

4.7 A 409GOPS/W Adaptive and Resilient Domino Register File in 22nm Tri-Gate CMOS Featuring In-Situ Timing Margin and Error Detection for Tolerance to Within-Die Variation, Voltage Droop, Temperature and Aging

Jaydeep P. Kulkarni, Carlos Tokunaga, Paolo Aseron, Trang Nguyen Jr, Charles Augustine, James Tschanz, Vivek De

Intel, Hillsboro, OR

8-transistor (8T) cell 1-read/1-write (1R1W) register files (RF) with domino read and static differential write are critical performance-limiting building blocks in high-performance microprocessor datapaths. The RF operating voltage (V) and frequency (F) are limited by the delay of the precharge-evaluate read critical path. Traditionally, the operating V/F is set to ensure no read timing error across all data access patterns in the RF array in the presence of within-die (WID) parameter (P) variations, and worst-case voltage droops, temperature (T) changes and transistor-aging-induced delay degradations. However, many of these worst-case conditions and events are rare during normal operation. Therefore, these V/F guardbands can severely limit the best-achievable performance and energy efficiency in scaled CMOS process.

Timing-error detection (TED) via double sampling of critical logic path delays with error-detection sequentials (EDS) or critical path delay monitoring with calibrated tunable replica circuit (TRC) have been used along with replay-based error recovery and V/F adaptation via recovery overhead tracking to minimize variation guardbands in flipflop-based static CMOS logic units [1-3]. These error detection techniques, however, cannot be directly used for 2-phase precharge-evaluate domino read critical paths in high-performance RF arrays. In this paper, we demonstrate, in-situ timing margin detector (TMD) and timing error detector (TED) circuits for domino read critical paths in 8T-cell 1R1W RF arrays (Fig. 4.7.1) to enable: 1) V/F adaptation to temperature and aging, and for excessive persistent timing errors produced by certain data access patterns, and 2) resiliency to occasional timing errors triggered by local high-frequency V droops and nominally random data access in the presence of WID delay variations. TMDa detects data arrival within margin detection window 2 (MDW2) set by the maximum-possible read delay change, as dictated by slow dynamic variations such as T, during the adaptive V/F response. The V/F adaptation controller then either lowers F or increases V until data arrival within the margin detection window 1 (MDW1), set by the maximum-possible read delay change in one clock cycle, is detected by TMDb. If data arrives ahead of MDW1+MDW2, the V/F adaptation controller increases F or lowers V until data arrival within MDW1 is again detected by TMDb. TEd detects timing errors triggered by fast V droops and data access patterns. The error response controller repeats the read operation at F/2 or higher V for error recovery. It also drives the V/F adaptation controller to lower F or increase V if the error count accumulated by the error rate tracker (ERTe) during the sampling period exceeds a threshold.

A 14KB RF array with embedded TMD, TED and error compaction (EC) circuits (Figs. 4.7.2 and 4.7.7) for the domino read critical paths is implemented in 22nm tri-gate CMOS [4]. The array is organized as 28 128-entry 4Kb sub-arrays with 32b data-out and 16 cells per local bitline (LBL) for read. A 2-way NAND merges 2 LBL's into the global bitline (GBL). 2 GBLs drive a set-dominant latch (SDL) for read-data capture (SDLOUT), which is sampled by a TMD with programmable MDW delay and a TED for timing margin and error detection, respectively. The EC combines 32 TMD/TED outputs into a single error-compact detection signal using a wide-domino NOR gate. TMD detects if the read delay transition occurs within the programmable MDW by sampling SDLOUT and DEL_SDLOUT with CLK, and comparing the sampled values (Fig. 4.7.3). Unlike DSTB-EDS-based double-sampling of MDW-delayed critical paths for timing margin detection [2], the effectiveness of this TMD is not limited by a small sampling window and it can detect transitions anywhere within a large MDW.

TED is implemented via double sampling of SDLOUT by CLK to detect transitions in the 3H clock phase. The sampling window of 50% cycle time is larger than that for the DSTB-EDS deployed in flipflop-based static CMOS logic units [2], where the maximum size of the frequency-independent fixed sampling window must be carefully restricted to avoid min-delay failures. The larger TED sampling window for the domino-read-evaluate path enables more aggressive

V/F push, until timing errors become undetectable. The maximum V/F push is limited by worsening LBL delay, and potential LBL transition failure, as well as the proportionally shrinking sampling window. Correct read data is captured in SDLOUT even when a read timing error is detected and latched in DOUT-LAT at the start of 3L clock phase. Conditional delayed precharge utilizes charge sharing between adjacent LBL/GBLs via the bitline equalization transistors EQ1 and EQ2. If NAOUT (SDLOUT) evaluates to 1, bitline charge sharing is initiated conditionally to speed-up precharge. If NAOUT (SDLOUT) does not evaluate to 1, it may be due to a slow LBL evaluate transition. The programmable-delayed LBL/GBL precharge allows more time for the LBL/GBL read evaluate transition to propagate to the SDL. In this case, P1, P2 and EQ1 are turned ON simultaneously for accelerated precharge completion within the reduced precharge time window. Thus, operating frequency dictated by the LBL/GBL evaluate delay is maximized while precharge and differential static write are designed not to be frequency limiters under worst-case conditions.

