

Some Limits of Power Delivery in the Multicore Era

Runjie Zhang
University of Virginia
Charlottesville, VA, USA
Runjie@virginia.edu

Brett H. Meyer
McGill University
Montréal, Québec, Canada
brett.meyer@mcgill.ca

Wei Huang
IBM Austin Research
Laboratory
Austin, TX, USA
huangwe@us.ibm.com

Kevin Skadron
University of Virginia
Charlottesville, VA, USA
skadron@cs.virginia.edu

Mircea R. Stan
University of Virginia
Charlottesville, VA, USA
mircea@virginia.edu

ABSTRACT

The ability to scale down threshold and hence supply voltages can no longer keep up with device density as technology scales. Microprocessor power density is therefore increasing. At the same time, the total number of C4s is predicted to be constant for the foreseeable future, according to ITRS 2011. As a result, more and more of the C4 pads are dedicated to power delivery, at the expense of off-chip I/O signals, impeding I/O throughput scaling—even though core counts and hence bandwidth requirements are increasing exponentially. It therefore becomes important to consider the *power delivery network* (PDN) as early as possible in the design process, both to ensure enough I/O pads and because a later redesign due to power delivery issues is costly. In this paper, we propose and validate a steady-state *architecture-level* PDN model, called VoltSpot, and explore the impact of the power delivery constraint for future technology nodes. Our results, based on a scaled multicore processor, indicate that worst-case on-chip IR drop at 16nm will be at least three times larger than that at 45nm. We propose a first-order optimization algorithm to derive the number and placement of C4 pads for by power delivery to achieve a specific IR-drop target. When optimizing to satisfy an IR-drop constraint of 5%, power delivery requires so many pads that multicore processors at 16nm will not be able to maintain constant per-core I/O bandwidth.

1. INTRODUCTION

In future CMOS technology nodes, threshold and supply voltages are not scaling down as fast as device density is increasing. Even if power supply or cooling limits limit total chip power, localized power densities will still increase. Continuing reductions in voltage, although slowing, further increase local current density, because current density is power density divided by the scaled voltage. Higher current density and total current place greater demands on the power delivery network (PDN); current-related chip phenomena such as electromigration (EM), resistive current (IR) drop, and inductive transient current (Ldi/dt) noise all get worse with higher current and larger current swings.

Electromigration refers to the gradual migration of ions in metal conductors due to high density current flow. EM happens mostly in the PDN where the current flow tends to be *uni-directional*, exacerbating EM effects. EM can cause open or short circuits in metal wires and eventually failure

of the entire chip.

IR drop comes from the resistivity of PDN wires, pads and pins, and describes the *voltage droop* from the power supply to the circuits in silicon, as well as the *ground bounce* from silicon to true ground. Large IR drops reduce the available circuit voltage headroom, hence increasing circuit delay and degrading circuit performance. It can also lead to timing errors if the IR drop exceeds the worst case design specifications.

The Ldi/dt effect is a dynamic noise effect and is caused by large and fast current swings in the intrinsic inductances of the PDN. In this paper, we focus on IR drop and electromigration and leave the extension to transient Ldi/dt as future work.

One major challenge in designing a PDN that scales well as current increases is the slow scaling of resources such as on-chip C4 pads. As a matter of fact, the total number of C4 pads for a fixed-area processor is predicted to remain constant for the foreseeable future, according to the latest ITRS roadmap [9]. In addition, C4 pads are not only used for power delivery, but also for I/O signals. Obviously, delivering higher current through a constant number of C4 pads creates significant design challenges. In order to better address these challenges, it is important to analyze power delivery trends for future technology nodes and take PDN issues into consideration early in the design process, *e.g.*, at the architecture level.

Among all available C4 pads, some are dedicated to the PDN, but others must be dedicated to off-chip I/O signals to communicate with memory and other chips. The limitation of available C4 pads creates an important tradeoff between I/O bandwidth and power delivery quality. However, it is impractical to explore this tradeoff space with high-resolution, post-RTL PDN simulation, because PDNs in modern microprocessors usually contain millions of nodes and take a significant amount of time to simulate, let alone the physical design turnaround time and cost if any changes are made. For these reasons, it is preferable to have an architecture-level pre-RTL PDN model to allocate and place on-chip resources to jointly mitigate issues thermal, reliability, power delivery, and I/O bandwidth constraints.

