The Role of Adaptation and Resiliency in Computation and Power Management

Special Session Paper

Arijit Raychowdhury, Saad Bin Nasir, Samantak Gangopadhyay Georgia Institute of Technology, GA, USA Contact Email: arijit.raychowdhury@ece.gatech.edu

I. Introduction

Variations, both static and dynamic in nature, impact power-efficiency the performance and of microprocessors and other digital integrated circuits. While static process variations can often be mitigated through binning or by post-silicon tuning, dynamic variations change as a function of time and environment and cannot be reduced by static tuning. Examples of these types of dynamic variations include power supply (V_{CC}) droop, temperature change, and device aging over time. The droop magnitude, frequency and duration depend on the complex interaction of capacitive and inductive parasitics at the board, package, and die levels and change with the workload and current demand [1-3]. V_{CC} droops contain high-frequency (i.e., fast changing) and lowfrequency (i.e., slow changing) components and occur locally and/or globally across the die. Temperature variations occur at a relatively slower time scale with local hot spots on the die, depending on environmental and workload conditions as well as the heat-removal capability of the package. Further, transistor aging slowly degrades the drive current over time as a function of gate bias and temperature conditions. Conventional designs build in V_{CC} guardband (V_{CC} GB) at a target frequency to ensure correct functionality even in the presence of worst-case dynamic variations. In the digital core, this GB accounts for the delay increase on critical paths when dynamic variations occur and guarantees that timing constraints are met even under worst-case conditions. In the storage arrays (register files and caches) this GB is added to the static minimum operating voltage (V_{MIN}) to obtain the resultant operating V_{MIN} and results in an increase of both leakage and dynamic array power. The key design aspect of adaptive and resilient designs is to remove a part of this GB, such that logic timing and memory functionality can be met at a lower V_{CC}. This improves energy efficiency and improves battery life; a key consideration in mobile and handheld platforms. The mitigation of the design GB can be achieved in two fundamentally different ways. For dynamic variations with time scales in the order of hundreds of us to ms (such as temperature variations or transistor aging), a reactive scheme where any functional error is avoided is applicable. For example, embedded temperature sensors can be used to trigger changes in the clock frequency or changes in the supply voltage, such that at low temperatures (when transistor currents are higher) a faster clock (at iso- V_{CC}) or a lower supply (at iso-clock frequency) can be employed [4]. For dynamic changes which occur in time scales of ns to us, such reactive schemes are impossible by design. Hence, an alternative scheme can be employed where timing errors (in logic) or readwrite errors (in memory) are allowed to happen. Once errors occur, they are detected, and the errant instruction is replayed with proper care such that data integrity is maintained. A combination of such adaptive schemes (Fig. 1) result in the mitigation of a major part of the design GB, resulting in higher energy efficiency.

On a similar note, power management blocks in microprocessors and systems on chip (SoCs) are typically designed for the worst case conditions. As an example embedded voltage regulators are mostly designed to provide stability at the worst case load conditions. Although it provides correct operation



Fig. 1. Adaptation for error free operation in logic and memory in response to slow changing environmental variables (like temperature). Resiliency for error detection and correction in response to voltage droops.

978-1-4799-6278-5/14/\$31.00 ©2014 IEEE



Fig. 2. (a) Embedded sensors allow close loop control of the operating voltage and frequency in processors. (b) Measurements illustrate dynamic frequency control in response to temperature changes [4].

across a wide dynamic range, the resultant design is less efficient across the entire range. Hence research has started in earnest to provide design solutions in embedded power management blocks, such that with changing load currents, the loop dynamics in the embedded voltage regulators can be adapted for more consistent performance and higher energy efficiency across a wide range of load currents. In this paper we provide an overview of design ideas and solutions that enable adaptation and resiliency in logic, memory and power management blocks.

