Walking Pads: Managing C4 Placement for Transient Voltage Noise Minimization

Ke Wang[†], Brett H. Meyer[*], Runjie Zhang[†], Mircea R. Stan[‡], Kevin Skadron[†]

[†]Dept. of Computer Science
University of Virginia
Charlottesville, VA, USA

[*]Dept. of Elec. & Comp. Eng.
McGill University
Montréal, QC, Canada

[‡]Dept. of Elec. & Comp. Eng.
University of Virginia
Charlottesville, VA, USA

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

ABSTRACT

Transient voltage noise, including resistive and reactive noise, causes timing errors at runtime. We introduce a heuristic framework—Walking Pads—to minimize transient voltage violations by optimizing power supply pad placement. We show that the steady-state optimal design point differs from the transient optimum, and further noise reduction can be achieved with transient optimization. Our methodology significantly reduces voltage violations by balancing the average transient voltage noise of the four branches at each pad site. When we optimize pad placement using a representative stressmark, voltage violations are reduced 46-80% across 11 Parsec benchmarks with respect to the results from IR-drop-optimized pad placement. We also show that the allocation of on-chip decoupling capacitance significantly influences the optimal locations of pads.

Categories and Subject Descriptors

B.7.2 [**Design Aids**]: Layout, Placement and routing

General Terms

Algorithms, Design and Reliability

Keywords

Power pad allocation, Power distribution network, Voltage noise

1. INTRODUCTION

The exponential growth in on-chip power and current density due to CMOS scaling leads to an increasing challenge in delivering a stable voltage supply to processor functional blocks. When the supply voltage deviates too significantly from the nominal, timing errors can occur.

Voltage supply noise occurs due to the resistance and inductance of the components of the power delivery network (PDN), such as the package, the controlled-collapse-chip-connection (C4) pads that connect the package to the PDN, and the PDN wires themselves. Recent studies show that transient inductive noise, proportional to LdI/dt, is expected to represent a larger proportion of total noise in future process technologies [1, 12].

In this paper, we explore the effect of C4 power supply pad allocation on transient voltage noise. As the interface between the package and the chip, C4 pads play a crucial role in determining the impedance of the whole power delivery system. C4 pads are used for both power delivery and I/O; while allocating additional C4 pads for power delivery can minimize voltage noise, doing so may reduce I/O bandwidth. Optimizing power pad count and placement, beyond improving reliability or performance by reducing voltage noise, also exposes opportunities to increase I/O bandwidth, a critical bottleneck in modern SoC design.

Prior work has targeted resistive voltage noise (IR drop) and optimized pad location and number to minimize worst-case IR drop [13, 14, 16, 18, 19]. All these state-of-the-art pad placement techniques focus on steady-state analysis and VDD pads only. While previous work [16] suggests reducing transient noise with IR-based-optimization, we observe that such optimization has limited benefit.

Pad placement optimization for transient noise mitigation is characterized by an enormous design space combined with costly design evaluation. First, computationally complex architectural modeling is required to derive the voltage violations needed to determine optimal pad location. VoltSpot [17], for instance, calculates grid node voltage at a sub-cycle granularity in order to achieve the fidelity needed for accurate, transient PDN behavior modeling. Second, the combinatorial design space of pad placement is huge for modern chips, consisting of over 1,000 candidate pad locations, of which typically 50% or more are used for power delivery. For example, the search space for the case study in this paper—a 16nm, 16-core Intel Penryn-like multiprocessor—is larger than 10^{489}.

To make transient noise mitigation tractable, we extend Walking Pads (WP), a heuristic optimization framework for fast IR-drop-optimized power pad placement [14]. WP converts the global pad placement optimization problem into a local virtual-force balance problem allowing simultaneous movement of all pads, reducing algorithm complexity significantly over the simulated annealing (SA) and mixed integer linear program (MILP) approaches in the literature [18, 19]. The computational efficiency of WP makes VDD and GND pad placement optimization for transient violation suppression feasible. To ensure WP selects a pad placement suitable across many benchmarks, and thus, many PDN behaviors, we propose optimizing placement using a benchmark representative of worst-case transient noise, a "stressmark."