Our main contributions in this paper are as follows:

- We propose VoltSpot, an architecture-level model of the on-chip PDN, including C4 pads, with a simple interface for use in other architecture-level tools. Only

high-level parameters such as chip size and metal pitch (given by ITRS or process-specific design rules) are required from users. We validate our model against an IBM power grid benchmark and find that it models pad current with less than 4% error on average.

- To the best of our knowledge, we are the first to study the tradeoff between signal I/O pads and power pads using architecture-level modeling of the power delivery network and resulting pad requirements.
- We also present a scaling analysis down to 16nm, investigating IR drop and the number of available I/O pads under IR-drop constraints. We observe that IR drop more than triples from 45nm to 16nm, becoming a more severe constraint than electromigration due to pad current. Under an assumption of 5% IR-drop tolerance, there will not be enough pads for I/O signals to keep per core bandwidth constant at 16nm.

The paper is organized as follows: Section 2 describes our PDN modeling methodology and Section 3 presents our validation results. The scaling scenario and other simulations are described in Section 4. Section 5 discusses our results and Section 6 reviews prior research on PDN modeling. Section 7 gives a summary and a discussion of future work of the paper.

2. ARCHITECTURE-LEVEL PDN MODELING METHODOLOGY

The power delivery system for modern microprocessors consists of voltage regulators, connectors and metal traces on PCB, loadline resistance, chip package and on-chip metal layers. The on-chip PDN starts at the power and ground C4 pads, and usually spans multiple layers of parallel metal wires. Within these layers, interleaved power and ground supply lines provide the required current to the chip. Depending on the design requirement, the on-chip PDN may consist of a single global power grid, or a coarser global grid to which local power grids connect (for power gating or design modularity). Regardless of their hierarchy, on-chip PDNs were designed to keep the on-chip voltage as *spatially uniform* and *temporally steady* as possible. C4 pads, a 2-D array of solder balls distributed between the silicon die and the package substrate, solder on-chip metal wires and the electrical package together and serve as both signal I/O channels and aqueducts for current. In this paper, we assume a single global on-chip PDN which consists of only one VDD grid and one GND grid.

We use a compact model of the on-chip PDN’s physical structure, and only require that the user specify (a) top-layer metal pitch and cross-section area, (b) chip dimensions, (c) V_{DD} /Ground C4 pad locations, (d) chip floorplan, and (e) chip power map. Given these inputs, VoltSpot solves for the voltage and current at each V_{DD} /Ground C4 pad and internal node in the resulting on-chip power delivery network.

The regularity of the on-chip PDN’s physical structure makes compact PDN modeling feasible. A well accepted methodology models the multi-layer V_{DD} and ground nets as separate regular 2-D circuit meshes [2, 7, 8]. Under steady-state assumptions, both meshes contain only resistors. C4 pads are modeled as individual resistors attached to on-chip

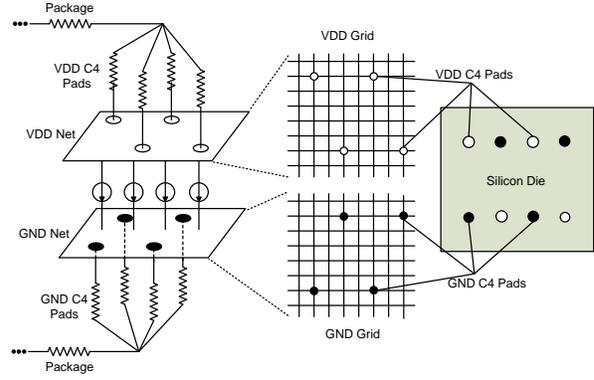


Figure 1: On-chip PDN model

grid nodes and the relative locations of those connection points in the grid represent the actual locations of the C4 pads on the silicon die. Ideal current sources are used to model the load (*i.e.* the switching transistors). Finally, off-chip components like the package or PCB board are lumped into single resistors. We adopt this methodology and build the model skeleton as in Figure 1. Since our main focus is on-chip PDN, we assume the PCB board represents an ideal power supply, and therefore the only off-chip parts in our implementation are the lumped package resistors.

Both grid size and grid resistance are determined by the shape and number of the power/ground lines in the top two metal layers. These parameters are independent of the chip floorplan. VoltSpot automatically calculates PDN grid parameters based on the top-layer metal pitch and chip dimensions. For example, the numbers of columns and rows in the grid equal to the number of longitude and latitude wires in top two layers. These numbers are derived by dividing chip width/height by metal pitch. The grid resistance was calculated by the metal resistance equation: $R = \rho * l/A$, where A is metal’s cross-sectional area and l is length. This feature makes modeling chips of arbitrary size easy and sets architects free from the electrical engineering details. We validate the accuracy of this method in Section 3.