Measurements of read timing margin and read timing errors are performed using a high-frequency membrane probe card on the RF array and the TMD successfully detects read evaluate transitions within varying MDW values across a wide V/F range (Fig. 4.7.3). Timing errors become predominantly multi-bit as V/F is pushed by only 5%, for both 16b and 32b data out (Fig. 4.7.4). Thus, parity protection, typically used in register files, or more complex ECC schemes are not as effective for reliable detection of timing errors with aggressive V/F push, while incurring minimal performance/power and area overheads. Measurements show that delayed precharge can improve the read timing error capture rate by up to 10x by allowing more time for the LBL evaluate transition.

Figure 4.7.5 demonstrates detection of a timing error induced by fast local 10% V droop, the predominant source of supply noise in high performance microprocessors – especially with fully-integrated voltage regulators [5], and successful error recovery via replay at F/2, as well as 10% higher V. Error rate measurements at 440mV show higher F at higher temperature, as expected from the inverse temperature dependence of drive current at low voltages. The operating F would be set to 860MHz to account for worst-case PVT variations and aging guardbands for the conventional design (target failure rate below $1e^{-6}$). With the TMD/TED approach, as F is pushed higher, throughput first increases proportionally, and then peaks when the error rate and the corresponding recovery overheads become too large. The peak throughput at 440mV can be improved by 36-39% using TMD/TED based V/F adaptation and error recovery via replay, depending on the number of recovery overhead (RO) cycles. The maximum achievable throughput (GOPS) can be improved by 21% using replay at F/2 or higher V (Fig. 4.7.6). For a target throughput of 0.9 GOPS, GOPS/W can be improved by 67% (65%) with replay at F/2 (10% higher V). The area overheads of the embedded TMD/TED circuits in the 14KB RF array are 6.4/12.8%, and can be further reduced when larger array designs are considered. The power overheads are minimal (0.2-0.3%).

Acknowledgements:

The authors thank K. Ikeda, L. Avery, and D. Jenkins for chip design, R. Forand & M. Haycock for help and support. This research was, in part, funded by the U.S. Government (DARPA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

References:

- [1] D. Ernst, *et al.*, "Razor: A Low-Power Pipeline Based on Circuit-Level Timing Speculation," *IEEE/ACM MICRO*, pp. 7-18, 2003.
- [2] K. A. Bowman, *et al.*, "A 45nm Resilient Microprocessor Core for Dynamic Variation Tolerance," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 194-204, 2011.
- [3] J. Tschanz, *et al.*, "Tunable Replica Circuits and Adaptive Voltage-Frequency Techniques for Dynamic Voltage, Temperature, and Aging Variation Tolerance," *IEEE Symp. on VLSI circuits*, pp. 112-113, 2009.
- [4] C.-H. Jan, *et al.*, "A 22nm SoC Platform Technology Featuring 3-D Tri-Gate and High-k/Metal Gate, Optimized for Ultra-Low Power, High Performance and High Density SoC Applications," *IEDM Dig. Tech. Papers*, pp. 4-7, 2012.
- [5] N. Kurd, *et al.*, "Haswell: A Family of IA 22nm Processors," *ISSCC Dig. Tech. Papers*, pp. 112-113, Feb. 2014.

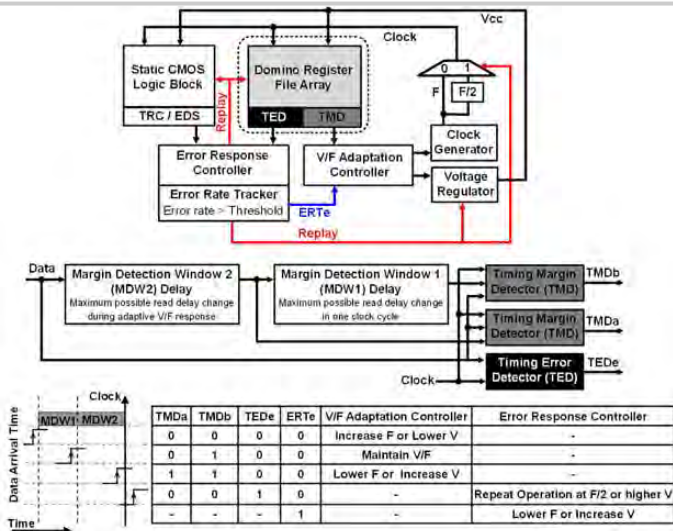


Figure 4.7.1: TMD and TED-based V/F adaptation/recovery mechanism with MDW settings and its relation with data arrival timing.

