

Architecture Implications of Pads as a Scarce Resource

Runjie Zhang[†] Ke Wang[†] Brett H. Meyer^{*} Mircea R. Stan[†] Kevin Skadron[†]
University of Virginia[†], McGill University^{*}

{runjie, kewang, mircea, skadron}@virginia.edu brett.meyer@mcgill.ca

Abstract

Due to non-ideal technology scaling, delivering a stable supply voltage is increasingly challenging. Furthermore, competition for limited chip interface resources (i.e., C4 pads) between power supply and I/O, and the loss of such resources to electromigration, means that constructing a power delivery network (PDN) that satisfies noise margins without compromising performance is and will remain a critical problem for architects and circuit designers alike. Simple guardbanding will no longer work, as the consequent performance penalty will grow with technology scaling.

In this paper, we develop a pre-RTL PDN model, VoltSpot, for the purpose of studying the performance and noise trade-offs among power supply and I/O pad allocation, the effectiveness of noise mitigation techniques, and the consequent implications of electromigration-induced PDN pad failure. Our simulations demonstrate that, despite their integral role in the PDN, power/ground pads can be aggressively reduced (by conversion into I/O pads) to their electromigration limit with minimal performance impact from extra voltage noise – provided the system implements a suitable noise-mitigation strategy. The key observation is that even though reducing power/ground pads significantly increases the number of voltage emergencies, the average noise amplitude increase is small. Overall, we can triple I/O bandwidth while maintaining target lifetimes and incurring only 1.5% slowdown.

1. Introduction

While CMOS technology scaling has resulted in exponentially greater transistor densities, threshold and supply voltages no longer decrease fast enough to prevent exponential growth in on-chip power and current density. As a result, delivering a stable voltage supply to switching transistors is increasingly challenging. The optimization of the power delivery network's (PDN's) physical structure (e.g., metal layer geometry [19], decoupling capacitor distribution [29] and controlled collapse chip connection (C4) pad location [35]) helps to reduce supply noise, and to date, noise can generally be handled with modest guardbands (i.e., slight under-clocking and/or over-volting). However, power delivery difficulties arise and will get worse for several reasons. First, despite advances that have produced sophisticated on-chip PDNs, the intrinsic resistance and inductance of the PDN circuit cannot be fully controlled with reasonable cost. This makes voltage fluctuation on the supply rails inevitable, threatening correct processor functionality with noise-induced short-term timing errors. Second,

the PDN also suffers from long-term reliability threats such as electromigration (EM). EM results in permanent failures and directly affects chip lifetime and voltage stability. Finally, there is a contention between power delivery needs and processor computation needs: the only connections between a silicon chip and the outside world are the C4 pads, required by both power supply and chip I/O signal links. The demand for I/O is expected to grow exponentially with on-chip processor counts, while the demand for power supply pads also increases as current consumption scales up. Unfortunately, C4 density is expected to remain flat in the foreseeable future [17].

We therefore hypothesize that C4 pads are a scarce resource valuable to both circuit designers and chip architects, and provisioning of C4 pads should be exposed for architectural exploration. In fact, we find that, despite their integral role in the PDN, power-ground (P/G) pads can be aggressively reduced to their electromigration limit with minimal performance impact due to voltage noise—but *only* with a suitable noise-mitigation strategy. The reason is that even though the number of voltage-noise events increases significantly, both due to technology scaling and further exacerbated by fewer P/G pads, the change in noise magnitude is small, and can be addressed with modest changes in noise mitigation and guardbanding. Making such analysis and related design choices available to designers in turn requires a pre-RTL PDN voltage model that is detailed enough to capture the effect of design-time pad placement and run-time noise mitigation strategies.

The major contributions of this paper are:

- We design and validate *VoltSpot*, a pre-RTL PDN model for architecture-level PDN noise and reliability evaluation. *VoltSpot* utilizes a fine-grained grid model capable of capturing the relationship between PDN design details (e.g., C4 pad count and placement) and supply-voltage noise. Combined with other architecture-level tools, *VoltSpot* provides a versatile platform for investigating the effect of application- and time-dependent noise, evaluating design- and run-time noise mitigation techniques, and estimating vulnerability to lifetime-reliability problems such as electromigration. Due to its pre-RTL nature, *VoltSpot* enables multi-dimensional design space explorations that include I/O-pad allocation. *VoltSpot* is publicly available at <http://lava.cs.virginia.edu/VoltSpot>.
- We explore the impact of C4 pad configuration on power-supply noise with the assistance of cycle-accurate performance simulation and power modeling. Our results indicate that replacing some power pads with I/O pads only

marginally increases the *amplitude* of supply noise, even though the *number* of noise events increases dramatically.

- We compare different state-of-the-art run-time noise mitigation techniques and observe that a hybrid mechanism that combines dynamic margin adaptation and noise-induced error recovery is the most robust to technology scaling and worst-case noise-inducing power viruses. With this hybrid technique, noise introduced by reducing the number of P/G pads can be mitigated with negligible overhead (1.5% slow-down). The key insight is that a very small increase in the default timing guardband eliminates problems due to the much greater frequency of small/medium noise events.
- We also study the effect EM-induced PDN pad failure has on transient noise. With reduced P/G pad count, earlier and more frequent PDN pad failures are expected. However, when appropriate noise mitigation strategies are available, performance degrades gracefully. EM ultimately limits the extent to which P/G pads can be reduced in favor of increasing I/O pads, because reducing the number of P/G pads increases the remaining pads' current.

These findings in turn enable a reduction in the number of P/G pads, with negligible performance overhead, in favor of a substantial increase in I/O bandwidth.

2. Background and Related Work

Due to the power delivery network's resistance, capacitance and inductance, the supply voltage will drop or fluctuate in response to current traveling through the PDN. Since transistor delay is directly related to the voltage between its source and drain [32], any voltage variation beyond the assumed design margin can cause a timing error, threatening program correctness. The main design goal for the power delivery system for a modern microprocessor is therefore to ensure that the on-chip supply voltage is as *spatially uniform* and *temporally stable* as possible. To achieve this, a modern PDN usually consists of several levels of voltage regulator modules (VRM) and decoupling capacitors. Metal traces on both the printed circuit board (PCB) and in the chip package deliver supply current from off-board or off-chip VRMs¹ to C4 pads, and on-chip metal layers further distribute current to the transistors [29].

As mentioned, considerable work has explored how to optimize the entire PDN, so that modest guardbanding—essentially under-clocking relative to the ideal voltage-frequency relationship—suffices today to cope with voltage noise. However, voltage droops² remain unavoidable and are worsening. This is because increasing current density exacerbates both localized LdI/dt and global LC resonance, while decreasing supply voltage reduces voltage-fluctuation tolerance: simple guardbanding will become more expensive.

¹In this work, we only consider off-chip VRMs. The research regarding on-chip VRMs is still in flux, so we leave pad-allocation and voltage-noise issues with on-chip VRMs for future work. However, VoltSpot can be easily extended to support such modeling.

²In this paper, we call transient noise 'droop' and static noise 'drop'.




Figure 1: (a) Side view of a chip showing its connection to the package through C4 pads. (b) Top view of a C4 array showing an example allocation of power pads and I/O pads (real processors have many more C4 pads than illustrated here).