II. Adaptation and Resiliency in Logic

There are two main approaches to reducing the impact of guardbands for dynamic variations: error avoidance and error correction. Error avoidance typically uses sensors to monitor the variation source (dc voltage, temperature, or aging) and a control circuit which attempts to adapt the circuit by adjusting supply voltage, body bias, or frequency. This technique works well for slow-changing variations such as temperature that also do not have significant spatial variation across the die or the area under control. Such example designs have been reported in [4]. Fig. 2a illustrates a typical design where embedded sensors have been used to trigger a dynamic voltage-frequency (DVF) block which can be used to control the resultant frequency of operation. It can be noted in Fig. 2b how an increase in temperature can be used to reduce the clock frequency such that an adaptive and dynamic design can be realized.

Although slow changing variations can be mitigated by adaptive and reactionary designs, fast-changing variations such as voltage droop do not provide sufficient time for the sensing, communication, and adaptation to occur before the variation results in circuit failure. Under nominal conditions, when worstcase variations are not present, there is sufficient time for data to propagate down each critical path and be latched into the receiving state element, and no error occurs. However, under the presence of a dynamic variation the critical path delay increases, resulting in incorrect data latched into the receiving latch or flipflop. This condition must be detected by an errordetection circuit.

1) Error-Detection Sequential (EDS): The first method for error detection involves replacing the critical-path flip-flops with error-detection sequentials (EDS). These circuits (Fig. 3) effectively double-sample the input data to the sequential – once on the rising clock edge and again on the falling clock edge [5-12]. If these two samples differ, data has arrived late to the flip-flop and this is assumed to be a timing error. The high phase of the clock therefore represents



Fig 3. Error-detection sequential (EDS). When mode=1, this circuit functions as a standard flip-flop without error detection [11].



Fig 4. Tunable replica circuit (TRC). Delay is tuned during test to match the delay of a pipeline critical path [10, 11].



Fig 5. Measured throughput (TP), as normalized to the conventional maximum TP, and recovery cycles, as a percentage of total cycles, versus clock frequency (F_{CLK}) for a benchmark program in a resilient and adaptive microprocessor [11].

the "error-detection window" – the period of time in which a timing error can be detected. Under normal operation, the data should never arrive in this window, meaning that even short (min-delay) paths must arrive in the low phase of the clock. Therefore there is a trade-off in the design of this timing window – a wider error-detection window allows larger variations to be detected, but results in a more stringent min-delay requirement.

2) Tunable Replica Circuits (TRC): An alternate form of error detection uses a replica of a critical path, rather than the critical paths themselves, to detect timing errors [8-10]. In this approach, a tunable replica circuit (TRC) is provided which can be tuned post-silicon to match the delay of the critical paths on the die (Fig. 4). An EDS at the end of this path reports a timing error if the delay of this path exceeds the clock period. It is assumed that if the TRC fails, an actual critical path on the die may have failed as well, and error correction is performed. Thus, the replica circuit is used as a predictor of errors for the processor pipeline. While technique requires additional this post-silicon calibration and does not detect path-level within-die dynamic variations as the EDS technique does, it results in less area and power overhead, and is easier to implement.

III. Adaptation and Resiliency in Memory Arrays Reactive mechanisms such as dynamic voltage and frequency adjustment, as has been described above, can be used in memory arrays as well. One such design has been reported in [13] where an adaptive word-line voltage has been used for adaptive boosting. However, error detection under dynamic V_{CC} droops in memory arrays poses design challenges. Details of such a design for 8T cell arrays have been shown in [14] and an overview is presented here.

The 8T SRAM cell [13,14] used in the present design offers fast read (RD) and write (WR), dual-port capability, and generally supports lower minimum V_{CC} (or V_{MIN}) than the 6T cell. By using a decoupled single-ended RD port with domino-style hierarchical RD bit-line, the 8T cell features fast RD evaluation path without causing access disturbance that limits RD V_{MIN} in the 6T cell. Using the 8T cell in a half-selectfree architecture eliminates pseudo-reads during partial writes, hence enabling WR V_{MIN} optimization independent of RD. As power limits are aggressively reduced, thus demanding a lower V_{MIN} of the 8T cell, both within-die (WID) and die-to-die (D2D) device parameter variations are getting worse with feature size scaling. The typical approach of sizing up the 8T RD and WR ports to mitigate process variation has limited



Fig 6. Circuit schematic for TRBRD.