One major novelty of our model is that we can expose C4 pads as an architectural resource, and expose power delivery as an architectural constraint. This allows architects to explore the tradeoff between chip I/O bandwidth and power delivery quality, for example, and better evaluate the benefits of various architectural choices that might affect power delivery (such as placement of high-power-density units) or I/O bandwidth (such as data compression or novel I/O signaling technologies). Designers are able to specify the number of pads as well as their locations via a simple interface and VoltSpot maps those pads onto the PDN grid. During the mapping process, the tool takes pad size and pitch into consideration and aligns the positions of pads in the grid so that all the pads’ locations fit in regular array. VoltSpot also provides an extensible framework for implementing pad optimization algorithms.

As an architecture-level tool, VoltSpot also takes a processor floorplan and power map as inputs. This is helpful to study the spatial variation of voltage or current within one die. To achieve that, we divide our PDN grid into blocks according to the processor’s floorplan and assign power consumption values at the granularity of functional blocks. Since

switching silicon is represented by ideal current sources between the power plane and the ground plane, we assign uniform values to current sources within each function block. According to the equation $Power = Voltage \times Current$, we divide power by supply voltage to get current source values. VoltSpot—like other pre-RTL tools that involve power—assumes that power density is uniform within a block. (If not, the block can be subdivided as necessary). This means that when a block covers more than one node in the power delivery network, the total current required by that block is divided equally. Specification of the blocks and power values uses the same input interface as HotSpot [14] and leverages a new, pre-RTL architecture-level floorplanning tool for rapid specification of floorplans [6]. More details can be found in Section 5.

To solve PDN voltage and current for a given floorplan and powermap, VoltSpot first maps blocks’ power to current sources. It then traverses each grid node as well as two package nodes to update voltage information based on its neighbour’s voltage or current using Kirchhoff’s Current Law. By iteratively traversing the entire circuit, the difference between two iterations (Δ) decreases and the solver stops as soon as Δ becomes smaller than a certain threshold. In our implementation, Δ was conservatively set to a value of 1.0×10^{-7} , which is several orders of magnitude smaller than the differences we are trying to observe (on-chip voltage and IR drop).

3. VALIDATION

To understand our model’s accuracy in predicting C4 pad current and on-chip IR drop, we validated VoltSpot against a power grid analysis benchmark suite released by IBM [12].

The benchmark suite consists of detailed PDN structural information for six chips with different die sizes, silicon design and number of metal layers. The PDN structure is given in SPICE format and the SPICE files provide each and every metal wire’s geometric information and resistance value. Other information like C4 pad placement or via location between metal layers can also be extracted from the SPICE file. Similar to what we assume in our model (see Section 2), the load is also modeled as ideal current sources. Besides the PDN structure, this benchmark suite also provides a steady-state power map for each test case as well as SPICE simulation results for the voltage at each PDN node.

We parsed the SPICE files and extracted PDN grid size and resistance value as well as C4 pad location information for all the six test cases. Since the benchmark directly provides top-layer metal grid size and resistance, there is no need to calculate it from pitch and metal size. Then we ran VoltSpot to simulate each case with those values and the power maps provided by the suite.

To compare our results, we chose C4 pad current as our metric for two reasons. First, we want to study the impact of different architectures on C4 pad currents, since electromigration in C4 pads is one of the significant challenges in PDN design. Second, since IR drop across a section of wire is directly proportional to the current through that wire, current results can be directly translated into IR drop results. For this reason, the estimated current can also directly provide an estimation of IR drop error. Table 3 shows the characteristics of each benchmark and validation results.

We use two error metrics to compare our simulation results to the data provided by IBM provided. The average

Name	# of Elements	Metal Levels	# of Pads	Average Error(%)	Top Error(%)
PG1	55K	2	100	6.2	9.6
PG2	0.25M	5	120	5.2	3.3
PG3	1.60M	5	461	3.3	3.7
PG4	1.84M	6	312	2.9	1.6
PG5	2.16M	3	177	3.7	3.7
PG6	3.25M	3	132	2.7	2.8

Table 1: Validation results. Except for PG1, which has smallest size and least regular metal structure, most of the benchmarks give less than or close to 5% pad current error. “Top error” shows the average error rate for the pads within top 5% current value. Both average error and top error tend to be lower for PDNs with either more metal layers or more wires (*i.e.*, more elements).