Figure 2: Voltage-emergency maps for different pad configurations for a 16nm, 16-core chip. A total of 1914 C4 locations are available. All three graphs have the same color scale; areas with warmer colors indicate more voltage violations.

Figure 1 illustrates how electrical current and I/O signals are delivered to a chip. C4 pads play a crucial role in power delivery. Due to growing power density, chips manufactured in future technology nodes (e.g., 16 nm) will require 150-200 A or more. However, the maximum feasible C4 pad density does not scale with CMOS technology [17], and pad composition and hence material properties (e.g., resistivity and maximum allowed current density) are improving slowly, if at all. If I/O pad count increases to keep pace with core count, there will be increasing competition between I/O and power-ground needs for scarce C4 pad sites.

Unfortunately, existing pre-RTL, architecture-level PDN studies often make simple assumptions about C4 pads, and none have considered the power/ground vs. I/O tradeoff. They either use lumped PDN models that collapse all pads into a single resistor-inductor pair [8, 10, 30], or use coarse-grained models that are not able to accurately capture the effect of pad count and location [9]. Fig. 2 illustrates the type of effect that pad count and location can have on on-chip voltage noise. For this figure, we simulate a 16-core processor with our fine-grained model running a PDN-stressing workload for 100K cycles. We count the number of cycles in which a voltage emergency (defined for this figure as a cycle with an average voltage droop larger than 5% V_{dd}) occurs on the chip.

Fig. 2a and 2b illustrate configurations with the same number of power supply pads, but only the latter case has optimized pad locations. Due to the sub-optimal allocation of pads, Fig. 2a experiences 6x more emergency cycles compared with Fig. 2b. Fig. 2c illustrates a configuration with optimized pad locations, but 40% fewer pad count than in Fig. 2b. Although the optimized locations in Fig. 2c prevent extreme voltage-noise hotspots, the system still experiences up to 3x more emergency cycles than Fig. 2c. Clearly, both pad count and pad locations have a large effect on on-chip voltage noise.

Other researchers have studied the effect of C4 optimization during pre-RTL design. For example, Wang et al. [35] evaluated the impact of pad placement on chip IR drop and proposed an optimization algorithm to minimize global IR drop. Zhang et al. [39] also used IR drop as a figure of merit and examined the impact of C4 pad count on technology scaling. Although the PDN models used by these works are fine-grained enough to precisely model C4 pads, they cannot simulate a PDN’s transient behavior. We will show in Sec. 5 that evaluating only steady-state IR drop severely underestimates supply voltage noise and performance effects. In fact, we also show that IR drop is only a small component of runtime voltage noise.

Runtime noise mitigation techniques have been proposed to either avoid excessive voltage noise [8, 21, 30], or recover from noise-induced errors [10]. To reduce the energy overhead of guarding against worst case, several proposals [11, 22] dynamically adjust circuit timing margins to save energy during average-case execution while guaranteeing functionality in the worst case. We will evaluate both error recovery and margin adaptation methods and show the necessity of combining these into a hybrid strategy.

3. VoltSpot

VoltSpot takes as input a processor floorplan, described at the level of architectural units, pad locations, and a per-unit power trace, and calculates the transient current and voltage observed at each pad and within each architectural block; its model structure is illustrated in Figure 3. The typical regularity of the on-chip PDN’s physical structure makes compact on-chip PDN modeling feasible. We adopt a well-accepted methodology [9, 13] that models the Vdd and ground nets as separate regular 2D circuit meshes. C4 pads are modeled as individual resistor-inductor branches attached to on-chip grid nodes, and on-chip decoupling capacitors as distributed capacitors connecting the Vdd and ground grids. Ideal current sources model the load (i.e., the power of the switching transistors and associated leakage), and the current values are calculated as $I = \frac{Power}{SupplyVoltage}$. Compared with on-chip wires, vias have much lower impedance due to their size. For this reason, VoltSpot ignores their resistance and inductance. We note that this abstraction significantly reduces the problem size and thus enables application-level noise simulation.

Since our main focus is the on-chip PDN, we model off-chip components such as the package with lumped RLC elements




Figure 3: VoltSpot model structure

and assume the PCB provides an ideal power supply. VoltSpot therefore captures both mid- (package) and high-frequency (on-chip) responses of the PDN [29]. Like other pre-RTL tools, VoltSpot assumes that power density is uniform within each architectural block. This assumption can be relaxed by simply providing a higher-resolution power map.

VoltSpot strikes a careful balance between modeling precision and simulation speed. Consequently, VoltSpot makes it possible to perform detailed, performance-simulator driven evaluations of transient PDN noise as well as the effect of design- and run-time mitigation techniques. Such explorations are not possible with physical level modeling due to complexity; higher-level models (e.g., lumped models) are too inaccurate. As a result, VoltSpot makes it possible to accurately simulate application-specific noise and related phenomena (e.g., electromigration), evaluate mitigation techniques, and in turn give architects the opportunity to make critical, early design trade-offs (e.g., power supply pads vs. I/O pads). VoltSpot is implemented as a C library and is designed to be compatible with a variety of architectural modeling tools and is publicly available as an open-source tool.

3.1. Improvements Beyond State-of-the-Art Models

VoltSpot makes two key improvements over alternative models in the literature: (1) transient PDN grid modeling at the granularity of pad pitch or smaller, and (2) multi-layer-metal PDN modeling using multiple, parallel RL branches.

Due to their coarse modeling granularity, previous pre-RTL PDN models are incapable of either modeling pads in detail, or precisely capturing localized noise. Alternatively, VoltSpot

makes the on-chip grid size a function of C4 pad array size. For example, for a flip-chip design with 2,500 C4 bumps (e.g., distributed as a 50x50 array), the minimum modeling grid size will also be 50x50. To calculate the resistance of on-chip grid, VoltSpot uses the metal resistance equation ($R = \rho * l/A$, where ρ is the resistivity per unit area of the metal, A is cross-sectional area, and l is wire length). For on-chip inductance calculation, VoltSpot adopts the equation from [19]:

$$L_{eff} = \frac{\mu_0 l}{N\pi} \left[\ln \left(\frac{w+s}{w+t} \right) + \frac{3}{2} + \ln \left(\frac{2}{\pi} \right) \right] \quad (1)$$

where N , μ_0 , l , w , t , s are the number of power and ground pairs, permeability of the vacuum, length, width and thickness of a single wire, and the spacing between wires. Our results show that, in some cases, even the finest grained on-chip grid from previous models (e.g., 12x12 [9]) underestimates the localized voltage noise amplitude by 20% and voltage emergency count by 3x. We also find that, due to the limitation of the granularity of our power model, increasing PDN modeling granularity beyond an average of four grid nodes per C4 pad (i.e., using a 100x100 grid to model a chip with a 50x50 C4 array) improves precision by less than 3% (in terms of noise amplitude). For these reasons, we set the grid-node-to-pad-ratio to 4:1.

To further improve accuracy, VoltSpot explicitly models the multi-layer structure of physical PDN. Previous work suggests that the electrical properties of on-chip metal layers heavily depend on their geometry (width, pitch, etc.) [26]. Different RL properties respond to different time constants as current is changing. A single RL circuit thus cannot accurately describe the entire stack. We replace the single RL pair in on-chip grid with multiple parallel RL branches (as shown in Fig. 3c) and calculate RL values based on the different metal layers' geometry. Our results show that the single RL pair model using values extracted from the top metal layers overestimates PDN inductance and reports a voltage noise amplitude 30% larger than a multi-pair-branch model that considers six layers of PDN metal.

