On-Chip Measurement of Deep Metastability in Synchronizers

Jun Zhou, *Student Member, IEEE*, David J. Kinniment, *Life Member, IEEE*, Charles E. Dike, *Member, IEEE*, Gordon Russell, *Member, IEEE*, and Alexandre V. Yakovlev, *Member, IEEE*

***Abstract*—A deep metastability measurement scheme has been implemented on chip using digital circuits with 0.18 μm technology. Compared with previous off-chip implementations using analog circuits, the on-chip implementation allows integration of both the synchronizer circuits and the measurement method, and eliminates high-speed off-chip paths which are a source of inaccuracy. It also makes control at the picosecond level easier because of the inherent stability of digital integrating counters and digital delay lines. Our results show that the digital delay line used to adjust the data to clock times is controllable to an increment of 0.1 ps, and the input time distribution is 5.2 ps compared with 7.6 ps for the analog version. Because of the use of high and low counters, we can control the ratio of high to low outputs so that the actual input distribution can be measured to within better than 1 ps. The metastability time constant τ has been measured down to 10^{-17} s which corresponds to an mean time between failures (MTBF) of 100 seconds in an experimental time of 10 minutes and can be extended to a lower level by increasing the measurement time. Our results also show that a new synchronizer circuit designed for robustness to variation in V_{dd} performed at least as well as the Jamb Latch at all values of V_{dd}, and is more than 20% faster when V_{dd} was reduced by 25%.**

***Index Terms*—Metastability, mean time between failures (MTBF), NoC, synchronizer.**

I. INTRODUCTION

As THE SIZE of systems on chip (SoC) have increased, it has become difficult or impossible to accurately distribute a single global clock across the entire system [1], [2]. Future systems are likely therefore to consist of many independently or semi-independently clocked regions, with a need for synchronization of the data passing between them [3]. Consequently, there will be many more synchronizers whose reliability is crucial to the reliability of the entire system on chip. Synchronizer outputs are assumed to be stable after a fixed time interval, usually a clock cycle, therefore to know how reliable a synchronizer circuit actually is, we should measure how often the output changes after the clock cycle time. This is difficult because we are looking for failure rates as long as some months or years, therefore the reliability is projected from simulation results or measurements that only measure failures over a few hours. The mean time between failures (MTBF) is used to represent the reliability of synchronizers. Normally, we obtain an input time and output time relationship first and then the corresponding MTBF can be computed. Simulators such as SPICE [4] and MATLAB [5] have been used to estimate the MTBF of synchronizers, but SPICE is not sufficiently accurate for long metastability time prediction because some devices exhibit variations in τ at this point. Traditional off-chip measurement techniques do not allow MTBF to be measured beyond the point where any initial switching transient has died away sufficiently to make accurate projections for long term reliability. This is what we might call the deep metastability region.

To overcome the drawbacks of simulation methods and traditional measurement techniques, a new measurement method has been proposed [6] to enables the measurements to be carried out further into the deep metastability region. However, the previous work in [6] was implemented by using off-chip analog variable delay lines and an operational amplifier *RC* integrator as components in a delay locked loop. Because of the instability of the analog components, it is not easy to control the operation of the delay lines or to characterize the actual synchronizer input stimuli time distribution. On-chip implementation of the deep metastability measurement using digital variable delay elements and counters allows integration of both the synchronizer circuits and the measurement method, and eliminates high-speed off-chip paths which are a source of inaccuracy. It also makes control at the picosecond level easier because of the inherent stability of digital integrating counters and digital delay lines.

This paper shows how the deep metastability measurement method can be implemented on chip using digital variable delay elements and counters. Section II reviews the metastability theory and metastability measurement methods which include the traditional measurement method and the deep metastability measurement method. In Section III, the details of the on-chip implementation are described. In Section IV, the measurement results are shown and comparison is made between the simulation results and the measurement results demonstrating that the on-chip measurement circuit works as expected.

II. MEASURING METASTABILITY

A. Arbiters and Synchronizers

Metastability occurs in both arbiters [7] and synchronizers. Arbiters are used to allow multiple processes to access to a single shared resource while synchronizers are used to retime the data passing between different clock regions in networks on chip (NOC), though they have similar structure. A basic arbiter is made from two cross-coupled NAND gates as in [7].

Manuscript received February 6, 2007; revised September 20, 2007. This work was supported in part by Intel Corporation, the U.K. EPSRC Research Grant EP/C007298/1 and equipment support from Agilent Technologies.

J. Zhou, D. J. Kinniment, G. Russell, and A. V. Yakovlev are with the University of Newcastle, Newcastle upon Tyne, U.K. NE1 7RU (e-mail: jun.zhou@ncl.ac.uk).

C. E. Dike is with the Intel Corporation, Hillsboro, OR 97124 USA (e-mail: charles.e.dike@intel.com).

Digital Object Identifier 10.1109/JSSC.2007.913160

Metastability happens if the two requests arrive very close together. A D-type latch which is the basic element of a conventional master–slave synchronizer, may also be made from two cross-coupled NAND gates. Similarly, metastability happens if the data and clock arrive very close together. For synchronizers, we hope that the metastability can resolve as fast as possible. The disadvantage of using NAND gates is that they are relatively slow when compared with inverters, and thus metastability is slow to resolve in a latch made from NAND gates. In order to achieve a faster resolution time, cross-coupled inverters with set and reset transistors have been used in the Jamb latch [4] circuit which is commonly used as a synchronizer because of its good performance. The details of the Jamb latch will be described in Section III-B.