error rate is calculated by averaging the absolute error rate across all pads, and the top error rate is the average error value of the top 5% of all C4 pad sorted by their current. We chose the top 5% because for both pad current and on-chip IR drop, we are most interested in the worst case. Except for the PG1 case, almost all the other five test cases give less than 5% average error and the top error is lower than that. According to the results from these test cases, VoltSpot has higher accuracy when modeling PDNs with more metal layers, or with more elements. PG1 not only has the lowest number of elements, number of metal levels and number of pads, but also has metal layers that are not organized in grid, and thus it does not map well to our PDN model—these are the reasons why PG1 has a higher error. It is worth mentioning that PDNs of modern high-performance processor chips usually contain multiple layers of regular metal traces. So for our study, PG1 is less representative than the other cases.

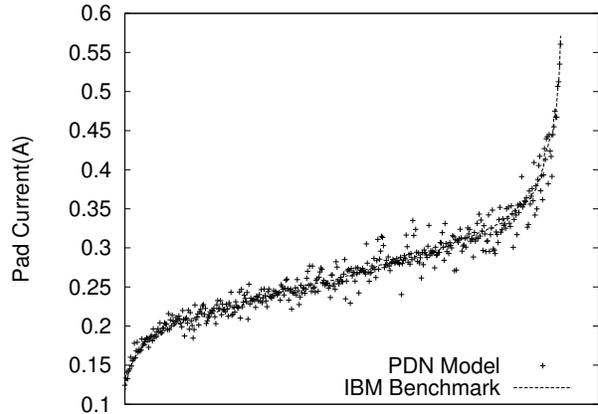


Figure 2: Alternative error representation for PG3. Pad current comparison results are sorted by original (IBM) current value and each data point’s X-axis value is its rank among all the pads. Although the top error rate for PG3 is even higher than average error, this graph shows that VoltSpot is still accurate at estimating worst case IR drop.

To better understand the accuracy of our model, we considered yet another error representation, presented in Figure 2 for PG3. The figure plots the current for all pads in

PG3, with pads sorted by the current they carry as reported by IBM. To show the validation error, we simply match pads from our model to those in the sorted list of pads in the IBM results. Although this representation loses spatial information in error distribution, it gives a better view of pad current distribution as well as error distribution in terms of pad current. Figure 2 illustrates that the error for pads with high current is lower than for pads with low current—this is important, since we are most concerned with accurately modeling those pads that deliver the highest current.

4. SIMULATION SETUP

To study the effect of technology trends on PDN noise in the near future and to explore the architectural trade-off space subject to PDN limitations, we integrate VoltSpot with an architecture-level power model and chip floorplanner. Using a 45nm Intel Penryn-like out-of-order core as a baseline, we create a series of scaled multicore processors down to 16nm and study the resulting PDN noise.

4.1 Multicore Scaling

We chose an Intel 45nm Penryn-like processor [5] as our baseline design. It has two 32-bit 4-way out-of-order cores and each core contains a 32kB L1 instruction cache and a 32kB L1 data cache. The core runs at 3.7GHz. Unified L2 caches are private to each core and 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. We also assume that L2 cache is always private. We use mesh-based network-on-chip (NoC) structure across all technology nodes.

4.2 Power Modeling and Chip Floorplanning

To get chip-wide power consumption data for all the technology nodes, we use McPAT [11], an integrated power and area model. Table 2 shows the area and peak power (including leakage power) results for our Penryn-like multicore designs in each technology.

Tech Node(nm)	45	32	22	16
# of Cores	2	4	8	16
Area(mm ²)	116.44	124.78	131.48	149.25
Supply Voltage(V)	1.0	0.9	0.8	0.7
Peak Total Power(W)	74.62	100.48	116.76	148.49
Peak Total Current(A)	74.62	111.64	145.95	212.13

Table 2: Area and power of multicore processors with Penryn-like cores

To estimate the worst-case power consumption for each system, we conducted performance simulations and activity factor analyses to extract an empirical reasonable worst-case switching activity. Based on these simulations, we use 80% of McPAT’s theoretical peak power as our best estimate for chip practical peak power consumption. Previous work on stressmark generation [10] suggests similar ratios between realistic peak power and theoretical peak power. McPAT calculates theoretical peak power by assuming maximum switching activity, corresponding to functional blocks being fully active every cycle. For most of the structures like L2 cache or NoC, this is neither achievable nor sustainable.