VoltSpot models PDNs with up to tens of thousands of RLC components. To solve such large-scale circuits efficiently and accurately at each time step, we choose the implicit trapezoidal numerical method, an A-stable method with 2nd-order accuracy [5]. This method is the default ordinary differential equation solver in SPICE and it is also widely used in circuit and PDN simulation [31]. Since solver error monotonically increases with simulation time step, we set our time step to one fifth of a cycle at 3.7GHz (around 50ps) to keep the numerical error of node voltage below 10^{-5} V. We use an open-source sparse matrix solver, SuperLU [24], and optimize memory efficiency with multiple minimum-degree reorderings, significantly reducing fill-ins in sparse LU decomposition.

3.2. Validation

We validated VoltSpot using an IBM PDN analysis benchmark suite [27]. The suite consists of detailed PDN structural infor-

mation for six chips with different die sizes, silicon designs and metal layer counts. The PDN structure is given in SPICE format and includes metal wires' geometric properties and resistance. Other information, such as C4 pads placement, on-chip decap distribution and workload pattern, can also be extracted from the SPICE file. Besides the PDN structure, this benchmark suite also provides both steady-state and transient SPICE simulation results for all or selected on-chip nodes.

The focus of VoltSpot is capturing within die current and voltage variation. We therefore evaluate VoltSpot's accuracy by comparing both simulated static C4 currents and transient on-chip voltage droops with those derived from the reference SPICE netlists accompanying the IBM benchmarks. For each benchmark (Bench), Table 1 shows the number of circuit nodes (# Nodes), metal layers (# of Layers) and power supply pads (# of Pads) and the validations results. The table also shows whether the resistance of vias is ignored in the SPICE model. Even though the variation of pad current within the same chip could be as large as 5x (observed in PG3), VoltSpot still accurately captures all pads' current with an average error of 5.2% (Pad Current Error). In transient validation results, Voltage Error Average shows the average node voltage mismatch across all simulated time steps and all given on-chip nodes, while Voltage Error Max Droop compares the max droops observed during the entire transient simulation. Both metrics give low error even for the benchmarks that include detailed via information. This demonstrates that VoltSpot provides high-quality estimation for on-chip voltage fluctuation and that vias can be safely omitted from the model.

We observe in the IBM benchmark suite that PDNs with more metal layers or elements usually have a more regular structure than smaller PDNs. Since VoltSpot assumes a regular on-chip metal stack, it is most accurate when modeling large-scale PDNs. PG1 not only has the fewest elements, but also employs asymmetric grids that do not map well to our PDN model. Since the PDNs of modern high-performance processors usually contain multiple layers of regular metal traces, PG1 is less representative. We therefore exclude PG1 from our validation.

4. Simulation Setup

To study the effect of technology scaling trends on PDN noise in the near future and explore the resulting trade-offs in architectural design choices due to power-delivery voltage-noise limitations, we create a series of multicore processor configurations scaled down to 16nm. For a reasonably modern configuration focused on single-thread performance, we chose a 3.7GHz 45nm Intel Penryn-like processor [7] as the baseline. It has two 32-bit 4-way out-of-order cores. Each core contains a 32kB L1 instruction cache and a 32kB L1 data cache. Unified L2 caches private to each core are each 3MB. For each technology node, we hold the processor architecture constant but assume that the number of cores (and therefore the number of L2s) doubles. For scalability and consistency, we assume a

| Bench | # of Nodes | # of Layers | Ignores Via R | # of Pads | Current Range(mA) | Pad Current Error (%) | Voltage Error: Average (%Vdd) | Voltage Error: Max Droop (%Vdd) | Voltage Error: Correlation (R^2) |
|-------|------------|-------------|---------------|-----------|-------------------|-----------------------|-------------------------------|---------------------------------|--------------------------------------|
| PG2 | 0.25M | 5 | No | 120 | 620-1530 | 5.2 | 0.21 | 0.86 | 0.968 |
| PG3 | 1.60M | 5 | No | 461 | 116-571 | 3.3 | 0.11 | 0.46 | 0.977 |
| PG4 | 1.84M | 6 | No | 312 | 13-24 | 2.9 | 0.04 | 0.06 | 0.967 |
| PG5 | 2.16M | 3 | Yes | 177 | 60-110 | 3.7 | 0.08 | 0.11 | 0.983 |
| PG6 | 3.25M | 3 | Yes | 132 | 210-410 | 2.7 | 0.11 | 0.54 | 0.966 |

Table 1: Static and transient validation results against IBM benchmark




Figure 4: Floorplan of the Penryn-like 16core processor.

| Tech Node (nm) | 45 | 32 | 22 | 16 |
|-------------------------|-------|-------|-------|-------|
| # of Cores | 2 | 4 | 8 | 16 |
| Area (mm ²) | 115.9 | 124.1 | 134.4 | 159.4 |
| Total C4 Pads | 1369 | 1521 | 1600 | 1914 |
| Supply Voltage (V) | 1.0 | 0.9 | 0.8 | 0.7 |
| Peak Total Power (W) | 73.7 | 98.5 | 117.8 | 151.7 |

Table 2: Characteristics of Penryn-like Multicore Processors

mesh network-on-chip (NoC). Fig. 4 shows the floorplan of our 16nm, 16-core processor.

To get chip-wide and application-specific power consumption and area for all technology nodes, we use McPAT [23], an architecture-level power model, and integrate it with Gem5 [2], a multi-core performance simulator. Table 2 shows the area and peak power (including leakage power) results for our designs. We use ArchFP [6] to generate our floorplans.

Besides tolerating voltage noise, design margins also deal with application-induced thermal variation, process variation, aging, and test inaccuracy [22]. In this paper, we only address the portion of margin needed for voltage noise and omit the margin for other sources of variations.

4.1. Power Trace Sampling and Stressmark

To accelerate simulation, we borrow the idea of statistical sampling [12, 38]. The idea behind statistical sampling is to simulate in detail only short samples from much longer applications and estimate the behavior of the whole application based on the sampled segments. In performance studies, the accuracy of sampling heavily depends on how structures are “warmed up” before each sample, since the state of many architectural blocks (e.g., cache, branch predictor) accumulates over a long history and significantly affects performance [12]. Fortunately, warming up a PDN simulation is much simpler: the only fac-

tor that might affect results accuracy is the electrical charge in decoupling capacitors, which have short time-constants. Our study shows that 1000 cycles of warm-up at 3.7GHz is sufficient.

According to [38], a total of 2 million instructions (including instructions for warm up) sampled at equal intervals is sufficient to estimate a multi-billion-cycle application’s IPC with $\pm 3\%$ error and 99.7% confidence. For this reason, we take 1000 samples at equal intervals from the end-to-end execution of simmedium inputs for 11 benchmarks in the Parsec 2.0 benchmark suite [1]. Two benchmarks (facesim and canneal) were incompatible with our performance simulation infrastructure and were omitted. Each sample contains 2000 cycles of per-cycle power information, of which the first 1000 cycles in each sample are used for warm-up.

