A 409 GOPS/W Adaptive and Resilient Domino Register File in 22 nm 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, *Senior Member, IEEE*, Carlos Tokunaga, *Member, IEEE*, Paolo A. Aseron, *Member, IEEE*, Trang Nguyen, Charles Augustine, *Member, IEEE*, James W. Tschanz, *Member, IEEE*, and Vivek De, *Fellow, IEEE*

***Abstract*—This paper presents an adaptive and resilient domino register file design featuring *in-situ* timing margin and error detection for the performance-critical domino read path. Voltage/frequency is adapted for slow-changing variations such as low-frequency supply noise, temperature fluctuation, and aging-induced degradation. Dynamic adaptation is combined with error detection and recovery for fast voltage droops and random data access patterns in the presence of within-die process variations. Throughput and energy efficiency gains are higher than the replica/canary based critical path approach. Timing margin is tracked by double-sampling the read output and its delayed version at the same clock edge. Timing errors are detected by double-sampling and comparing the read output within a clock window. The sensing errors in the precharge/evaluate domino read path are converted into timing errors using a conditional delayed-bitline precharge technique that does not impact the subsequent precharge operation. The proposed techniques incur 6–13% area overhead and 0.2–0.3% power overhead for a 4 Kb sub-array. The measurement results from a 22 nm tri-gate CMOS testchip demonstrate 21% throughput and 67% energy efficiency improvement with a peak energy efficiency of 409 GOPS/W.**

***Index Terms*—Adaptation, domino read operation, error compaction, error detection sequentials, register file, resiliency, timing error detection, timing margin detection, tunable replica circuits.**

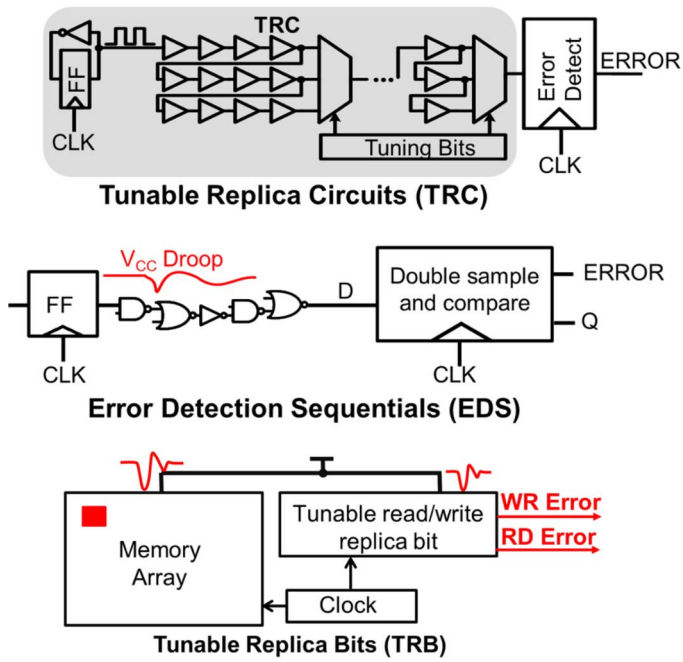

Fig. 1. Previous work: Tunable replica circuits (TRC) and in-situ error detection sequential (EDS) for adaptive and resilient static CMOS datapath, tunable replica bits (TRB) for adaptive memory array design.

I. Introduction

In MODERN microprocessor design, supply voltage (Vcc) scaling is the primary driver to reduce energy consumption [1]. The minimum operating supply voltage (Vmin) of a core consisting of static CMOS datapath logic and high-performance domino register file (RF) arrays is typically limited by process variations in the RF array bitcells that use minimum-sized transistors. The RF operating voltage (V) and frequency (F) are limited by the delay of the precharge-evaluate read critical path. Furthermore, additional V/F guardbands are applied to account for the worst case dynamic variations such as voltage droops, temperature fluctuations, and aging-induced degradation. However, since most systems usually operate at nominal conditions, the fixed V/F guardbands for infrequent dynamic variations significantly limit the best achievable performance and energy efficiency [2]. Previously reported guardband reduction techniques fall into two categories based on whether or not they allow actual timing errors to occur. Dynamic adaptation techniques that monitor the available timing margin in the design and modulate the V/F dynamically do not allow errors to occur. Resilient techniques, on the other hand, are more aggressive in V/F push. In this case, the errors are allowed to happen, they are detected,

Manuscript received May 05, 2015; revised July 10, 2015; accepted July 19, 2015. Date of publication September 22, 2015; date of current version December 30, 2015. This paper was approved by Guest Editor Jinuk Luke Shin. This research was supported in part 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.

The authors are with the Circuit Research Lab, Intel Corporation, Hillsboro, OR 97124 USA (e-mail: jaydeep.p.kulkarni@intel.com).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2015.2463083

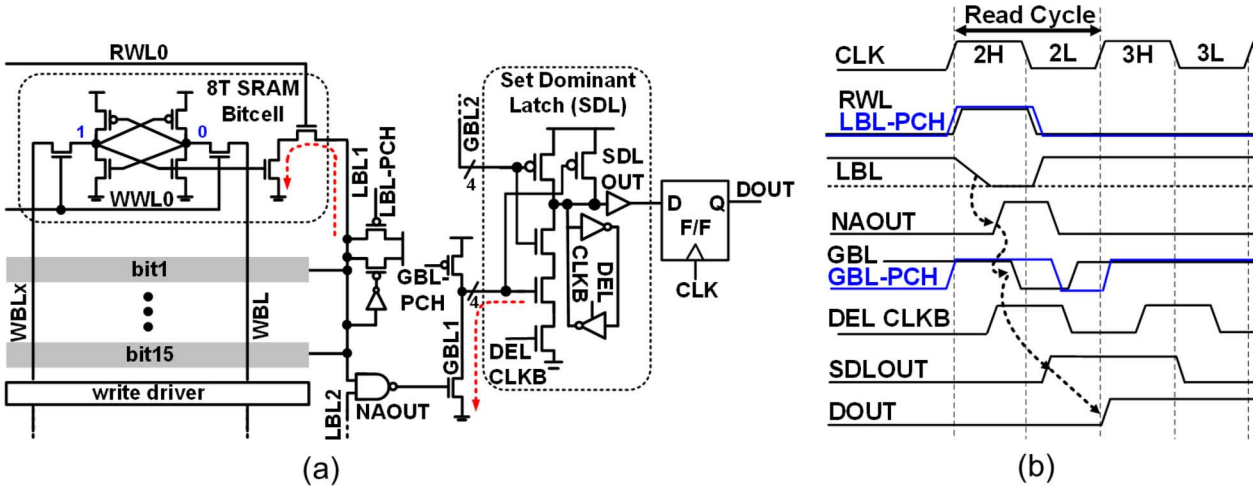

Fig. 2. (a) Read path schematics for the baseline domino register file array; (b) precharge/evaluate read timing diagram for the baseline design.

and then corrected using appropriate replay mechanisms. For error detection in flip-flop based static CMOS logic units, a replica-based approach such as tunable replica circuits (TRC) has been proposed [3]. In this approach, a set of replica circuits are calibrated to match the critical path pipeline stage delay and timing errors are detected by double-sampling the TRC outputs (Fig. 1). The key requirement is that the TRC must always fail before the critical path fails. The TRC is an area-efficient and non-intrusive technique, but it cannot leverage the probabilities of critical path activation, multiple simultaneous switching at inputs of complex gates, or worst case coupling from neighboring signal lines. The alternative in-situ approach for timing error detection uses error detection sequentials (EDS) in the critical paths of the pipeline stage. Timing errors are detected by a double-sampling mechanism using a flip-flop and a latch (Fig. 1) [4], [5]. Errors are corrected by performing a replay operation at higher V or lower F. V/F can also be adapted by monitoring the error rate and accounting for error recovery overheads. These error detection techniques, however, cannot be directly used for 2-phase precharge-evaluate domino read critical paths in high-performance RF arrays since the data outputs are valid only during the evaluate phase. For error detection in RF arrays, replica-based techniques such as tunable replica bits (TRB) has been proposed [6]. In this approach, a set of replica memory bits are tuned at test time so that in the presence of dynamic variations the TRBs fail before the worst case memory bit fails (Fig. 1). This approach is very similar to the TRC approach in the logic circuits. In addition, a double sensing technique for differential sensing in 6T SRAMs has been proposed [7]. The 6T SRAMs used in last level caches are optimized for highest density and require higher Vmin for successful operation. These arrays operate on a separate and relatively stable voltage supply.