B. Metastability Theory

In a synchronizer, if the data input goes high sufficiently far in advance of the clock edge, the synchronizer output will always go high and if it is significantly after the clock it will always go low. If the two edges are close enough, the high or low outcome is affected by circuit noise and is non-deterministic. Here we will define the separation between data and clock which gives an exactly equal probability of a high or low outcome as the balance point. In the absence of noise, an input exactly at the balance point would take an infinite time to resolve. The synchronizer response from metastability is usually exponential [6]. Thus, for inputs a time $\Delta t_{\rm in}$ away from the balance point, where $\Delta t_{\rm in}$ is less than the metastability window, $\mathbf{T}_{\rm w}$, the relationship between resolution time t and $\Delta t_{\rm in}$ is given by

$$t = \tau \cdot \ln \frac{T_w}{\Delta t_{\rm in}} \tag{1}$$

Customarily, t is measured from the normal propagation delay, $\mathrm{T_d}$. A small change in the input time will therefore cause a change in the output time:

$$dt = -d\Delta t_{\rm in} \cdot \frac{\tau.}{\Delta t_{\rm in}}. \tag{2}$$

If the resolution time is longer than the time allowed for synchronization, the synchronizer may fail as a result of an undefined output level. The number of failure events caused by the data edge occurring less than $\Delta t_{\rm in}$ from the balance point in a total time T depends on the clock rate and the data rate and is given by

$$\mathrm{Number}_{\rm failure\_events} = T \cdot \Delta t_{\rm in} \cdot f_c f_d. \tag{3}$$

From (3), the mean time between each failure event is

$$\mathrm{MTBF} = \frac{1}{\Delta t_{\rm in} \cdot f_c f_d}. \tag{4}$$

Using (1) and (4), MTBF can also be expressed in terms of known system and circuit parameters:

$$\mathrm{MTBF} = \frac{e^{\frac{t}{\tau}}}{T_w f_c f_d}. \tag{5}$$

The values of the circuit parameters are found in practice by plotting ln(MTBF) against t. From (5), the slope of this graph is $1/\tau$. T_w can be found from projecting the graph back to the ln(MTBF) axis where $t = 0$, since at that point

$$T_w = \frac{1}{\mathrm{MTBF_0}.f_c f_d}. \tag{2}$$

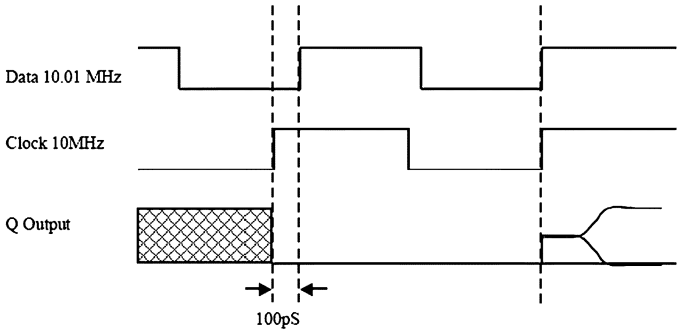

Fig. 1. Traditional measurement method using two oscillators.

C. Deep Metastability Measurement Method

As shown in Fig. 1, the traditional measurement method uses two oscillators with a similar frequency to provide data and clock for the synchronizer. In this way the overlap time between data and clock is evenly distributed over a range of times equal to the difference between the clock and data periods. The drawback of this method is that there are relatively few deep metastability events as these events are produced by very small overlap times and they have very small probability. This makes it difficult to measure τ in the deep metastability region. Measurements or simulation of the early deterministic region can give a falsely optimistic result [4], [6].

Recently, a new measurement method called deep metastability measurement was introduced by Kinniment *et al.* in [6].

As can be seen in Fig. 2, the deep metastability measurement method uses only one oscillator and two delay lines to provide data and clock for the synchronizer. One delay line is fixed and the other one is variable. The output of the synchronizer is used to control the variable delay line so that the loop settles at a point where the number of high output events is the same as the number of low output events. We call this the balance point. When the loop has settled the distribution of data input times is small, and close to a normal distribution. In this way the synchronizer is forced into metastability on almost every clock cycle and more deep metastability events can be observed. The measurement can then be conducted in the deep metastability region, which gives a more reliable result for the synchronizer performance. Our measurement is made by comparing the distribution of input events with the distribution of output events. In this experiment we count the number of input events where the data is ahead of the balance point by times between 0 and $t_{\rm in}$ and then count the output events between infinity and a time $t_{\rm out}$. The value of $t_{\rm out}$ that gives the same output count as the input count given by $t_{\rm in}$ establishes the correspondence between $t_{\rm in}$ and $t_{\rm out}$. The method is described in detail in [6], and allows us to construct the input time versus output time from the input distribution and output distributions recorded by the oscilloscope. One problem is that the input time distribution is obscured by measurement noise. In [6], it is also shown how this noise can

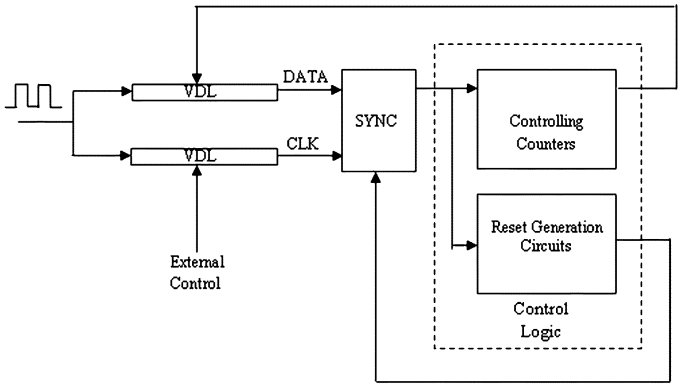

Fig. 2. Deep metastability measurement.

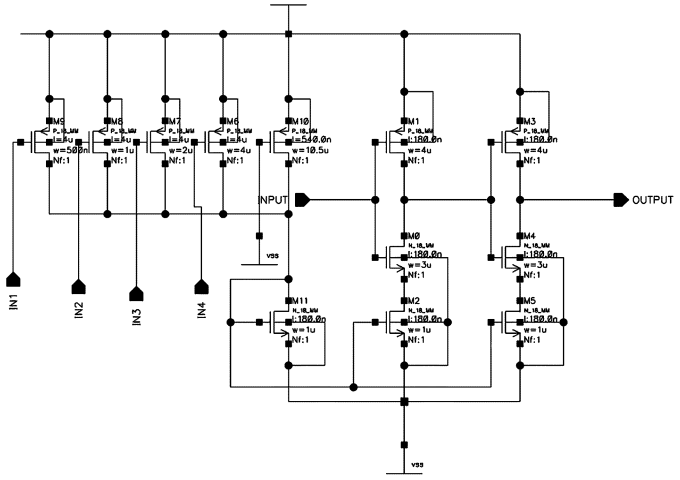

Fig. 3. Variable delay line.