We take two further steps to ensure that we simulate worst-case behavior. First, the sampled power traces are taken from 2-core simulations; for technology nodes with more than two cores, we replicate the 2-core power trace to 4, 8 or 16 cores. In this way, transient current fluctuation representative of that in Parsec benchmarks occurs simultaneously in each pair of cores, increasing the stress on the PDN. Second, we construct a stressmark to simulate a voltage noise virus by selecting the most noisy (in terms of noise amplitude) power trace among all samples and replicating it 1000 times. We illustrate the stressmark’s noise pattern in Fig. 5.

4.2. PDN Parameters and Pad Location Optimization

Table 3 lists the major physical PDN parameters we use with VoltSpot. For on-chip metal, we use copper and adopt a metal layer structure similar to an Intel 45nm metal stack [36]. Table 3 reports the width (W), pitch (P), and thickness (T) of the global, intermediate, and local layers. For on-chip decoupling capacitors, we use deep trench capacitors, which have the highest capacitance density [28]. The die area allocated to on-chip decap is a design parameter and will be discussed in Sec. 6. We assume SnPb is the primary material for C4 pads. Typical pad diameter and resistivity are derived from [37]. Pad spacing was selected so that our pad density matches ITRS [17] projections. Package resistance, inductance and capacitance come from [14]. We performed sensitivity studies on package and pad parameters (e.g., SnAg pads) and observe that the effect of pad allocation (Sec. 5) or wear-out (Sec. 7) on voltage noise

| | |
|---|-----------------|
| On-Chip Metal Resistivity (ρ) | 1.68e-8 |
| Global PDN Layers W/P/T (μm) | 10 / 30 / 3.5 |
| Intermediate PDN Layers W/P/T (μm) | 400 / 810 / 720 |
| Local PDN Layers W/P/T (μm) | 120 / 240 / 216 |
| On-Chip De-cap Density (nF/mm^2) | 100 |
| C4 Pad Diameter/Pitch (μm) | 100 / 285 |
| C4 Pad Resistance/Inductance ($m\Omega/pH$) | 10 / 7.2 |
| Package Resistance (R_{pkg_s}) ($m\Omega$) | 0.015 |
| Package Inductance (L_{pkg_s}) (pH) | 3 |
| Package Resistance (R_{pkg_p}) ($m\Omega$) | 0.5415 |
| Package Inductance (L_{pkg_p}) (pH) | 4.61 |
| Package Capacitance (C_{pkg_p}) (μF) | 26.4 |

Table 3: PDN Parameters

is insensitive to these variables: the impedance of the PDN’s pad layer mostly depends on pad configuration.

As mentioned before in Sec. 2, the number and locations of C4 pads have a significant effect on power delivery quality. To avoid sub-optimal pad allocation, we adopt the simulated-annealing algorithm described in [35] and extend it to jointly optimize both Vdd and ground pad locations.

5. Transient Voltage Noise: Scaling and Effects of C4 Pad Configuration

Power supply noise comes from three major sources: static IR drop, LdI/dt droop, and LC resonance. IR drop is the consequence of PDN resistance at any given time irrespective of processor behavior sequence. Instead, dynamic noise is triggered by certain chip behaviors such as a sudden change in power consumption, which will create large LdI/dt noise, or reoccurring power consumption patterns at or near an LC resonance frequency. Fig. 5 compares an on-chip node’s resistive drop (a function of instantaneous current and resistance only) and transient noise (from all three noise sources) over 1000 clock cycles. We observe that IR drop only constitutes a small fraction of the overall noise, suggesting that considering IR drop alone (as have all prior studies of C4 pads) is insufficient in PDN design and optimization. For this reason, for the remainder of the paper we will use aggregate transient voltage droop as the key metric for evaluating the quality of a PDN; IR drop will not be distinguished as a separate issue.

5.1. Scaling Trends of Supply-Voltage Noise

As manufacturing process technologies scale, current density grows as power density increases and supply voltage decreases. We compared maximum voltage droop and the number of voltage noise events across process technology nodes. Here we define a transient voltage droop greater than a certain threshold as a voltage-droop violation. We assumed a fixed PDN metal stack structure across technologies. This is because in our study, we focus on the upper layers of metal, where wires are bulky and thus less affected by technology. Our sensitivity studies show that the impact of metal width on voltage noise




Figure 5: Comparison of transient voltage noise and static IR drop over a 1K cycle window in the benchmark ferret. Periodic oscillation implies that LC resonance is the major cause. We build our noise-inducing stressmark based on this segment.

| Tech Node (nm) | 45 | 32 | 22 | 16 |
|---------------------------|------|------|------|-------|
| Maximum Noise (% Vdd) | 7.96 | 8.91 | 9.49 | 11.87 |
| Violations (8% Threshold) | 0 | 3 | 37 | 598 |
| Violations (5% Threshold) | 1515 | 2288 | 2881 | 6668 |

Table 4: Voltage Noise Scaling Trend, Ideal (all pads allocated to power/ground), in fluidanimate benchmark

amplitude is small ($\pm 50\%$ metal-width changes max noise-amplitude by less than 0.5% Vdd). Other parameters such as decap density and C4 pad density are also kept constant for fair comparisons between technologies. For this scaling limit study, we allocate all available C4 pads to power, to study the upper bound of PDN quality. The impact of C4 pad placement will be discussed in detail in Sec. 5.2. We simulate fluidanimate, one of the most noisy applications in the suite. The results of our experiment are summarized in Table 4.

First, we observe that the magnitude of voltage droop increases as feature size decreases: the maximum on-chip voltage droop increases by approximately 4% of Vdd from 45nm to 16nm. Second, voltage-noise violation events occur more frequently with scaling. We evaluated noise event frequency for two different thresholds (5% and 8% of Vdd); our results show rapid growth rate of violation count in each case.

This is a best-case scenario, with all pads allocated to power/ground, and a realistic benchmark. With a more realistic pad configuration and our stressmark, we observe that the maximum noise is actually 13% at 16nm. We use this value as our static safety margin in the rest of the paper, because a static guardband must countenance worst-case behavior.

5.2. Voltage Noise and Pad Configuration Effects

Our primary interest in this paper is the contention between power supply and chip I/O; it is therefore important to quantify the effect of C4 pad allocation on power supply noise. Since power noise increases rapidly with technology scaling, we henceforth focus on the 16nm node and use a 16-core Penryn-like processor as the platform for our analysis. To better illustrate the trade-off between I/O pad count and processor




Figure 6: Voltage noise change across different pad configurations. Each MC needs 30 I/O pads.

off-chip memory bandwidth, we assume all on-chip memory controllers (MC) are single-channel and convert the number of available I/O pads into the number of supported MCs (different MC channels cannot share C4 pads [18]); more MCs means fewer Vdd and GND pads, and vice versa.

Based on the pad requirement breakdowns for three different commercial processors [15, 16, 33], we assume that our 16-core chip includes four inter-chip links (85 pads each) and a total of 85 miscellaneous pads (including clock, dynamic voltage scaling control, sensing, debug, testing, etc.). We assume all MCs are FBDIMM interfaces (as in the Xeon E7-8800) to represent the trend toward narrower, more serial interfaces; FBDIMM requires far fewer pads (about 30 per MC channel) than DDR3 (80-120 pads per channel). The total number of C4 pads for our 16nm chip is 1914; all pads not used for I/O are dedicated to power supply.

