

Wide-Range Many-Core SoC Design in Scaled CMOS: Challenges and Opportunities

Sriram Vangal¹, Senior Member, IEEE, Somnath Paul¹, Member, IEEE, Steven Hsu, Member, IEEE, Amit Agarwal¹, Member, IEEE, Saurabh Kumar¹, Ram Krishnamurthy¹, Fellow, IEEE, Harish Krishnamurthy¹, Member, IEEE, James Tschanz¹, Member, IEEE, Vivek De¹, Fellow, IEEE, and Chris H. Kim¹, Fellow, IEEE

Abstract—The system-on-chip (SoC) designs for future Internet of Things (IoT) systems, spanning client platforms to cloud datacenters, need to deliver uncompromising and scalable performance with extreme energy efficiency for diverse workloads and applications, while satisfying a wide range of energy budgets, as well as platform cooling and power delivery constraints. Low-latency, burst-mode responsiveness, and scalable high-throughput performance must be delivered on demand for a range of thread-parallel, task-parallel, and data-parallel workloads covering traditional and emerging applications. This article discusses the challenges and opportunities for many-core SoC design in scaled CMOS process operating over a wide voltage-frequency range including near-threshold-voltage (NTV) that can meet the compute demands of the future at scale, flexibly, and efficiently. This article covers: 1) circuit design techniques for NTV cores; 2) mitigation techniques for within-die parameter variations via multivoltage frequency schemes; 3) digital integrated voltage regulators (VRs) for fine-grain and wide-range voltage modulation; and 4) radiation-induced soft error rate (SER) characterization and mitigation techniques to enable reliable operation at NTV. Silicon prototype examples will be used to illustrate the different techniques and highlight future research directions.

Index Terms—Digital low dropout (LDO), low-power, low-voltage SRAM memory circuits, minimum-energy design, near-threshold-voltage (NTV) computing, power-performance, resilient adaptive computing, soft error rate (SER) characterization, variation-aware many-core dynamic-voltage-frequency scaling (DVFS).

I. INTRODUCTION

NEAR-THRESHOLD computing promises dramatic improvements in energy efficiency. For many CMOS designs, the energy consumption reaches an absolute minimum in the near-threshold-voltage (NTV) regime that is of the order of magnitude improvement over super-threshold operation [1]–[3]. However, frequency degradation due to aggressive voltage scaling may not be acceptable across all single-threaded or performance-constrained applications. The

key challenge is to lock-in this excellent energy efficiency benefit at NTV, while addressing the impacts of: 1) loss in silicon frequency; 2) increased performance variations; and 3) higher functional failure rates in memory and logic circuits. Enabling digital designs to operate over a wide voltage range is key to achieving the best energy efficiency [2] while satisfying varying application performance demands. To tap the full latent potential of NTV, multilayered co-optimization approaches that crosscut architecture, devices, design, circuits, tool flows and methodologies, and coupled with fine-grain power management techniques are mandatory to realize NTV circuits and systems in scaled CMOS processes.

In this article, we present several multivoltage designs across four technology nodes, featuring many-core system-on-chip (SoC) building blocks. The silicon prototypes, designed at Intel Labs, Hillsboro, OR, USA, and in partnership with academia, include an 80-core TeraFLOPS processor [4] in 65-nm technology [5], a NTV-enhanced Pentium class IA-32 CPU [6] built using second generation 32-nm high- k /metal gate transistors [7], two 22-nm designs featuring 3-D trigate and high- k /metal gate devices [8]—a single instruction, multiple data (SIMD) vector permute engine [9], and a resilient 2×2 , 2-D mesh network-on-chip (NoC) fabric [10]. Key data from four more test-chips designed using a 14-nm second generation trigate CMOS process [11]—a mm-scale, 0.79-mm² NTV IA-32 Quark microcontroller (MCU) [12], [13] for μ W wireless sensor nodes (WSNs), a Quad-core Atom processor with digital low dropout (LDO) linear regulators [14] and irradiation soft error rate (SER) test structures [15], [16] to accurately analyze memory and logic SER are highlighted. The NTV-optimized designs demonstrate a wide dynamic power-performance range, including reliable NTV regime operation for maximum energy efficiency.

This article is organized as follows. Section II describes various NTV design techniques for SRAM and logic circuits. Architecture-driven adaptive mechanisms to address higher functional failure rates and variation-tolerant resiliency at NTV for SoC fabrics are also presented, including the tools, flows, and recipes for wide-dynamic range design. Moderation techniques for within-die (WID) parameter variation are discussed in Section III. Section IV presents digitally controlled voltage regulators (VRs) that enable per-core

Manuscript received October 4, 2020; accepted November 28, 2020.
(Corresponding author: Sriram Vangal.)

Sriram Vangal, Somnath Paul, Steven Hsu, Amit Agarwal, Saurabh Kumar, Ram Krishnamurthy, Harish Krishnamurthy, James Tschanz, and Vivek De are with Circuit Research, Intel Labs, Intel Corporation, Hillsboro, OR 97124 USA (e-mail: sriram.r.vangal@intel.com).

Chris H. Kim is with the Department of Electrical Engineering and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TVLSI.2021.3061649>.

Digital Object Identifier 10.1109/TVLSI.2021.3061649

dynamic-voltage-frequency scaling (DVFS) to significantly improve the overall energy efficiency of multicore processors. In addition, in Section V, results from radiation-induced SER characterization circuits on memory and logic circuits operating down to NTV are presented and analyzed, before concluding this article in Section VI.

II. NTV CIRCUIT DESIGN METHODOLOGY

The most common limit to voltage scaling is the failure of SRAM and logic circuits. SRAM cells fail at low voltage because device mismatches degrade the stability of the bit-cell for read, write, or data retention. SRAM cells typically use the smallest transistors. Also, they are the most abundant among all circuit types on a die. Therefore, the V_{\min} of the SRAM cell array limits V_{\min} of the entire chip. Logic circuits, clocking, and sequentials fail at low voltage because of noise and process variations. Alpha and cosmic ray-induced soft errors cause transient failure of memory, sequentials, and logic at NTV. Frequency starts degrading exponentially as the supply voltage approaches the device threshold voltage V_T . This also sets a limit on the minimum operational voltage (V_{\min}). This limit can be alleviated to some extent by trigate transistors. Since they have a steeper subthreshold swing, they can provide a lower V_T for the same leakage current target. Aging degradations cause the failure of SRAM cells at low voltages since different transistors in the cell undergo different amounts of V_T shift under voltage–temperature stress and thus worsen device mismatches in the bit-cells. All these effects degrade and limit V_{\min} . Sections II-A–II-F describe low-voltage design techniques used for SRAM memory, combinational cells, sequentials, and voltage-level shifter circuits.

A. SRAM Memory and Register File Optimizations

An 8-T SRAM cell [Fig. 1(a)] is commonly used in single- V_{DD} microprocessor cores, particularly in performance critical low-level caches and multiported register-file (RF) arrays. The 8-T cell offers fast simultaneous read and write, dual-port capability, and generally lower V_{\min} than the 6-T cell. With independent read and write ports in the 8-T cell, significantly improved read noise margins can be realized over the traditional 6-T SRAM cell, at an additional area expense. The noise margin improvement is due to the elimination of the read-disturb condition of the internal memory node by the introduction of a separate read port in the SRAM cell. As a result, variability tolerance is greatly enhanced, making it a desirable design choice for ultralow power (ULP) SRAM memory operating at lower supply voltages down to NTV and energy-optimum points.

The 8-T bit-cell is still prone to write failures due to write contention between a strong PMOS pull-up and a weak NMOS transfer device across process, voltage, and temperature (PVT) variation. This contention becomes worse as V_{DD} is lowered, limiting V_{\min} . A variation-tolerant dual-ended transmission gate (DETG) cell is implemented on the 22-nm NTV-SIMD RF array by replacing the NMOS transfer devices with full transmission gates (TGs) [Fig. 1(b)]. This design enables a strong “1” and “0” write on both sides of the cross-coupled inverter pair. The DETG cell always has two NMOS or two

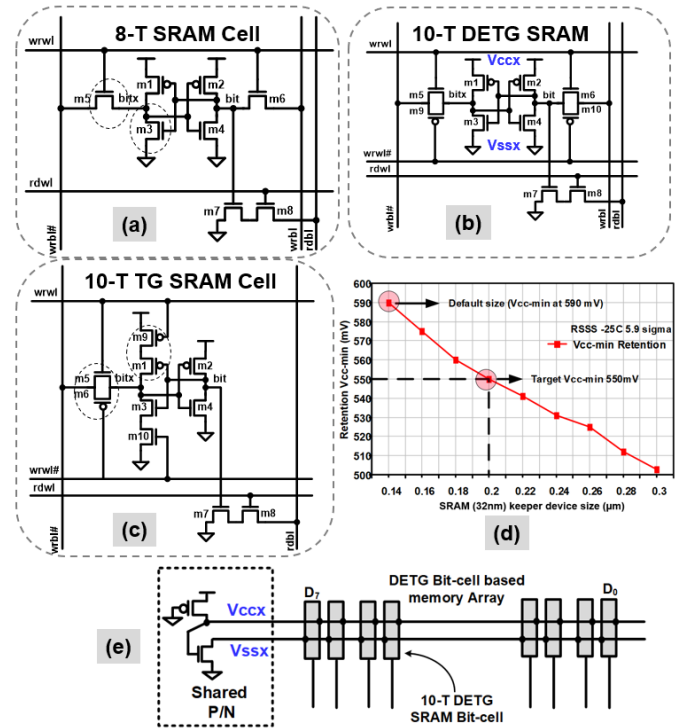


Fig. 1. Prototypes use variability-tolerant SRAM bit-cells. (a) 8-T SRAM bit-cell used in the NTV-MCU [13]. (b) SIMD engine uses a 10-T DETG SRAM topology [9]. (c) Alternate 10-T TG SRAM bit-cell used in the NTV-CPU [6]. (d) Simulated retention voltage simulations for the 10-T TG SRAM in 32 nm, as a function of keeper device size (m9, m10) in the presence of random variations (5.9σ , slow skew, -25°C). (e) Shared PMOS/NMOS on the virtual supplies improve memory write V_{\min} by 125 mV in the 22-nm DETG-based memory array.