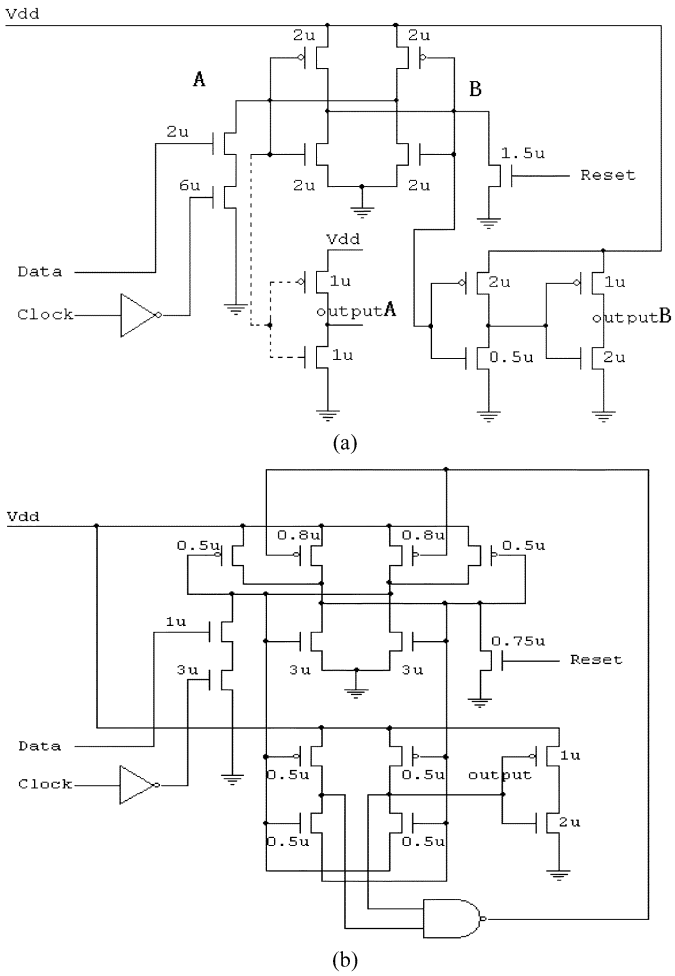

Fig. 4. Devices under test.

be removed by adjusting the ratio of high output events and low output events.

III. On-Chip Implementation

As shown in Fig. 2, the on-chip measurement circuit is composed of three parts: variable delay lines (VDLs), devices under test (DUTs), and control logic. Together they form a feedback loop to adjust the input time of the synchronizer. The details of each part are described separately below.

A. Variable Delay Lines

There are two VDLs in the on-chip measurement circuit. One is used to vary the delay in the DATA path and the other is used to vary the delay in the CLK path. The VDL in the DATA path is controlled by a 16-bit counter on the chip. The VDL in the CLK path is controlled externally. Fig. 3 shows the architecture of the VDL. The VDL is based on a current mirror structure and has been proposed by Maymandi-Nejad and Sachdev [8]. Compared with traditional VDL, its advantage is that the delay behavior is monotonic.

As can be seen in Fig. 3, a current starved buffer, M0–M5, is the main element. The current through this buffer is controlled by a current mirror circuit composed of transistors M2 and M11. Because of the current mirror structure, the controlling transistors do not have to be placed below the main N-type transistor, so the charge sharing effect is reduced and the delay behavior of the VDL is monotonic [8]. An appropriate current through M11 can be adjusted by turning on controlling transistors M6–M9, while transistor M10 is always on as a base transistor. Here the W/L of controlling transistors M6–M9 are arranged in a binary fashion so that the number of controlling transistors can be minimized. In order to get a very small incremental delay and large delay range, each VDL includes 4 cascaded stages similar to Fig. 3. The maximum delay of each stage is different and is designed to achieve an incremental delay of 0.1 ps and a delay range of 0–500 ps.

B. Devices Under Test

Three different synchronizers have been incorporated on the chip for measurement and comparison, namely Jamb Latch A, Jamb Latch B and a robust synchronizer. A single latch from each synchronizer is shown in Fig. 4.

As can be seen in Fig. 4(a), Jamb Latch A and Jamb Latch B have the same structure but different output configuration (output A for Jamb Latch A and output B for Jamb Latch B). They have been reported to have different characteristics in the deterministic region [9] and we want to confirm the effect. Jamb latch is commonly used as a synchronizer because of its relatively good performance. Here, the flip-flop is reset by pulling

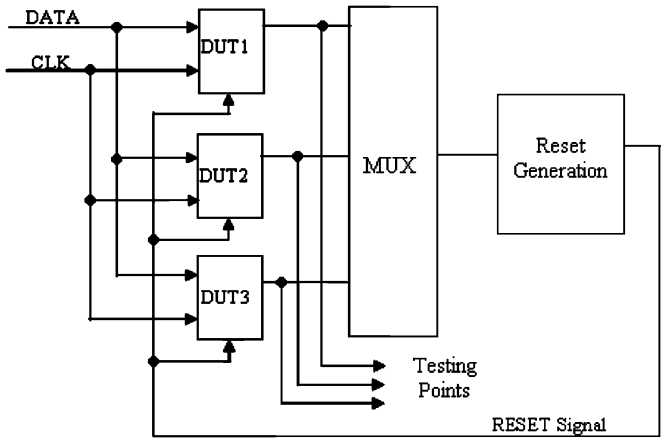

Fig. 5. Multiplexer circuit for DUTs.