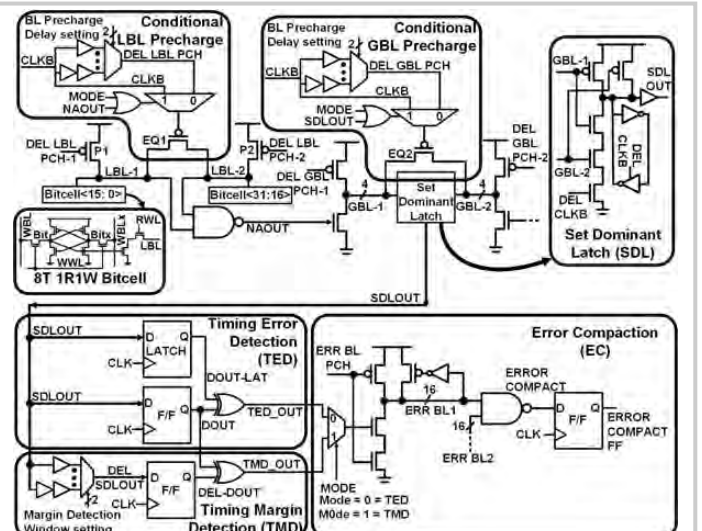


Figure 4.7.2: Resilient domino register file read path featuring delayed bitline precharge, TMD, TED and EC circuits.

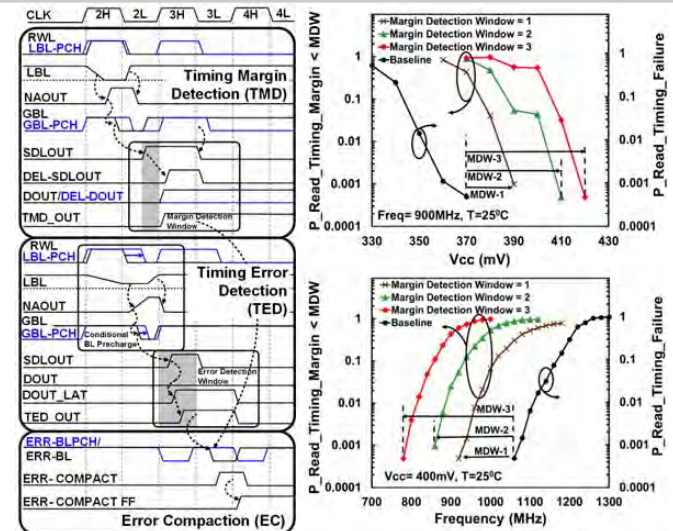


Figure 4.7.3: Resilient register file read timing diagram; measured read timing margin errors and read timing failures iso-V_{cc}/frequency.

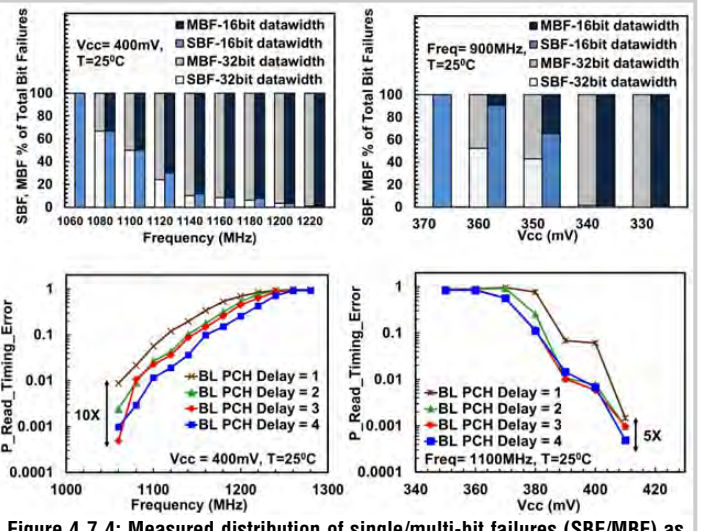


Figure 4.7.4: Measured distribution of single/multi-bit failures (SBF/MBF) as a function of voltage/frequency; sensitivity of bitline precharge (BL PCH) delay to read timing errors captured by TED.