DAC'14, June 01–05, 2014, San Francisco, CA, USA.

http://dx.doi.org/10.1145/2593069.2593243.

Our paper makes four principal contributions:

1. We extend Walking Pads to be capable of transient voltage violation suppression.
2. We show that the steady-state optimal pad location differs from the transient optimal, and that greater noise reduction can be achieved with the latter.
3. We observe that transient optimization against a voltage violation stressmark reduces the violation counts of whole benchmarks.
4. We show that the on-chip distribution of decoupling capacitance (decap) has a significant effect on pad placement optimization for transient voltage noise.

2. RELATED WORK

The state-of-the-art voltage-noise-minimizing pad allocation techniques are all based on IR drop analysis. Successive Pad Assignment (SPA) [13] and a mixed integer linear program (MILP) technique [18] were proposed for power pad assignment in a pad ring structure. Both of these two methods suffer from scalability issues, and are not tractable for large scale 2D C4 arrays. Zhong and Wong proposed a simulated-annealing-based, fast power-pad placement optimization algorithm [19]. Yu and Wong designed a more efficient pad allocation algorithm to reduce global IR drop [16]. Our earlier work on Walking Pads efficiently optimized pad location by leveraging virtual-force-directed, simultaneous, pad movement. Other work minimizes statistical timing margin based on steady-state analysis [6]. However, these approaches ignore reactive circuit elements in the package, power-pad impedance, on-chip inductance and decoupling capacitors (decap), making them inappropriate for suppressing transient noise. In this work, we model PDN with all resistive noise resources considered by previous works, as well as the more detailed structures necessary to capture transient voltage drop, including package RLC, pad impedance, on-chip inductance and decap. Based on this PDN model, we extend WP to consider transient noise and propose a stressmark-based method to optimize pad placement such that transient noise is suppressed.

3. PROBLEM FORMULATION

The objective of power supply pad location optimization is, given a) an architectural-block-level system floorplan, b) a cycle-by-cycle system power trace for each architectural block, c) the number of power supply pads to place, and d) the voltage violation threshold (any cycle with a larger voltage drop is considered to be a violation cycle), to identify the set of candidate pad locations at which to place power supply pads in order to minimize the number of transient violation cycles.

3.1 Power Delivery Network Modeling

To capture the transient noise at architectural blocks in the PDN system, we utilize VoltSpot, an architectural level PDN model [17]. As shown in Fig. 1a, this model includes two major parts: a model of the on-chip power delivery grid and a model of package. The VDD and GND nets are modeled as separate regular 2D circuit meshes. To accurately model the number of C4 pads and their locations, we use a 4-to-1 grid-node-to-pad ratio (i.e., using a 100×100 grid to model a PDN with 50×50 pads) [17]. We use multiple RL pairs to model multiple on-chip metal layers [11]. C4

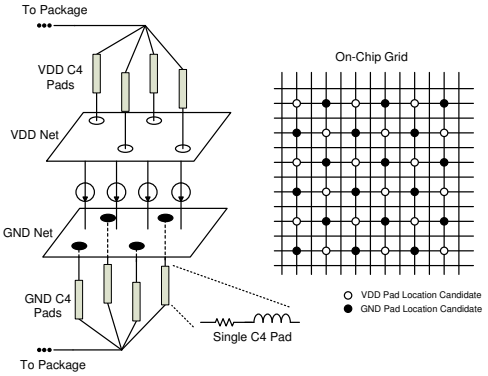

(a) On-chip power delivery network model.

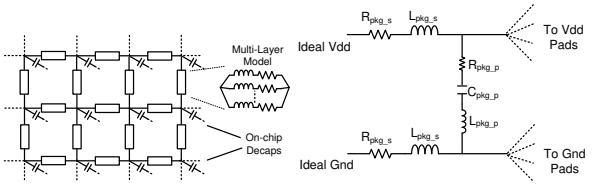

(b) On-chip grid model (c) Package model

Figure 1: **PDN model structure [17]**

pads are modeled as individual resistor-inductor branches attached to on-chip grid nodes. On-chip decoupling capacitors connect to each VDD and GND grid node. Ideal current sources are used to model the load (i.e., the power of the switching transistors and associated leakage). We use a lumped RLC model for off-chip components such as the package or PCB board. Our WP algorithms are compatible with detailed package models—but support for such models is the subject of future work. The VoltSpot PDN model has been validated with IBM transient power grid benchmarks with a maximum error of 0.8%VDD [17].