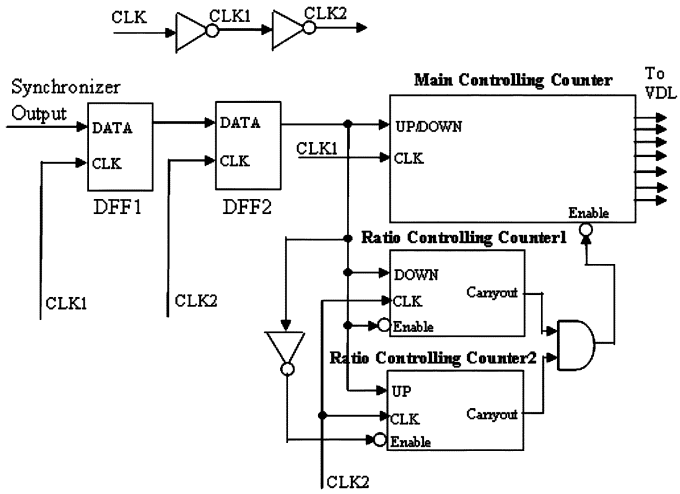

Fig. 6. Controlling counters.

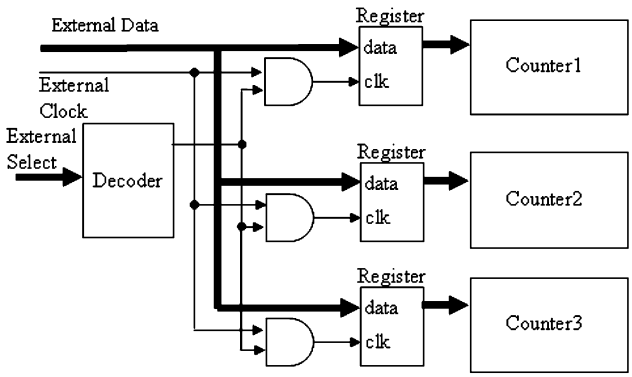

Fig. 7. Loading circuit for controlling counters.

node B to ground, and then set if data is high and clock is low, by pulling node A to ground. Metastability occurs if the overlap of data and clock is at a critical value which causes node A to be pulled down, and node B to be pulled up to the metastable level. A problem of the Jamb Latch is that metastability time constant τ is sensitive to V_{dd} and temperature variation.

To overcome this problem, a robust synchronizer shown in Fig. 4(b) has been proposed [10]. By increasing the latch current only during metastability, the circuit can be made relatively insensitive to V_{dd} and temperature variation without significantly increasing the power. The structure of this circuit is shown in Fig. 4(b). The power dissipation in the robust synchronizer is greater than in the conventional Jamb latch during metastability, but because metastability only lasts on average for a very small proportion of the clock cycle time, typically less than 100 ps of a 1 nS clock cycle, the average power dissipation is comparable. More importantly, the latch current which determines the resolution speed of metastability is greatly increased during metastability in the robust synchronizer. This allows the same power budget to be used in a more effective way, and leads to a faster resolution time especially at low V_{dd}. It is shown in [10] that with similar power consumption the robust synchronizer is faster than Jamb latch at 1.8 V and becomes much faster as V_{dd} decreases. Each synchronizer described above is made from two latches similar to Fig. 4. As shown in Fig. 5, all the synchronizers share the same DATA, CLK and RESET signal. There is a multiplexer for selecting different synchronizers on the chip for measurement. When one of the synchronizers is selected, its output goes through the multiplexer and the reset generation circuits to generate the RESET signal for all the synchronizers. The multiplexer here is used to make the testing circuitry identical for all designs, but it can introduce a relatively large delay. In order to obtain the accurate output time, measurement points are placed before the multiplexer.

C. Control Logic

The control logic is divided into two parts: controlling counters and reset generation circuits.

The controlling counters include one 16-bit main controlling counter and two 8-bit ratio controlling counters. The outputs of the main controlling counter are used to control the VDL in the DATA path. The two ratio controlling counters are used to adjust the ratio of high output events and low output events. Fig. 6 shows the details of the controlling counters.

As shown in Fig. 6, two D-flip-flops are used to detect output events from the synchronizer. The ratio controlling counter 1 counts only when there is a low output event from the synchronizer and it always counts down. The ratio controlling counter 2 counts only when there is high output event from the synchronizer and it always counts up. The main counter counts only when there is a carryout from either ratio controlling counter 1 or ratio controlling counter 2 and it can count up for counter 1 or down for counter 2, depending on the output event detected.

All the controlling counters must be loaded with initial values at the beginning of test. Because of the limitation of the number of pins, a multiplexed loading circuit is used for loading different controlling counters. Fig. 7 shows the architecture of the multiplexer circuit.

As shown in Fig. 7, some registers are used for holding the loaded values for different controlling counters. The clock signals of the registers are generated by ANDing the external clock and the outputs of a decoder which is controlled by an external select signal. If one controlling counter is selected, the corresponding output of the decoder goes high and thus the external clock can go through the AND gate to latch the data into the registers of the counter.

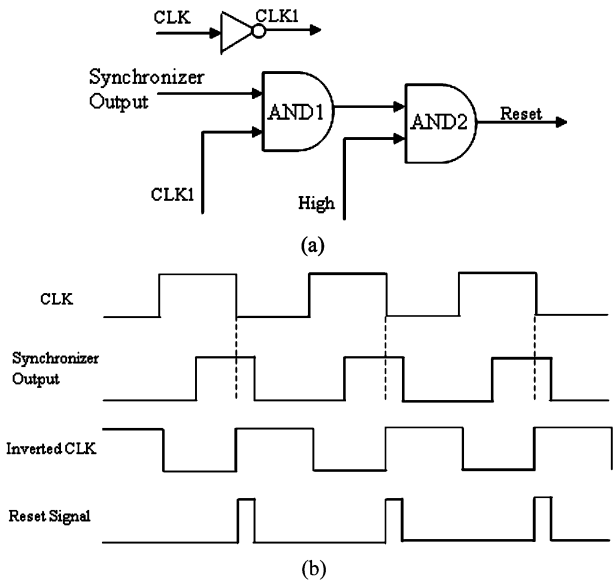

Fig. 8. Generation of RESET signal.