In this paper, we propose and demonstrate, for the first time, in-situ timing margin detector (TMD) and timing error detector (TED) circuits for domino read paths in 8T 1R1W bitcell large signal sensing RF arrays. The proposed TMD and TED circuits enable:

1) V/F adaptation to low-frequency voltage variations, temperature and aging, as well as to excessive persistent timing errors produced by certain data access patterns, and

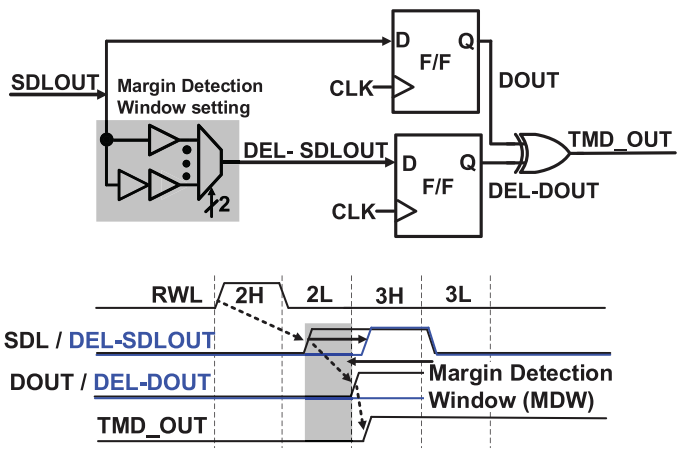

Fig. 3. Timing margin detection (TMD) by double sampling the SDLOUT and its delayed version at the same clock edge. Programmable margin detection window (MDW) setting for margin detection across wide range.

2) resiliency to timing errors triggered by local high-frequency voltage droops and nominally random data access in the presence of within-die (WID) delay variations.

The paper is organized as follows. Section II describes the design of the baseline domino RF array, TMD, TED, conditional bitline precharge, and error compaction (EC) circuits. Section III discusses the usage modes of the TMD and TED, along with various data arrival scenarios. Section IV presents the measurement results from a 22 nm tri-gate CMOS testchip. Section V concludes the paper with the key findings. Section VI presents an analytical treatment of the operating frequency improvement in the resilient design (in the Appendix).

II. Circuit Design

A. Baseline Design

Each baseline sub-array is organized as a 128-entry 4 Kb sub-array with 32 bit data-out and 16 cells per local bitline (LBL) for read. A 2-way NAND merges 2 LBLs into the global bitline (GBL). Two GBLs drive a set-dominant latch (SDL) for read data capture (SDLOUT), and the data output is captured

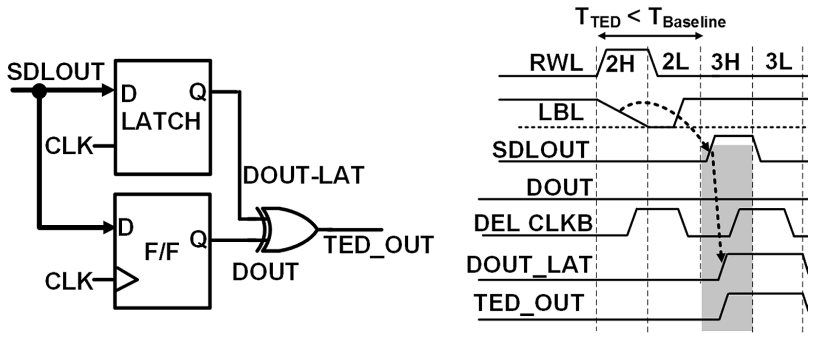

Fig. 4. Timing error detection (TED) by double sampling SDLOUT with a flip-flop and a latch combination and corresponding timing diagram.

in a flip-flop, thus completing the read operation in one cycle [Fig. 2(a)]. The read operation consists of two-phase precharge/evaluate domino logic. In 2 H clock phase, when a selected read wordline (RWL) signal is asserted, the LBL is discharged depending on the bitcell data [Fig. 2(b)]. A falling transition on the local bitline triggers the merge-NAND gate that merges two local bitlines. The NAOUT node in turn evaluates a GBL which feeds the SDL. A falling transition on GBL is captured by SDL and latched by a minimum-sized cross-coupled inverter pair. The SDL acts as an interface between the domino and static logic paths. SDL functionality is controlled by delayed inverted clock (DEL CLKB) and sets SDLOUT to 0 if both GBLs are pre-charged.

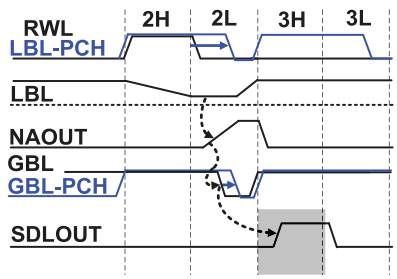

Fig. 5. Timing diagram for the delayed bitline precharge used to detect the domino read path errors by converting sensing errors into timing errors by allowing more time for the bitline evaluate transitions to propagate to the SDLOUT.

B. Timing Margin Detection (TMD)

Timing margin detection (TMD) is used to adapt the V/F to slow variations such as low-frequency supply noise, temperature fluctuation, and aging-induced degradation. TMD detects if the read delay transition occurs within the programmable margin detection window (MDW) by sampling SDLOUT and DEL_SDLOUT with CLK, and comparing the sampled values at the same clock edge (Fig. 3). Unlike double-sampling-time-borrowing (DSTB) EDS-based double sampling of MDW-delayed critical paths for timing margin detection [5], the effectiveness of this TMD is not limited by a small sampling window and it can detect transitions anywhere within a large MDW. In the DSTB-EDS scheme, the MDW is limited by min-delay constraints. Since the DEL-SDLOUTs are flopped with the same clock, it is susceptible to the metastability event as DEL-SDLOUT can arrive anywhere within the setup-hold time window of the flip-flop. The metastability probability can be alleviated by using multiple TMDs with varying MDW and multiple sampling + XOR comparators. In a more sophisticated TMD design, a thermometer code of MDWs can be generated to adapt the V/F at much finer granularity. With a thermometer coded TMD, the metastability issue can be mitigated by observing the 3 bit window around $0 \rightarrow 1$ transition and adapting V/F accordingly at the expense of increased TMD area, extra load on CLK and SDLOUT nodes, and a complex adaptation algorithm [8].

However, in the proposed approach, TMDs with only two MDWs are used (usage model will be described in detail in Section III). With a small number of TMDs, there may be a scenario where metastability event in DEL-SDLOUT flip-flops can cause adaptation to set lower V and/or higher F than acceptable for error-free operation for a given voltage/temperature/aging condition. This timing margin mis-estimation is not catastrophic since it will be detected as a timing error by the TED and recovered by an appropriate replay mechanism (details in next section). If the TMD metastability event is persistent and if the TED error rate (due to TMD metastability mis-estimation) exceeds a certain threshold, the error rate tracker can initiate the adaption loop to back off V/F, thus mitigating TMD metastability. Therefore, the TED + replay mechanism can be used to alleviate metastability issues for small number of TMDs without increasing the design overheads.

C. Timing Error Detection (TED)

The TED circuit is implemented to enable resilient operation in the presence of fast variations such as high-frequency voltage droop and random bitcell access patterns in the presence of WID variations. The key mechanism is to double sample SDLOUT with an additional latch to detect transitions in the 3H clock phase (Fig. 4). The outputs of the flop and latch are compared with an XOR gate to detect the timing error [4]. 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. Note that the sampling window of 50% cycle time is larger than that for the DSTB-EDS in-situ timing error detection in logic circuits [5] where the maximum size of the frequency-independent fixed sampling window must be carefully restricted to avoid min-delay failures. In the baseline design, the probability of a metastability event is around the rising edge of 3H clock phase, while for TED based resilient design, it

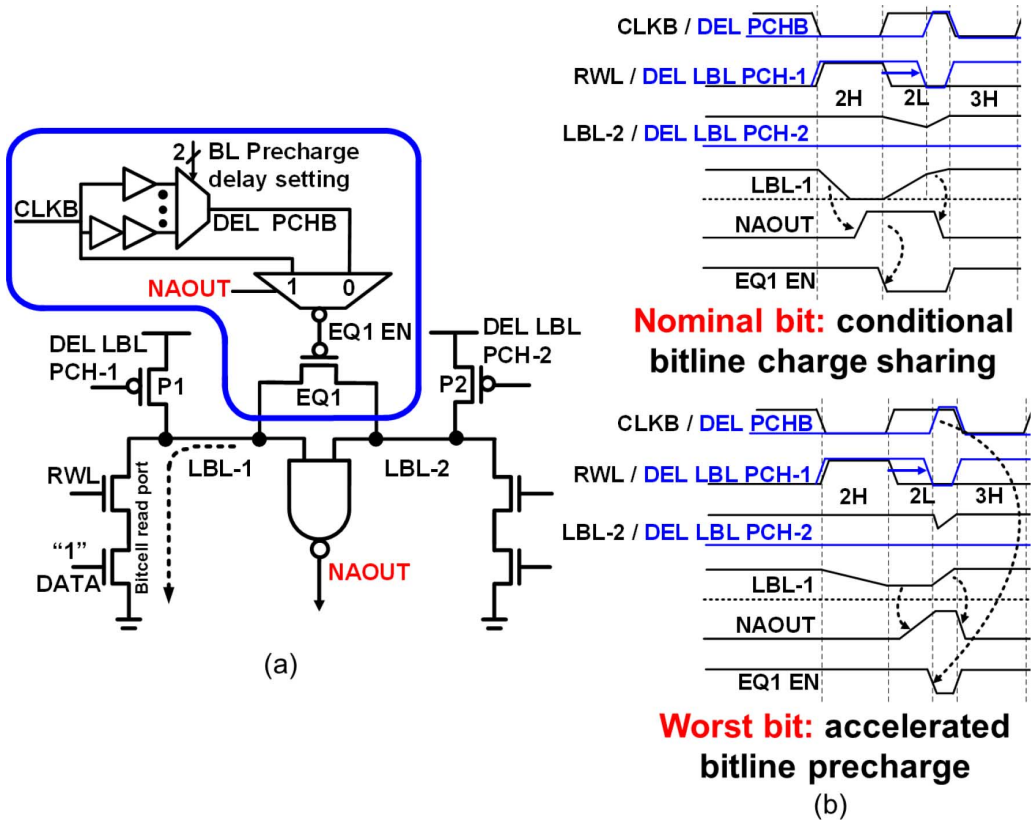

Fig. 6. (a) Conditional bitline precharge circuit with bitline charge sharing for accelerated precharge; (b) timing diagram for the nominal bit and the worst case bit read operation.