PMOS devices to write a “1” or “0,” on nodes bit and bitx. This inherent redundancy averages the random variation effect across the transistors, improving both contention and write-completion. Moreover, the cell is symmetric with respect to PMOS and NMOS skew, which reduces the effect of systematic variation. DETG cell simulations show a 24% improvement in write delay, allowing a 150-mV reduction in write V_{\min} . However, the DETG cell is contention limited at its write V_{\min} , which can be reduced by the shared P/N circuits. An always “ON” PMOS and NMOS is shared across the virtual supplies of eight DETG cells [Fig. 1(e)]. The shared P/N circuit limits the strength of the cross-coupled inverters across variations reducing write contention by 22%. This circuit optimization results in an additional 125-mV write reduction compared to DETG, enabling an overall 275-mV write V_{\min} reduction when compared to the 8-T SRAM cell.

Caches in the 32-nm NTV-CPU use a modified, single-ended, and fully interruptible 10-T TG SRAM bit-cell [Fig. 1(c)], which allows for contention-free write operations. This topology enables a 250-mV improvement in write V_{\min} over an 8-T bit-cell. With this improvement, bit-cell retention now becomes a key V_{DD} limiter. The simulated retention voltage data for 32-nm 10-T TG SRAM, as a function of keeper device size (m9, m10) and in the presence of random variations (5.9σ , slow skew, -25°C) is shown in Fig. 1(d). Clearly, larger keeper devices lower the retention voltage. The keeper device is increased from 140 to 200 nm to realize

TABLE I
COMPARISON BETWEEN 8T SRAMS BUILT
WITH ULP AND SP BIT-CELLS

8T SRAM device type	Gate pitch	Normalized frequency (0.5V)	Normalized Leakage (0.5V, 25C)	14-nm bit-cell area (μm^2)
Standard performance (SP)	70nm	5×	26×	0.100 μm^2
Ultralow power (ULP, NTV MCU Memory)	84nm	1×	1×	0.155 μm^2 (1.55×

a 550-mV retention V_{\min} target. For a reliable read operation, bit-lines incorporate a scan-controlled, programmable stacked keeper, which can be configured to three or four PMOS device stacks to reduce read contention and improve read V_{\min} , across a wide operating voltage/frequency (V/F) range.

To achieve low standby power in the WSN, all on-die memories and caches on the 14-nm NTV-MCU use a custom 8-T [Fig. 1(a)], 0.155- μm^2 bit-cell, built using 84-nm gate pitch ULP transistors [13]. The 8-T bit-cell provides a well-balanced tradeoff in V_{\min} and area over the 6-T and 10-T SRAM cells. The ULP transistor-optimized memory arrays are designed to provide low standby leakage. However, as summarized in Table I, a 5× performance slowdown is estimated over standard performance (SP) transistor 8T memory at 500 mV but is still fast enough for edge computing applications. Context-aware power-gating of each 2-kB array is supported for further leakage reduction with no state retention. The ULP array also enables 26× lower leakage (at 500-mV supply) and has a 55% area cost over an SP-based 8T memory array, drawn on a 70-nm gate pitch. The ULP memory leakage scales from 114 pA at 1-V voltage down to 8.28 pA per bit at the retention limit of 308 mV, as measured at room temperature (25 °C).

PVT and aging adaptive on-die boosting of read word-line (RWL) and write word-line (WWL) as a common circuit assist technique for further lowering SRAM V_{\min} is described in [17] and [18]. Boosting RWL enables a larger read “ON” current without forcing a larger PMOS keeper. Boosting WWL helps write V_{\min} for two reasons—it improves contention without upsizing NMOS pass device size (or lowering its V_T) and improving write completion by writing a “1” from the other side. At the iso-array area, on-die WL boosting achieves twice as much V_{\min} reduction over bit-cell upsizing [17]. However, word-line boosting requires an integrated charge-pump or another method for generating a boosted voltage on die.

B. Combinational Cells Design Criteria

Circuits are optimized for robust and reliable ultralow voltage operation. A variation-aware pruning is performed on the standard cell library to eliminate the circuits which exhibit dc failures or extreme delay degradation at NTV due to reduced transistor ON/OFF current ratios and increased sensitivity to process variations. Simulated 32-nm normalized gate delays (y-axis), as a function of VDD for logic devices in

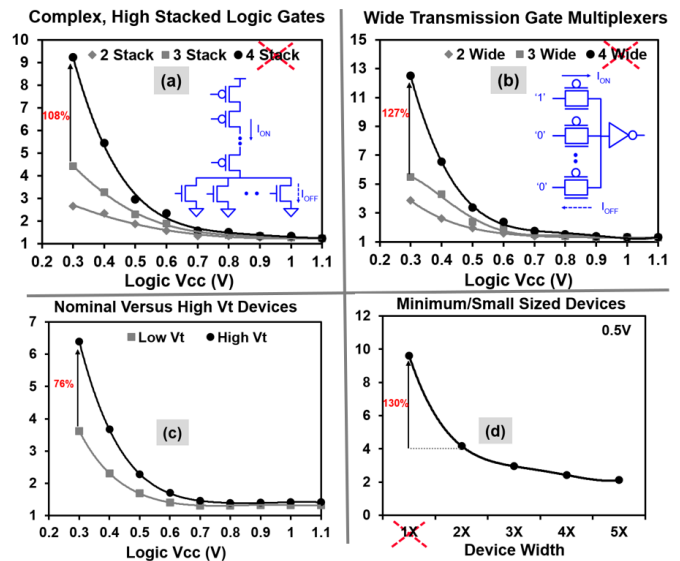


Fig. 2. Simulated 32-nm normalized gate delays (y-axis) versus supply voltage for logic devices in the presence of random variations (6σ). To limit excessive gate delays at NTV, the data indicate that (a) transistor stack sizes need to be limited to three, including (b) wide pass-gate multiplexers, (c) high VT devices have a 76% higher delay penalty over nominal VT flavors due to variations, and (d) minimum width (1×, Z_{\min}) devices show 130% higher delay at 500 mV, requiring restricted use.

the presence of random variations (6σ) is presented in Fig. 2. Complex logic gates with four or more stacked devices and wide TG multiplexers with four or more inputs are pruned from the library because they exhibit more than 108% and 127% delay degradation compared to three-stack or three-wide multiplexers, respectively [Fig. 2(a) and (b)]. Critical timing paths are designed using low VT devices because high VT devices indicate a 76% higher delay penalty at 300-mV supply, in the presence of variation [Fig. 2(c)]. All minimum-sized gates with transistor widths less than 2× of the process-allowed minimum (Z_{\min}) are filtered from the library due to 130% higher variation impact [Fig. 2(d)], and the use of single fin-width devices is limited in logic design in technology nodes below 22 nm.

C. Sequential Circuit Optimizations

At lower supply voltages, degradation in transistor $I_{\text{ON}}/I_{\text{OFF}}$ ratio, random and systematic process variations, affect the stability of storage nodes in flip-flops. Conventional TG-based master-slave flip-flop circuits typically have weak keepers for state nodes and larger TGs. During the state retention phase, the ON-current of the weak keeper contends with the OFF-current of the strong TG affecting state node stability. Additionally, charge-sharing between the internal master and slave nodes (write-back glitch) can result in state bit-flip due to reduced noise margins at low VDD. The NTV-CPU employs custom sequential circuits to ensure robust operation at lower voltages under process variations. A clocked CMOS-style flip-flop design (Fig. 3) replaces master and slave TGs with clocked inverters, thereby eliminating the risk of data write-back through the pass gates. In addition, keepers are upsized to improve state node retention and are made fully




Fig. 3. Low-voltage pass-gate free sequential circuit. (a) Clocked-CMOS-style flip-flop implementation. (b) Clocked inverter logic gate.

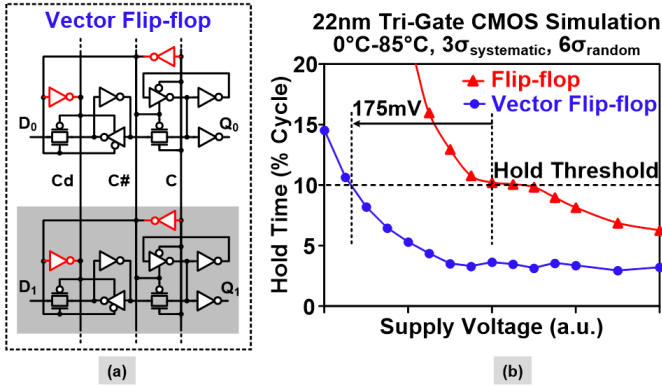


Fig. 4. Averaging impact of device parameter variations at NTV. (a) Vector flip-flop topology with shared clock input drivers. (b) Simulated hold time improvement in a 22-nm node in the presence of random variations (6σ).

interruptible to avoid contention during the write phase of the clock, thus improving V_{\min} .

Clock edges also degrade severely at low voltages, since local clock drivers inside flip-flops are small (Z_{\min}) and prone to process variations. This can cause hold-time degradations and min-delay failures. The NTV-SIMD engine employs vector flip-flops, where the clock outputs of local drivers are shared across the flip-flops. As shown in Fig. 4(a), the drivers on nodes C, C# and Cd are shared across multiple flip-flops. This helps average the impact of device parameter variations. In addition, vector flip-flops enable lower clock power since they present a reduced input capacitance to the clock tree drivers. Under worst case variation, if one of the local clock inverters becomes weak, the other shared clock inverter will compensate for the reduced drive strength. Vector flip-flop simulations (22 nm) across two adjacent cells with shared local min-sized clock inverters [Fig. 4(b)] show better hold time violations at NTV and improved hold time margins by 175 mV. Stacked min-delay buffers limit variation-induced transistor speedup, further improving hold time margins at NTV by 7%–30%.