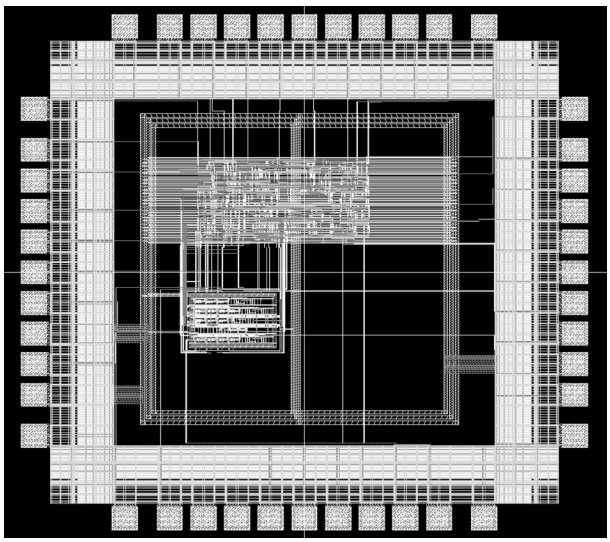

Fig. 9. Layout of on-chip measurement circuit.

To ensure that the measurements we make are consistent, we always reset the devices under test before the data changes. Fig. 8(a) shows RESET generation circuits. The RESET signal is generated by ANDing the synchronizer output and the back edge of the clock. In order to hold the RESET signal for some time, another AND gate is used to add the delay. Fig. 8(b) shows the generation of the RESET signal.

D. Layout of On-Chip Measurement Circuit

The on-chip measurement circuit has been fabricated using UMC 0.18 μm technology and its layout is shown in Fig. 9. The control circuits are laid out using a standard cell library and occupy the larger block in Fig. 9. The variable delay lines and the devices under test are in the smaller custom designed block. The power supply of the devices under test can be varied separately from all other power supplies.

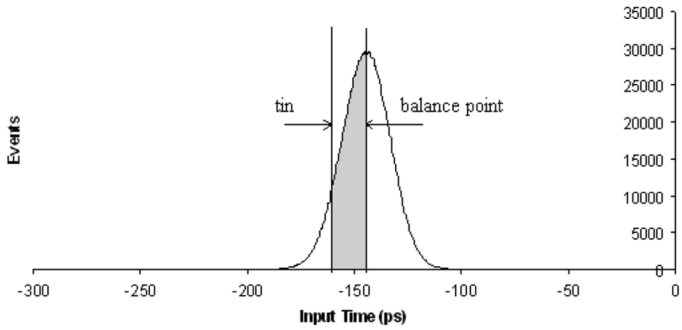

Fig. 10. Input histogram of Jamb latch A.

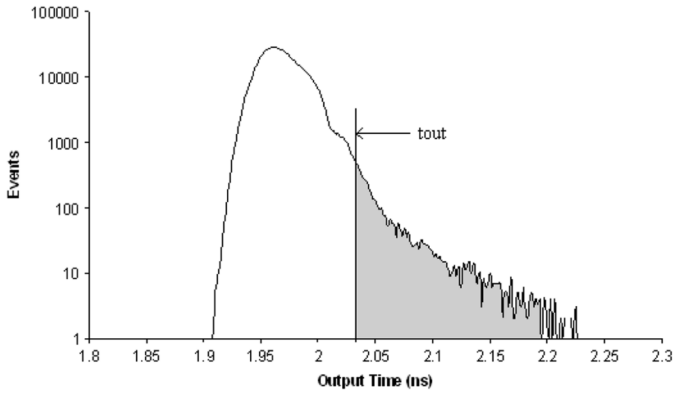

Fig. 11. Output histogram of Jamb Latch A.

IV. Results

A. Input Histogram

Fig. 10 shows the input histogram for the Jamb Latch A synchronizer at a V_{dd} of 1.8 V. The clock is used as trigger to observe the data. As shown in Fig. 10, the data time is held within a very small range around the balance point of the synchronizer demonstrating that the delay locked loop is stable. The standard deviation of the distribution of data including oscilloscope measurement noise is about 11 ps.

B. Output Histogram

Fig. 11 shows the output histogram of Jamb Latch A at a V_{dd} of 1.8 V. Again, the clock is used as trigger to observe the output of the synchronizer. For this experiment, we used a ratio of 1:1 between high and low outputs. To achieve this, the values of the two ratio controlling counters were both set to 1. By using a digital histogramming oscilloscope the total number of high output events and low output events can be recorded and displayed. We recorded the output events over a period of time and plotted the total number of high output events against that of low output events in Fig. 12, where it can be seen that the ratio of high to low output events is approximately 1:1, which demonstrates that the ratio is held very constant over the time of the measurement.

By setting different values for the two ratio controlling counters the proportion of high to low output events can be changed and the median value of the input distribution can be shifted. The relationship between the shift and the percentage of high output events are plotted later in Fig. 13, which also demonstrates that

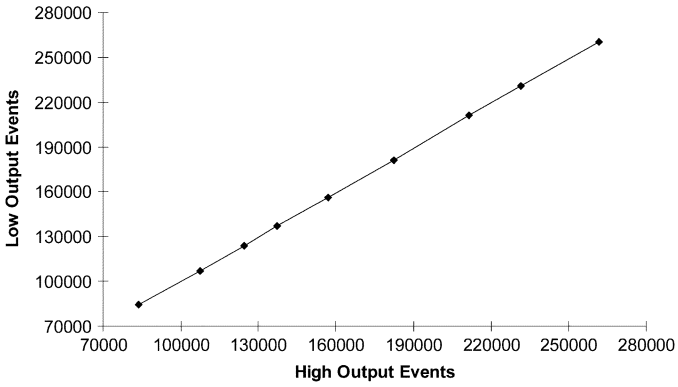

Fig. 12. High output events versus low output events.