 V_{MIN} returns. The use of Tunable Replica Bits (TRBs) can potentially reduce or eliminate a significant part of the V_{CC} GB in 8T SRAM arrays. The usage model for the TRB involves the following three steps:

(a) Experimentally determining the nominal (i.e., static, under no dynamic variation) V_{MIN} of the 8T SRAM array for both RD and WR (V_{MINRD} and V_{MINWR}),

(b) Separately calibrating a read TRB (TRBRD) circuit and a write TRB (TRBWR) circuit to fail once V_{CC} falls below the static RD and WR V_{MIN} , plus a small TRB margin with the resultant voltages called V_{TRBRD} and V_{TRBWR} , respectively, as given by:

 $V_{TRBRD} = V_{MINRD} + TRB Margin for RD$ ------ [1a]

 $V_{TRBWR} = V_{MINWR} + TRB Margin for WR$ ------ [1b]

The TRB margin arises from the ability to regulate the TRB voltages using an embedded calibration circuit and the output resolution of such a circuit.

(c) Operating the array and TRBs at the higher of V_{TRBRD} and V_{TRBWR}

Operating Array $V_{MIN} = Max(V_{TRBRD}, V_{TRBWR}) -- [2]$ The TRB provides a monitor which shows if an access error has occurred in the memory under a dynamic event. This monitor in conjunction with a system level recovery technique can ensure correct functionality even under such events. These techniques include, but are not limited to (1) replaying the instruction at the same CLK frequency, or (2) replaying at slower (e.g., half) CLK frequency. Care needs to be ensured that at valid data is never lost in the process of a WR failure in the array. This can be done by providing an adequate data buffer or retaining a golden copy of the data in the main memory. The circuit implementation of the TRBs is shown in Fig. 6 and 7.

The schematic diagram of the TRBRD is shown in Fig. 6. It includes (a) a timing error detection circuit for RD-"1" and (b) a noise detector for RD-"0". The timing error detector comprises of (1) an accessed 8T bit cell (storing a "1"), whose word-line (RDWL) is weakly turned on using a tuning circuit, and (2) 7 unaccessed cells (storing "0") on a common RDBL whose RDWLs are connected to V_{SS} . When a V_{CC} droop occurs, the RD port of the accessed cell is weakened and it fails to sufficiently discharge the local BL in the given cycle time which leads to a timinginduced RD failure. Similarly, the noise detector comprises of (1) an accessed cell (storing a "0") whose RDWL is fully ON and (2) 7 unaccessed cells (storing "1") on a common RDBL whose RDWL is weakly turned on by the tuning circuit. In the event of a V_{CC} overshoot, the RD ports of the unaccessed cells

become strong and can inadvertently discharge the local BL, thereby causing a noise induced failure. The timing and noise-induced RD failures are OR-ed to obtain the RD Error signal. The tuning circuit for the timing and the noise failure monitors can be individually calibrated. The tuning circuit comprises of two programmable stages: (1) a weak RD voltage generator (generated by a scan-programmable transistor-based voltage divider), and (2) a scan-programmable delay (which mimics the delay in the decoder circuits). A combination of these two design elements ensure that both the timing and noise detectors produce successful RDs for $V_{CC} \ge V_{TRBRD}$ and that one or both of the detectors fails below V_{TRBRD} .

The TRBWR is designed on the principle of weak write in the array. It is designed using a low swing WRBL voltage that causes a successful WR only for $V_{CC} \ge V_{TRBWR}$. Each detection cycle (i.e., the weak WR cycle where the TRBWR cell detects the presence of a dynamic event) is followed by three extra (or setup) cycles when the opposite data is strongly written (at



Fig 8. Measured number of single bit failures in the 16KB array with and without V_{CC} droop (A_{DROOP}=13%). A 9.5% static guardband (GB) is required [14].