We employ the implicit trapezoidal numerical method to solve such large-scale circuits efficiently and accurately [4, 17]. We set our time step to one fifth of a cycle at 3.7GHz to keep the numerical error of node voltage below 10^{-5}V.

3.2 Power Supply Pad Placement

While steady-state voltage noise optimization can focus on VDD pads [6, 16, 19], transient voltage violation minimization requires the joint modeling/optimization of VDD and GND pads to correctly simulate the package-pad-mesh-pad-package circuit. To simplify the problem, we divide the pad grid into "white" and "black" (like a checkerboard) for VDD and GND pad candidate positions respectively, as illustrated in Fig. 1a. Each possible allocation of VDD and GND pads to grid locations is called a *configuration*. The total number of configurations is larger than 10^{489} in the case studies we explore. We note that checkerboard VDD-GND pad layout is not a limitation of the Walking Pads optimization framework. WP works as long as VDD and GND pad candidate positions are uniformly distributed.

4. WALKING PADS

WP is based on the hypothesis that, when all pads are at their locally noise-balanced positions, with equal voltage noise along wires in two opposite directions from each pad, global voltage noise will be minimized. In other words, if we move a noise-balanced pad in any direction, it will

cause more noise in the opposite direction. The total current and total dI/dt are both invariant in any given cycle, independent of pad placement: balancing noise by moving the pad toward higher noise will trade less noise in one area for more noise in another. However, if this balancing reduces the magnitude of the noise events such that they more often fall below the provisioned threshold, without pushing events elsewhere above the threshold, then the total noise is greatly decreased. Consequently, we designed the WP framework to convert a global optimization problem, the placement of n VDD and n GND pads given $m \geq 2n$ candidate locations, into a local balance problem, the placement of individual pads with respect to various nearby voltage noise demands. To determine how to place pads, we have adopted a virtual-force-directed pad movement strategy; this virtual force is defined by a vector in the direction of the combined voltage noise observed by a given pad.

4.1 Algorithm Framework

Walking Pads adopts a general three-step algorithm framework to incrementally move all pads toward their noise-balanced positions: (1) determine per-grid-node voltage levels for each cycle; (2) calculate virtual forces and determine the direction and distance of movement for each pad based on the virtual force; and, (3) move all pads. In step 1, transient PDN simulation is performed by simulating with power traces derived from activity factors calculated by a performance simulator. In step 2, for each pad, the transient voltage gradient along each wire from that pad (the rate of voltage change along the wire from the pad to an adjacent grid node) is averaged over all simulated cycles. The virtual force is then calculated by vector summation of the transient voltage gradient across the four wires connected to each pad. Noise in opposite directions in the grid from the pad cancel each other; when local noise is balanced, this vector is $\vec{0}$.

When averaging the per-wire voltage gradient, we tried three different strategies: median, mean, and root mean square (RMS) of the voltage gradient. We observe that the RMS results in the most effective noise minimization among the three approaches, because RMS can also capture the AC component of voltage noise. We adopt RMS in this paper.

Our method achieves a significant performance improvement over other methods like SA by (a) employing a deterministic approach to pad movement in step 2, and (b) allowing all pads to move simultaneously in step 3. As all pads move toward their estimated balanced positions in each iteration, fewer iterations are needed than in approaches that move fewer pads at once.

4.2 Practical and Deep Optimization

We propose two algorithm variants—"practical" optimization (PO) and "deep" optimization (DO)—for different trade-offs between optimization time and result quality.