is moved to the falling edge of 3L clock phase. The metastability in the data path can be avoided by swapping the flop and latch in this design; similar to [5]. In this case, the metastability will be contained only within the error path. The larger TED sampling window for the domino read evaluate path enables more aggressive V/F push, until timing errors become undetectable. The timing error detectability is limited by worsening LBL delay, and potential LBL transition failure resulting in a sensing failure. These sensing failures cannot be captured using the double sampling mechanism and limit the maximum frequency achieved by the resilient design.

D. Conditional Delayed Precharge

In order to convert sensing failures into timing failures, the precharge of the LBL/GBL is delayed to allow more time for the bitline-evaluate transitions to propagate to the SDL output and be captured by the double sampling mechanism (Fig. 5). Due to reduced precharge time, the subsequent precharge operation can become a performance limiter. To mitigate this issue, we use conditional delayed precharge which utilizes charge sharing between adjacent bitlines via the bitline equalization transistor (EQ1), as shown in Fig. 6(a). For a nominal bit read scenario, the LBL is evaluated completely within the 2H clock phase and triggers the merge-NAND output (NAOUT). Then, there is no need to delay the subsequent precharge operation. The NAOUT controls a multiplexer which selects the inverted clock (CLKB) signal or delayed precharge signal (DEL PCH). If NAOUT evaluates to 1, bitline charge sharing is initiated conditionally to speed-up precharge by activating the equalization transistor (EQ1) in the 2L clock phase (Fig. 6(b)). If NAOUT does not evaluate to 1, it may be due to slow local bitline (LBL) evaluation (weak-1 read) or due to read-0 operation. For read-0, the bitlines will be maintained at Vcc with the help of keeper transistors and there is no timing requirement for the subsequent precharge operation. For a weak-1 read scenario, the programmable delayed LBL precharge allows more time for the LBL 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 by asserting the delayed-precharge signal (DEL PCH) (Fig. 6(b)). Note that the LBL is partially discharged in this case, and would not require full-swing bitline precharge as in the nominal bit read scenario. Thus, the operating frequency dictated by the bitline evaluate delay is maximized while precharge and differential static write are designed not to be frequency limiters under worst case process, voltage, temperature, and aging conditions.

E. Error Compaction (EC)

The next step in the RF read path is the error compaction phase which consists of a wide domino NOR gate combining the 32 TMD/TED outputs from every bitslice into a single error compact signal (Fig. 7). This operation evaluates in the 3L clock phase and generates an error-detect signal for every sub-array at the beginning of the 4H clock phase.

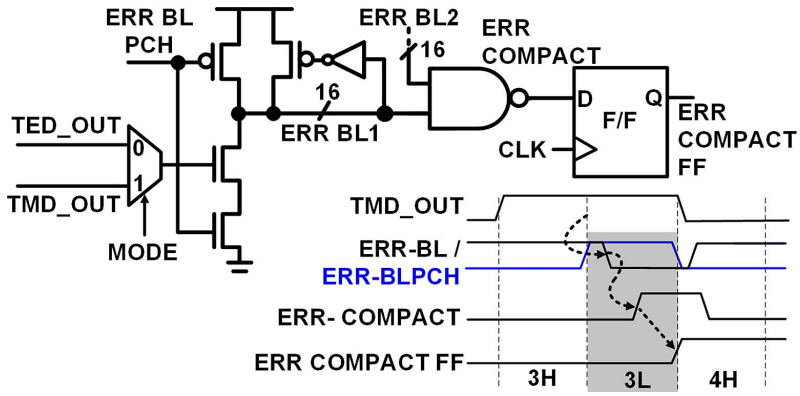

Fig. 7. Error compaction using wide NOR gate and corresponding timing diagram.

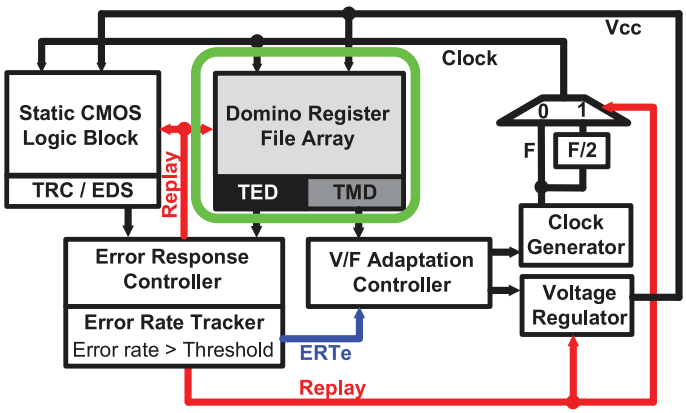

Fig. 8. Unified framework with timing margin detection (TMD) and timing error detection (TED) based V/F adaptation and error response.

III. Discussion

Fig. 8 shows the unified framework for the adaptive and resilient static CMOS logic block along with the domino register file operating at the same Vcc and frequency. The TRC + EDS address the dynamic variations in the static CMOS logic block. For the domino RF array, TMD addresses the slow variations and generates an output to the V/F adaptation controller. The adaptation controller modulates the clock generator and/or voltage regulator to adapt the voltage and/or frequency applied to the design. On the other hand, TED addresses the fast dynamic variations that occur too quickly to be mitigated by V/F adaptation. The TED circuits give an output to the error response controller which initiates a replay mechanism at increased supply voltage or reduced clock frequency. An error rate tracker is also included to monitor the error rate. If the error rate exceeds a certain threshold, a signal (ERTe) is sent to the adaptation controller which adapts voltage and frequency to minimize the overheads due to frequent replay mechanisms.

Fig. 9 shows various scenarios for data arrival times. The read operation is 1 cycle long (2H and 2L clock phase) and consists of worst case delay due to WID variations + guardbands for voltage droop, temperature fluctuations, and aging-induced degradation + setup time of the capturing flip flop. TMD is operational just before the 3H rising clock edge. It consist of at least two margin detection windows (MDW). TMDa detects the data arrival time within MDW2, set by the maximum possible read delay change, as dictated by slow dynamic variations during the adaptive V/F response. The 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 adaptation controller increases F or lowers V until data arrival within MDW1 is again detected by TMDb. TED detects timing errors triggered by fast voltage droops and data access patterns in the 3H clock phase. The error response controller repeats the read operation at F/2 or higher V for error recovery. It also drives the 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. TED also captures the TMD mis-estimations (such as TMD metastability event) and corrects the error by initiating a replay operation.

The proposed in-situ adaptive and resilient technique is applied only for the register file read operation. Write operation should be designed to work correctly even in the presence of dynamic variations, since it is difficult to perform a speculative write operation. For example, if the write operation is performed speculatively (removing all PVT and aging guardbands), then a subsequent read operation needs to be performed to confirm whether the write was successful. This read operation needs to be tolerant to dynamic variations. If a write error is detected then a subsequent replay write operation at higher V or lower F needs to be performed. Hence, the performance degradation for in-situ write error detection can be severe. An alternative approach for write error prevention skews the bitcell to guarantee correct write operation under fast variations, with optional use of tunable replica bits (TRB) for mitigating guardbands for slow-changing variations.