D. Level Shifter Circuit Optimizations

NTV designs, operating at low supply voltages require level shifters to communicate with circuits at the higher voltages (e.g., I/O). Similar to RF writes, conventional cascode voltage switch logic (CVSL) level shifters are inherently contention circuits. The need for wide range, ultralow voltage level shifter to a high supply voltage further exacerbates this contention. The ultralow voltage split-output, or ULVS, level shifter decouples the CVSL stage from the output driver stage




Fig. 5. ULVS level shifter. (a) Circuit diagram with critical devices circled. (b) Simulated 22-nm V_{\min} benefit of 125-mV node in the presence of random variations (6σ).




Fig. 6. Simulated 22-nm SIMD engine RF and logic optimization benefits across 0 °C–85 °C, 3σ systematic, 6σ random variation.

and interrupts the contention devices, thus improving V_{\min} by 125 mV (Fig. 5). Full interruption of contention devices occurs for voltages $V_{in} \geq V_{out}$, while for voltages $V_{in} < V_{out}$, the contention devices are only partially interrupted but still are beneficial at low voltages. For equal fan-in and fan-out, the ULVS level shifter weakens contention devices, thereby reducing power by 25%–32%.

Fig. 6 summarizes improvements achieved from applying multiple circuit techniques for both the RF and logic circuits across 0 °C–85 °C in the 22-nm SIMD engine. The static RF read circuits and shared P/N DETG write SRAM bit-cells improve overall RF V_{\min} by 250 mV. Shared gates, ULVS level shifters, and vector flip-flops improve overall logic V_{\min} by 150 mV.

E. Architecture Driven NTV Resilient NoC Fabrics

Architectural techniques can help regain some of the performance loss from engaging aggressive VDD reduction. The limits to NTV-based parallelism to reclaim performance have been discussed in [19]. Dynamic adaptation techniques have been shown to monitor the available timing margin and guard bands in the design and dynamically modulate the V/F, thus preventing the occurrence of timing errors [20]. Architecture-assisted resilient techniques, on the other hand, are more aggressive with the V/F push. In this case, the errors are allowed to happen, they are detected and then corrected using appropriate replay mechanisms.

Replica path-based methods such as tunable replica circuits (TRC) have been proposed [21] for error detection in




Fig. 7. NTV-NoC. (a) Two-stage router data path and control logic indicating critical timing path. (b) Pipe stage 2 is enhanced with an EDS circuit to detect failures in critical timing paths down to NTV.

flip-flop-based static CMOS logic blocks. 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. The key requirement is that the TRC must always fail before the critical path fails. The TRC is an area-efficient and nonintrusive technique, but it cannot leverage the probabilities of critical path activation, multiple simultaneous switching at inputs of complex gates, or worst case coupling from adjacent signal lines. An 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. 7(b)] [22]. Errors are corrected by performing a replay operation at higher V_{DD} or lower F . The V_{DD}/F can also be adapted by monitoring the error rate and accounting for error recovery overheads.

NoCs have rapidly become the accepted method for connecting a large number of on-chip components. Packet-switched routers are key building blocks of NoCs [23]. Margins for operating V_{DD}/F are used to guarantee error-free operation limit achievable energy efficiency and performance at the minimum-energy voltage optimum (V_{OPT}). While error-correction codes (ECC) have been previously used to mitigate transient failures in routers [24], the associated performance and energy overheads can be significant for the detection and correction of multibit failures. Timing error detection using EDS has been used for processor pipelines with minimal overhead [22]. An NTV router, designed in a 22-nm node and enhanced with error detection sequentials (EDS) and a flow-control unit (FLIT) replay scheme, provides resilience to multibit timing failures for on-die communication. The goal is to evaluate the performance and energy benefits of the single-error correction double-error detection (SECDED) ECC method over an EDS-based approach, from nominal V_{DD} down to NTV.

Resilient Router Architecture and Design: The six-port packet-switched router in the 2×2 , 2-D mesh NoC fabric communicates with the traffic generator (TG) via two local ports and with neighboring routers using four bidirectional, 36-bit 1.5-mm-long on-die links [Fig. 1(c)]. Inbound router FLITs are buffered in a 16-entry 36-bit wide FIFO [Fig. 7(a)]. The most critical timing path in the router consists of request generation, lane, and port arbitration, FIFO read, followed by a fully nonblocking crossbar (XBAR) traversal. Any failure in this timing path is detected by the EDS circuit [Fig. 7(b)] embedded in the output pipe stage (STG 2). The two-cycle EDS enhanced router can be run in two modes, with and without error detection. The TG contains SECDED logic

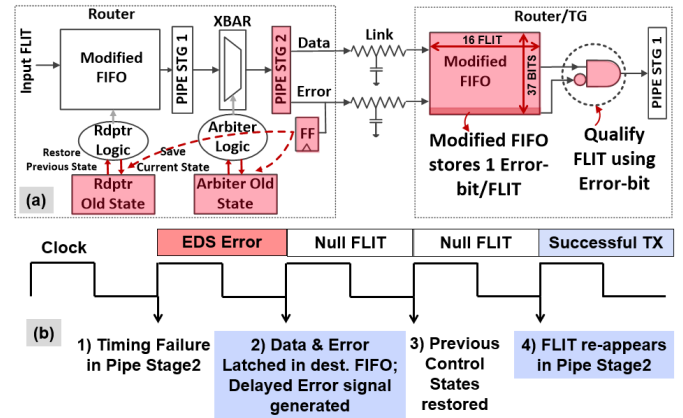


Fig. 8. Internal router architecture. (a) Modifications to enable FLIT replay and forward error correction. (b) Clock cycle diagram showing stages of timing failure detection, replay, and recovery.

which appends or retrieves nine ECC bits from a packet's tail FLIT, thus allowing end-to-end detection and correction of errors in the payload. A programmable noise injector [21] is introduced at each node on the VNoC supply to induce noise events during packet transmission.

The router control logic recovers from timing failures by saving critical states for the last two FLIT transmissions [Fig. 8(a)]. In the event of a timing failure, the Error signal generated by the EDS circuit in STG 2 is captured along with the erroneous FLIT in the recipient's FIFO, modified to accommodate an additional error bit as shown. Forward error correction is achieved by qualifying the FIFO output with the Error flag. In the router with the timing failure, the Error signal is latched to mitigate metastability. This synchronized Error flag is then used to roll-back the arbiters and FIFO read pointers to the previous functionally correct state. The current FLIT is again forwarded as part of the replay. Error synchronization and roll-back incur two clock cycles of delay between an error event and successful recovery [Fig. 8(b)]. To avoid min-delay failures at STG 2, a clock with scan-tunable duty cycle control is implemented for the EDS latches. Additional min-delay buffers are inserted in the XBAR data path for added hold margin at a 2.4% area cost. In addition, the resilient router incurs the following overheads: 1) about 2.5% of router sequentials are converted to EDS; 2) enabling replay causes a 10.5% increase in sequential count with 1.6% area overhead; and 3) the power overhead for the entire router is 8.7% with a 2.8% area cost.

Silicon measurements are performed at 25 °C for a representative NoC traffic pattern with FLIT injection at each router port every clock cycle at 10% data activity. The measured NoC silicon logic analyzer trace (Fig. 9) shows a supply noise-induced timing failure on the control bits of the packet header FLIT, followed by two cycles of bubble (null) FLITs and persistent retransmission (replay) of the FLIT until successful recovery. As shown, timing error synchronization and roll-back incur a two-cycle delay between an error event and successful recovery.

Fig. 10 plots the measured bandwidth (BW) for the resilient router at 400 mV, in the presence of a 10% VNoC droop induced by the on-die noise injectors. The number of




Fig. 9. Silicon logic analyzer trace showing successful recovery of FLITs from timing failures.




Fig. 10. Improvement in resilient router BW at V_{OPT} (400 mV) over a nonresilient version.

erroneous FLIT increases exponentially with F_{CLK} . To account for such droop, a nonresilient router must operate with 28% (700 mV) and 63% (400 mV) F_{CLK} margins, respectively, thus limiting F_{MAX} . The resilient router reclaims these margins and offers near-ideal BW improvement until higher error rates and FLIT replay overheads limit overall BW gains. Past the point-of-first failure (PoFF), both control and data bits are corrupted. While ECC can identify data bit failures, control bit failures can invalidate the entire FLIT, rendering any ECC scheme ineffective. If control paths are designed with enough timing margins such that the control bits do not fail, the F_{CLK} gain from SECEDED ECC is only 7% beyond PoFF, since several data bits fail simultaneously. In contrast, at 400 mV, the EDS scheme provides tolerance to multibit failures over a $9\times$ wider F_{CLK} range, past PoFF. Compared to a conventional router implementation, the resilient router offers 28% higher BW for 5.7% energy overhead at 700 mV and 63% higher BW with 14.6% energy improvement at 400 mV.

F. Designing for Wide-Dynamic Range: Tools, Flows and Methodologies

Device optimizations need to work in concert with automated CAD design flows for optimal results. The 14-nm NTV-WSN design uses high performance (HP), standard performance (SP), Ultra-low performance (ULP) and thick-gate (TG) devices—all four transistor families in a 14-nm second-generation trigate SoC platform technology [11]. To minimize variation-induced skews, the clock distribution is completely designed using HP devices. The lower threshold voltage (V_T) of the HP devices allows improved delay predictability on the clock paths at NTV. SP devices are used for 100% of logic cells to achieve sufficient speeds during the active mode of operation, with memory using ULP transistors for low standby power. The bidirectional CMOS IO circuits are designed using high voltage (1.8 V) TG transistors.




Fig. 11. NTV-CPU. (a) Optimizations for wide range design convergence. (b) Design criteria varies widely at NTV (0.5 V) versus 1.05-V corner.

The optimized cell library for a wide operational range is characterized at 0.5, 0.75, and 1.05 V V_{DD} corners for design synthesis and timing convergence and are optimized for robust and reliable ultralow voltage operation. Statistical static timing analysis (SSTA) is employed—a method that replaces the normal deterministic timing of gates and interconnects with probability distributions and provides a distribution of possible circuit outcomes [25]. As discussed earlier, a variation-aware SSTA study is performed on the standard cell library to eliminate the circuits which exhibit dc failures or extreme delay degradation due to reduced transistor ON/OFF current ratios and increased sensitivity to process variations. As a result, the standard cell library was conservatively constrained for use in the NTV optimized designs.