Practical optimization, which trades lower results quality for faster optimization, is shown in Algorithm 1. PO moves pads in the direction defined by the normalized virtual force $\vec{F}/\left\|\vec{F}\right\|$. $\vec{F}$ is a function of G_{dir}, the average voltage noise in direction *dir*. To force pads to stop at approximately balanced positions, we introduce a "freezing" process, which gradually decreases the move distance D_i of each pad at a constant freezing rate γ. PO terminates when pads no longer move. Since D_i is a real number, pads move to the nearest unoccupied candidate position of the same type. The initial large-step stage of PO helps pads to jump out of local minima, while the later small-step stage helps pads gradually converge into balanced positions. One round of optimization of either GND pads or VDD pads is called an iteration, as each requires the most time-consuming part of the algorithm: solving the transient voltage equation. Pads from the GND and VDD grids move in alternating iterations. VDD and GND pads can be moved simultaneously, halving the required iterations, but resulting in up to 3% more violations in our case studies. *is_better()* checks if the current pad placement improves upon all previous placements. We use three criteria to compare different pad placements. First, the placement with the fewest violation cycles is the best (better global violation-count optimization). In the event of a tie, the placement with the fewest violation cycles in the grid node with the maximum violation cycles is the best (better local violation-count optimization). In the event of a tie, the placement with the lowest maximum noise amplitude is the best (better global violation-amplitude optimization). Using these criteria, the best placement is returned to the designer after PO converges.

```
Set: initial move distance D_0, freezing rate γ
repeat
    foreach grid ∈ {GND, VDD} do
        Solve transient equation;
        foreach pad do
            F⃗ = (G_north − G_south, G_east − G_west);
            D⃗isp = F⃗/‖F⃗‖ ∗ D_i;
            Move;
        end
        D_{i+1} = D_i ∗ γ
        is_better();
    end
until is_converged() == True;
```

Algorithm 1: Practical Optimization algorithm.

All initial pad locations for PO are generated randomly, and the initial location has little effect on the resulting quality of PO. Using criteria described in the literature [14], we choose an initial move distance $D_0 = 3\sqrt{2} * pad\_pitch$ and freezing rate $\gamma = 0.99$ for our case studies. Automatic or adaptive parameter selection is the subject of future work.

Deep optimization, which sacrifices performance for better violation suppression, is shown in Algorithm 2. Simultaneous pad movements (as in PO) reduce the quality of the solution to some extent because the forces on one pad may change when other pads move. To address this, DO performs a local greedy search: it moves one pad in each iteration and only accepts movements that improve solution quality. We hypothesize that moving the pads nearest to the grid node with the highest violation count and amplitude (the *hot spot*) is the best way to quickly improve solution quality. DO sorts the VDD and GND pads by their distances to the *hot spot* and lets nearby pads (whether VDD or GND) move first. When the *hot spot* moves, DO re-sorts the pads and continues. DO terminates when no pad movement improves design quality as measured by *is_better()*. To reduce search time, DO is only used after PO (noted PO+DO) when high-quality optimization is required.

5. EXPERIMENTAL SETUP

To evaluate our pad placement optimization approach, we conducted a series of simulations. First, we trained

```
repeat
    Sort power pads by distance to noise hot spot → PadList;
    foreach pad ∈ PadList do
        F⃗ = (G_north − G_south, G_east − G_west);
        Move pad to the neighboring candidate position in
        the direction of F⃗/‖F⃗‖;
        Solve transient equation, identify hot spot;
        if is_better() then
            accept the movement; break;
        else
            undo the movement;
        end
    end
until is_converged() == True;
```

Algorithm 2: Deep Optimization algorithm.

the pad configuration using a variety of power trace segments, including a stressmark we have derived for the purpose. Next, we evaluated the pad configurations using the same trace segments as well as entire benchmarks, measuring the number of transient voltage emergencies that occur. In each case, we compare our transient-based optimization techniques with pad placements optimized to minimize steady-state IR drop.

To optimize pad location against worst IR-drop, we adopt the simulated annealing framework described in [19] and extend it to jointly optimize both VDD and GND pad locations. We use a very slow cooling speed—0.98 [14]—to get high-quality results for IR-drop optimization within a reasonable time. All our experiments were conducted on an Intel Xeon E5-1650 3.20GHz CPU with 32GB memory.