IV. Measurement Results

The proposed timing margin detection (TMD) and timing error detection (TED) circuits for the adaptive and resilient domino register file have been implemented on a 22 nm tri-gate CMOS testchip [9]. It consists of a 7 KB baseline array, a 7 KB resilient array, scan chain, clock source, and noise injector (Fig. 10). The 7 KB arrays are organized as 14 sub-arrays; each 4 Kb sub-array is organized as 128 entries X 32 bits (Fig. 10). Each 128 bit slice is organized hierarchically with 16 bits per local bitline (LBL) and 4 LBLs per global bitline. The area overhead of the embedded TMD/TED circuits in the 4 Kb RF array is 6.4/12.8%, and can be further reduced when larger

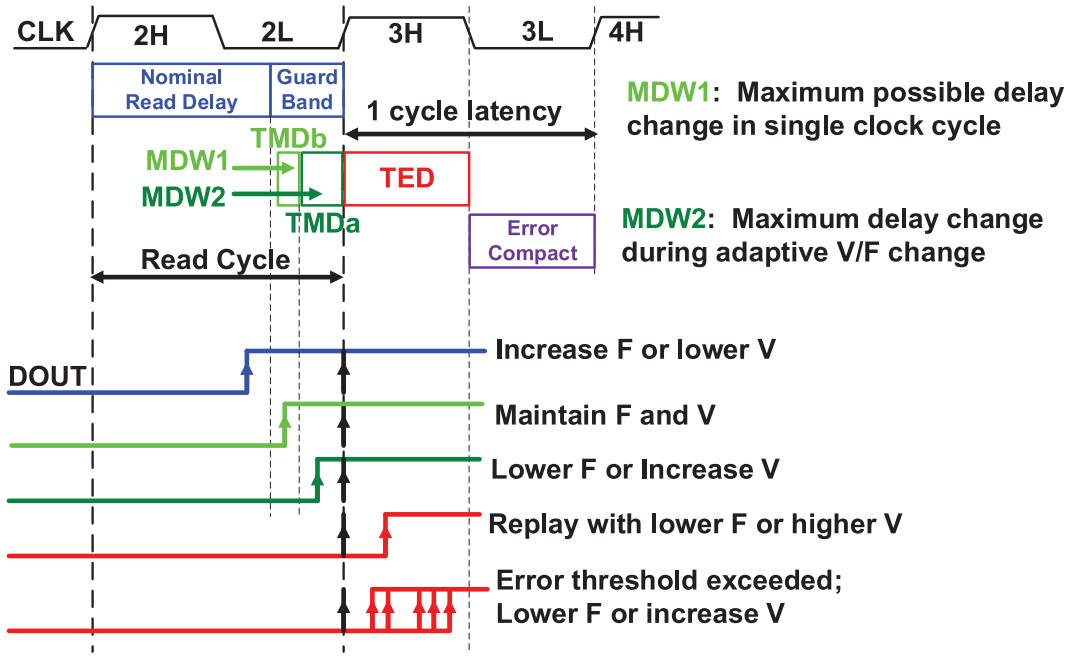

Fig. 9. Various data arrival scenarios and corresponding adaption/recovery mechanisms.

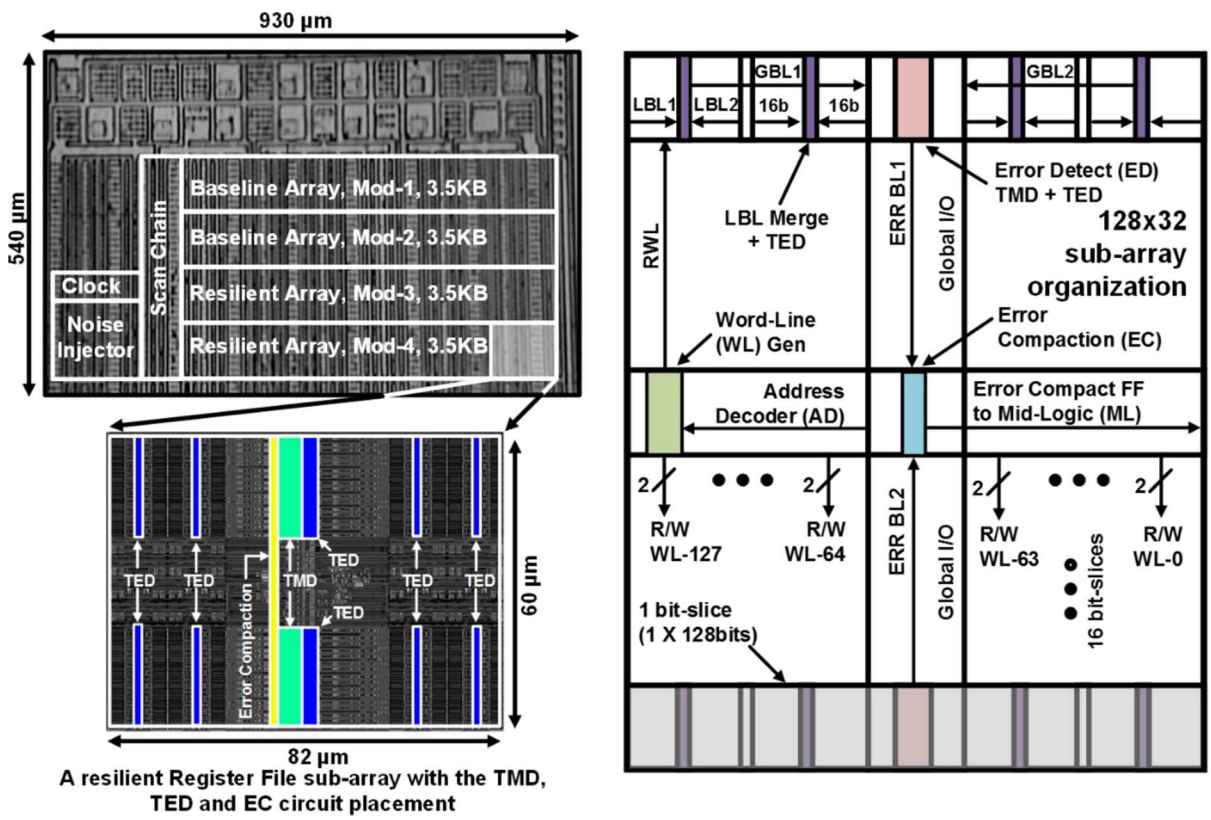

Fig. 10. Chip micrograph with sub-array layout and organization annotating TMD, TED, and EC block placements.

array designs are considered. The power overheads at iso-Vcc and iso-frequency are minimal (0.2–0.3%). Other testchip details are listed in Table I.

Read timing margin and read timing error measurements on the RF arrays are performed using a high-frequency membrane probe card. Supply voltage is aggressively lowered to induce random bitcell failures in the small size array. Fig. 11 shows measured failure probability of the read timing margin being less than the margin detection window (MDW) for different MDW setting, as well as the read timing failure probability as functions of supply voltage and frequency. The TMD successfully detects read evaluate transitions within varying MDW

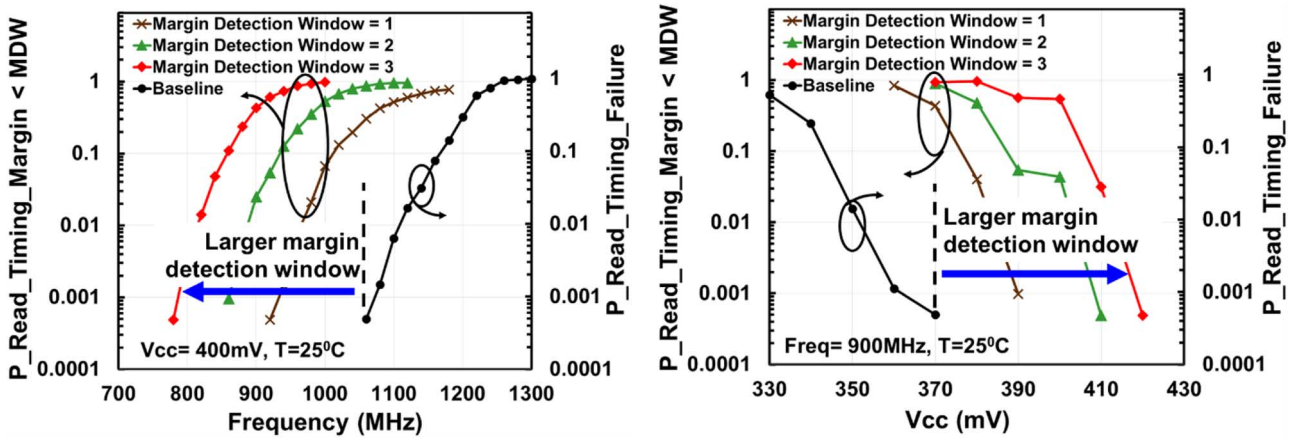

Fig. 11. Measured read timing margin errors (margin < detection window) and actual read timing failures for different margin detection windows (MDW) at iso-Vcc and iso-frequency.

