had the slave been metastable (about 5.2ns) is 1.6ns. This additional 1.6ns delay is clearly observable in back edge measurements similar to Figure 14 1.6ns is equivalent to over 10τ, and will significantly affect the reliability of the synchronizer

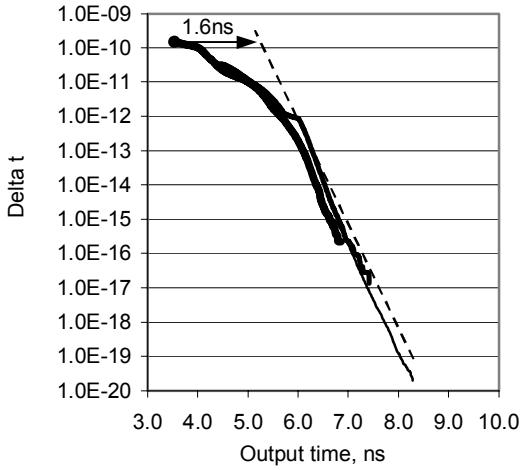

Figure 16 Delay change in 74F5074

5. On Chip implementation

The measurements described here were made by using off chip analog variable delay lines and an operational amplifier RC integrator as components in a

delay locked loop, however implementation on chip, in a digital form should be relatively straightforward. Each open collector delay element in Figure 3 can be replaced by a current starved inverter variable delay element shown in Figure 17.

Simulations of a DLL using these elements [12] with 0.35µ process parameters show that delay increments of 1ps for each p-type transistor switched by the control lines should be achievable, though several delay elements in series may be needed to give sufficient range of input delay. Since each control line switches just one transistor a thermometer code input in which a 1ps delay increase is represented by a change from n out of m control lines high to n + 1 high.

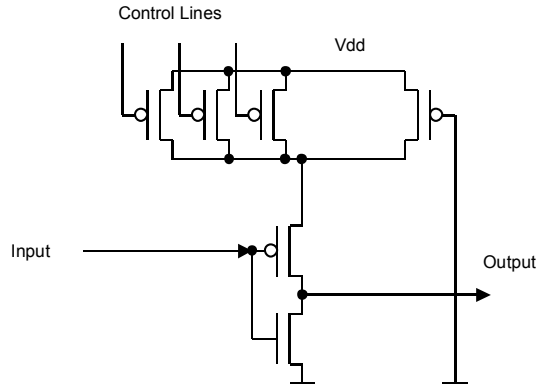

Figure 17 Variable delay element (VDE)

The analog integrator can be replaced by an up/down counter. The count required to produce a given delay change the variable delay would define the effective time constant, for example the 1.5ms time constant of the analog integrator in Figure 3 changes the delay by around +0.06ps for every high output, and -0.06ps for low output. A similar effect would result from a 4 bit initial count, followed by further bits decoded into a thermometer output to drive the control lines of the delays as shown in Figure 18.

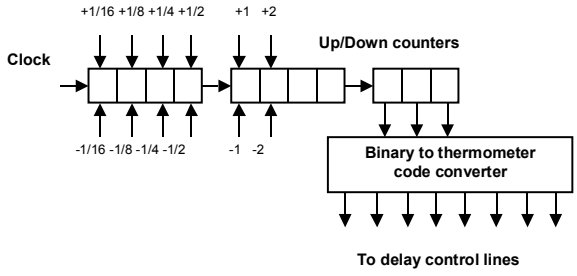

Figure 18 Digital integrator

In this schematic 16 high or low inputs of value +1 or -1 would be needed to change the most significant bit of the 4-bit counter, and hence alter the delay by 1ps.

One advantage of a digital integrator is that it is relatively simple to alter the ratio of high to low Q outputs needed to set the delays. Instead of counting +1 for a high output and -1 for a low output it would be possible to add extra bits on to the least significant end of the counter to get effective counts of +1/2, or +1/4 for a high output, and to bypass some lower order bits to get a count of -2 for a low output. Characterising the actual flip-flop input distribution now becomes easier than with the analog integrator, since it is just a matter of selecting the correct inputs on the counter to get a precise ratio of high to low outputs.

It is interesting to note that a very long integrator time constant (i.e. a large counter) will require many clock periods to alter the relative D to Clock delay by 1ps, so the distribution of input times is then quite narrow at less than 10ps. On the other hand, a much shorter time constant (fewer counter bits), together with the random nature of the output from the flip-flop under test gives a much greater variation in input times. This could provide a mechanism for controlling the input time distribution.