5.1 Model Parameters

We select a 16-core, Intel Penryn-like multiprocessor scaled to 16nm technology as our evaluation platform. We assume a supply voltage of 0.7V; architectural floorplans were generated using an architecture-level floorplanner, ArchFP [5]. The area of each functional block is calculated with McPAT [10]. Because of the size of C4 bumps, the rows and columns of the C4 pad array maintain a minimum spacing, or pad pitch, of $285\mu m$, based on ITRS projections for the pad density [7]. This results in a grid with 1914 pad candidate locations for our 16-core system. Other PDN model parameters are adopted from the literature [17].

5.2 Power Trace Generation and Sampling

To get chip-wide, application-specific power consumption data, we used Gem5 [3], a multi-core performance simulator, and McPAT. During benchmark execution, Gem5 generates activity factors at the granularity of architectural blocks, which McPAT uses to derive cycle-level power traces for each block. To model steady-state worst case for IR drop based optimization, we assume that each architectural unit dissipates 85% of its maximum power [8, 14].

To accelerate transient noise simulation, we borrow the idea of statistical sampling from processor performance research [15]. We take 1000 samples at equal intervals from the end-to-end execution of sim-medium inputs for 11 benchmarks in the Parsec 2.0 benchmark suite (the remaining two benchmarks are currently incompatible with our performance simulation infrastructure) [2]. Each sample (power trace segment) contains 2000 cycles of per-cycle power information; the first 1000 cycles in each sample are used to "warm-up" the reactive elements in the PDN. These benchmarks are diverse enough to be considered representative of general-purpose workloads [2].

To train and evaluate our pad placements under worst-case behavior, we extracted power traces for each benchmark from 2-core simulations and then made eight copies of the 2-core power trace simulation on the 16-cores platform. In this way, transient current fluctuation representative of that in Parsec benchmarks occurs simultaneously in each pair of cores, increasing the pressure on the PDN. To identify voltage violation, we assume a 5% voltage drop threshold.

6. RESULTS

6.1 Comparison Techniques

We compare our algorithms with steady-state IR-drop-optimizing simulated annealing (SA), because transient optimization using SA is impractical (too many costly iterations). To evaluate the feasibility of SA-based transient optimization, we used noise violations observed during one sample trace (the sample with highest violation amplitude across all benchmarks, denoted "stressmark") to place 267 VDD and 267 GND pads over a 16-core model PDN described above. One PO/DO iteration or one SA trial movement causes one round of PDN transient simulation, requiring one minute of computation. PDN transient simulation is the most time consuming operation in pad placement optimization, requiring over 99.99% of the total optimization time across all evaluated methods.

We extended the IR-targeted SA technique [19] for transient noise optimization by adopting a cooling rate of 0.90 and using violation count as the objective function; it executed 13884 iterations, requiring 9.6 *days*, to reduce the violation count from 522 (randomly initialized pad location) to 305. Our PO approach, on the other hand, required only 178 iterations, 178 minutes, to reach the violation count of 286 from the same initial pad placement. PO+DO reached 279 violations in 284 minutes. To achieve the same quality as WP methods, SA needs to further slow its cooling speed, requiring significantly more computation (estimated at 5× or 48 days [14]). Increases in pad count exacerbate this performance penalty. Consequently, all remaining comparisons are made against the more practical IR-drop-based SA.

6.2 Transient- vs. IR-Optimized Placement

Our first simulations compare the effectiveness of IR-drop-optimizing pad placement method and our Walking Pads transient optimization method on transient noise control. In these simulations we assume a uniform distribution of on-chip decoupling capacitance. We randomly selected four samples among all the samples from the 11 Parsec benchmarks: A, B, C, D. We also identified the sample with highest violation amplitude across all benchmarks, and denote this sample the stressmark S.

The violation cycle counts are 26, 97, 159, 271, and 326, for these five samples when 267 VDD and 267 GND pads are optimized for IR-drop. Figure 2 shows reduction in violation cycle count achieved by WP across the five samples with respect to the IR-drop-optimized placement. "PO on itself" indicates the effectiveness of training and evaluating PO using the sample in question. "PO on stressmark" indicates the effectiveness of training PO using S and then evaluating PO using the sample in question. PO can reduce violations by up to 80% after 178 iterations.