Fig. 6 shows the relationship between supply voltage noise and I/O configuration. We evaluate voltage noise with two different metrics: noise violation count and noise amplitude. The violation rate bars assumes a 5% voltage-droop violation threshold and report the number of noise violations (averaged across all 1000 samples). The lines in Fig. 6 illustrate the maximum observed voltage noise (averaged across all samples). Independent of the application, increasing the number of MCs worsens voltage noise in terms of both violation event rate and noise amplitude, though noise amplitude increases only marginally (up to 1.5% Vdd).

Violation count increases rapidly as power supply pads decrease, because of the sensitivity of violation count to noise amplitude. The 2D array of C4 pads allows current to be directly delivered to an arbitrary on-chip point without traveling through long, high impedance, on-chip wires. As we reduce the number of power supply pads, we effectively increase the average physical distance between power supply pads and loads, increasing impedance and therefore noise. Since we optimize the location of power supply pads, the increment of average pad-to-load distances across different pad configurations is small; consequently, noise amplitude increases insignificantly. However, even a small change in noise amplitude can result in a large number of new violations, as many different nodes may already be close to the threshold.

6. Run-time Voltage Noise Mitigation

VoltSpot makes it possible to evaluate run-time noise mitigation strategies in the context of architectural design decisions. In this section, we evaluate the performance overhead of several run-time mitigation techniques, and the effect of the number of integrated memory controllers. Increasing the number of memory controllers improves performance by reducing memory access latency. However, pads for I/O must be taken from those budgeted for power supply, increasing the demands on (and performance overhead of) noise mitigation. For fair comparisons between (i) dynamic margin adaptation and (ii) error recovery techniques, we present all results in terms of speedup. We assume a constant supply voltage and adjust the timing margin only. Since there is a roughly linear relationship between supply voltage droop and circuit delay within a reasonable range [32], we assume a voltage droop of $X\%$ Vdd increases circuit delay by $X\%$.

6.1. Dynamic Margin Adaptation

As an alternative to wasteful, fixed margins, Lefurgy et al. [22] propose a technique to detect available timing margin with critical path monitors (CPM) at run-time and use fast digital phase lock loops (DPLL), capable of reducing frequency by 7% within 5ns, to increase timing margin until the integral frequency control loop catches up and raises voltage as needed. Note that this emergency response to voltage droop is a one-shot control, immediately dropping frequency by the maximum (7%) needed to preserve correct operation in the presence of the worst possible voltage droop. However, voltage could continue to drop rapidly before the DPLL can adjust, so the processor must always underclock (relative to the current voltage) by enough to preserve correct behavior given the worst-case voltage slew rate in 5ns. The magnitude of this guardband has not been made public.

In this paper, we only focus on performance. We adapt the above methodology by assuming that the voltage remains fixed at the maximum, and only clock speed is reduced in response to voltage droop.³ We assume ideal voltage sensing in each

³For lower DVFS settings, performance is less critical, and margin can be more generous, so here we focus on the highest performance state.

core, and per-core DPLLs to respond to per-core voltage-droop behavior. We observe that most applications exhibit phases of low and high voltage droop. Our integral loop tracks worst-case voltage droop over a monitoring period (one sample in our simulation methodology) and then sets the clock speed accordingly for the next sample. Within a sample, any voltage droop in excess of the margin set by the integral loop initiates a one-shot safety response that lowers the clock speed by 7%, or to the worst-case margin of 13% (whichever is smaller, given the last setting from the integral loop), in order to protect against a worst-case, rapid droop. Frequency is reset (the one-shot change is removed) at the next integral loop update.

One-shot control is not enough to protect against rapid voltage droop during the DPLL change. An extra safety margin must always be maintained to allow for the worst-case voltage slew rate during the 5ns DPLL update. This means that if the integral loop currently allows a voltage droop of $X\%$, the clock frequency must be reduced by an additional $S\%$. In this case, before the one-shot control engages, the clock speed is $X + S\%$ below nominal. An $X\%$ droop is the trigger; if a voltage droop exceeds $X\%$, the one-shot control engages, reducing frequency to $X + S + 7\%$ (or 13%, whichever is smaller) below nominal. S , in other words, accommodates the worst-case voltage droop that could occur while the DPLL is changing.

We determine the necessary safety margin (S) as a function of technology node (using a brute-force search). Both S and the worst-case margin grow significantly from 45nm to 16nm, as shown in Table 5. We observe that the required safety margin increases by almost 2%, significantly reducing the performance benefits of margin adaptation: on average, the average portion of the 13% worst-case margin that can be removed for performance improvement shrinks from 27% to 9%. The problem is that margin adaptation must be very conservative to guard against potential (but rare) worst-case voltage droops. Since margin adaptation only removes margin during low-noise program phases, we choose fluidanimate instead of our stressmark (Sec. 4.1) for this analysis, otherwise the margin controller could not reduce margin at all due to the stressmark’s constantly noisy behavior.

In these simulations, we hold the PDN design (power supply pads density, on-chip decap density, metal structure, etc.) constant across technologies. It is possible to reduce voltage droop slew rate and margin adaptation safety margins by adding more on-chip decap as technology scales. However, our design space exploration study shows that, to keep the 16 nm chip’s performance overhead on a par with that of 45 nm chip, at least 15% more die area must be allocated to decap, a cost equivalent to two cores.

6.2. Recovering From Noise-Induced Errors

An alternative to dynamic margin adaptation is to roll back and recover when a timing error is detected [10]. Such techniques address a key weakness of dynamic margin adaptation: false positives where noise reduces timing margin but would not

| Tech Node (nm) | 45 | 32 | 22 | 16 |
|-----------------------------|------|------|------|-----|
| Safety Margin (S , %Vdd) | 2.5 | 2.9 | 3.1 | 4.3 |
| % of Margin Removed | 26.9 | 23.6 | 20.9 | 8.6 |

Table 5: Dynamic Margin Adaptation and Scaling

ultimately cause an error. To analyze the overhead of noise recovery, we first simulate benchmarks to completion and collect noise amplitude data. Then, we perform post-processing to determine, given an allowed voltage droop and recovery overhead in clock cycles, the total performance overhead in cycles. Fig. 7 shows the performance effect of different timing margin settings. Using a design with constant margin of 13% as our baseline (no recovery needed), we evaluated recovery-based methods’ performance with different benchmarks on our 16 nm, 16-core chip with 24 memory controllers.

We observe that as we remove timing margin, the processor runs faster but also experiences more errors. As a result, the recovery penalty associated with removing too much margin overwhelms the benefit of increased clock frequency. In extreme cases (e.g., fluidanimate with only 5% margin), aggressive margin settings introduce so many errors that it significantly hurts processor performance. We assume here that each error recovery requires 30 cycles (rollback 10 cycles and replay at half frequency [10]) and observe that on average, 8% timing margin gives the best performance.

6.3. A Hybrid Technique

Preventive margin adaptation will perform poorly in future technologies because it has to preserve a large safety margin to prevent timing errors. However, by adding the protection of error *recovery*, the margin controller no longer needs to prevent all errors and can operate with a much lower safety margin. We combine the above two methods, so that the processor can both adjust margin at run-time and recover from errors that exceed the margin. The margin controller in this hybrid technique monitors voltage noise. When a voltage emergency is detected, the controller records the amplitude of that violation and triggers recovery. After the recovery, the controller increases timing margin to match the observed noise amplitude. Both the hybrid technique and recovery-only technique react after noise events happen. The advantage of the hybrid technique is the ability to adjust the noise tolerance threshold via frequency scaling.