Achieving the performance targets across the entire voltage range is challenging since critical path characteristics change considerably due to nonlinear scaling of device delay and a disproportionate scaling of device versus interconnect (wire) delay. It is critical to identify an optimal design point such that the targeted power and performance are achieved at a given corner without a significant compromise at the other corner. Synthesis corner evaluations for the NTV-CPU [Fig. 11(a)] suggest that 0.5 V, 80 MHz synthesis achieves the target frequency at both 0.5 (80 MHz) and 1.05 V (650 MHz). In comparison, it is observed that 1.05-V synthesis does not sufficiently size up the device dominated data paths which become critical at lower voltages, resulting in 40% lower performance at 0.5 V. Although 1.05-V synthesis achieves lower leakage and better design area, the 0.5-V corner was selected for final design synthesis of the NTV prototypes, considering its low voltage performance benefits and promise for wide operational range. Performance, area, and power metrics at the two extreme design corners in a 32-nm node are presented in Fig. 11(b). For subsequent NTV prototypes, a multicorner design performance verification (PV) methodology that simultaneously co-optimizes timing slack across all three performance corners was developed. This PV approach ensures that performance targets are met across the wide voltage operational range. The method accounts for nonlinear scaling of device delays in the critical path versus interconnect delay scaling across wide VDD. At low voltages, severe effects of process variations result in path delay uncertainties and may cause setup (max) or hold (min) violations. Setup violations can be corrected by frequency binning. However, hold violations can cause critical functional failures. The design timing

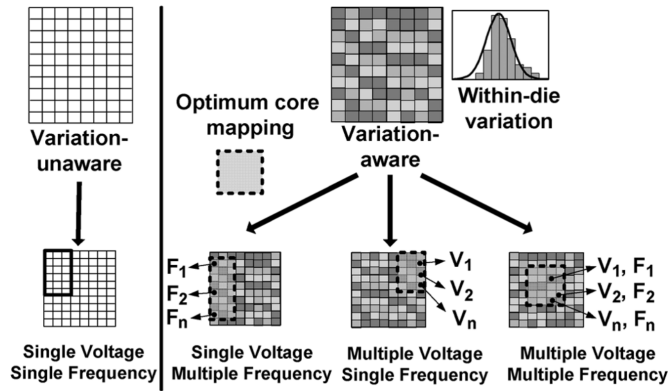


Fig. 12. Conventional DVFS chip versus variation-aware DVFS and optimal core allocation.

convergence methodology is enhanced to consider the effect of random variations and provide enough variation-aware hold margin guard-bands for robust NTV operation.

III. WID VARIATION-AWARE DVFS METHODS

Many-core processors with on-die NoC interconnects are emerging as viable architectures for energy-efficient high-performance computing (HPC) [26], [27]. Aggressive supply voltage scaling results in a dramatic improvement in energy efficiency at the expense of performance [3], [28]. Many-core processors can regain this performance by parallelizing workloads and employing more cores. Future trends for energy-efficient architectures: 1) more small cores integrated on a single die; 2) larger die sizes for increased parallel performance; and 3) lower operating voltages for increased energy efficiency; when coupled with worsening WID variations due to device scaling will result in increased vulnerability to core-to-core maximum clock frequency (F_{max}) and leakage variations [29]. However, these effects can be mitigated by a combination of variation-aware software, architecture, and design approaches. Many-core processors with DVFS can optimize power performance of parallel workloads by varying V/F levels of independent voltage frequency islands (VFI). A combination of DVFS with dynamic core-count scaling (or DVFCs) has been proposed to further improve performance and energy efficiency across varying workloads [30]. Chips with single-voltage/single-frequency (SVSF) for all cores running homogeneous threads [4] as well as multiple-voltage/multiple-frequency (MVMF), running heterogeneous applications and using independent V/F control for each core [26], have been reported. In a conventional DVFS design all cores are treated identical and the DVFS operating points are determined by the slowest core on the chip (Fig. 12 variation-unaware). With technology scaling, however, core-to-core variations in F_{max} and leakage due to device parameter variations have become significant, thus motivating the need for variation-aware DVFCs. Fig. 12 also illustrates a variation-aware DVFCs approach with per-core F_{max} and leakage profiles and mapping of an application to an optimal set of cores [31].

Silicon measurements are performed on the 80-core TeraFLOPS processor [4], an NoC architecture that contains

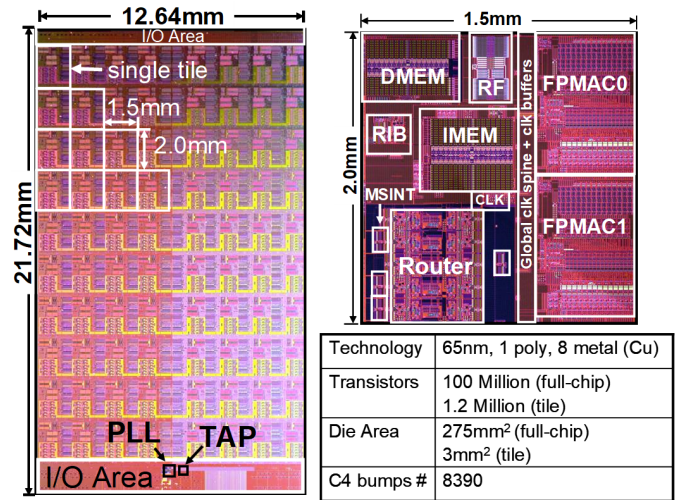


Fig. 13. 80-core processor tile micrograph and characteristics.

80 tiles arranged as a 10×8 2-D mesh network and implemented in a 65-nm [5] process technology (Fig. 13). The full chip has 100 M transistors and 1.2 M in each tile. The full-chip die area is 275 mm² with an individual tile area of 3 mm². Each tile consists of a processing engine connected to a five-port router for intertile communication. The processing engine contains two independent single-precision floating point multiply-accumulator (FPMAC) units, 3-kB single-cycle instruction memory (IMEM), and 2-kB data memory (DMEM). A five-port two-lane pipelined packet-switched router core with phase-tolerant mesochronous links forms the key communication fabric [23]. The 15 fan-out-of-4 (FO4) design uses a balanced core and router pipeline with critical stages employing performance-setting semidynamic flip-flops. Several application kernels have been mapped to the design to study performance scalability and the percentage of peak performance achievable [32].

The optimum operating point for a workload is defined by the number of active cores, their physical locations on the die, and individual V/F values, which results in minimum energy per operation, while still meeting the execution time target.

Clock and leakage power management features implemented on the chip, combined with WID core-to-core variation profiles and workload characteristics influence the optimum settings and are governed by balances among clock and data switching energy, intercore communication energy, and leakage energy of active and idle cores. The maximum clock frequency (F_{max}) and leakage (active standby) power were measured as a function of V_{DD} for each of the 80 cores on several dies, along with switched capacitances for clock, data activity, and intercore communications through the embedded per-core routers and the on-die NoC interconnect. Parameterized energy and performance models are populated by these silicon measurements and assume each tile to be an independent VFI and the entire 2-D-mesh as one single VFI. Application-specific attributes, the number of floating point operations (FLOP), switching activity, intertile communication activity (number of flit transfers), and execution cycle penalty resulting from communication cycles that cannot be

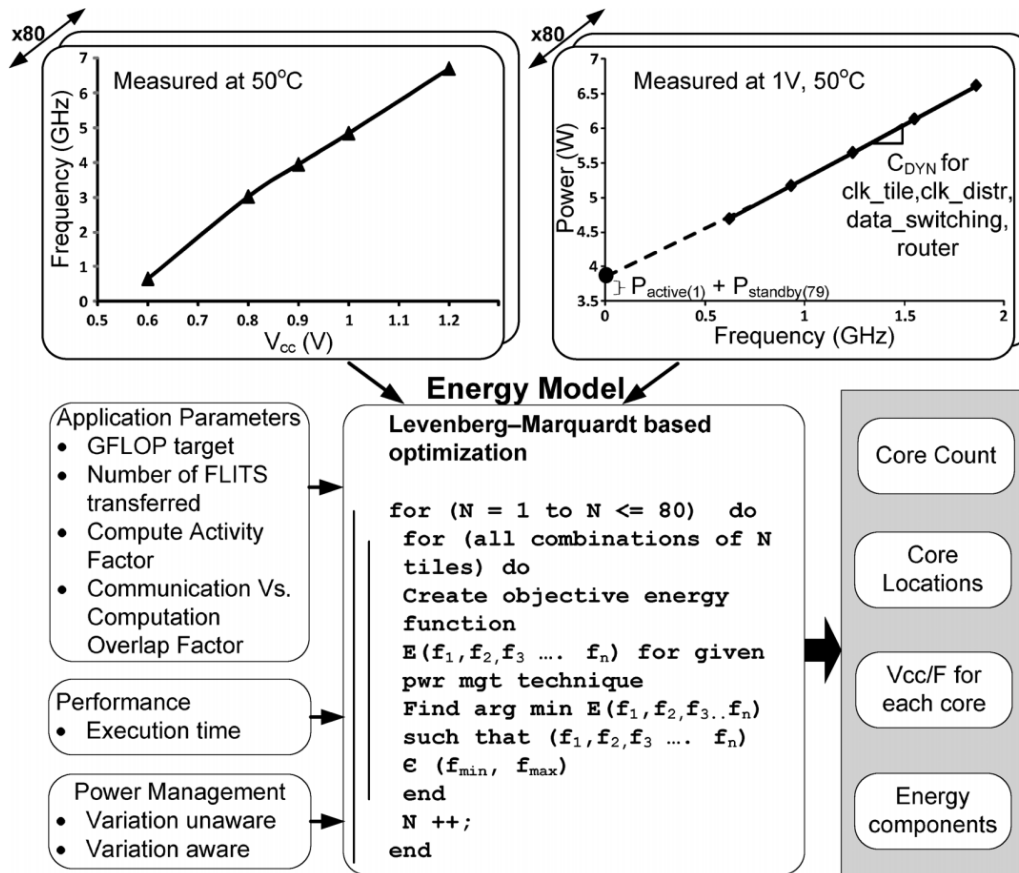


Fig. 14. Illustration of the procedure to obtain minimum energy point for an application.

overlapped with compute cycles, are calibrated to real applications like communication-intensive 2-D-fast Fourier transform (FFT) and compute-intensive partial differential equation (PDE) solvers. They are used by a multidimensional global optimizer (Fig. 14) to determine the optimal V/F values and core allocation that minimizes energy/FLOP while meeting a target execution time, for various scenarios. The input to the optimizer is the energy model of the chip determined by the voltage and frequency of each core, number of flits transferred between cores, average hop distance for each flit, execution time to complete the application, of computing one FLOP and of transferring one flit over one hop of the mesh. For example, to accurately model the communication energy for the PDE solver workload, the energy model was populated with the number of flits transferred between cores and the average hop distance that each flit traverses for this workload. The compute energy is calculated from the number of FLOP's in the workload. The optimizer aims to find the minima of the objective energy function by using the Levenberg Marquardt algorithm. The optimization variables are the frequency of each core and are bound by the minimum performance requirement of the workload (F_{min}) and F_{max} for that core. The initial start value is the midfrequency between F_{min} and F_{max} . The model outputs the core count, location of the cores, V/F for each of 80 cores and individual energy components at the optimal operating point.

