

25.8 A Near-Threshold-Voltage Network-on-Chip with a Metastability Error Detection and Correction Technique for Supporting a Quad-Voltage/Frequency-Domain Ultra-Low-Power System-on-a-Chip

Chuxiong Lin¹, Weifeng He¹, Yanan Sun¹, Bingxi Pei¹, Zhigang Mao¹, Mingoo Seok²

¹Shanghai Jiao Tong University, Shanghai, China

²Columbia University, New York, NY

Emerging embedded systems, such as autonomous robots/vehicles, demand a new system-on-a-chip (SoC) that is ultra-low power (mW or even sub-mW level) but highly robust. Such an SoC typically integrates heterogeneous building blocks for supporting a range of features, each ideally operating in an independent voltage and frequency (V/F) domain [1]. In such an architecture, a network-on-chip (NoC) has played a key role to enable high-speed and energy-efficient networking. However, it is increasingly challenging to meet a robustness target since each V/F domain uses a significantly different voltage, e.g., from nominal 1V to near-threshold voltage (NTV), and clock frequency, e.g., from hundreds of MHz to sub-MHz. Furthermore, any two clocks may have uncertain and time-varying phase and frequency relationships. These properties significantly worsen robustness, particularly metastability, in an NoC.

The synchronizer is an effective technique to counter metastability, but it neither guarantees the correctness of data nor provides long-term metastability mitigation. [2-3] investigated metastability characterization and on-chip measurement of synchronizers, but focused on test structures. [4] proposed metastability-immune circuits based on timing-error detection and correction (EDAC) [5] but in a single-clock system. [6] demonstrated an NoC that detects and corrects timing errors (yet not those arising from metastability) in the ultra-low voltage region. [7] presented a self-synchronization technique for a two-clock system, but only considered an unknown phase relationship of the same clock frequency.

In this work, we experimentally demonstrate a globally-asynchronous-locally-synchronous (GALS) NoC test chip in a 65nm low-power (LP) process. The NoC contains four independent V/F domains from 0.4V/7.3MHz to 1V/175MHz. We propose a metastability error detection and correction (MEDAC) technique and employ it in the dual-clock FIFOs of the routers in the NoC. It supports any clock-phase difference with several clock frequency ratios from 0.25 to 4. It can detect and correct the necessary conditions of metastability, i.e. data arrival in the metastability window of a receiving clock, and mitigate the chance of entering such conditions in future. The proposed technique significantly reduces the risk of triggering the necessary conditions of metastability by one to three orders of magnitude.

Figure 25.8.1 shows the proposed 2-by-2 GALS NoC architecture, consisting of four routers (R_1 to R_4) and four processing elements (PE_1 to PE_4). The PE and the router in the same V/F domain share a local ring-oscillator-based clock generator. A PE generates a data packet and checks the integrity of a received packet. If a PE missed a packet, it sends a request packet to the source PE, which then retransmits the missing packet. The router has two channels, each of which supports bilateral data transfer to/from the neighboring V/F domain. Each channel has voltage-level converters and two dual-port asynchronous FIFOs (Fig. 25.8.1 top).

In the FIFO, we replace the conventional synchronizer with our MEDAC unit. The unit consists of a metastability detector and corrector. The detector has the main and the shadow synchronizers (Fig. 25.8.2 top). The main synchronizer receives Rx_clk_d , which is the delayed version of the shadow synchronizer clock Rx_clk by the interval window (W). Suppose W_M is the metastability window of a flip-flop (Fig. 25.8.2 bottom left), a bit of Tx_data arrives in one of the five regions in time (Fig. 25.8.2 bottom right). If a bit arrives far from both the Rx_clk and Rx_clk_d edges (i.e. Regions 1 and 5), all the flip-flops in the synchronizers stably capture the correct bit. However, if a bit arrives near either of the clock edges (Regions 2 and 4), the output of the first flip-flop of the synchronizers (R_{m0} , R_{s0}) may enter a metastable state. The bits on R_{m0} and R_{s0} may or may not escape from metastability in one clock cycle – this is difficult to detect. Instead, the MEDAC unit compares the outputs of two synchronizers and flags this as a metastability condition. Note that if W is too large, Region 3 is created (Fig. 25.8.2 bottom right). If a bit arrives in Region 3, the detector makes a false alarm. To minimize this, we calibrate W to roughly equal W_M . The detection window (W_D), then, becomes $2 \cdot W_M$ ($= W_M + W$).