Fig. 9. Demonstration of resilient microprocessor while executing the *edgedetect* benchmark at an F_{CLK} of 1.5GHz and a V_{CC} of 1.0V. (a) Input bitmap image. (b) Correct output of edge-detected image with resilient circuits enabled. (c) Output of edge-detected image when resilient circuits are disabled during processing [11].

half the frequency) and WR drivers are properly setup (Fig. 7). Consequently, each TRBWR cell provides data every four cycles. Hence four TRBWR cells work in unison in a time-interleaved fashion, such that the TRBWR generates an Error/Pass signal at the end of each CLK cycle. The weak WRBL voltage is provided by a scan programmable tuning circuit, similar in design to the tuning circuit of the TRBRD cell. The resolution of the tuning circuit depends on the granularity of the delay elements and the transistor based voltage generator network. This is further aggravated under die-to-die and within-die variations. In experimental setup, we have noted a minimum resolution of 20mV as repeatable and consistent across several dies. Fig. 8 illustrates how the number of single bit failures (SBF) can be lowered by emplying resiliency and adaptation in cache arrays.

The use of adaptation and resiliency in logic and memory results in mitigation of a major portion of the design guardband. Once a timing error has been detected through either the EDS or TRC technique for the core or the TRB technique in the cache, error correction must be performed. Error correction in the core must ensure that (1) the offending instruction, which contains erroneous data, does not commit this result to memory or change the state of the system, and (2) the erroneous instruction and all following instructions must be repeated until an error-free result is obtained. In a microprocessor this can be achieved in a similar way as a branch misprediction: the pipeline is flushed and instruction replay begins from the offending instruction. To ensure correct execution, it is also possible to reduce the clock frequency during instruction replay. Error Correction in the cache would also include replaying the instruction and ensuring that an errant WB does not corrupt the correct data. Adaptation and resiliency features allow a microprocessor to operate in conditions which would otherwise be impossible. Fig. 9 illustrates resilient and adaptive design in a microprocessor and how resiliency allows correct processing.

IV. Adaptation in Embedded Power Management Just as error avoidance (by adaptation) and error detection and correction are design principles geared towards improving the system energy efficiency, adaptive designs in embedded power management can also provide efficient voltage regulation under a wide dynamic range of operation. Here we provide such a design choice in digital low dropout regulators. A digital linear regulator [16, 17] and the corresponding linearized control model are shown in Fig. 10. In this design a comparator compares the regulated output voltage (V_{OUT}) with the reference voltage. Depending on whether V_{OUT} is more or less than the reference, power MOSFETS are either turned on or turned off till regulation is achieved. From the control model of such a regulator, as has been described in [17], it is noted that the sampling clock frequency determines the location of the closed loop system poles. This is illustrated in Fig.



Fig. 10. (a) Circuit schematic and (b) Control model of a digital linear regulator.



Fig. 11: SPICE simulations showing the change of loop behavior as the load condition changes from 3.5mA to 350uA at iso- sampling frequency in response to a step load.



Fig. 12. Power efficiency of the adaptive controller illustrating that a severe loss of efficiency at lighter load conditions can be compensated for by employing adaptive loop control. The $F_{SAMPLING}$ corresponding to 50uA, 500uA and 5mA load currents are F_{LOW} , $F_{NOMINAL}$ (=3X F_{LOW}) and F_{HIGH} (=3X $F_{NOMINAL}$). The case of "no adaptation" illustrates the design for worst case load. The theoretical maximum for power efficiency is 70% (V_{OUT} =0.7V and V_{DD} =1V).

11 where with iso-sampling frequency, as the load current changes, the system moves from an overdamped to an unstable or oscillatory step response. This necessitates the use of adaptation in the control loop of the regulator to achieve a stable and consistent step response. This is achieved in [17] by sensing the location of the output pole (due to the load resistor) and adjusting the sampling frequency accordingly. It can be also shown that the resultant design mitigates a part of the design guardband that exists in case of light load conditions. Hence, the overall design improves the energy efficiency by as much as 50% as shown in Fig. 12.

V. Conclusions

This article provides an overview of adaptation and resiliency as design parameters in energy efficient systems. By allowing embedded sensor based adaptation in logic, memory and embedded voltage regulators, we can mitigate a part of the design guardband enabling lower operating power across a wide dynamic range.