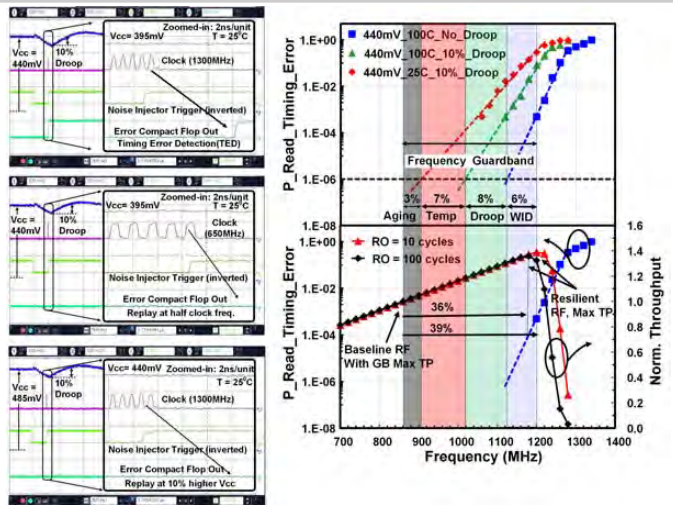


Figure 4.7.5: TED demonstration for 10% V_{cc} droop, error recovery via replay at F/2 and 1.1X V_{cc}; measured frequency guard-bands at iso-V_{cc}, with throughput gain for various recovery-overhead (RO) cycles.

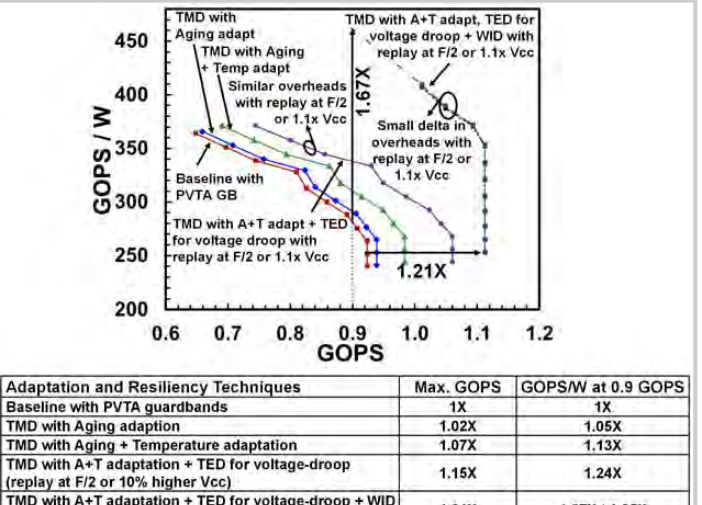


Figure 4.7.6: Measured GOPS/W vs. GOPS comparison for TMD-based adaptation and TED-based recovery schemes with max. 409 GOPS/W.

Adaptation and Resiliency Techniques	Max. GOPS	GOPS/W at 0.9 GOPS
Baseline with PVTa guardbands	1X	1X
TMD with Aging adaption	1.02X	1.05X
TMD with Aging + Temperature adaptation	1.07X	1.13X
TMD with A+T adaptation + TED for voltage-droop (replay at F/2 or 10% higher Vcc)	1.15X	1.24X
TMD with A+T adaptation + TED for voltage-droop + WID (replay at F/2 or 10% higher Vcc)	1.21X	1.67X / 1.65X

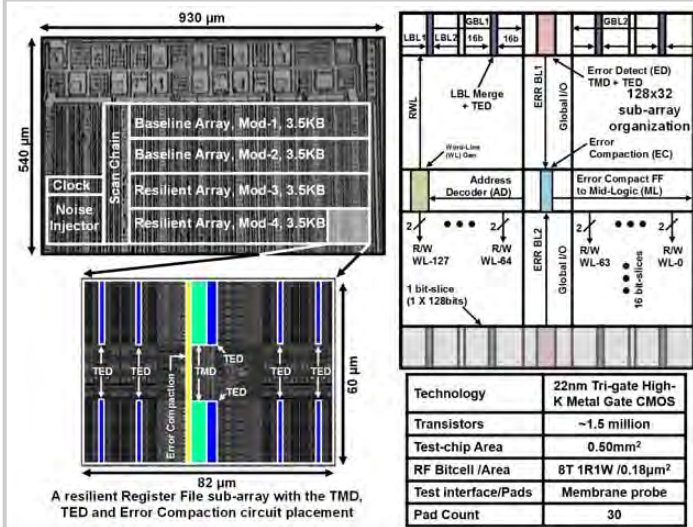


Figure 4.7.7: Chip micrograph with sub-array organization and implementation summary table.