Fig. 8 depicts a performance comparison between the techniques discussed above. The chip evaluated here is also a 16-core, 24 MC processor. “Ideal” bars represent the performance gain achieved by an oracle margin controller that always maintains the minimum required margin without causing any timing errors. As a sensitivity study, we explored three different rollback penalties for the recovery technique. Using the analysis in Fig. 7, we select the optimal timing margin setting for each rollback penalty assumption. As we expected,




Figure 7: Speed up of recovery-based techniques with different timing margin settings. We evaluate our 16nm, 16core processor with 24MCs here. The baseline case enforces a 13% timing margin and thus guarantees timing is error-free.

the margin-adaptation-only technique has lower speedup compared with the recovery-based technique. We also observe that the recovery-based techniques’ performance is minimally sensitive to rollback penalty. This is because, with proper margin setting, errors happen rarely, thus spending more cycles on each error would not significantly degrade overall performance. In contrast, the performance of the hybrid technique is much more sensitive to error recovery overhead, because it relies on the occurrence of errors to trigger margin adjustment, and therefore experiences more errors than a well-tuned recovery technique. The average results from the Parsec benchmark suite show that the hybrid technique only barely outperforms the recovery-only technique when recovery cost is low.

While evaluating mitigation techniques’ performance with our stressmark, we observe that recovery-only experiences significant slowdown (right most bars in Fig. 8; note that the stressmark’s performance was excluded from Parsec average calculation). This is because, in order to achieve optimal performance with normal applications, we relax timing margin toward the average case. However, since the stressmark constantly excites PDN resonance, it will experience frequent errors (12 per 1000 cycles) under this tight margin setting, and thus makes the recovery technique suffer from frequent recovery penalties. In contrast, the hybrid technique can quickly adapt the margin to the required level at the beginning of the stressmark and avoid all remaining errors. In summary, although recovery-only performs better with typical workloads even when the recovery cost is high, the hybrid technique is more robust to worst-case behavior. If guarding against worst-case noise is the priority, designers should choose hybrid over recovery-only techniques.

6.4. Trading Power Pads for Performance

Our primary interest is the tradeoff between power supply pads and I/O pads. Sec. 5.2 points out that trading power/ground pads for I/O connections will degrade power delivery quality. Based on the preceding discussions, we choose the hybrid noise mitigation technique and assume a pessimistic per-error recovery cost of 50 cycles. Fig. 9 shows the pad-induced noise mitigation overhead for different applications. Each

application uses its own performance with the 8 MC case as the baseline. While increasing MC count proportionally improves performance, it also exacerbates supply voltage noise. Our results indicate that the performance penalty to mitigate the extra noise is fairly low, even if we aggressively increase chip MC count from 8 to 32, reducing power/ground pad allocations from 1254 to 534. The reason is that even though the number of voltage-noise events increases significantly, the change in amplitude of most of these events is small and modest changes in noise mitigation and guardbanding can address these more frequent noise events.

One side-effect of increasing the number of I/O channels is more complicated package I/O routing: lateral I/O wires could “cut” through package metal layers and split power delivery planes into separate islands. As a result, the impedance of package PDN will increase. We performed a first-order analysis of this effect by increasing the impedance of the package’s serial portion (R_{pkg_s} and L_{pkg_s} in Fig. 3b). Our analysis shows that, although larger R_{pkg_s} increases the PDN’s impedance at low frequency, it also helps to reduce its impedance at resonant frequencies due to damping. Overall, the observed maximum noise amplitude is insensitive to the change of R_{pkg_s} and L_{pkg_s} (doubling RL values only increases noise amplitude by 0.15% Vdd). With the extension of distributed package modeling, VoltSpot could be used to perform more detailed analysis of I/O routing’s impact on PDN impedance—an interesting area for future work.

7. Chip I/O Bandwidth and C4 EM Lifetime

VoltSpot also makes it possible to evaluate the effect of EM-induced PDN pad failure on transient voltage noise. EM refers to gradual mass transport in metal conductors induced by momentum transfer from electrons to atoms, is characterized by a median time to failure (MTTF) that is inversely related to current density [3]. As current density increases with technology scaling, EM pressure on C4 pads is expected to increase. In this section, we explore the relationship between EM-induced PDN pad failure, voltage noise, memory controller count, and the effectiveness of the mitigation techniques explored above.

Electromigration failures are the result of continuous, high-




Figure 8: Performance comparison between different noise mitigation techniques. Numbers after recover/hybrid in legend represent the cost (in cycles) of each recover, from error. The baseline case enforces a 13% constant timing margin.




Figure 9: Performance penalty of mitigating extra voltage noise caused by reduced power/ground pads. We use the hybrid technique with a conservative assumption of 50-cycle rollback-and-replay penalty per error. Each benchmark uses its own performance with 8MC case as baseline.

density current flow. Previous work [34] suggests that the lifetime of metal under high-frequency AC current stress is determined by the DC component of the stressing current alone. We therefore focus on DC stress only and use the processor’s peak power consumption for worst-case analysis. To be more specific, we use 85% (a ratio suggested by [39]) of theoretical peak power reported by McPAT as a stressmark power input. Although I/O pads are also vulnerable to EM, changing the number of power/ground pads would not affect per-I/O pad current. For this reason, we do not evaluate EM’s effect on I/O pads.

7.1. Whole-Chip Electromigration MTTF Calculation

For a single C4 pad, EM-induced failure times follow a log-normal distribution (with $\sigma = 0.5$ [25]). The median time to failure is determined by Black’s equation [3] and adjusted to consider current crowding and Joule heating [4]:

$$t_{50} = A(cJ)^{-n} \exp\left(\frac{Q}{k(T + \Delta T)}\right) \quad (2)$$

where J is current density, n and Q are material-specific constants (for SnPb solder bump, $n = 1.8$, $Q = 0.8eV$ [20]), k is Boltzmann’s constant, $c = 10$, $\Delta T = 40^\circ C$ [4], A is an empirical constant, and T is temperature in Kelvin. Given the MTTF, we can calculate the probability of failure, $F(t)$, for a single pad after any time t .

To quantitatively evaluate the chip’s EM lifetime, we derive a new cumulative distribution function that describes the distribution of time t when the first PDN pad failure occurs:

$$P(t) = 1 - \prod_i (1 - F_i(t)) \quad (3)$$

where $F_i(t)$ is the failure probability of the i th pad after time t and is calculated based on the individual pad current density given by VoltSpot. By definition, the median value of the above distribution is the time where $P(t) = 0.5$; it represents the median time to first PDN pad failure in the whole chip, considering all pads. We refer to it as MTTFF.

Table 6 shows the scaling trend of average on-chip current density, single-pad worst current and both MTTF and MTTFF for C4 pads. MTTF/MTTFF results are normalized to the 45nm MTTFF value. When we examine the whole chip’s robustness against EM, MTTFF is much worse than the worst single element’s expected lifetime. For example, if every single pad in a 45nm chip were designed to have a 10 year MTTF under the worst case, the median time to first PDN pad failure in the entire chip would be around 3.4 years. As technology scales, a power delivery system designed to work 10 years at 45nm would only be EM-failure-free for about 2.4 years at 16nm. All our calculations assume a temperature of $100^\circ C$ to represent the worst-case scenario.