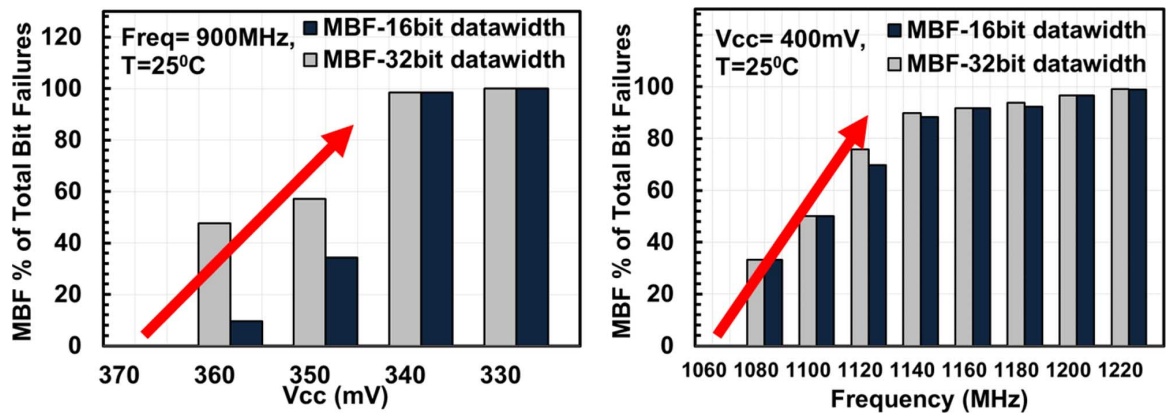

Fig. 12. Measured distribution of multi-bit failures (MBF) at iso-voltage and iso-frequency.

TABLE I
TESTCHIP IMPLEMENTATION SUMMARY

Technology	22nm FinFET High-K Metal Gate CMOS
Transistors	~1.5 million
Test-chip Area	0.50mm^2
RF Bitcell /Area	8T 1R1W /0.18μm^2
Test interface/Pads	Membrane probe
Pad Count	30

values across a wide V/F range. This MDW-delayed double sampling scheme allows a larger delay tuning range than the TRC/EDS-based margin detection approach that is limited by a fixed margin detection window [3], [5]. Fig. 12 shows the distribution of multi-bit failures (MBFs) as a percentage of total bit failures for varying Vcc and frequency. Timing errors become predominantly multi-bit as V/F is pushed by only 5%, for both 16 bit and 32 bit data out. Thus, parity protection, typically used in register files, or more complex ECC schemes are prohibitive for reliable detection of timing errors with aggressive V/F push at low area and power cost. The latency/area cost of multi-bit ECC increases rapidly with increasing bit error detection capability [10]. Thus, SEC-DED or DEC-TED ECC schemes can provide very small V/F push, similar to the resilient router design in [11]. Resiliency mechanisms with timing error detection (TED) along with replay at higher V or lower F can be more power efficient for tolerating multi-bit failures, thus allowing more aggressive V/F push. Fig. 13 shows the effect of conditional bitline precharge on the read timing errors for varying V and F. The bitline precharge delay consists of a fixed delay (two multiplexer gate delay) and can be incremented in steps of two inverter delays. Conditional bitline precharge helps in two ways:

1) It allows the bitline sensing failures to be converted into timing failures by allowing LBL/GBL evaluate transitions to propagate to the SDL where the timing errors are detected using a double sampling mechanism.
2) It helps reduce sensing failures by allowing the bitcells that are on the verge of read failure to evaluate the bitlines for a longer duration. These bitline transitions are captured correctly at the SDL output at the rising edge of the 3 H clock phase without incurring a timing error. By modulating the conditional bitline precharge delay, the read timing errors can be reduced by 5–10X as Vcc and frequency is varied.

Fig. 14(a) demonstrates the TED operation captured on the oscilloscope. Ten percent (10%) Vcc noise is induced using the noise injector at 440 mV/1300 MHz operating point. The noise

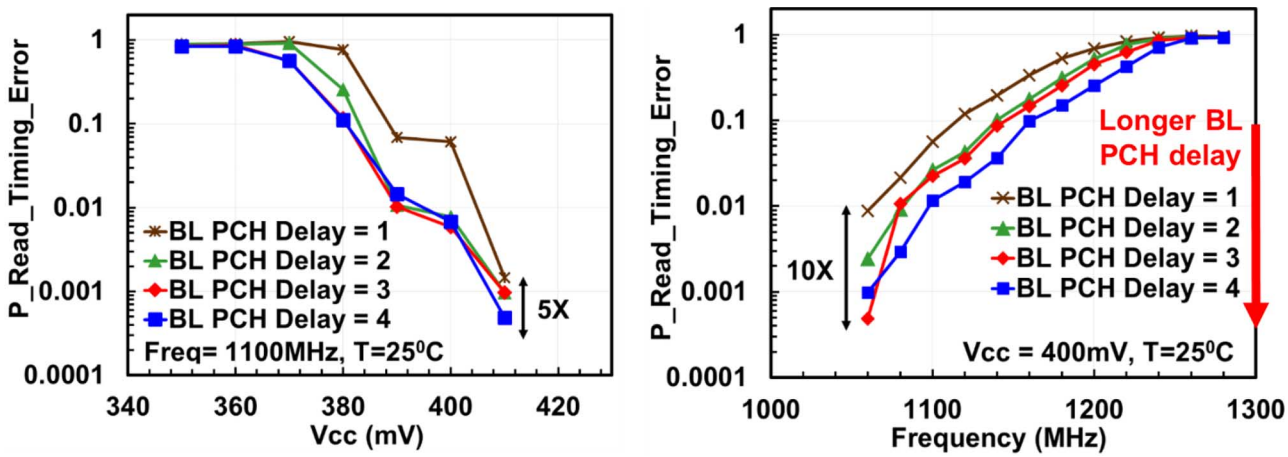

Fig. 13. Measured sensitivity of bitline precharge (BL PCH) delay to the read timing errors captured by timing error detector (TED) at iso-Vcc and iso-frequency. The BL PCH delay is incremented in steps of two inverter delays.

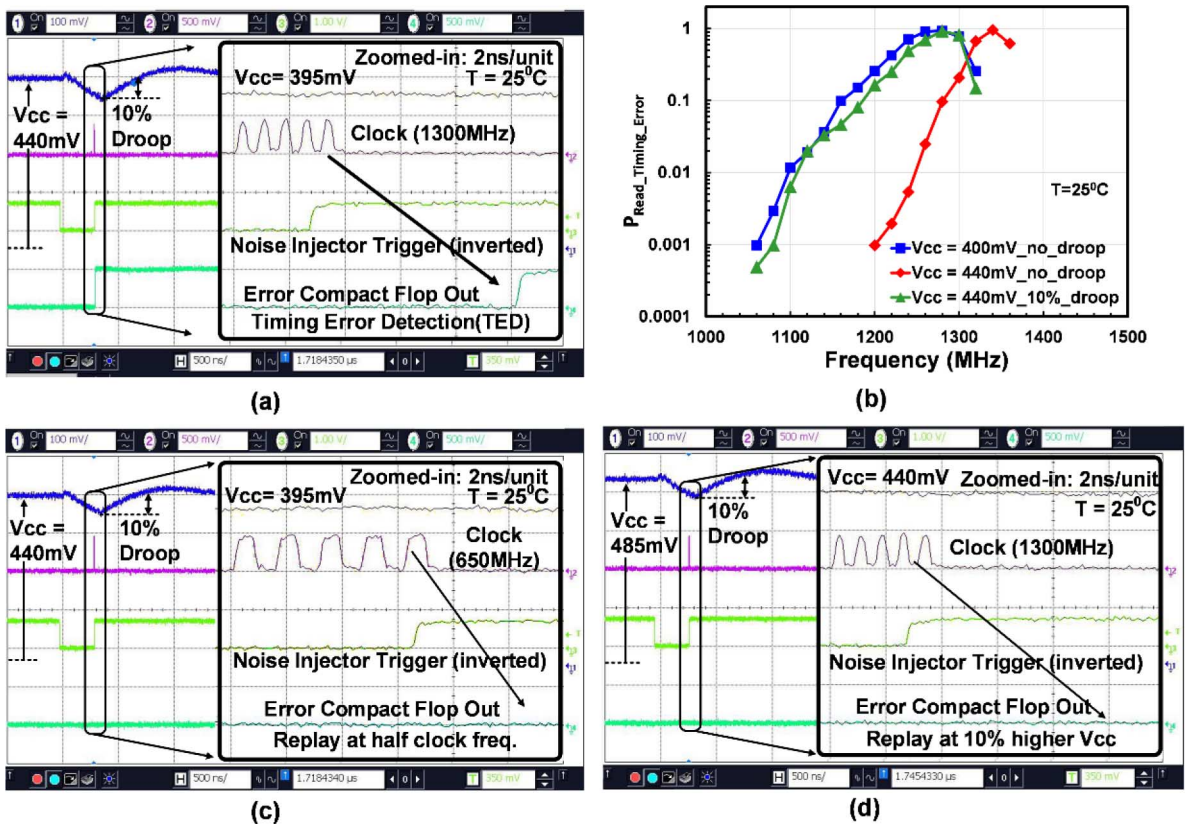

Fig. 14. (a) Demonstration of timing error detection (TED) circuit operation due to 10% Vcc droop; (b) read timing errors correlation of induced 10% voltage droop with 10% lower static Vcc; (c) error recovery via replay at F/2 and (d) at 10% higher Vcc.

injector trigger is adjusted to result in maximum droop during the read operation. The zoomed-in version shows Vcc drooping by ∼45 mV (10% of nominal Vcc) and an error-compact signal is triggered. The error rate statistics for 10% voltage droop at 440 mV correlate well with that at 400 mV without any droop [Fig. 14(b)]. The error can be recovered by via replay at F/2 (650 mV) [Fig. 14(c)] or at 10% higher Vcc [Fig. 14(d)].