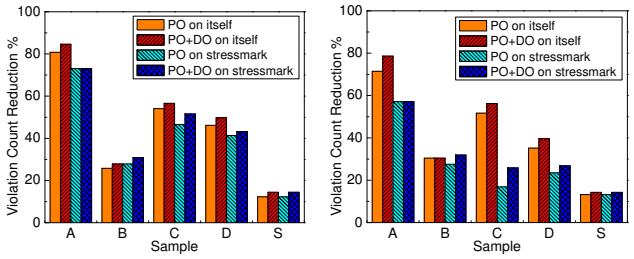

(a) 267 VDD + 267 GND (b) 387 VDD + 387 GND

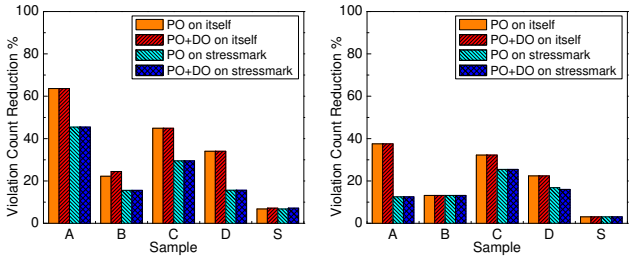

(c) 507 VDD + 507 GND (d) 627 VDD + 627 GND

Figure 2: **Effect of transient optimization on violation count. "on itself" means: optimize pads and evaluate violation count on this sample. "on stressmark" means: optimize pads on stressmark and evaluate violation count on this sample.**

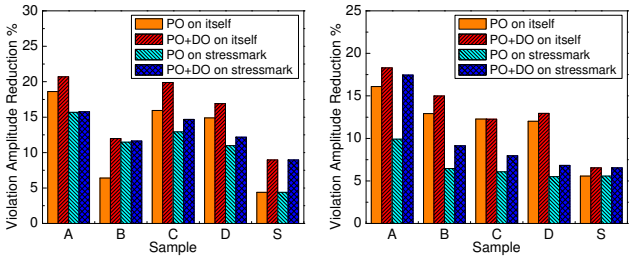

(a) 267 VDD + 267 GND (b) 387 VDD + 387 GND

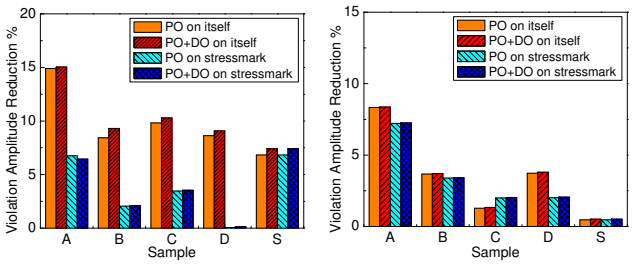

(c) 507 VDD + 507 GND (d) 627 VDD + 627 GND

Figure 3: **Effect of transient optimization on maximum violation amplitude.**

Under PO+DO, DO is terminated after $0.2 \cdot N$ iterations where N is the total number of power supply pads allocated, both VDD and GND. These additional iterations reduce violations up to 7% when compared with PO results.

Figure 3 shows the reduction in the maximum violation noise amplitude achieved by WP with respect to the IR-drop-optimized placement. PO and PO+DO reduce the maximum violation amplitude by up to 18.6% and 20.7%.

Optimizing pad placement for transient noise clearly reduces both the number of voltage emergencies and their magnitude, but it does so at the expense of steady-state IR drop: transient-optimized pad placement results in an increase of up to 1.3% of the worst IR drop on the IR-drop stressmark compared with IR-drop-optimized pad placement in our study cases. However, this is ultimately of little consequence: steady-state IR-drop is clearly a poor proxy for transient noise, the mitigation of which requires direct optimization. Note that transient noise combines the effects of all sources of noise, including IR drop.

We observe in Figures 2 and 3 that optimizing pad placement by training with the stressmark S significantly reduces the violation count and amplitude not only for S, but also for other samples. This indicates the feasibility of training on one representative sample to reduce noise violations across a variety of samples.