WID core-to-core F_{max} measurements for each of the 80 cores measured individually are shown in Fig. 16. Targeted tests were written to capture the behavior of critical cycle and phase paths in the tile. Presilicon timing analysis and silicon test measurements identified the single-cycle accumulate loop in the FPMAC to be the critical path and the maximum frequency limiter on the chip and the IMEM to multiport RF issue loop to be the worst phase path on the chip. Fewer logic gates between flops and deeper pipelining aggravate the impact of random WID variation. Measurements at 1.2 V, 50 °C demonstrate F_{max} of cores between 5.7 and 7.3 GHz. Both cycle and phase paths have similar F_{max} distributions (Fig. 9). F_{max} variation for a core is determined by (threshold voltage) and (effective channel length) fluctuations for devices in the critical path as opposed to leakage power, which is an aggregation for all devices in the core. Fig. 16 shows core leakage measurements correlate well with F_{max} data.

Many-core processors operating at lower voltages for increased energy efficiency are also exposed to greater variability due to increased sensitivity of (ON current) to fluctuations, as demonstrated by a frequency spread of 28% at 1.2 V versus 62% at 0.8 V between the fastest and slowest cores on the die. The delay distribution plot (Fig. 16) for 80 cores on a single die shows an of 5.93% at 1.2 V and increases by 45%–8.64% at 0.8 V. The data makes a strong case for variation-aware DVFS at lower operating voltages.

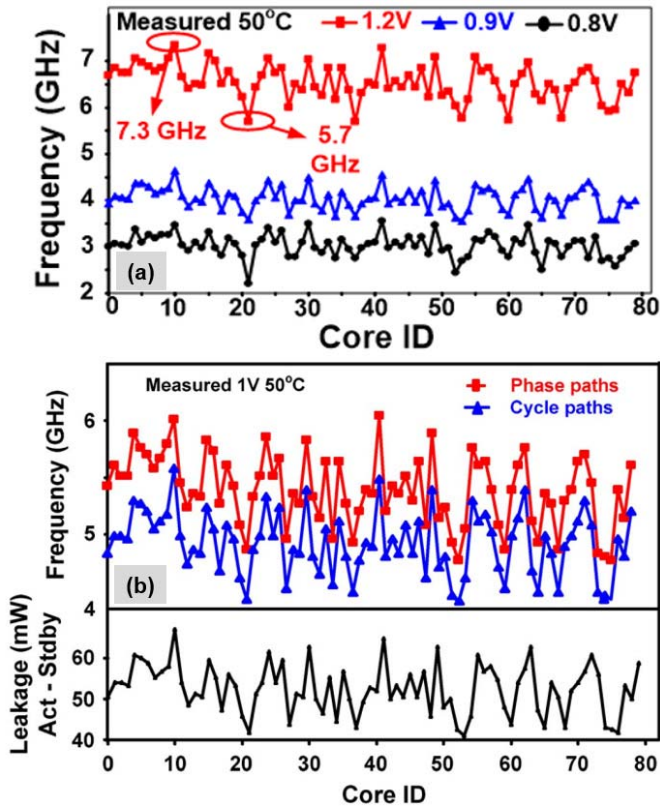


Fig. 15. Measured (a) core-to-core F_{max} variation for 80 cores on a single die and (b) core-to-core cycle and phase path frequency and leakage power correlation [31].

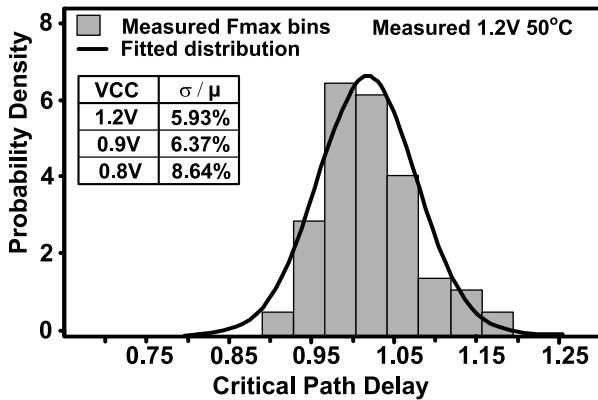


Fig. 16. Critical path delay distribution and σ/μ at different voltages (65 nm).

Fig. 17 shows measured F_{max} profiles for three dies. Since each die contains 80 separate cores, this data provides a statistical significance for measuring the impact of WID variations on core-to-core F_{max} and leakage variations. Lack of correlations among F_{max} profiles from the three different dies confirm the absence of any deterministic design issues behind the F_{max} variability and point to the need for per-die postsilicon F_{max} profiling for implementation of variation-aware schemes.

Fig. 18 compares variation-aware DVFCS for four applications across two performance targets and parallel workload characteristics. Energy gains, number of active cores, and number of independent voltage islands (VI) or frequency islands (FI) are shown for eight optimal operating points (four




Fig. 17. Measured F_{max} profiles for three different dies.

SCHEME	INDEX	2% of Peak			50% of Peak			2% of Peak			50% of Peak		
		Gain	pJ/FLOP	Core #	Gain	pJ/FLOP	Core #	Gain	pJ/FLOP	Core #	Gain	pJ/FLOP	Core #
LOW	SINGLE VCC MULTIPLE FREQ	9%	305	28	0%	633	56	7%	42	32	2%	85	68
	MULTIPLE VCC SINGLE FREQ	23.5%	256	16	6.32%	593	56	19%	36	24	7.5%	80	60
	MULTIPLE VCC AND FREQ	SAME AS MULTIPLE VCC SINGLE FREQ						SAME AS MULTIPLE VCC SINGLE FREQ					
HIGH	SINGLE VCC MULTIPLE FREQ	15%	364	16	0%	1345	48	13%	49	16	0.5%	169	52
	MULTIPLE VCC SINGLE FREQ	35%	278	12	7.16%	1249	48	30%	39	12	7%	158	48
	MULTIPLE VCC AND FREQ	SAME AS MULTIPLE VCC SINGLE FREQ						SAME AS MULTIPLE VCC SINGLE FREQ					

= Number of unique VCC or Freq domains

Fig. 18. Comparison of variation-aware DVFCS (SVMF, MVSF, and MVMF) across four workloads of different compute/communication activities and performance targets.

applications at two performance targets each). An application with low compute activity, high communication activity, and 2% peak performance target gains 15% on a single voltage multiple frequency (SVMF) and 35% on a multiple voltage single frequency (MVSF) design as compared to SVSF. However, results for 50% performance targets show no gains for SVMF and 7% for MVSF. SVMF offers improvement only when the leakage energy savings from idling active cores earlier exceeds the extra clock energy needed for higher frequency operation. Thus, minimal improvements are visible at high-performance targets where leakage power is not significant. MVSF offers the highest gains in energy/FLOP: 19%–35% at low performance targets and 6%–7% at higher performance levels, and also incurs minimal implementation overheads since no frequency domain crossing circuits are needed for homogeneous workloads.

Fine-grained voltage resolution control per VI is essential for closely tracking WID variations across the die in an MVSF design. Voltage resolution of 6–12 mV, available in typical VRs, preserves most of the energy/FLOP gains of MVSF. A 6.25-mV resolution has an energy penalty of 1% when compared to infinite voltage resolution. Fig. 19 illustrates increased sensitivity to voltage resolution at lower V_{DD} due to worsening WID variations and shows an 18.75-mV resolution that degrades efficiency by 5% at 0.64-V supply.

IV. DIGITALLY CONTROLLED VRs FOR WIDE-RANGE DVFS

Per-core DVFS can significantly improve the overall energy efficiency of multicore processors by enabling each core to

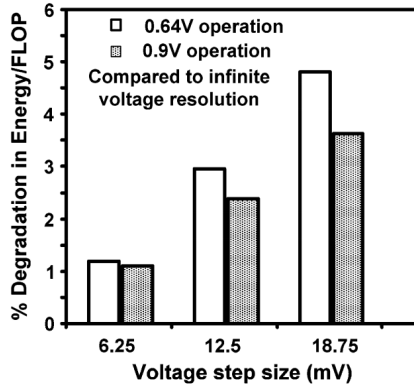


Fig. 19. Degradation in energy/FLOP with increasing voltage step-size.

operate independently across a wide voltage–frequency range. The local voltage domain of each core can be created and controlled using integrated digital LDO linear regulators. While using integrated inductor or capacitor-based switching regulators can enable higher power efficiency over a wide voltage range, they incur significant process complexity overheads for on-die or in-package integration of high power density passives. Digital LDOs implemented using embedded power gates (PGs) can avoid the minimum dropouts required by their analog counterparts and allow easier local integration on die. In addition, they enable better scalability and portability across scaled process nodes. Furthermore, many of the digital LDO building blocks can often be synthesized, thus improving design productivity.