Chip power consumption is directly related to a workload’s dynamic activity. Depending on the magnitude and

duration of peak power instances, transient local voltage drop could be filtered out by on-chip decoupling capacitance. It also could be magnified by dI/dt effect, if the intrinsic inductance is large enough. Since VoltSpot currently focuses on steady-state effects, we assume that the peak power consumption will last long enough to ignore transient effects. We leave the study of dynamic behaviours as important future work.

We use a floorplanner developed in [6] to draw all our chip floorplans. The chip floorplan is another important input because we want to examine both global and local PDN noise. Figure 3 shows the floorplan of our Penryn-like core (L2 cache is not shown in this graph). The area of each functional block is calculated by McPAT. According to our scaling assumption, chips at different technology nodes share the same single core structure—we therefore build our multicore floorplans based on the core shown in Figure 3 and add NoCs and memory controllers.

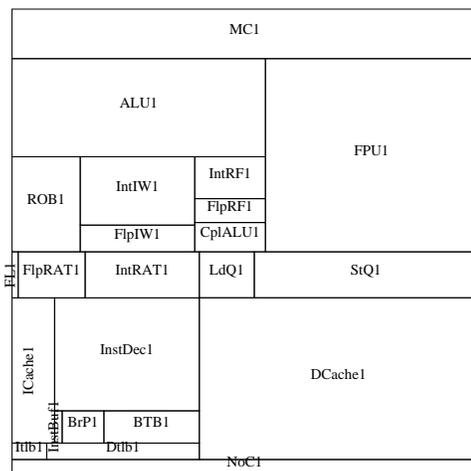


Figure 3: 45nm baseline Penryn-like core

4.3 PDN Parameters

Table 3 lists the major PDN physical parameters we used with VoltSpot. For on-chip metal, we use copper and choose pitch, width and thickness to approximate an Intel 45nm metal stack [15]. For C4 pads we use SnPb; its resistivity can be found in [4]. Pad spacing was selected so that our pad density matches ITRS projections. Package resistance comes from [1]. According to our sensitivity study, the C4 pad diameter has a negligible impact on IR drop results because it only affects pad resistance, which is relatively small compared to on-chip metal resistance. On-chip resistance depends on metal cross-sectional area and metal pitch, and therefore these two parameters are the most sensitive ones. Section 5.4 provides more detail.

5. RESULTS

5.1 Electromigration on C4 Pads

EM is one of the major failure mechanisms that deserve designers’ attention. According to [16], aluminum and copper metal wires, commonly used for on-chip interconnections, can carry two orders of magnitude higher current density than solder joints. This suggests that C4 solder bumps

Top Layer Metal Pitch (μm)	30
Top Layer Metal Width (μm)	6
Top Layer Metal Thickness (μm)	5
Top Layer Metal Resistivity (ρ)	1.68e-8
C4 Pad Diameter (μm)	130
C4 Pad Pitch (μm)	285
C4 Pad Resistivity (ρ)	1.46e-7
Package Resistance ($m\Omega$)	0.03

Table 3: PDN parameters selected for scaling study

are more vulnerable to EM. For this reason, we calculate the max current density on C4 pads, illustrated by the line in Figure 4. In order to determine the upper bound of the PDN capacity (or the lower bound of PDN noise), we assume that all pads are used for power or ground (and that each type is distributed uniformly). While this is an unrealistic assumption for a real system, it allows us to determine the best-case trend in PDN behavior. In the event that the PDN imposes constraints on the rest of the design under this best case, clearly any design under more realistic assumptions will be constrained by the PDN as well.

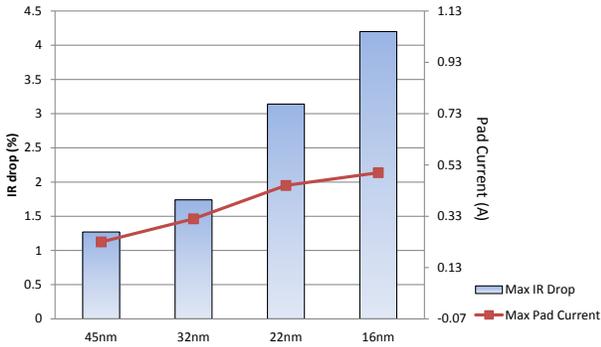


Figure 4: Maximum pad current and max on-chip IR drop at each technology node. The upper range of the right Y-axis is the threshold current value for EM (at 100°C). For IR drop, we do not set an explicit threshold value but a 3.8% IR drop could cause as high as 51% delay increase [13]. IR drop therefore poses a more significant risk to failure than EM.