Upon detection of metastability, the MEDAC unit asserts *Stall* to notify the router for data retransmission. In addition, it takes preventive action to reduce the probability of future metastability by modulating the phase difference (ΔP) of Tx_clk and Rx_clk . If the ratio of Rx_clk to Tx_clk ($k = Rx_clk/Tx_clk$) is rational (i.e. k is an integer or the inverse of an integer) and static, we can completely avoid the metastability condition by modulating ΔP and sufficiently separating the clock edges. However, if k is irrational and parasitically varying, ΔP would change over time and eventually bring two edges too close, causing metastability. Furthermore, the metastability condition may last for several continuous cycles in some cases. For example, if $k \approx 1$, the phase difference of two clocks may be smaller than W_M for several cycles, making data transmission between two V/F domains less reliable.

We therefore adjust the phase to maximally extend the time until the next metastability condition. The phase difference normalized to Rx_clk 's period, i.e. $\Delta P_{norm} = \Delta P/T(Rx_clk)$, can be classified to two classes, *Class-A*: infrequently and *Class-B*: frequently entering the metastability condition. Fig. 25.8.3 (left) shows the exemplary case with $k=0.4$. Our goal is to tune the phase of Rx_clk to one of the *Class-A* ΔP_{norm} s. To do so, at design time, we estimated all possible phase shift (PS) values for each k from 1/4 to 4 (Fig. 25.8.3 right). We found that there exist three PS values, 180, 90, and 54 degrees, which can adjust ΔP_{norm} to be in class *Class-A* for all considered k values. We define this minimum set of PSs to be *Class-A* PS (APS). Based on the above idea, we designed the metastability corrector in the MEDAC unit (Fig. 25.8.4 top left). The mean-time-between-metastability (MTBM) timer counts the number of cycles between two adjacent metastability occurrences and estimates if the current ΔP_{norm} is *Class-A* or *-B*. If it is *Class-B*, the phase shifter modulates the three delay lines to change ΔP_{norm} to be *Class-A*. We designed the variable delay lines to support the APS for the target k values.

We prototyped the NoC test chip (Fig. 25.8.7 top). Each router takes 0.44mm². The area overhead of the MEDAC is 4.4%. We mainly measured the communication between two V/F domains. The first domain uses $V_1=0.5V$ and $F_1=5$ MHz. We sweep the supply of the second domain (V_2) from 0.5V to 1V and the clock frequency (F_2) from 1.25MHz to 20MHz. Over the V/F conditions, the MEDAC can reliably detect the necessary conditions of metastability. At $V_2=0.5V$, the MEDAC-enabled router enters metastability by up to 1600x less often. At $V_2=1V$, the reduction becomes smaller, since metastability reduces exponentially with increasing supply voltage. The mitigation technique reduces the retransmission rate, thereby improving throughput and energy efficiency. At $V_2=0.5V$, the MEDAC-enabled router achieves up to 19.5% higher throughput (Fig. 25.8.5 top right) and 16.1% less energy consumption (Fig. 25.8.5 bottom). The test chip also includes the test structure (Fig. 25.8.6 top) to characterize the optimal W s across supply voltages (Fig. 25.8.6 bottom left) and the size of metastability window (W_M) (Fig. 25.8.6 bottom right). We compare our design to the recent NTV NoC work [6] (Fig. 25.8.7 bottom). The proposed MEDAC technique will benefit the design of future multi-V/F near-threshold NoCs demanding both ultra-low power and extreme robustness.

Acknowledgements:

This work was supported by the National Key Research and Development Program of China (No.2018YFB2202004), NSFC (Grant No. 61774104), Semiconductor Research Corporation (TxACE 2712.012), and US National Science Foundation (CCF-1453142).

References:

- [1] J. P. Cerqueira et al., "Catena: A 0.5-V Sub-0.4-mW 16-Core Spatial Array Accelerator for Mobile and Embedded Computing," *IEEE Symp. VLSI Circuits*, 2019.
- [2] J. Zhou et al., "On-Chip Measurement of Deep Metastability in Synchronizers," *IEEE JSSC*, vol. 43, no. 2, pp. 550-557, Feb. 2008.
- [3] L. Clemenz et al., "Metastability in CMOS Library Elements in Reduced Supply and Technology Scaled Applications," *IEEE JSSC*, vol. 30, no. 1, pp. 39-46, Jan. 1995.
- [4] K. Bowman et al., "Energy-Efficient and Metastability-Immune Timing Error Detection and Instruction-Replay-Based Recovery Circuits for Dynamic-Variation Tolerance," *ISSCC*, pp. 402-403, Feb. 2008.
- [5] S. Kim, M. Seok, "Variation-Tolerant, Ultra-Low-Voltage Microprocessor with a Low-Overhead, Within-a-Cycle In-Situ Timing-Error Detection and Correction Technique," *IEEE JSSC*, vol. 50, no. 6, pp. 1478-1490, June 2015.
- [6] S. Paul et al., "A 3.6GB/s 1.3mW 400mV 0.051mm² Near-Threshold Voltage Resilient Router in 22nm Tri-gate CMOS," *IEEE Symp. VLSI Circuits*, 2013.
- [7] F. Mu et al., "Self-Tested Self-Synchronization Circuit for Mesochronous Clocking," *IEEE TCAS-II*, vol. 48, no. 2, pp. 129-140, 2001.