Per-core DVFS using distributed digital linear microregulators has been previously demonstrated to improve overall energy efficiency for multicore processors [33], [34]. A hybrid combination of a digital LDO at lower dropouts and an integrated switched capacitor VR at larger dropouts has been used for DVFS of a graphics core to further improve power efficiency over a wide voltage range [35]. The digitally controlled LDO [14] in 14-nm trigate CMOS powers an Atom core with embedded PGs to enable per-core DVFS over a wide voltage–frequency range. The multimode digital controller features a classic linear Type-II control along with nonlinear mode and adaptive gain for robust and fast transient response over wide voltage and current ranges. A code roaming algorithm is also implemented to complement the main controller for improved reliability by eliminating issues due to self-heating and electromigration across a wide range of voltages and load currents.

A processor with four Atom cores is implemented as 2-core modules, each with a dedicated bus interface unit (BIU) (Fig. 20). Each LDO powers its respective core and BIU and receives a programmable external supply voltage whose levels are varied based on the workload. When the cores are at iso-performance levels, the LDO controller is powered OFF, the PG is bypassed, and the external variable rail drives the same optimal voltage to all the cores for maximum energy efficiency (Fig. 21).

When the cores are not in the iso-performance mode, the core running at a higher frequency uses its LDO in bypass mode, thus maximizing its own energy efficiency,

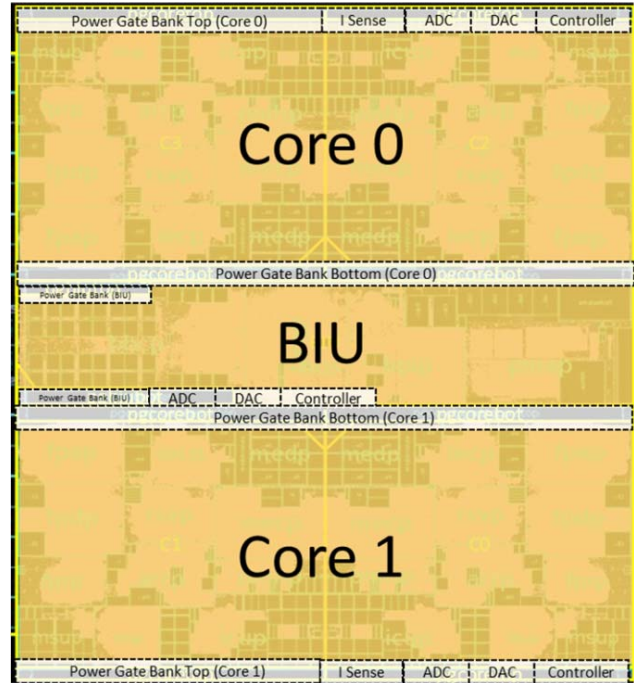


Fig. 20. Quad-core Atom processor layout floorplan with digital LDO using embedded PGs for per-core DVFS.




Fig. 21. Digital LDO supports both iso and noniso performance modes of the processor.

while the core running at a lower frequency engages the LDO in regulation mode to drop the input voltage to its required optimal value. This scheme enables high density and cost-effective power delivery while improving overall energy efficiency. The Atom processor presents a wide dynamic load with a wide range of output voltages. The load current range of 0.2–2.5 A across an output voltage range from 0.55 to 1.15 V is powered with the digital LDO from an input voltage range of 0.565–1.165 using embedded PGs as shown in the layout floor plan (Fig. 20) of the core module with integrated LDO controller and PGs. The active area of silicon to deliver this power is 0.115 mm², which represents a total power density of 26.1 W/mm².

A step load transition of 1.2–2.4 A in <1 ns at an output voltage of 1.05 V is demonstrated (Fig. 23). The nonlinear




Fig. 22. Measured step load transitions waveform showing the effect of the controller modes (linear, nonlinear, and adaptive gain) [14].




Fig. 23. Combinational versus sequential SER illustration.

controller with the adaptive gain improves the settling time by $10\times$ to <20 ns as compared to a conventional linear controller. Measurements also demonstrate a 50% improvement in droop to <100 mV for the multimode controller with a combination of linear, nonlinear, and adaptive gain modes working in tandem. Seamless transition between modes is also achieved by the use of digital control. The peak current efficiency of the LDO is 99.6% at the maximum load of 2.5 A.

As summarized in the comparison Table II, the digital LDO achieves high power density with high current efficiency and the best minimum dropout voltage with the lowest figure of merit (FOM) as defined in [37].

V. RADIATION-INDUCED SER IMPACTS AT NTV

Soft errors remain a critical reliability concern in scaled technologies due to the increasing transistor count per chip and the reduced operating voltage [38], [39] down to NTV. To tackle radiation effects in integrated circuits, designers need to accurately analyze circuit dependencies that impact SER. Fig. 23 illustrates a combinational SER event occurring in a logic gate chain and a sequential SER event occurring in a latch storage node, respectively. As the radiation particle penetrates the silicon, an ion track is generated along its strike path. This results in charge generation which is collected by nearby diffusion junctions. The collected charge, if large

TABLE II
COMPARISON BETWEEN VRs

	Unit	[33]	[34]	[35]	[36]	Work in [14]
Technology	-	45-nm	22-nm SOI	22-nm	14-nm	14-nm
Input V_{DD}	V	1.18-1.62	1.1	0.65-1.05	1-1.15	0.565-1.165
Output V_{DD}	V	0.9-1.1	0.61-1.05	0.38-0.92	0.5-1.12	0.55-1.15
Min. Dropout	mV	85	50	130	30	<15
Max. Load Current	A	0.04	11.9	NA	1.5	2.5
Peak current efficiency	%	77.5	96.9	NA	99.3	99.6
Power Density	W/mm ²	0.616	34.5	NA	16.2	26.1
FOM	ps	86.4	NA	NA	18.4	14.8

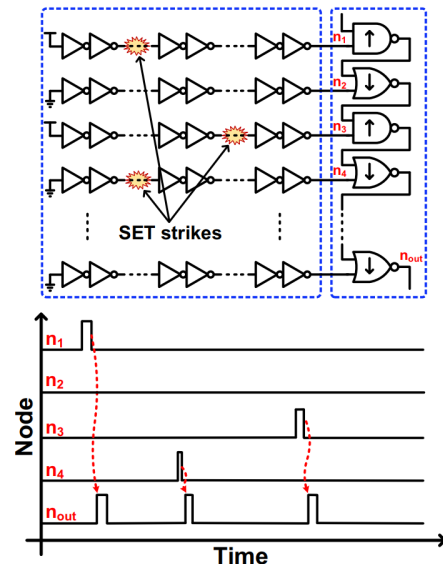


Fig. 24. SET pulse collection using skewed NAND-NOR readout chain.

enough, induces a transient pulse in case of a combination logic path or a bit flip in case of a memory storage node. The former is referred to as single event transient (SET) while the latter is called single event upset (SEU).

Process parameters such as fin geometry, doping profile, doping density, and junction geometry affect the charge collection efficiency. In addition, circuit parameters such as transistor width, threshold voltage (V_T), fan-out, chain length, supply voltage, and P/N ratio have a strong impact on SER. Understanding the contribution of various circuit parameters on SER is a critical aspect in developing a radiation hardening strategy.

A. SER in Combinational Circuits in 14-nm Trigate CMOS

High-density SER characterization circuits using skewed NAND-NOR standard logic gates have been proposed [16]. Fig. 24 illustrates the operation of the NAND-NOR readout chain. All detection chains are connected in parallel to the readout chain. This enables funneling and propagation of

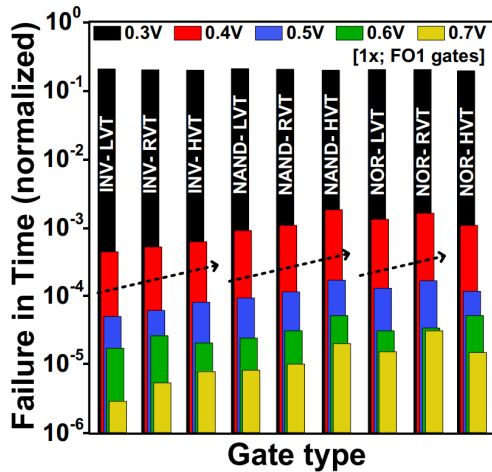


Fig. 25. Measured FIT rates for various gates with different V_T and V_{DD} .

all SET pulses in the readout chain while expanding the pulsewidth to ensure they reach a final 6-bit triple modular redundant (TMR) counter. The skewed NAND/NOR gates allow for pulse expansion once the SET pulses enter the readout chain. This ensures that all pulses injected by the detection chains reach the final output node without disappearing halfway. The dedicated SER test structures, designed in 14-nm trigate technology [11] were irradiated under a neutron beam at the Los Alamos National Laboratory, Washington, DC, USA, collecting large amount of statistical data for studying SER dependence on various circuit parameters.

Measured radiation data (Fig. 25) shows normalized failure in time (FIT) rates for gates with different V_T across 0.3–0.7-V supply voltages. FIT is defined as the number of soft errors per gate in a billion hours of circuit operation. Low V_T (LVT) gates exhibit lower SER over regular V_T (RVT) and high V_T (HVT) gates. Supply voltage scaling also has a strong impact on SER. For a lower V_{DD} , SER increases exponentially. There is about a $10^5 \times$ increase in FIT rate when V_{DD} is scaled from a nominal value of 0.7 V to an NTV value of 0.3 V. Critical charge (Q_{crit}) is typically defined as the minimum collected charge at a circuit node that can cause a SET or SEU event. At NTV, the Q_{crit} becomes extremely low, hence almost all particles that induce charge collection, result in a soft error. This is reflected in the black bars in Fig. 7. That is, V_T does not have a dominating impact on the FIT rate at 0.3 V resulting in constant FIT rates irrespective of the V_T type. The data also shows that FIT rates for NAND and NOR gates were higher compared to inverters due to two sensitive internal nodes susceptible to SER.

As the device size (W) increases, $I_{restore}$ gets stronger, junction area becomes larger, and gate capacitances get bigger. These cumulative effects result in much higher Q_{crit} , thereby, leading to a much lower SER across all V_{DD} points (Fig. 26). Unlike the case of varying V_T , here noticeable variation can be seen in FIT even at 0.3 V. Although the contribution from $I_{restore}$ is small at lower V_{DD} , variation in node capacitance can still dominate the final SER values even at lower voltages. This confirms that increasing the capacitance of a sensitive circuit




Fig. 26. Measured FIT rates for various gates with device size (W) versus V_{DD} .