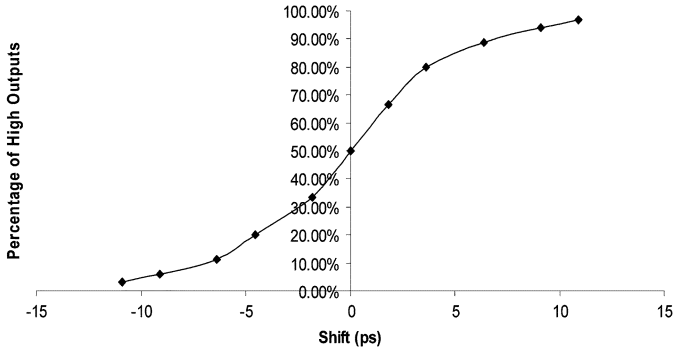

Fig. 13. Measurement of actual input distribution.

the measurement of input events can be made to an accuracy of around 1 ps.

Finally, the output histogram in Fig. 11 is very similar to the one given in Fig. 6 in [4], where the window of input events has been shrunk to 10 ps, which demonstrates that the on-chip measurement correlates to previously published results.

C. Corrected Input Histogram

The input histogram recorded on the oscilloscope also contains the measurement noise from the oscilloscope itself which is typically 9.2 ps according to the specification of the oscilloscope. Because of this relatively large measurement noise component we cannot reliably use Fig. 10 to assign input times to output times. We eliminate the measurement noise and find the real density of inputs around the balance point by altering the ratio of high to low outcomes, and hence shifting the balance point. If we measure the time shift and we know the average number of trajectories that have changed from high to low, then we know the number of inputs within the time represented by the shift despite any oscilloscope noise. In this way we can plot the actual average number of trajectories against time. The correction of the input histogram can be done by adjusting the values of the two sub-counters to produce different proportions of high and low outputs from the synchronizer. This causes the balance point which is the peak of the input histogram to shift in order to achieve the proportions set by the sub-counters, and the proportion of high outputs can be plotted against the shift as in the method described in [6]. The shifts required to give different probabilities of high outputs are plotted in Fig. 13.

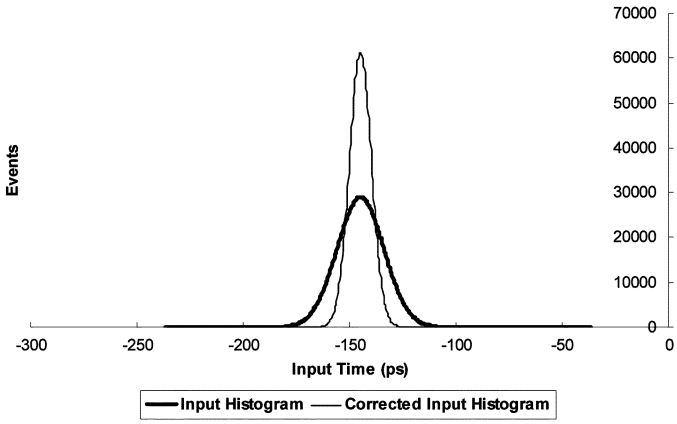

Fig. 14. Corrected input histogram of Jamb latch A.

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 Fig. 13 represents the cumulative probability of a high output for a random input time deviation of 5.2 ps, so we can conclude that the actual distribution has a deviation of this value. Furthermore, the observed input deviation of 11 ps is close to the square root of the sum of the squares of 9.2 ps and 5.5 ps. The corrected input histogram is shown in Fig. 14.

As can be seen from Fig. 14, the corrected input histogram has a standard deviation of about 5.2 ps. This result demonstrates that the feedback loop holds the delay difference between data and clock to within very close limits, and that the distribution of the delay difference is nearly random.

D. Input Time Versus Output Time

After getting the input and output histogram of the synchronizer, the input time and output time relationship can be plotted using the mapping method described in [6].

Fig. 15 shows the measured input time versus output time for Jamb A, Jamb B and robust synchronizer at 1.8 V. In order to avoid the problem of long connections on the chip and compare the metastability characteristics of the three synchronizers, for all measurements, output times are computed relative to the time when the largest number of events are recorded. The reciprocal of the slope of the curves in Fig. 15 represents the metastability time constant τ. As can be seen from the figure outputs with input time down to 10^{-17} s can be plotted, which corresponds to an MTBF of 100 seconds given that both the clock frequency and data rates are 30 MHz, thus τ can be measured in the deep metastability region.

Fig. 16 shows the simulation results of input time versus output time for the three tested synchronizers. Again output times are made start from zero where the normal propagation delay is. By comparing Figs. 15 and 16 it can be found that the slope of the curves between 10^{-11} s and 10^{-14} s in the two figures are similar. In the simulation results only the input time down to 10^{-14} s are plotted because the minimum simulation step used is 10^{-14} s. Below this value the accuracy of the simulation results is not reliable. Figs. 15 and 16 show that the robust synchronizer can be more reliable than both Jamb latch

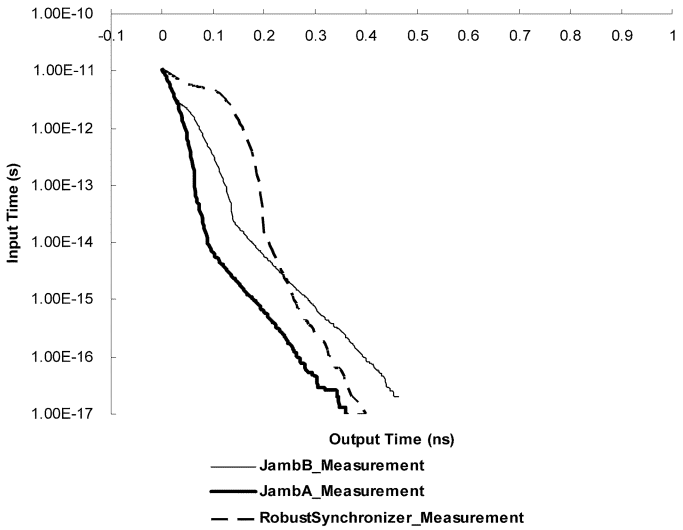

Fig. 15. Measured input time (s) versus output time (ns).