Fig. 15 shows the read timing error measurements at 440 mV under voltage and temperature fluctuations. The read failure measurements are extrapolated to 1e-6 to estimate the required frequency guardbands for 1 Mb array size. With 10% voltage droop, the operating frequency is lowered by 8%. With temperature variation from 100°C to 25°C, the frequency further reduces by 7%. Note that inverse temperature dependence is observed at low operating voltage, similar to the combinational static CMOS logic [12]. The aging guardband is estimated to be 3% based on a representative ring oscillator delay degradation [13]. Therefore a baseline design with PVT and aging guardband would operate at 860 MHz at 440 mV. The corresponding throughput is normalized to 1 and used as reference for subse-

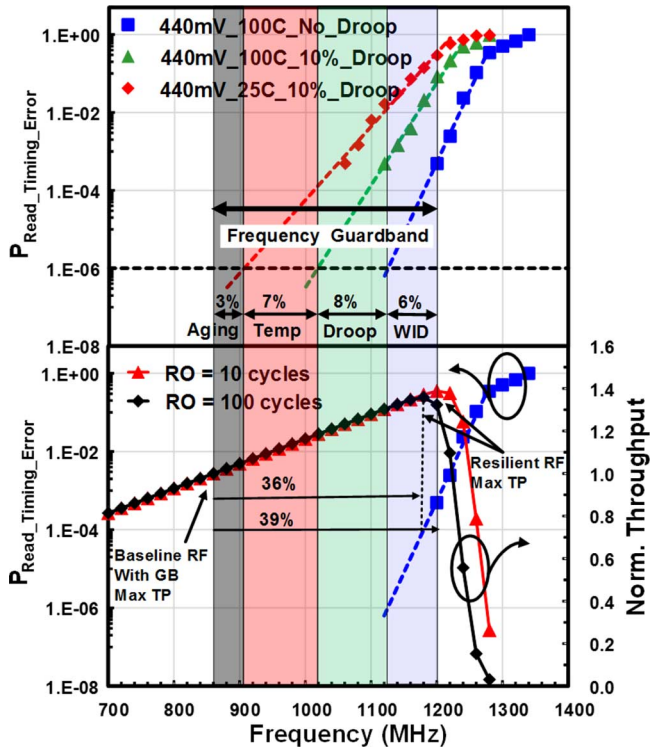

Fig. 15. Measured frequency guardbands for temperature, voltage droop and WID variations, corresponding throughput gain for various recovery overhead (RO) cycles.

quent throughput comparisons. As frequency is pushed higher using the TMD + TED scheme, throughput first increases proportionally, and then peaks when the error rate and the corresponding recovery overheads (RO) become too large. As the recovery overhead increases, throughput improvement further reduces. The peak throughput at 440 mV can be improved by 36–39% using the TMD/TED based V/F adaptation and error recovery via replay, depending on the number of recovery overhead (RO) cycles. With TED, frequency can be pushed by 6%, thus compensating for WID variation impacts across random bitcell access patterns. As the frequency is pushed higher, the read failures increase substantially, thereby increasing the recovery overheads and worsening overall throughput.

Fig. 16 shows the measured energy efficiency (in GOPS/W) vs. throughput (in GOPS). For slow variations, uniform aging-induced degradation and temperature change during the design lifetime is assumed. For aging case, the adaptive RF with TMD will operate at 3% higher frequency at beginning of life and trend towards the same frequency as the baseline design at end of life. If this 3% frequency gain is distributed equally over the design lifetime in 10 steps, for example, the average frequency gain by eliminating the aging guardband using TMD-based adaptation will be 1.65%. Similarly, for the temperature guardband, initially the adaptive RF with TMD will operate at 7% higher frequency. Over the design lifetime, the average frequency gain due to elimination of the temperature guardbands would be 4.4%. The probability of occurrence of 10% voltage droop on average is assumed to be 1e-6 as observed in [14]. Due to very low probability of occurrence of a fast voltage droop event, the recovery overheads due to replay at F/2 or 10% higher

TABLE II
THROUGHPUT IMPROVEMENT AT 250 GOPS/W AND ENERGY EFFICIENCY IMPROVEMENT AT 0.9 GOPS FOR TMD AND/OR TED SCHEMES

Adaptation and Resiliency Techniques	Thoroughput	Energy Efficiency
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

Vcc are similar. However in the case of TED for WID variations, higher dynamic switching capacitance (C_{DYN}) cost is incurred for replay at 10% higher Vcc (quadratic relationship with dynamic power) compared to the replay at F/2 (linear relationship with dynamic power). For a target throughput of 0.9 GOPS, energy efficiency (GOPS/W) can be improved by 67% with replay at F/2 and by 65% with replay at 10% higher Vcc. The maximum achievable throughout (GOPS) can be improved by 21% using TMD for aging and temperature adaption and TED for voltage droops and WID with replay at F/2 or higher V. The peak measured energy efficiency is 409 GOPS/W. When the proposed adaptive and resilient register file is embedded with the logic core, the power and frequency gains will be determined by the Vmin of the entire core, total switching capacitance (Cdyn), V/F guardbands for the core logic, and the error rate increase of the core logic with V/F push. Table II describes the throughput and energy efficiency improvement using TMD and/or TED for each guardband component.

V. CONCLUSIONS

We have proposed an adaptive and resilient domino register file design featuring in-situ timing margin and error detection for the performance-critical domino read path. Voltage and frequency are adapted to slow-changing variations such as low-frequency voltage droop, temperature fluctuation, and aging-induced degradation combined with error detection and recovery for fast voltage droops and random data access patterns in the presence of within-die process variations. The timing margin is detected by double-sampling the read output and its delayed version at the same clock edge. The timing errors are detected by double-sampling the read output with a flop + latch comparison. The sensing errors in the precharge/evaluate domino read path are converted into timing errors using conditional delayed bitline precharge without affecting the subsequent precharge operation. The metastability in TMD can be addressed by multiple MDWs or with the help of TED. The metastability in TED can be mitigated by using a latch in the datapath and confining the metastability to the error path. The proposed techniques incur 6–13% area overhead and 0.2–0.3% power overhead for a 4 Kb sub-array. The measurement results from 22 nm tri-gate CMOS testchip demonstrate TMD operation across larger MDWs; larger than DSTB-EDS approach in which timing margin detection window is governed

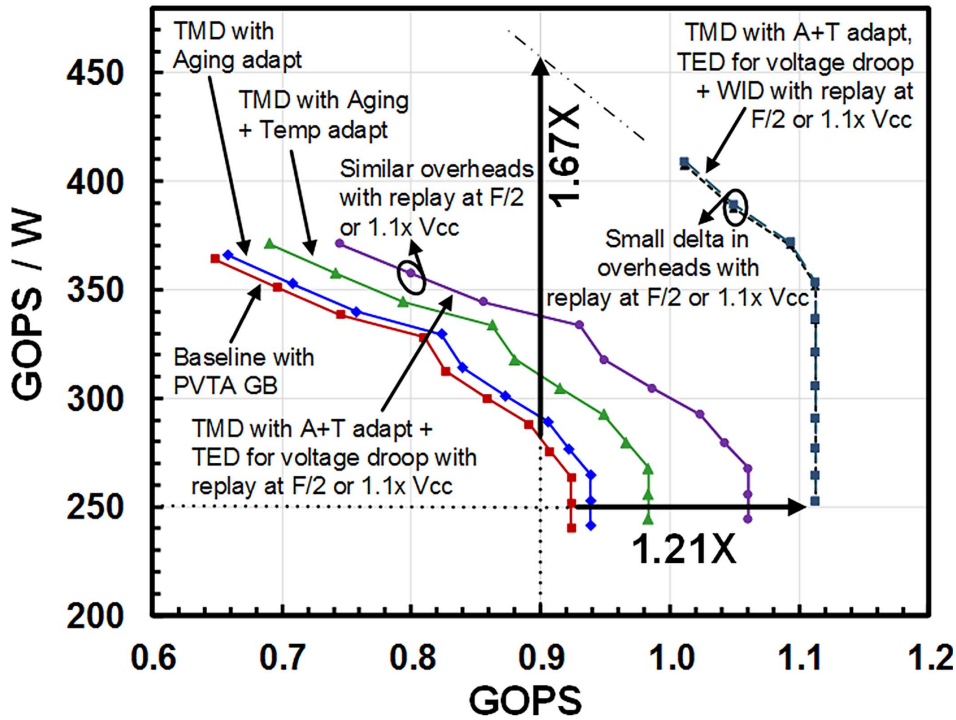

Fig. 16. GOPS/W vs. GOPS comparison of TMD (adaptation) and TED (resilient) schemes.

by min-delay constraints. For aggressive V/F push, the resilient approach of error detection and recovery is more effective in handling multi-bit failures compared to conventional ECC schemes. Conditional bitline precharge helps TED in two ways by converting sensing failures into timing failures and also by lowering sensing failures by allowing more time for the bitline evaluation. With the resilient approach, the throughput gain is eventually limited by the replay overheads. Maximum achievable throughput gain observed is 21%. For a given throughput (0.9 GOPS) the energy efficiency is improved by 65–67% with a peak energy efficiency 409 GOPS/W.