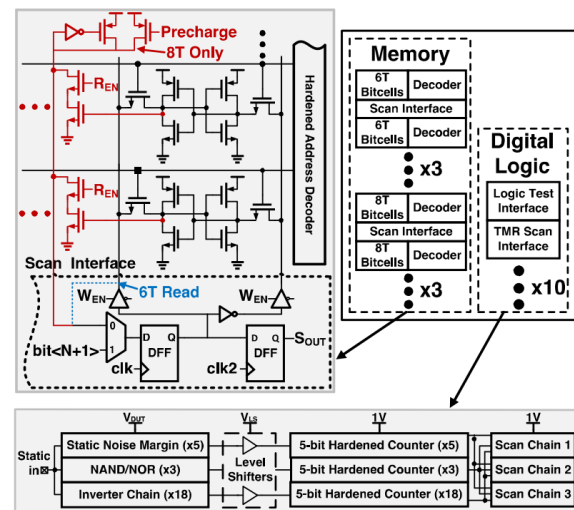


Fig. 27. Characterization setup for radiation-induced SRAM errors [15].

node is a viable option for reducing the SER across different supply voltages.

B. Multibit Upsets in SRAM Circuits in 65-nm CMOS

Custom test circuits have been developed and tested under neutron and alpha radiation, with memory measurement results that show SER dependence on V_{DD} [15]. Neutron-induced SER in SRAM increases by 18% in 90-nm SRAM SER for every 10% decrease in V_{DD} down to 0.7 V [40]. A $7.8 \times$ increase in SER is reported for a 10-T subthreshold SRAM in 65-nm when V_{DD} is reduced from 1.0 to 0.3 V [41]. A variety of SRAM, RF, and digital logic test structures (Fig. 27) are included in a 65-nm test chip to provide a comprehensive assessment of circuit sensitivities to radiation at low V_{DD} [15].

Fig. 28 shows measured SER versus V_{DD} data for 6-T and 8-T memory arrays. When V_{DD} reduces from 1 V to 330 mV, SER increases on average by $6.45 \times$ and $2.5 \times$ for accelerated neutron and alpha radiation, respectively. A sharp increase in SER is observed once the threshold voltage is crossed ($V_T \sim 0.35$ V), as the transistors enter the weak-inversion region, where I_{ON} has an exponential response to any change in gate voltage. The neutron-induced multibit upset (MBU)

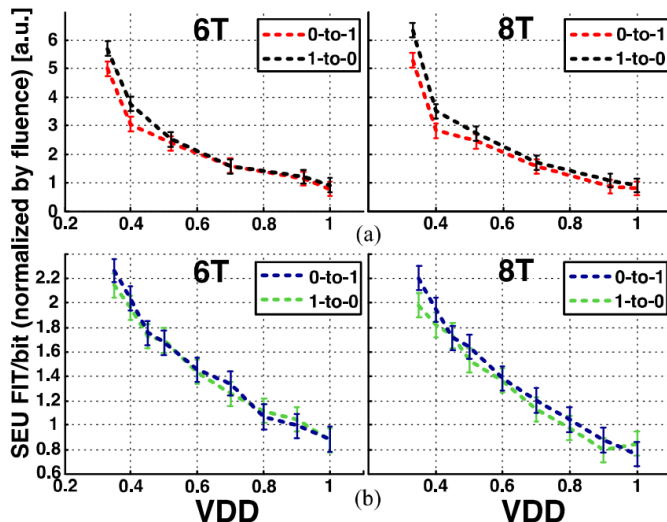


Fig. 28. SRAM/RF SER versus V_{DD} for (a) neutron and (b) alpha radiation.

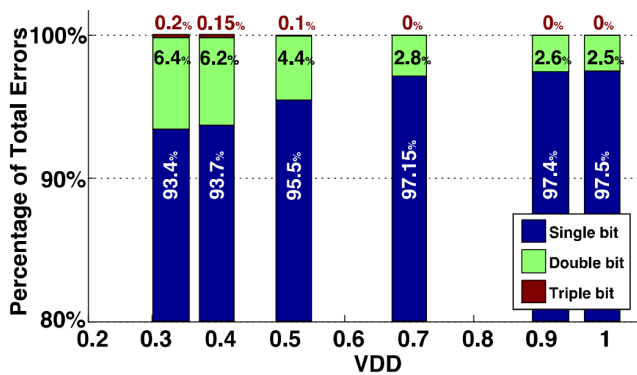


Fig. 29. Measured SRAM MBU rate versus V_{DD} under neutron irradiation.

rate versus V_{DD} (Fig. 29) shows a $2.6\times$ increase in errors, when V_{DD} scales from 1.0 to 0.33 V, with triple-bit upsets occurring below 0.5 V.

NTV circuits show strong sensitivity to soft errors, dictating aggressive SER protection over a design operating at nominal V_{DD} . Mitigating soft errors at NTV requires judicious tradeoffs across area, speed, and power metrics. Silicon measurements confirm that using larger devices and increasing the capacitance of sensitive circuit nodes can be effective in reducing SER at NTV for logic, while memory systems can be protected using information redundancy techniques, such as ECCs.

VI. CONCLUSION

The SoC designs for future Internet of Things (IoT) systems need to deliver uncompromising and scalable performance with extreme energy efficiency. Many-core designs with resilient and efficient NoC for intercore communications, with each core operating across a wide DVFS range, including NTV, are essential. WID variation-aware DVFS with fine-grain per-core voltage control using digital LDOs is needed for mitigating variation impacts and enabling fine-grain power management. Effective mitigation techniques of radiation-

induced soft errors in logic and memory circuits are needed for robust operation at NTV.

ACKNOWLEDGMENT

The authors thank S. Jain, S. Khare, V. Honkote, M. Abbott, M. Anders, T. Majumder, C. Tokunaga, M. Cho, P. Aseron, T. Nguyen, M. Khellah, and K. Ravichandran at Intel Labs, R. Pawlowski, R. Muthukaruppan, D. Mallik, and V. Grossnickle at Intel, Hillsboro, OR, USA, for technical contributions and encouragement, the support from Intel's ATTD Team with package assembly, and the faculty at Electrical Engineering Departments of the University of Minnesota, Minneapolis, MN, USA, and Oregon State University, Corvallis, OR, USA.

REFERENCES

- [1] R. G. Dreslinski, M. Wiecekowsk, D. Blaauw, D. Sylvester, and T. Mudge, "Near-threshold computing: Reclaiming Moore's law through energy efficient integrated circuits," *Proc. IEEE*, vol. 98, no. 2, pp. 253–266, Feb. 2010, doi: [10.1109/JPROC.2009.2034764](https://doi.org/10.1109/JPROC.2009.2034764).
- [2] V. De, S. Vangal, and R. Krishnamurthy, "Near threshold voltage (NTV) computing: Computing in the dark silicon era," *IEEE Des. Test.*, vol. 34, no. 2, pp. 24–30, Apr. 2017, doi: [10.1109/MDAT.2016.2573593](https://doi.org/10.1109/MDAT.2016.2573593).
- [3] S. Hanson *et al.*, "Ultralow-voltage, minimum-energy CMOS," *IBM J. Res. Develop.*, vol. 50, nos. 4–5, pp. 469–490, Jul. 2006.
- [4] S. R. Vangal *et al.*, "An 80-tile sub-100-W TeraFLOPS processor in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 29–41, Jan. 2008, doi: [10.1109/JSSC.2007.910957](https://doi.org/10.1109/JSSC.2007.910957).
- [5] P. Bai *et al.*, "A 65 nm logic technology featuring 35 nm gate lengths, enhanced channel strain, 8 Cu interconnect layers, low-k ILD and 0.57 μm^2 SRAM cell," in *IEDM Tech. Dig.*, San Francisco, CA, USA, Dec. 2004, pp. 657–660, doi: [10.1109/IEDM.2004.1419253](https://doi.org/10.1109/IEDM.2004.1419253).
- [6] S. Jain *et al.*, "A 280 mV-to-1.2 V wide-operating-range IA-32 processor in 32 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2012, pp. 66–68, doi: [10.1109/ISSCC.2012.6176932](https://doi.org/10.1109/ISSCC.2012.6176932).
- [7] C.-H. Jan *et al.*, "A 32 nm SoC platform technology with 2nd generation high-k/metal gate transistors optimized for ultra low power, high performance, and high density product applications," in *IEDM Tech. Dig.*, Baltimore, MD, USA, Dec. 2009, pp. 1–4, doi: [10.1109/IEDM.2009.5424258](https://doi.org/10.1109/IEDM.2009.5424258).
- [8] 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 *IEDM Tech. Dig.*, San Francisco, CA, USA, Dec. 2012, pp. 3.1.1–3.1.4, doi: [10.1109/IEDM.2012.6478969](https://doi.org/10.1109/IEDM.2012.6478969).
- [9] S. Hsu *et al.*, "A 280 mV-to-1.1 V 256 b reconfigurable SIMD vector permutation engine with 2-dimensional shuffle in 22 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2012, pp. 178–180, doi: [10.1109/ISSCC.2012.6176966](https://doi.org/10.1109/ISSCC.2012.6176966).
- [10] 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 *Proc. Symp. VLSI Circuits*, Kyoto, Japan, Jun. 2013, pp. C30–C31.
- [11] C.-H. Jan *et al.*, "A 14 nm SoC platform technology featuring 2nd generation tri-gate transistors, 70 nm gate pitch, 52 nm metal pitch, and 0.0499 μm^2 SRAM cells, optimized for low power, high performance and high density SoC products," in *Proc. Symp. VLSI Technol.*, Kyoto, Japan, Jun. 2015, pp. T12–T13, doi: [10.1109/VLSIC.2015.7231380](https://doi.org/10.1109/VLSIC.2015.7231380).
- [12] Intel Corporation. *Intel Quark Processors*. Accessed: Sep. 30, 2020. [Online]. Available: <http://www.intel.com/content/www/us/en/embedded/products/quark/overview.html>
- [13] S. Paul *et al.*, "A sub- cm^3 energy-harvesting stacked wireless sensor node featuring a near-threshold voltage IA-32 microcontroller in 14-nm tri-gate CMOS for always-ON always-sensing applications," *IEEE J. Solid-State Circuits*, vol. 52, no. 4, pp. 961–971, Apr. 2017, doi: [10.1109/JSSC.2016.2638465](https://doi.org/10.1109/JSSC.2016.2638465).
- [14] R. Muthukaruppan *et al.*, "A digitally controlled linear regulator for per-core wide-range DVFS of atom cores in 14 nm tri-gate CMOS featuring non-linear control, adaptive gain and code roaming," in *Proc. 43rd IEEE Eur. Solid State Circuits Conf. (ESSCIRC)*, Leuven, Belgium, Sep. 2017, pp. 275–278, doi: [10.1109/ESSCIRC.2017.8094579](https://doi.org/10.1109/ESSCIRC.2017.8094579).