7.2. Tolerating Pad Failures with Runtime Mitigation

The challenge when a pad fails is that it introduces more voltage droop events and increases their amplitude in the neighborhood of that pad. However, the same solution as we employed before to allow increased I/O pads can be used again: by improving voltage-droop mitigation, we can tolerate more PDN pad failures. In fact, we can tolerate multiple PDN pad failures while still converting some power/ground pads into I/O pads.

| Tech Node (nm) | 45 | 32 | 22 | 16 |
|-----------------------------------|------|------|------|------|
| Chip current density (A/mm^2) | 0.54 | 0.75 | 0.93 | 1.16 |
| Worst single pad current (A) | 0.22 | 0.29 | 0.43 | 0.50 |
| Normalized single pad MTTF | 2.94 | 1.71 | 0.87 | 0.70 |
| Normalized whole chip MTTF | 1.00 | 0.63 | 0.29 | 0.24 |

Table 6: C4 Pad EM Lifetime Scaling Trend




Figure 10: PDN pad failure’s effect on noise mitigation performance overhead and expected EM lifetime. F indicates the number of failed pads. The baseline is an 8 MC chip with no failed pads. For *Recovery*, the performance overhead with 32 MC goes off-chart to 15% if 20 pads fail (25% if 60 pads fail).

To evaluate PDN pad failures’ effect on voltage noise, we remove power pads from the previously studied 16nm, 16-core chip and simulate the EM-damaged chips with the benchmark fluidanimate. Since EM-induced failure is a stochastic event, any power supply pad could fail at any time. Thus the number of possible pad configurations is huge (e.g., for a chip with 1200 power supply pads, there are more than 10^{43} ways to have 20 failures). It is impossible to traverse all failure combinations, so as a “practical worst case” estimation, we choose to fail the pads with the highest current density. There are three reasons behind this methodology. First, a single pad’s MTTF is inversely related to current density, so pads with the highest current density tend to fail first. Second, pads carrying high current are closer to blocks with high power consumption and those blocks (e.g., the ALU) are more likely to produce large voltage noise. Therefore, removing supporting pads for those blocks introduces more voltage fluctuation and thus gives a better estimation for the worst consequences of PDN pad failure. Third, although PDN pad failures happen sequentially, and any PDN pad failure could change on-chip current distribution and thus change the distribution of current among pads, such phenomena will not reduce the failure risk of initially high-current pads: EM is an effect that accumulates over time.

The lines in Fig. 10 show the noise mitigation performance overhead of different pad-failure-tolerances F and different I/O configurations. We examine both recovery-only and hybrid techniques with a conservative recovery cost of 50 cycles per error. The baseline is the performance of recovery-only

with an 8 MC chip and no PDN pad failures. For the same reason noted in Sec. 5.2, noise amplitude only increases mildly despite the significant increase in voltage violation rate that results from a limited number of PDN pad failures. As a result, the recovery technique, which enforces a constant timing margin, suffers from more frequent rollbacks as pads fail. The hybrid technique, however, can avoid this penalty by dynamically adjusting the timing margin. Even with a relatively high recovery cost, the hybrid mechanism more effectively tolerates PDN pad failures, especially in wide-I/O chips, and in particular as technology scales.

Regardless of noise mitigation technique selection, we observe that for chips with abundant power supply pads (e.g., in the 8 MC case, where the total number of power supply pads is 1254), the performance overhead when $F = 60$ is fairly small. For chips with relatively lower power supply pad count (and thus higher I/O pad count), the performance overhead required to tolerate the same amount of PDN pad failures is higher. With the hybrid technique, a 24 MC chip can tolerate 60 pad failures with less than 1.5% performance overhead.

7.3. EM Lifetime Considering Pad Failure Tolerance

We have shown that a small amount of PDN pad failures can be tolerated with noise mitigation techniques; we therefore must adjust how MTTF is calculated to properly account for these failures (earlier MTTF values calculated the expected EM-failure-free lifetime). As mentioned before, the combinational space for allowing tens of pads to fail is enormous, making the derivation of an analytical solution for expected lifetime with PDN pad failure tolerance impractical. Fortunately, the times at which individual pads fail follow a known probability distribution. We have therefore used Monte Carlo Simulation to estimate MTTF under multiple PDN pad failures.

The bars in Fig. 10 show the normalized expected lifetime of different pad-failure tolerance levels F across different chip I/O configurations. If we do not allow any PDN pad failures ($F = 0$), increasing the memory controller count from 8 to 24 reduces EM lifetime by half with modest noise-mitigation performance overhead. However, by tolerating a small number of PDN pad failures, the whole chip’s expected lifetime is extended significantly. For example, if 40 pads are allowed to fail, increasing processor’s MC count from 8 to 24 would not hurt the expected system lifetime. According to our performance results, the penalty of tolerating 40 PDN pad failures with a 24 MC chip is negligible (1%). However, there are limits to how far this approach can extend lifetime. Going beyond 24 MCs and giving up more power pads to accommodate 32 MCs will place too much pressure on the rest of the power pads and even PDN pad failure tolerance cannot extend chip lifetime enough to match the baseline case. Thus we conclude that the power and I/O pad tradeoff is ultimately limited by C4 EM lifetime—to 24 MCs in the scenarios we evaluate.

8. Conclusions and Future Work

Power delivery quality is becoming a limiting factor in multi-core processor design. In this paper, we introduce VoltSpot, a pre-RTL, C4-aware PDN modeling tool, and combine it with other architecture-level tools to quantitatively evaluate voltage droop and electromigration (EM) lifetime changes subject to power/ground and I/O C4 pad tradeoffs. Our results suggest that with noise-aware pad allocation, replacing power-supply pads with I/O pads only mildly increases supply-voltage noise amplitude. By evaluating the performance of different run-time voltage noise mitigation schemes, we conclude that the penalty of mitigating extra noise caused by reducing power-supply pads is negligible, as long as run-time noise control is carefully designed. We also discover that, with the help of dynamic noise mitigation, chips can tolerate a small number of C4 pad EM failures with low performance overhead. Combined with a detailed EM lifetime study, we show that with a performance overhead of just 1%, chip I/O bandwidth can be tripled without sacrificing EM lifetime. More sensitivity and limit studies are available in a technical report [40].

VoltSpot enables a number of directions for future work. Combined with a thermal model, VoltSpot closes the loop for reliability research related to temperature, EM and transient voltage noise. Additional architecture-level studies can be supported with some extensions. For example, the recent industry trend of moving towards tighter in-package integration (e.g., stacked DRAM) alleviates constraints such as off-chip memory bandwidth. However, such integration along the third dimension exacerbates the challenge of power delivery, with increased current draw and inter-layer voltage noise propagation. VoltSpot can be easily extended to model a variety of 3D organizations, including microbumps and on-chip VRMs. Such extensions will be helpful for identifying bottlenecks in near-future processors' PDNs and to design and evaluate solutions to improve PDN reliability.

Acknowledgments

This work is supported by NSF grant CNS-0916908 and MCDA-0903471, and DARPA MTO under contract no.HR0011-13-C-0022.