Appendix
Resilient Design Frequency Improvement

In this analysis, conditional delayed precharge is used only for LBL stage and GBL and SDL and the capture flop-setup time assumed to be the same in both cases for simplicity. For a conventional domino read path, total time available for read operation as shown in Fig. 17 (T_{CONV}):

$$T_{CONV} = T_{BL\_CONV} + T_{NA\_CONV} + T_{GBL} + T_{SDL} + T_{GB\_CONV} + T_{SETUP}$$

where $T_{BL\_CONV} = 0.5 * T_{CONV}$ (worst case bitline evaluation time with 50% clock duty cycle); T_{NA_CONV} = merge-NAND delay; T_{GBL} = global bitline evaluation delay; T_{SDL} = set dominant latch delay; T_{GB_CONV} = timing guardband; T_{SETUP} = setup time of the capturing flip-flop. Let us assume: $T_{GBL} + T_{SDL} + T_{SETUP} = \eta * T_{CONV}$ and $T_{GB\_CONV} = \alpha_{CONV} * T_{CONV}$; where α_{CONV} is the portion of the total cycle time allotted for the guardbands.

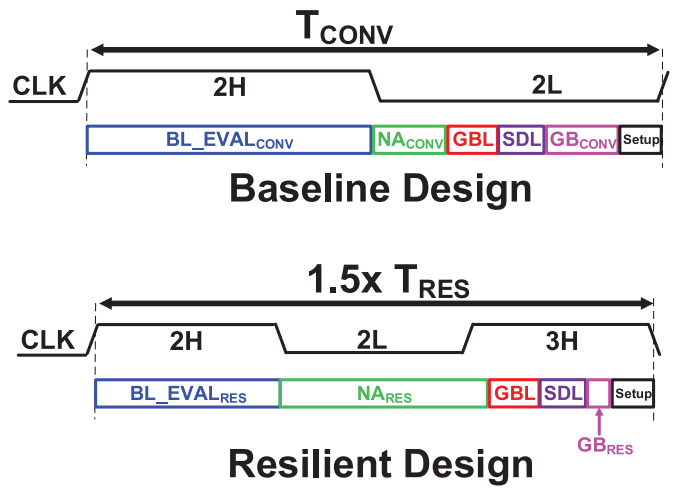

Fig. 17. Baseline design and resilient design: read cycle time break-up.

$$\therefore T_{NA\_CONV} = (0.5 - \eta - \alpha_{CONV}) * T_{CONV}. \quad (1)$$

In the conventional design during the evaluate phase, let us assume that LBL takes the entire evaluate phase to discharge to Vss and the NAND gate PMOS is triggered at this time with a gate overdrive voltage of $V_{CC} - V_T$. The NAND delay variation due to input slew rate (LBL transition in this case) is not accounted for the first order delay analysis.

$$T_{NA\_CONV} = \frac{(C_{NA} * V_{CC})}{(V_{CC} - V_T)} \quad (2)$$

$$\gamma = \frac{(\eta + \delta * \alpha_{CONV} + \beta) \pm \sqrt{(\eta + \delta * \alpha_{CONV} + \beta)^2 - 4 * [\alpha_{CONV}(1 - \beta + \beta * \delta) + \eta - 0.5(1 - \beta)]}}{2}$$

where C_{NA} = output load on merge-NAND, and V_T = threshold voltage of the PMOS in NAND. For a resilient domino read path, total time available for read operation as shown in Fig. 17:

$$1.5 * T_{RES} = T_{BL\_RES} + T_{NA\_RES} + T_{GBL} + T_{SDL} + T_{GB\_RES} + T_{SETUP}$$

where $T_{BL\_RES} = 0.5 * T_{RES}$ (worst case bitline evaluation time with 50% clock duty cycle). Let us assume GBL, SDL and capturing flip-flop setup time are the same as conventional design:

$$\therefore T_{GBL} + T_{SDL} + T_{SETUP} = \eta * T_{CONV}.$$

With the resilient approach, the reduced timing guardband is $T_{GB\_RES} = \alpha_{RES} * T_{CONV}$

$$\therefore T_{NA\_RES} = T_{RES} - (\eta + \alpha_{RES}) * T_{CONV}. \quad (3)$$

As the cycle time in the resilient design is smaller than in the conventional design, the LBL may not discharge completely and bitline swing will be reduced by the ratio of T_{RES}/T_{CONV}.

$$\Delta V_{BL} = \frac{T_{RES}}{T_{CONV}} * V_{CC} = \gamma * V_{CC}$$

where $T_{RES} = \gamma * T_{CONV}$, where γ = cycle time reduction ratio due to resilient operation

$$\therefore T_{NA\_RES} = \frac{(C_{NA} * V_{CC})}{(\gamma * V_{CC} - V_T)}. \quad (4)$$

Taking ratio of (4) and (2)

$$\therefore \frac{T_{NA\_RES}}{T_{NA\_CONV}} = \frac{(V_{CC} - V_T)}{(\gamma * V_{CC} - V_T)} = \frac{(1 - \beta)}{(\gamma - \beta)} \quad (5)$$

where $V_T = \beta * V_{CC}$, where β = ratio of the threshold voltage to the supply voltage.

Substituting (1) and (3) in (5)

$$\frac{T_{RES} - (\eta + \alpha_{RES}) * T_{CONV}}{(0.5 - \eta - \alpha_{CONV}) T_{CONV}} = \frac{(1 - \beta)}{(\gamma - \beta)}$$
$$\rightarrow \quad \frac{(\gamma - \eta - \alpha_{RES})}{(0.5 - \eta - \alpha_{CONV})} = \frac{(1 - \beta)}{(\gamma - \beta)}. \quad (6)$$

Let us denote $\alpha_{RES} = \delta * \alpha_{CONV}$, where δ = the ratio of the guardband reduction due to resilient design. Equation (6) can be rearranged as:

$$\gamma^2 - \gamma * (\eta + \delta * \alpha_{CONV} + \beta) + [\alpha_{CONV}(1 - \beta + \beta * \delta) + \eta - 0.5(1 - \beta)] = 0. \quad (7)$$

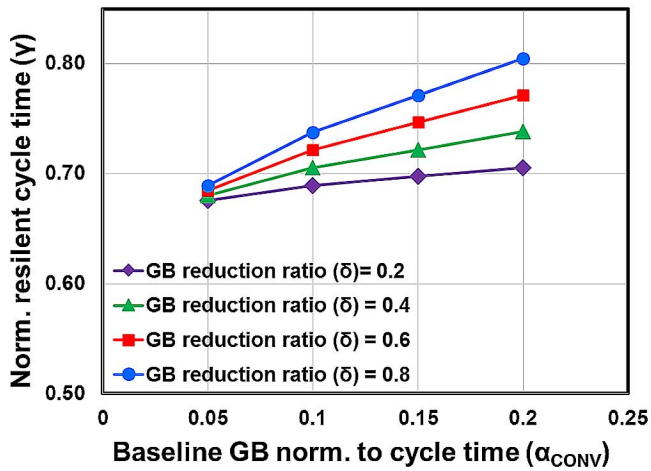

Fig. 18. Normalized resilient cycle time (γ) as a function of the ratio baseline guardband to the baseline cycle time (α_{CONV}) for different guardband reduction ratio (δ).

Equation (7) can be solved for γ to quantify the cycle time improvement in resilient design, as shown in the equation at the top of the page.Fig. 18 shows the cycle time improvement in the resilient design (γ ratio) as a function of guardband in the baseline design (α_{CONV}) for various guardband reduction ratios (δ ratio) for $\beta = 0.3$ (V_T/V_{CC}) and for $\eta = 0.1$ ($T_{GBL} + T_{SDL} + T_{SETUP}/T_{CONV}$). The trends indicate that:

1) when the guardband reduction ratio (δ) is smaller, most of the guardband is recovered in the resilient design and hence the resilient cycle time is smaller (higher frequency improvement);
2) for higher original guardband, if only a small portion of the guardband is recovered by the resilient design (higher δ ratios), cycle time improvement is lower. In this case, the cycle time improvement is compensated by the merge-NAND gate delay push out due to partial bitline evaluation.

ACKNOWLEDGMENT

The authors thank K. Ikeda, L. Avery, and D. Jenkins for chip design, and R. Forand and M. Haycock for help and support.