- [15] R. Pawlowski *et al.*, "Characterization of radiation-induced SRAM and logic soft errors from 0.33 V to 1.0 V in 65 nm CMOS," in *Proc. IEEE Custom Integr. Circuits Conf.*, San Jose, CA, USA, Sep. 2014, pp. 1–4, doi: [10.1109/CICC.2014.6946138](https://doi.org/10.1109/CICC.2014.6946138).
- [16] S. Kumar *et al.*, "An ultra-dense irradiation test structure with a NAND/NOR readout chain for characterizing soft error rates of 14 nm combinational logic circuits," in *IEDM Tech. Dig.*, San Francisco, CA, USA, Dec. 2017, pp. 39.3.1–39.3.4, doi: [10.1109/IEDM.2017.8268521](https://doi.org/10.1109/IEDM.2017.8268521).
- [17] A. Raychowdhury *et al.*, "PVT-and-aging adaptive wordline boosting for 8T SRAM power reduction," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2010, pp. 352–353, doi: [10.1109/ISSCC.2010.5433815](https://doi.org/10.1109/ISSCC.2010.5433815).
- [18] J. Kulkarni, B. Geuskens, T. Karnik, M. Khellah, J. Tschanz, and V. De, "Capacitive-coupling wordline boosting with self-induced V_{CC} collapse for write V_{MIN} reduction in 22-nm 8T SRAM," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2012, pp. 234–236, doi: [10.1109/ISSCC.2012.6176990](https://doi.org/10.1109/ISSCC.2012.6176990).
- [19] N. Pinckney *et al.*, "Assessing the performance limits of parallelized near-threshold computing," in *Proc. 49th Annu. Design Autom. Conf. (DAC)*, San Francisco, CA, USA, 2012, pp. 1143–1148, doi: [10.1145/2228360.2228571](https://doi.org/10.1145/2228360.2228571).
- [20] J. Tschanz *et al.*, "Adaptive frequency and biasing techniques for tolerance to dynamic temperature-voltage variations and aging," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2007, pp. 292–604, doi: [10.1109/ISSCC.2007.373409](https://doi.org/10.1109/ISSCC.2007.373409).
- [21] 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 *Proc. Symp. VLSI Circuits*, Kyoto, Japan, Jun. 2009, pp. 112–113.
- [22] K. A. Bowman *et al.*, "Energy-efficient and metastability-immune resilient circuits for dynamic variation tolerance," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 49–63, Jan. 2009, doi: [10.1109/JSSC.2008.2007148](https://doi.org/10.1109/JSSC.2008.2007148).
- [23] S. Vangal, A. Singh, J. Howard, S. Dighe, N. Borkar, and A. Alvandpour, "A 5.1 GHz 0.34 mm² router for network-on-chip applications," in *Proc. IEEE Symp. VLSI Circuits*, Kyoto, Japan, Jun. 2007, pp. 42–43, doi: [10.1109/VLSIC.2007.4342758](https://doi.org/10.1109/VLSIC.2007.4342758).
- [24] D. Rossi, A. K. Nieuwland, A. Katoch, and C. Metra, "New ECC for crosstalk impact minimization," *IEEE Design Test Comput.*, vol. 22, no. 4, pp. 340–348, Apr. 2005, doi: [10.1109/MDT.2005.91](https://doi.org/10.1109/MDT.2005.91).
- [25] C. S. Amin *et al.*, "Statistical static timing analysis: How simple can we get?" in *Proc. 42nd Annu. Conf. Design Autom. (DAC)*, Anaheim, CA, USA, 2005, pp. 652–657, doi: [10.1145/1065579.1065751](https://doi.org/10.1145/1065579.1065751).
- [26] D. N. Truong *et al.*, "A 167-processor computational platform in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1130–1144, Apr. 2009, doi: [10.1109/JSSC.2009.2013772](https://doi.org/10.1109/JSSC.2009.2013772).
- [27] L. Benini and G. De Micheli, "Networks on chips: A new SoC paradigm," *Computer*, vol. 35, no. 1, pp. 70–78, 2002, doi: [10.1109/2.976921](https://doi.org/10.1109/2.976921).
- [28] B. H. Calhoun and A. Chandrakasan, "Characterizing and modeling minimum energy operation for subthreshold circuits," in *Proc. Int. Symp. Low Power Electron. Design*, Newport Beach, CA, USA, Aug. 2004, pp. 90–95, doi: [10.1109/LPE.2004.240808](https://doi.org/10.1109/LPE.2004.240808).
- [29] K. A. Bowman, S. G. Duvall, and J. D. Meindl, "Impact of die-to-die and within-die parameter fluctuations on the maximum clock frequency distribution for gigascale integration," *IEEE J. Solid-State Circuits*, vol. 37, no. 2, pp. 183–190, Feb. 2002, doi: [10.1109/4.982424](https://doi.org/10.1109/4.982424).
- [30] J. Li and J. F. Martinez, "Dynamic power-performance adaptation of parallel computation on chip multiprocessors," in *Proc. 12th Int. Symp. High-Perform. Comput. Archit.*, Feb. 2006, Austin, TX, USA, Feb. 2006, pp. 77–87, doi: [10.1109/HPCA.2006.1598114](https://doi.org/10.1109/HPCA.2006.1598114).
- [31] S. Dighe *et al.*, "Within-die variation-aware dynamic-voltage-frequency-scaling with optimal core allocation and thread hopping for the 80-core TeraFLOPS processor," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 184–193, Jan. 2011, doi: [10.1109/JSSC.2010.2080550](https://doi.org/10.1109/JSSC.2010.2080550).
- [32] T. G. Mattson, R. Van der Wijngaart, and M. Frumkin, "Programming the intel 80-core network-on-a-chip terascale processor," in *Proc. ACM/IEEE Conf. Supercomput. (SC)*, Austin, TX, USA, Nov. 2008, pp. 1–11, doi: [10.1109/SC.2008.5213921](https://doi.org/10.1109/SC.2008.5213921).
- [33] J. F. Bulzaccelli *et al.*, "Dual-loop system of distributed microregulators with high DC accuracy, load response time below 500 ps, and 85-mV dropout voltage," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 863–874, Apr. 2012, doi: [10.1109/JSSC.2012.2185354](https://doi.org/10.1109/JSSC.2012.2185354).
- [34] Z. Toprak-Deniz *et al.*, "Distributed system of digitally controlled microregulators enabling per-core DVFS for the POWER8™ microprocessor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2014, pp. 98–99, doi: [10.1109/ISSCC.2014.6757354](https://doi.org/10.1109/ISSCC.2014.6757354).
- [35] S. T. Kim *et al.*, "Enabling wide autonomous DVFS in a 22 nm graphics execution core using a digitally controlled hybrid LDO/switched-capacitor VR with fast droop mitigation," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2015, pp. 1–3, doi: [10.1109/ISSCC.2015.7062972](https://doi.org/10.1109/ISSCC.2015.7062972).
- [36] T. Mahajan, R. Muthukaruppan, D. M. Shetty, S. Mangal, and H. K. Krishnamurthy, "Digitally controlled voltage regulator using oscillator-based ADC with fast-transient-response and wide dropout range in 14 nm CMOS," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, Austin, TX, USA, Apr. 2017, pp. 1–4, doi: [10.1109/CICC.2017.7993670](https://doi.org/10.1109/CICC.2017.7993670).
- [37] P. Hazucha *et al.*, "High voltage tolerant linear regulator with fast digital control for biasing of integrated DC-DC converters," *IEEE J. Solid-State Circuits*, vol. 42, no. 1, pp. 66–73, Jan. 2007, doi: [10.1109/JSSC.2006.885060](https://doi.org/10.1109/JSSC.2006.885060).
- [38] D. Rossi, J. M. Cazeaux, M. Omana, C. Metra, and A. Chatterjee, "Accurate linear model for SET critical charge estimation," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 17, no. 8, pp. 1161–1166, Aug. 2009, doi: [10.1109/TVLSI.2009.2020391](https://doi.org/10.1109/TVLSI.2009.2020391).
- [39] S. Kumar *et al.*, "Statistical characterization of radiation-induced pulse waveforms and flip-flop soft errors in 14 nm tri-gate CMOS using a back-sampling chain (BSC) technique," in *Proc. Symp. VLSI Technol.*, Kyoto, Japan, Jun. 2017, pp. C114–C115, doi: [10.23919/VLSIT.2017.7998134](https://doi.org/10.23919/VLSIT.2017.7998134).
- [40] P. Hazucha *et al.*, "Neutron soft error rate measurements in a 90-nm CMOS process and scaling trends in SRAM from 0.25- μ m to 90-nm generation," in *IEDM Tech. Dig.*, Washington, DC, USA, Dec. 2003, pp. 21.5.1–21.5.4, doi: [10.1109/IEDM.2003.1269336](https://doi.org/10.1109/IEDM.2003.1269336).
- [41] H. Fuketa, M. Hashimoto, Y. Mitsuyama, and T. Onoye, "Neutron-induced soft errors and multiple cell upsets in 65-nm 10T subthreshold SRAM," *IEEE Trans. Nucl. Sci.*, vol. 58, no. 4, pp. 2097–2102, Aug. 2011, doi: [10.1109/TNS.2011.2159993](https://doi.org/10.1109/TNS.2011.2159993).