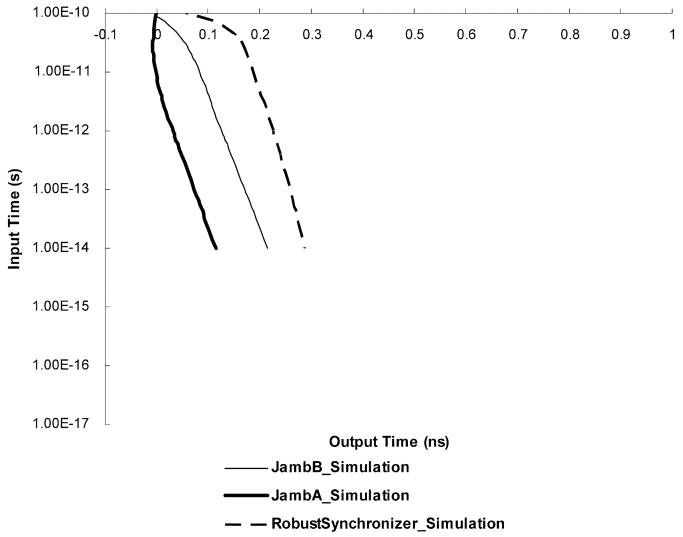

Fig. 16. Simulated input time (s) versus output time (ns).

A and B, that is, it has smaller τ, which leads to faster resolution of metastability and thus fewer very long output times. This is because the latch current which determines the resolution speed of metastability is greater in the robust synchronizer than it is in the Jamb latch during metastability. This also makes the robust synchronizer much less sensitive to V_{dd} and temperature variation. Figs. 15 and 16 also show that the three tested synchronizers are slower in the deep metastability region than they are in the deterministic region, which indicates that the prediction of MTBF based on early simulation can lead to falsely optimistic results.

E. Tau Versus V_{dd}

In order to compare the sensitivity of the Jamb latch and the robust synchronizer to V_{dd} variation, we measured the value of τ for different values of V_{dd}. Table I shows the measurement results. Here the simulation results are also shown for comparison.

TABLE I
TAU VERSUS V_{dd} FOR JAMB B AND ROBUST SYNCHRONIZER

	Measurement Results(ps)				Simulation Results(ps)	
Vdd (v)	Jamb Latch B		Robust Synchronizer		Jamb Latch B	Robust Synchronizer
	$>10^{-14}$	$<10^{-14}$	$>10^{-14}$	$<10^{-14}$		
1.8	19.44	35.55	15.27	34.92	18.99	14.69
1.7	21.75	37.29	16.53	35.76	20.36	15.36
1.6	25.64	40.93	19.38	38.25	22.24	16.19
1.5	28.77	52.36	20.29	43.07	24.99	17.23
1.4	36.22	66.17	23.75	50.36	29.31	18.59
1.35	45.43	75.35	28.51	58.19	36.85	20.39

Table I shows that for both synchronizers, τ increases with V_{dd} decreasing. This is because the latch current which determines the resolution speed of metastability decreases with V_{dd} decreasing. As can be seen from Table I, the robust synchronizer circuit performed at least as well as the Jamb Latch at all values of V_{dd}, and was more than 20% faster when V_{dd} was reduced by 25%.

Table I also shows that the values of τ match the simulation results well when the input time is above 10^{-14} s which is the minimum simulation step used. Below 10^{-14} s the simulation results are not reliable, and Table I shows that the measured τ below 10^{-14} s is greater than above 10^{-14} s, which means that the tested devices are slower in the deep metastability region than they are in the deterministic region. For this reason, we cannot rely on simulation to predict MTBF at realistic synchronization times, and it is necessary to check the value of τ in deep metastability with accurate measurement.

V. CONCLUSION

A deep metastability measurement scheme has been implemented on chip by using digital circuits and fabricated with UMC 0.18 μm technology. Compared with the previous off-chip implementation using analog circuits, the on-chip implementation using digital circuits allows integration of both the synchronizer circuits and the measurement method, and eliminates high-speed off-chip paths which are a source of inaccuracy. It also makes control at the picosecond level easier because of the inherent stability of digital integrating counters and digital delay lines.

Our results show that the measurement method is stable and reliable. The digital delay line was controllable to an increment of 0.1 ps, and the input time distribution was 5.2 ps compared with 7.6 ps for the analog version. Through the use of high and low counters the actual input distribution could be measured to within better than 1 ps. τ was measured down to 10^{-17} s within an experimental time of 10 minutes, corresponding to an MTBF of 100 seconds. By extending the measurement time and filtering early responses in a future version of this chip, this could be extended to 1 000 000 seconds or approximately 10 days MTBF. The responses from the tested devices were shown to correspond with simulation down to 10^{-14} s, but values of τ for input times below that reachable by simulation were shown to be greater than simulation, which means that tested device is

slower in the deep metastability region than it is in the deterministic region. For this reason, we cannot rely on simulation to predict MTBF at realistic synchronization times, and it is necessary to check the value of τ in deep metastability with accurate measurement. We have also made a comparison between the Jamb latch and the robust synchronizer at different V_{dd}s, the results show that the robust synchronizer circuit performed at least as well as the Jamb Latch at all values of V_{dd}, and was more than 20% faster when V_{dd} was reduced by 25%.

REFERENCES

[1] N. A. Kurd, J. S. Barkatullah, R. O. Dizon, T. D. Fletcher, and P. D. Madland, "Multi-GHz clocking schemes for Intel Pentium 4 microprocessors," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2001, pp. 404–405.

[2] S. Tam, S. Rusu, U. N. Desai, R. Kim, J. Zhang, and I. Young, "Clock generation and distribution for the first IA-64 microprocessor," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1545–1552, Nov. 2000.