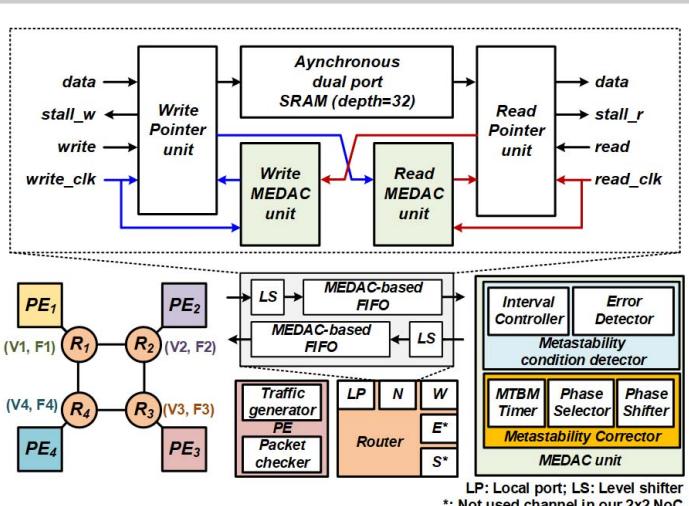


Figure 25.8.1: The 2x2 NoC architecture supporting four V/F domains featuring the proposed MEDAC-based FIFO architecture.

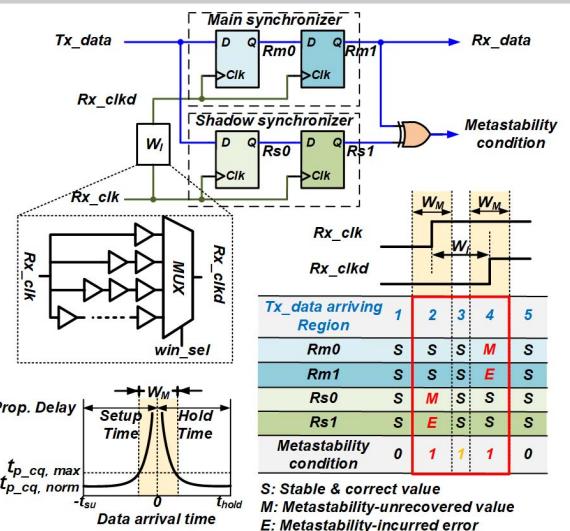


Figure 25.8.2: The circuits and mechanisms of the metastability condition detector.

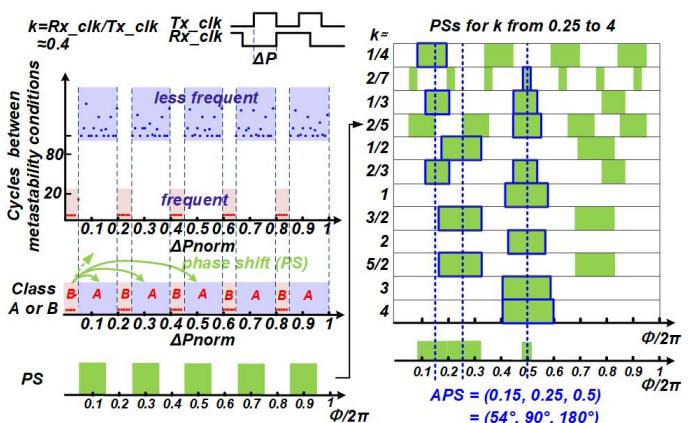


Figure 25.8.3: The metastability condition mitigation scheme. The phase difference of two clocks are classified into *Class-A* and *-B*. The mitigation scheme uses the set of *Class-A* phase-shift (APS) values to reduce the rate of entering metastability.