REFERENCES

[1] S. Borkar and A. A. Chien, "The future of microprocessors," *Commun. ACM*, vol. 54, no. 5, pp. 67–77, 2011.
[2] K. Bowman *et al.*, "Circuit techniques for dynamic variation tolerance," in *Proc. 46th Annu. Design Automation Conf.*, 2009, pp. 4–7.
[3] J. Tschanz, K. Bowman, S. Walstra, M. Agostinelli, T. Karnik, and V. De, "Tunable replica circuits and adaptive voltage-frequency techniques for dynamic voltage, temperature, and aging variation tolerance," in *Symp. VLSI Circuits Dig.*, 2009, pp. 112–113.

[4] D. Ernst *et al.*, "Razor: A low-power pipeline based on circuit-level timing speculation," *IEEE/ACM MICRO-36*, pp. 7–18, Dec. 2003.
[5] K. A. Bowman *et al.*, "A 45 nm resilient microprocessor core for dynamic variation tolerance," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 194–204, Jan. 2011.
[6] A. Raychowdhury *et al.*, "Tunable replica bits for dynamic variation tolerance in 8T SRAM arrays," *IEEE J. Solid-State Circuits*, vol. 46, no. 4, pp. 797–805, Apr. 2011.
[7] E. Karl, D. Sylvester, and D. Blaauw, "Timing error correction techniques for voltage-scalable on-chip memories," in *Proc. Int. Conf. Circuits Syst. (ISCAS)*, 2005, pp. 3563–3566.
[8] A. Drake *et al.*, "Distributed critical-path timing monitor for a 65 nm high-performance microprocessor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2007, pp. 398–399.
[9] C.-H. Jan *et al.*, "A 22 nm 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," in *Proc. Int. Electron Device Meeting (IEDM)*, 2012, pp. 4–7.
[10] S.-L. Lu, A. Alameldeen, K. Bowman, Z. Chishti, C. Wilkerson, and W. Wu, "Architectural-level error-tolerant techniques for low supply voltage cache operation," in *IEEE Int. Conf. IC Design and Technology*, 2011, pp. 1–5.
[11] S. Paul *et al.*, "A 3.6 GB/s 1.3 mW 400 mV 0.051 mm^2 near-threshold voltage resilient router in 22 nm tri-gate CMOS," in *Symp. VLSI Circuits Dig.*, 2013, pp. CC30–31.
[12] M. Cho, M. Khellah, K. Chae, K. Ahmed, J. Tschanz, and S. Mukhopadhyay, "Characterization of inverse temperature dependence in logic circuits," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC)*, 2012, pp. 1–4.
[13] S. Pae, J. Maiz, C. Prasad, and B. Woolery, "Effect of BTI degradation on transistor variability in advanced semiconductor technologies," *IEEE Trans. Device Mater. Rel.*, vol. 8, no. 3, pp. 519–525, Sep. 2008.
[14] V. J. Reddi, S. Kanev, S. Campanoni, M. D. Smith, G.-Y. Wei, and D. Brooks, "Voltage smoothing: Characterizing and mitigating voltage noise in production processors using software-guided thread scheduling," in *Proc. Annu. IEEE/ACM Int. Symp. Microarchitecture*, 2010, pp. 77–88.

Jaydeep P. Kulkarni (M'09–SM'15) received the B.E. degree from the University of Pune, Pune, India, in 2002, the M.Tech. degree from the Indian Institute of Science (IISc) Bangalore, India, in 2004, and the Ph.D. degree from Purdue University, West Lafayette, IN, USA, in 2009, all in electrical engineering.

During 2004–2005, he was with Cypress Semiconductors, Bangalore, India, where he was involved in low-power SRAM design. He is currently a staff research scientist in the Circuit Research Lab, Intel Corporation, Hillsboro, OR, USA. His research is focused on energy-efficient integrated circuit design and circuits/applications of emerging non-silicon technologies. He has filed 20 patents and published 50 papers.

Dr. Kulkarni has received the 2004 Best M.Tech. Student Award from IISc Bangalore, 2008 SRC Inventor Recognition Awards, 2008 ISLPED Design Contest Award, 2008 Intel Foundation Ph.D. Fellowship Award, Best Paper in Session Award at 2008 SRC TECHCON, 2010 Outstanding Doctoral Dissertation Award from Purdue School of ECE, 2012 Intel Patent Recognition Award, five Intel Divisional Recognition Awards, 2015 IEEE Circuits and Systems Society's TRANSACTIONS ON VLSI SYSTEMS Best Paper Award, and 2015 Semiconductor Research Corporation's (SRC) Outstanding Industrial Liaison Award. He has participated in Technical Program Committees of A-SSCC, ISLPED, and ASQED conferences and has been involved in IEEE Circuits and Systems Society's VLSI technical committee. He serves as an associate editor for IEEE TRANSACTIONS ON VLSI SYSTEMS, and as an industrial liaison at the SRC and SONIC STARnet research program.

Carlos Tokunaga (S'98–M'08) received the B.S. degree in electronics engineering from the University of Los Andes, Bogota, Colombia, in 2001, and the M.S. and Ph.D. degrees in electrical engineering from the University of Michigan, Ann Arbor, MI, USA, in 2005 and 2008, respectively.

He is currently a Research Scientist at the Circuit Research Lab, Intel, Hillsboro, OR, USA. His research interests include VLSI design with particular emphasis on energy-efficient resilient circuits and security-based circuit design.

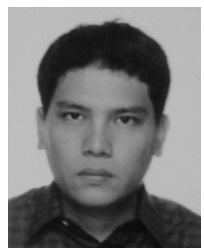

Paolo A. Aseron (M'07) received the B.S. degree in computer engineering from the University of the Philippines in 2001.

He has been with Intel Labs of Intel Corporation, Hillsboro, OR, USA, since 2006. Prior to Intel, he worked for Canon on Systems-on-a-Chip platform development from 2001 to 2003. His interests include high-performance low-power architecture and circuits, memory, and integrated power delivery.

Trang Nguyen received the A.A.S. degree in electronic engineering technology from Portland Community College, Portland, OR, USA, in 2002.

Since 2002, she has been with Intel Corporation, Hillsboro, OR, USA, as a Laboratory Technician at the Circuit Research Laboratories, under the ADR group.

Charles Augustine (S'08–M'11) received the Bachelor degree in electronics from BITS, Pilani, India, in 2004, and the Ph.D. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, USA, in 2011.

He is currently a Senior Research Scientist in the Circuit Research Lab (CRL) at Intel Corporation in Hillsboro, OR, USA. His primary research interests include ultra-low-power memory and logic circuits for GPUs, and circuit/architecture design for neuromorphic systems. He has held positions at Texas Instruments, ST Microelectronics, Philips Semiconductors, and Freescale Semiconductor, where he worked on CMOS digital integrated circuits and memories, including spin-torque based memories. He has published more than 50 papers in refereed journals and conferences and has filed nine patents (pending).

Dr. Augustine received 2015 Mahboob Khan Outstanding Industry Liaison Award from SRC, Best Paper Award in International Symposium on Low Power Electronics and Design (ISLPED) in 2012, Best Paper in Session Award at SRC Techcon in 2009, AMD Design Excellence Award from Purdue in 2008, nominated for Best Paper Award at International Symposium on Quality Electronic Design (ISQED) in 2009, and won the Bronze Medal for academic excellence from BITS, Pilani in 2004.

James W. Tschanz (M'99) received the B.S. degree in computer engineering and the M.S. degree in electrical engineering from the University of Illinois at Urbana-Champaign, IL, USA, in 1997 and 1999, respectively.

Since 1999, he has been a circuits researcher with the Intel Circuit Research Lab in Hillsboro, OR, USA, where his research interests include low-power digital circuits, design techniques, and methods for tolerating parameter variations. He also taught VLSI design for seven years as an adjunct faculty member at the Oregon Graduate Institute in Beaverton, OR, USA. He has published 53 conference and journal papers in this field, has authored three book chapters, and has over 41 issued patents.

Vivek De (M'89–SM'07–F'11) received the B.Tech. degree from the Indian Institute of Technology, Chennai, India, the M.S. degree from Duke University, Durham, NC, USA, and the Ph.D. degree from Rensselaer Polytechnic Institute, Troy, NY, USA, all in electrical engineering.

He is an Intel Fellow and Director of Circuit Technology Research in Intel Labs. He is responsible for providing strategic technical directions for long term research in future circuit technologies and leading energy efficiency research across the hardware stack. He has 231 publications in refereed international conferences and journals and 199 patents, with 30 more patents filed (pending).

Dr. De received an Intel Achievement Award for his contributions to an integrated voltage regulator technology. He received a Best Paper Award at the 1996 IEEE International ASIC Conference, and nominations for Best Paper Awards at the 2007 IEEE/ACM Design Automation Conference (DAC) and 2008 IEEE/ACM International Conference on Computer-Aided Design (ICCAD). One of his publications was recognized in the 2013 IEEE/ACM Design Automation Conference (DAC) as one of the "Top 10 Cited Papers in 50 Years of DAC". He has been a member of the ISSCC High Performance Digital Subcommittee since 2013. He was an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I during 2008–2010 and the IEEE TRANSACTIONS OF VLSI SYSTEMS during 2011–2015. He has been an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS since December 2014.