Acknowledgements: The author would like to acknowledge the contributions of J. Tschanz, K. Bowman, M. Khellah, T. Karnik and V. De of Intel Corporation. A part of the work was sponsored by the Semiconductor Research Corporation under grant 1836.140.

References:

- Muhtaroglu, G. Taylor, and T. R. Arabi, "On-Die Droop Detector for Analog Sensing of Power Supply Noise," *IEEE J. Solid-State Circuits*, pp. 651-660, Apr. 2004.
- [2] T. Fischer, J. Desai, B. Doyle, S. Naffziger, and B. Patella, "A 90-nm Variable Frequency Clock System for

a Power-Managed Itanium Architecture Processor," *IEEE J. Solid-State Circuits*, pp. 218-228, Jan. 2006.

- [3] R. McGowen, et al., "Power and Temperature Control on a 90-nm Itanium Family Processor," *IEEE J. Solid-State Circuits*, pp. 229-237, Jan. 2006.
- [4] J. Tschanz, et al., "Adaptive Frequency and Biasing Techniques for Tolerance to Dynamic Temperature-Voltage Variations and Aging," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 292-293.
- [5] D. Ernst, et al., "Razor: A Low-Power Pipeline Based on Circuit-Level Timing Speculation," in *Proc. IEEE/ACM Intl. Symp. Microarchitecture (MICRO-36)*, Dec. 2003, pp. 7-18.
- [6] S. Das, et al., "A Self-Tuning DVS Processor Using Delay-Error Detection and Correction," *IEEE J. Solid-State Circuits*, pp. 792-804, Apr. 2006.
- [7] S. Das, et al., "Razor II: In Situ Error Detection and Correction for PVT and SER Tolerance," *IEEE J. Solid-State Circuits*, pp. 32-48, Jan. 2009.
- [8] K. A. Bowman, et al., "Energy-Efficient and Metastability-Immune Resilient Circuits for Dynamic Variation Tolerance," *IEEE J. Solid-State Circuits*, pp. 49-63, Jan. 2009.
- [9] J. Tschanz, et al., "Tunable Replica Circuits and Adaptive Voltage-Frequency Techniques for Dynamic Voltage, Temperature, and Aging Variation Tolerance," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, June 2009.
- [10] J. Tschanz, et al., "A 45nm Resilient and Adaptive Microprocessor Core for Dynamic Variation Tolerance," in *Proc. IEEE ISSCC Dig. Tech. Papers*, Feb. 2010, pp. 282-283.
- [11] Bowman, et al., "A 45nm Resilient Microprocessor Core for Dynamic Variation Tolerance," in JSSC, Jan 2011, pp. 194-207
- [12] A. Raychowdhury et al.., "Error Detection and Correction in Microprocessor Core and Memory Due to Fast Dynamic Voltage Droops," Emerging and Selected Topics in Circuits and Systems, IEEE Journal on, vol.1, no.3, pp.208,217, Sept. 2011
- [13] A. Raychowdhury et al., "PVT-and-aging adaptive wordline boosting for 8T SRAM power reduction," Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International, vol., no., pp.352,353, 7-11 Feb. 2010
- [14] Raychowdhury et al., "Tunable replica bits for dynamic variation tolerance in 8T SRAM," in Proc of VLSI Circuit Symposium, June 2010, pp. 39-40
- [15] Wang et al., "Techniques to Extend Canary-Based Standby VDD Scaling for SRAMs to 45 nm and Beyond" *IEEE J. Solid-State Circuits*, pp. 2514, Nov. 2008.
- [16] Y. Okuma, et al., "0.5-V input digital LDO with 98.7% current efficiency and 2.7-uA quiescent current in 65nm CMOS," *IEEE Custom Integrated Circuits Conference*, pp. 1-4, Sept. 2010.
- [17] Gangopadhyay et al. ""Modeling and analysis of digital linear dropout regulators with adaptive control for high efficiency under wide dynamic range digital loads," Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014, vol., no., pp.1,6, 24-28 March 2014