6 Conclusions

By reducing typical input times from 100ns to around 10ps, a four order of magnitude increase in the probability of metastable events has been achieved, allowing much more of the event histogram of a synchronizer to be observed. Difficulties in measuring the distribution of input times in the presence of picosecond level noise in the measuring equipment can be overcome by the technique of measuring the shift in the distribution as a function of the proportion of high and low outputs. This technique also allows a range of time difference distributions to be created on chip by altering the integrator time constant, and to be accurately measured off chip in the presence of noise.

Limitations in the repetition rate of the measuring equipment can reduce the rate of collection of useful output data to as little as 1 data point collected in 1000 generated. This can be overcome by generating only events in the region of interest (deep metastability) so that the repetition rate required is low, and normal propagation time events do not obscure the others. Thus a total of seven orders of magnitude improvement over conventional methods of measuring metastability is possible using the methods described.

Our results show approximately double the useful range of input times, 10^{-10}s to 10^{-20}s, or over 25τ in the

case of the 74F5074, achieved with data collection times of only 1000 seconds. Increasing the range of output times has also revealed an additional final slope and has enabled the effect of the back edge of the clock to be observed for the first time. In the circuits we measured, the back edge increased the resolution time because metastability in the opaque mode of the slave of a master slave flip-flop response introduces more delay than the transparent mode. Our results show that as much as 10τ can be added to resolution times, reducing reliability by a factor of more than 1000. Both of these results could be important in safety critical applications.

The results are presented in a form, input time against output time, which is more meaningful than the usual events scale, and is easily converted to MTBF when the clock and data rates are known.

We aim now to show how these methods can be used on-chip version by replacing the analog variable delay elements by a digital version based on a current-starved inverter, [10][11], and the integrator with a limiting counter which can be incremented and decremented by programmable amounts. This will enable the proportion of high and low outputs to be varied and therefore the actual distribution of input times to be measured off chip.

7. Acknowledgements:

The authors would like to acknowledge equipment support from Agilent Technologies, discussions with Charles Dike of INTEL and funding from the UK EPSRC research grant EP/C007298/1, without which this work would not have been possible

8 References

[1] K. Y. Yun, and A. E. Dooply, "Pausible clocking based heterogeneous systems", IEEE trans, VLSI systems, Vol.7, No. 4, pp 482-487, Dec 1996,.

[2] R. Dobkin, R. Ginosar, and C. Sotirou, "Data Synchronization Issues in GALS SoCs", Proc ASYNC 2004 pp 170-179

[3] D.J.Kinniment, A. Bystrov, A.V. Yakovlev, "Synchronization Circuit Performance", IEEE Journal of Solid-State Circuits, 37(2), pp202-209, 2002.

[4] Y. Semiat, and R. Ginosar "Timing Measurements of Synchronization Circuits" Proceedings ASYNC2003, Vancouver, 12 – 16 May 2003, pp68-77.

[5] C.Dike and E.Burton. "Miller and Noise Effects in a Synchronizing Flip-Flop". IEEE Journal of Solid State Circuits Vol. 34 No. 6, pp.849-855, June 1999

[6] D.J.Kinniment and J.V.Woods, 'Synchronization and arbitration circuits in digital systems', Proc. IEE vol. 123, No. 10, October 1976

[7] Takayasu Sakurai "Optimisation of CMOS Arbiter and Synchronizer Circuits with Submicrometer MOSFET's" IEEE Journal of solid -state circuits vol 23 no 4 August 1988 pp901 -906

[8] C.H. van Berkel and C.E. Molnar,"Beware the 3-Way Arbiter," IEEE Journal of Solid-State Circuits, Vol.34, pp. 840-848 1999.

[9] O. Maevsky, D.J. Kinniment, A.Yakovlev, A. Bystrov "Analysis of the oscillation problem in tri-flops", Proc. ISCAS'02, Scottsdale, Arizona, May 2002, IEEE, volume I, pp.381-384.

[10] Piotr Dudek et al. "A high-resolution CMOS Time-to-Digital converter utilizing a Vernier delay line". *IEEE Journal of Solid-State Circuits*, 35(2):240.247, February 2000.

[11] M. Mota and J. Christiansen. "A four-channel selfcalibrating high-resolution time to digital converter". In *IEEE International Conference on Electronics, Circuits and Systems*, volume 1, pages 409.412. IEEE, 1998.

[12] Crescenzo D'Alessandro, Kelvin Gardiner, David J. Kinniment, Alex V. Yakovlev "On-chip Sub-Picosecond Phase Alignment", 2nd UK embedded Forum, 19/20 October 2005