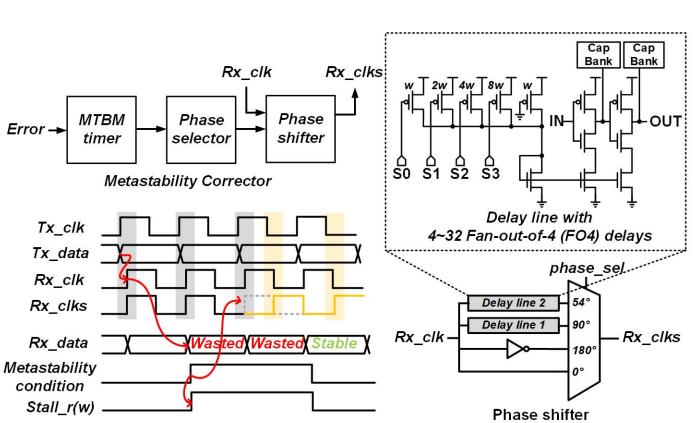


Figure 25.8.4: The metastability corrector architecture and timing diagram of metastability correction. Variable delay lines are used to support the pre-estimated APS.

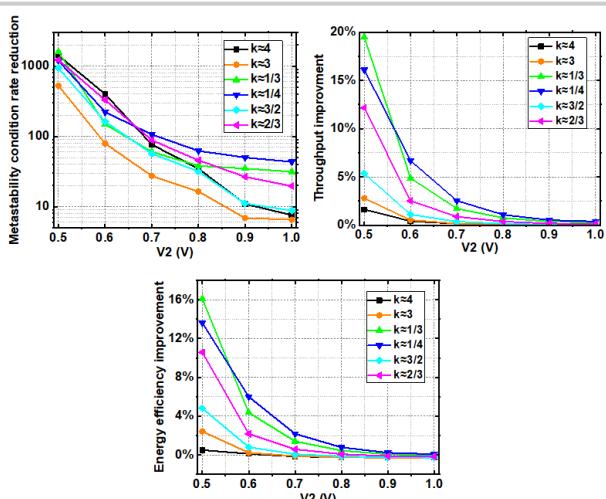


Figure 25.8.5: Measurement results. The MEDAC reduces the chance to enter the metastability condition and thus retransmission rate. The latter improves the throughput and energy-efficiency.

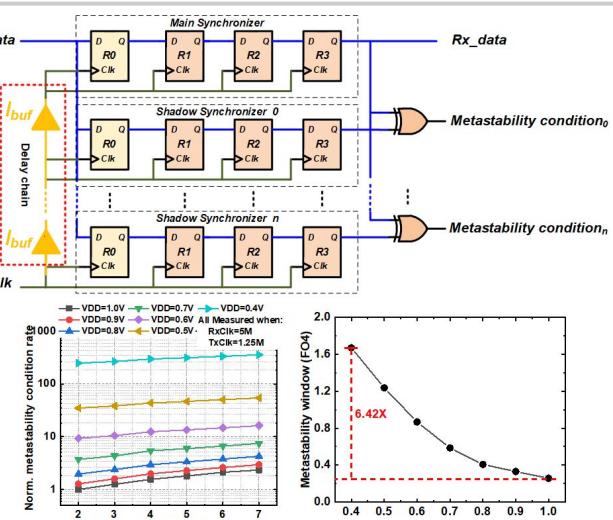


Figure 25.8.6: Test structure to characterize the metastability condition rate and the metastability window (W_M) across V_{DD} s.

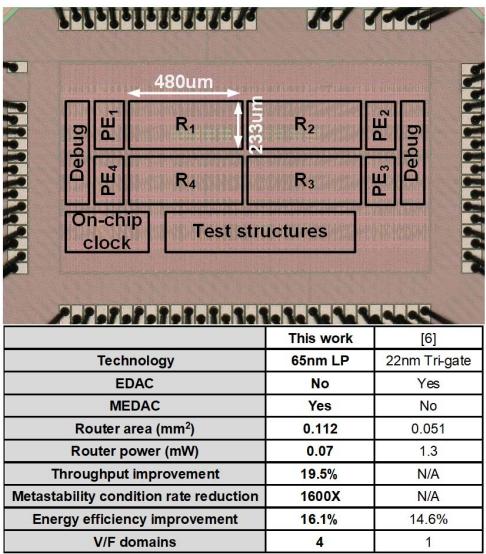


Figure 25.8.7: Testchip die photo and the comparison with the previous work.