References

- [1] C. Bienia, "Benchmarking modern multiprocessors," Ph.D. dissertation, Princeton University, 2011.
- [2] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill, and D. A. Wood, "The gem5 simulator," *SIGARCH Comput. Archit. News*, vol. 39, no. 2, 2011.
- [3] J. Black, "Electromigration: A brief survey and some recent results," *IEEE Trans. on Electron Devices*, vol. 16, no. 4, 1969.
- [4] W. Choi, E. Yeh, and K. Tu, "Mean-time-to-failure study of flip chip solder joints on Cu/Ni (V)/Al thin-film under-bump-metallization," *J. of Applied Physics*, vol. 94, no. 9, 2003.
- [5] I. J. Chung, *Modeling and Hybrid Simulation of On-chip Power Delivery Network Using an Unconditionally Stable Electromagnetic Field Solver*. ProQuest, 2007.
- [6] G. G. Faust, R. Zhang, K. Skadron, M. R. Stan, and B. H. Meyer, "ArchFP: Rapid prototyping of pre-rtl floorplans," in *VLSI-SoC*, 2012.
- [7] V. George, S. Jahagirdar, C. Tong, S. Ken, S. Damaraju, S. Scott, V. Naydenov, T. Khondker, S. Sarkar, and P. Singh, "Penryn: 45-nm next generation Intel core 2 processor," in *ASSCC*, 2007.
- [8] E. Grochowski, D. Ayers, and V. Tiwari, "Microarchitectural simulation and control of di/dt-induced power supply voltage variation," in *HPCA*, 2002.
- [9] M. S. Gupta, J. L. Oatley, R. Joseph, G. Y. Wei, and D. Brooks, "Understanding voltage variations in chip multiprocessors using a distributed power-delivery network," in *DATE*, 2007.
- [10] M. S. Gupta, K. K. Rangan, M. D. Smith, G. Y. Wei, and D. Brooks, "DeCoR: A delayed commit and rollback mechanism for handling inductive noise in processors," in *HPCA*, 2008.
- [11] M. S. Gupta, J. A. Rivers, P. Bose, G. Y. Wei, and D. Brooks, "Tribeca: design for PVT variations with local recovery and fine-grained adaptation," in *MICRO*, 2009.
- [12] J. W. Haskins Jr and K. Skadron, "Accelerated warmup for sampled microarchitecture simulation," *ACM TACO*, vol. 2, no. 1, 2005.
- [13] M. B. Healy, F. Mohamood, H. S. Lee, and S. Lim, "Integrated microarchitectural floorplanning and run-time controller for inductive noise mitigation," *ACM TODAES*, vol. 16, no. 4, 2011.
- [14] Intel, "Intel Pentium 4 processor in the 423 pin package / Intel 850 chipset platform," 2002.
- [15] —, "Intel Core i7-900 Desktop processor extreme edition series and Intel Core i7-900 desktop processor series on 32-nm process," 2011.
- [16] —, "Intel Xeon processor E7-8800/4800/2800 product families," 2011.
- [17] ITRS, 2011, <http://www.itrs.net>.
- [18] B. Jacob, S. Ng, and D. Wang, *Memory systems: cache, DRAM, disk*. Morgan Kaufmann, 2010.
- [19] R. Jakushokas and E. G. Friedman, "Multi-layer interdigitated power distribution networks," *IEEE TVLSI*, vol. 19, no. 5, 2011.
- [20] JEDEC, "Failure mechanisms and models for semiconductor devices," *JEDEC Publication JEP122G*, 2011.
- [21] R. Joseph, D. Brooks, and M. Martonosi, "Control techniques to eliminate voltage emergencies in high performance processors," in *HPCA*, 2003.
- [22] C. R. Lefurgy, A. J. Drake, M. S. Floyd, M. S. Allen-Ware, B. Brock, J. A. Tierno, and J. B. Carter, "Active management of timing guardband to save energy in POWER7," in *MICRO*, 2011.
- [23] S. Li, J. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *MICRO*, 2009.
- [24] X. S. Li, "An overview of SuperLU: Algorithms, implementation, and user interface," *ACM TOMS*, vol. 31, no. 3, 2005.
- [25] J. Lloyd, "On the log-normal distribution of electromigration lifetimes," *J. of Applied Physics*, vol. 50, no. 7, 1979.
- [26] A. V. Mezhiba and E. G. Friedman, "Electrical characteristics of multi-layer power distribution grids," in *ISCA*, 2003.
- [27] S. Nassif, "Power grid analysis benchmarks," in *ASP-DAC*, 2008.
- [28] C. Pei, R. Booth, H. Ho, N. Kusaba, X. Li, M. J. Brodsky, P. Parries, H. Shang, R. Divakaruni, and S. Iyer, "A novel, low-cost deep trench decoupling capacitor for high-performance, low-power bulk CMOS applications," in *ICSICT*, 2008.
- [29] M. Popovich, A. V. Mezhiba, and E. G. Friedman, *Power Distribution Networks with On-chip Decoupling Capacitors*. Springer, 2008.
- [30] M. D. Powell and T. Vijaykumar, "Exploiting resonant behavior to reduce inductive noise," in *ISCA*, 2004.
- [31] A. Ramachandran, "Methodologies for transient simulation of hybrid electromagnetic/circuit systems with multiple time scales," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2009.
- [32] M. Saint-Laurent and M. Swaminathan, "Impact of power-supply noise on timing in high-frequency microprocessors," *IEEE Trans. on Advanced Packaging*, vol. 27, no. 1, 2004.
- [33] B. Sinharoy, R. Kalla, W. Starke, H. Le, R. Cargnoni, J. Van Norstrand, B. Ronchetti, J. Stuecheli, J. Leenstra, G. Guthrie *et al.*, "IBM POWER7 multicore server processor," *IBM J. of Research & Development*, vol. 55, no. 3, 2011.
- [34] J. Tao, B. K. Liew, J. F. Chen, N. W. Cheung, and C. Hu, "Electromigration under time-varying current stress," *Microelectronics Reliability*, vol. 38, no. 3, 1998.
- [35] K. Wang, B. H. Meyer, R. Zhang, K. Skadron, and M. R. Stan, "Walking pads: Fast power-supply pad-placement optimization," in *ASP-DAC*, 2014.
- [36] N. H. Weste and D. M. Harris, *CMOS VLSI Design A Circuit and Systems Perspective*, 4th ed. Addison-Wesley, 2011.
- [37] S. Wright, R. Polastre, H. Gan, L. Buchwalter, R. Horton, P. Andry, E. Sprogis, C. Patel, C. Tsang, J. Knickerbocker, J. Lloyd, A. Sharma, and M. Sri-Jayantha, "Characterization of micro-bump C4 interconnects for Si-carrier SOP applications," in *ECTC*, 2006.
- [38] R. E. Wunderlich, T. F. Wenisch, B. Falsafi, and J. C. Hoe, "SMARTS: Accelerating microarchitecture simulation via rigorous statistical sampling," in *ISCA*, 2003.
- [39] R. Zhang, B. H. Meyer, W. Huang, K. Skadron, and M. R. Stan, "Some limits of power delivery in the multicore era," in *WEED*, June 2012.
- [40] R. Zhang, K. Wang, B. H. Meyer, M. R. Stan, and K. Skadron, "Architecture implications of pads as a scarce resource: Extended results," *University of Virginia, Tech. Rep. CS-2014-01*, May 2014.