6.3 Stressmark-based Optimization

We next explore the effect that training on such a stressmark has when evaluating the noise across complete benchmarks. Figure 4 shows the resulting violation count reduction for the 11 benchmarks in the Parsec suite. A 46% (*fluidanimate*) to 80% (*rtview*) reduction in violation count is achieved when pad placement is optimized based on our transient stressmark, compared with the results from IR-optimized placement. Clearly, transient pad optimization using a representative sample can help to reduce the violation count of a whole benchmark significantly. We further observe that PO+DO always performs better than PO, by 3% on average.

Table 1 details our optimization results for the benchmark *Blackscholes*. For each combination of pad count (Pads) and optimization algorithm (Alg.), Table 1 reports the number of violations (Vio. Count), average violation amplitude (Ave. Amp. %), and maximum violation amplitutde (Max Amp. %). We observe that our transient-based optimization reduces the violation count significantly for all tested pad counts. Furthermore, our transient-based optimization algorithms also reduce the average and maximum violation amplitude, which are used to perform voltage margin assignment for some runtime control techniques [9, 17].

6.4 Decap Effect on Optimal Pad Location

To show the effect of decoupling capacitance allocation on optimal pad location, we performed further simulations using a non-uniform decap allocation. As the ALU of each core is a transient noise hotspot, we divide the total on-chip decap by 16 and place 1/16th of the decap as close to the ALU of each of the 16 cores as possible. We then compare IR-drop and stressmark-based transient pad optimization. Figure 5 shows the reduction in violation counts in the stressmark produced by transient optimization for (a) uniform decap (Uni Decap) distribution, (b) non-uniform decap (Non-uni Decap), and (c) pad placement assuming a non-uniform distribution but applied to a scenario with uniform

Table 1: **Optimization Results for *Blackscholes***

Pads	Alg.	Vio. Count	Ave. Amp. %	Max Amp. %
267	IR	31234	6.404	8.961
267	PO+DO	9643	5.562	7.644
387	IR	11807	5.770	8.435
387	PO+DO	5160	5.221	7.620
507	IR	6291	5.406	7.982
507	PO+DO	4143	5.125	7.673
627	IR	4373	5.184	7.890
627	PO+DO	2940	4.925	7.580

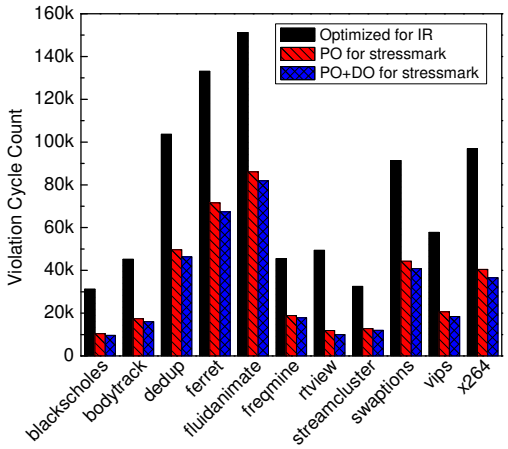

Figure 4: **Violation-cycle reduction by transient optimization. The pad count is 267 for VDD/GND. The pad locations are optimized on stressmark. The violation cycles for each benchmark are counted over all samples (1M cycles in total).**

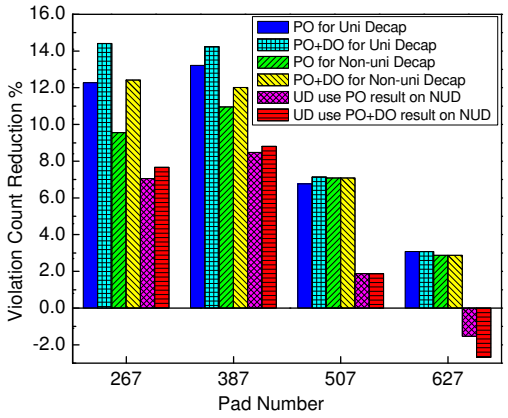

Figure 5: **Decap effect on optimal pad locations. The stressmark sample is used for this figure. "PO/PO+DO for Uni/Non-uni Decap" means we optimize and evaluate the results on the same decap configuration. "UD use PO/PO+DO on NUD" means we optimize pad location with non-uniform decap and evaluate the results with uniform decap.**

distribution. We observe that the distribution of decoupling capacitance affects pad placement: when pad placement is optimized for one case (e.g., non-uniform distribution) but then applied to another (e.g., uniform distribution), the reduction in violations shrinks significantly. The data for 627 pads show that violations may get worse under deep optimization on another decap configuration. In some cases, a change of decap allocation may even make the transient-optimized pad placement worse than the IR-optimized pad placement; joint pad placement and decap allocation optimization is clearly needed.