[3] K. Y. Yun and A. E. Dooply, "Pausible clocking based heterogeneous systems," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 7, no. 4, pp. 482–487, Dec. 1996.

[4] C. Dike and E. Burton, "Miller and noise effects in a synchronizing flip-flop," *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 849–855, Jun. 1999.

[5] S. Yang and M. Greenstreet, "Computing synchronizer failure probabilities," in *Proc. DATE'07*, Apr. 2007, pp. 1–6.

[6] D. Kinniment, K. Heron, and G. Russell, "Measuring deep metastability," in *Proc. ASYNC'06*, Grenoble, France, Mar. 2007, pp. 2–11.

[7] C. Mead and L. Conway, *Introduction to VLSI Systems*. Reading, MA: Addison-Wesley, 1979.

[8] M. Maymandi-Nejad and M. Sachdev, "A digitally programmable delay element: Design and analysis," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 12, no. 10, pp. 1126–1126, Oct. 2004.

[9] D. J. Kinniment, A. Bystrov, and A. V. Yakovlev, "Synchronization circuit performance," *IEEE J. Solid-State Circuits*, vol. 37, no. 2, pp. 202–209, Feb. 2002.

[10] J. Zhou, D. J. Kinniment, G. Russell, and A. Yakovlev, "A robust synchronizer circuit," in *Proc. ISVLSI'06*, Mar. 2006, pp. 442–443.

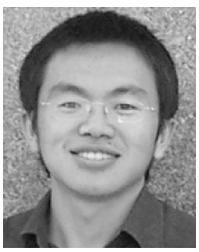

Jun Zhou (S'07) received the dual B.Sc. degree in communication engineering and in microelectronics from the University of Electronic Science and Technology of China in 2004. He is currently working toward the Ph.D. degree at the University of Newcastle upon Tyne, Newcastle, U.K.

His research interests include IC design and networks on chip.

David J. Kinniment (M'67–LM'06) received the M.Sc. degree in electrical engineering in 1963 and the Ph.D. degree in computer science in 1968 from Manchester University, Manchester, U.K., where he was part of the teams involved in the design of the asynchronous ATLAS and MU5 computers.

From 1964 to 1979, he was Assistant Lecturer, Lecturer, and Senior Lecturer in the Computer Science Department at Manchester University. In 1979, he was appointed to the Chair of Electronics at Newcastle University, Newcastle, U.K., and was Head of the Electrical and Electronic Engineering Department from 1982 to 1990, and 1996 to 1998. He is currently an Emeritus Professor in the University of Newcastle, where he works with the Microelectronic Systems Design research group at the School of Electrical, Electronic and Computer Engineering. His current interests include IC design, and asynchronous networks on chip.

Prof. Kinniment is a member of the IET and a Life Member of the IEEE.

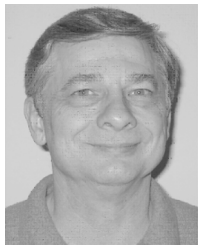

Charles E. Dike (M'95) received the B.S.E.E. and M.S.E.E. degrees from Brigham Young University, Provo, UT, in 1977 and 1984, respectively.

He was an Integrated Circuit Design Engineer with Signetics Corporation, Orem, UT, for 11 years. While there, he developed expertise in metastability theory and asynchronous arbitration. He worked in bipolar, BiCMOS, and CMOS technologies. He joined Intel Corporation, Beaverton, OR, in 1992 and presently works for the Design Automation and Collateral Department for the Ultra-Mobile Group in Hillsboro, OR, and is interested in asynchronous logic, ideas related to synchronization, metastability and networks on chip. He has received 14 patents and is the author or coauthor of several articles.

Gordon Russell (M'84) received the B.Sc. and Ph.D. degrees in electrical and electronic engineering from the University of Strathclyde, Glasgow, U.K., in 1970 and 1977, respectively.

From 1975 to 1977, he was Research Fellow in the Department of Computer Science and a Post Doctoral Fellow in the Department of Electrical Engineering at the University of Edinburgh, involved in the design of CAD tools for LSI circuit design and the computer modeling of charge coupled devices, respectively. In 1979, he joined the academic staff of the School of Electrical, Electronic and Computer Engineering, at the University of Newcastle upon Tyne, U.K. His research interests include concurrent error detection (CED) techniques, asynchronous design and test, BIST for SoC, on-chip time measurement circuits, and development of dynamic burn-in systems for circuits used in hazardous environments. He has written many technical articles on testing and design for testability and has been involved as co-author or co-editor of several books related to testing and other aspects of CAD for VLSI. He has given series of invited lectures, worldwide, on testing and design for testability.

Dr. Russell is a member of the IEE and the IEEE.

Alexandre V. (Alex) Yakovlev (M'98) was born in Russia in 1956. He received the D.Sc. degree from Newcastle University, Newcastle, U.K., in 2006. He received the M.Sc. and Ph.D. degrees from St. Petersburg Electrical Engineering Institute in 1979 and 1982, respectively, where he worked in the area of asynchronous and concurrent systems since 1980, and in the period between 1982 and 1990 held positions of Assistant and Associate Professor in the Computing Science Department. He first visited Newcastle as a Postdoc British Council Scholar in 1984/1985 for research in VLSI and design automation.

After coming back to Britain in 1990, he worked for one year at University of Glamorgan. Since 1991, he has been at Newcastle University, where he worked as a Lecturer, Reader and Professor in the Computing Science Department until 2002, and is now heading the Microelectronic Systems Design research group in the School of Electrical, Electronic and Computer Engineering. His current interests and publications are in the field of modelling and design of asynchronous, concurrent, real-time and dependable systems on a chip. He has published four monographs and more than 200 papers in academic journals and conferences, and has managed over 20 research contracts. He has chaired the program committees of several international conferences, and is currently a chairman of the steering committee of the Conference on Application of Concurrency to System Design.

Dr. Yakovlev is a member of the IEEE and the IEE.