7. CONCLUSIONS

In this paper, we present a power supply pad optimization framework—Walking Pads (WP)—that suppresses transient voltage noise by placing pads such that the average voltage gradient observed along each wire from the pad is balanced. Converting the global pad-placement problem into a local balance problem, WP allows all pads to move simultaneously, reducing algorithm complexity substantially compared with simulated annealing implementations, which prove to be impractical in this context. Our results clearly demonstrate that when transient pad placement optimization is performed, the resulting noise reduction is far superior to that achieved with state-to-the-art IR-drop-optimizing pad placement techniques. When we optimize pad placement using a representative stressmark, voltage violations are reduced 46-80% across 11 Parsec benchmarks. Combining WP and stressmark selection makes pad placement optimization for transient noise suppression tractable. We also observe that the distribution of decoupling capacitance influences the placement of pads. Consequently, we hypothesize that decap allocation and pad placement ought to be jointly optimized; this is a subject for future work.

8. ACKNOWLEDGMENTS

This work was supported in part by NSF grant CCF 1255907 and C-FAR, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

9. REFERENCES

[1] International technology roadmap for semiconductors, 2011.
[2] C. Bienia et al. The PARSEC benchmark suite: Characterization and architectural implications. In *PACT*, Oct. 2008.
[3] N. Binkert et al. The gem5 simulator. *SIGARCH Comput. Archit. News*, Aug. 2011.
[4] J. Chung. *Modeling and Hybrid Simulation of On-chip Power Delivery Network Using an Unconditionally Stable Electromagnetic Field Solver*. PhD thesis, UIUC, 2007.
[5] G. G. Faust et al. ArchFP: rapid prototyping of pre-RTL floorplans. In *VLSI-SoC*, Oct. 2012.
[6] K. Haghdad and M. Anis. Power supply pads assignment for maximum timing yield. *IEEE Trans. Circuits Syst. II, Exp. Briefs2*, 58(10):697–701, 2011.
[7] ITRS. 2011. http://www.itrs.net.
[8] A. Joshi et al. Automated microprocessor stressmark generation. In *HPCA*, Feb. 2008.
[9] C. R. Lefurgy et al. Active management of timing guardband to save energy in POWER7. In *MICRO*, Dec. 2011.
[10] S. Li et al. McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures. In *MICRO*, Dec. 2009.
[11] A. V. Mezhiba and E. G. Friedman. Electrical characteristics of multi-layer power distribution grids. In *ISCAS*, May 2003.
[12] M. Popovich et al. *Power distribution networks with on-chip decoupling capacitors*. Springer, New York, 2008.
[13] T. Sato, H. Onodera, and M. Hashimoto. Successive pad assignment algorithm to optimize number and location of power supply pad using incremental matrix inversion. In *ASP-DAC*, Jan. 2005.
[14] K. Wang et al. Walking pads: Fast power-supply pad-placement optimization. In *ASP-DAC*, Jan. 2014.
[15] R. E. Wunderlich et al. SMARTS: accelerating microarchitecture simulation via rigorous statistical sampling. In *ISCA*, Jun. 2003.
[16] T. Yu and M. Wong. A novel and efficient method for power pad placement optimization. In *ISQED*, Mar. 2013.
[17] R. Zhang et al. Architecture Implications of Pads as a Scarce Resource. In *ISCA*, Jun. 2014.
[18] M. Zhao et al. Optimal placement of power supply pads and pins. In *DAC*, Jun. 2004.
[19] Y. Zhong and M. D. F. Wong. Fast placement optimization of power supply pads. In *ASP-DAC*, Jan. 2007.