In [16], the author gives an EM threshold current density for SnPb solder. At 100°C, the maximum current density that a solder joint can carry without electromigration damage is $8.5 \times 10^3 A/cm^2$. Combined with our pad diameter assumption, we calculate the per pad current limit as 1.13A.

The max value of the right Y-axis in Figure 4 indicates the current limit; it is obvious that even though the maximum pad current increases as the technology scales, the absolute value is still far away from the electromigration threshold. This suggests that under ITRS’s projections for total pad count, there would be enough guard band for electromigration in C4 pads for at least the near future. For physical-design/package communities, this observation might not be novel but it is still important for architects to be aware of.

5.2 Steady-State IR Drop

IR drop is an important PDN metric because it is directly related to silicon delay increase and frequency degradation.

As technology scales, the impact of IR drop would increase due to higher currents. Similar to the previous section, we dedicate all potential pad locations to power and ground pads and no pads to I/O signals. We then use the model to find the maximum on-chip IR drop ratio for each technology. This gives a lower bound on IR drop and the results are shown in Figure 4. The reported IR drop value combines both voltage droop from power plane and ground bounce to ground plane.

IR drop, unlike electromigration, does not directly result in immediate failure when a threshold current has been crossed, but results in performance degradation instead. Previous work [13] suggests that a 0.05V voltage drop at $0.13\mu m$ with 1.35V power supply would cause a 15% average and up to 51% maximum delay increase. The bars in Figure 4 show that the IR drop increases as the power density increases with technology scaling, and that the IR-drop ratio value reaches above 4% at 16nm—result in non-trivial performance degradation. For a more realistic scenario where not all pads were dedicated to power and ground, the problem would be even worse.

5.3 I/O Pads vs. Power Supply Pads

Since both off-chip signal I/O channels and the power supply system use C4 pads as the interface between silicon die and outside world, our previous study of dedicating all possible locations to power supply pads does not show the impact of PDN noise on the number of signal I/O pads, and hence performance as a function of I/O bandwidth. To expose this, we propose an optimization algorithm that replaces power pads with I/O pads while keeping the worst on-chip IR drop below a given threshold.

Starting from an arbitrary power pad placement with a given chip floorplan and worst-case power map, our algorithm iteratively selects one of the two following actions until a termination condition is satisfied. One possible action is removing the power pad with lowest current; the other is adding a power pad to an adjacent vacant pad location near the worst IR-drop point. Optimization terminates when either: (1) the worst IR-drop point has no adjacent pad spot that is vacant; or (2) two adjacent steps add/remove the same pad, indicating that the max IR-drop spot is close to the pad with the lowest current. Once the algorithm terminates, all the remaining vacant pad locations are allocated to I/O signals.

Figure 5 shows the results of our optimization approach. Here we assume an IR-drop constraint of 5%. The total number of pads increases because the chip area increases (see Table 2). As technology scales, the available room for I/O pads gradually scales down because the increasing chip power density requires that more and more pad space be used for power delivery. If the memory bandwidth requirement is proportional to the number of cores, the available I/O pads will soon be insufficient to support multicore scaling. Furthermore, if we assume a more strict IR-drop constraint, the chip will require more power pads, further decreasing the available I/O bandwidth.

5.4 Sensitivity Study

Most of our physical PDN parameters were selected from published industrial data, but different designs are expected to present different design choices. We therefore conducted sensitivity studies on selected variables to test whether our

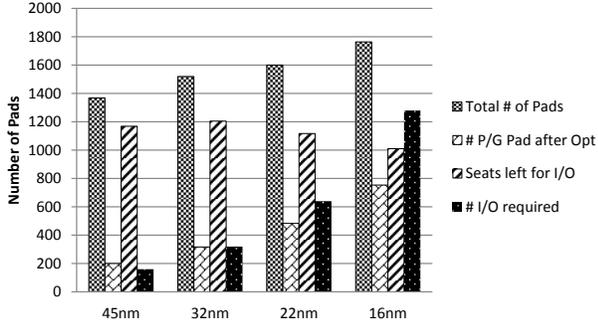


Figure 5: Number of required power pads and available pads for I/O. The number of I/O required is calculated under the assumption that the # of memory controllers is equal to the # of cores. At 16nm, the available I/O pads can no longer support the required bandwidth.

previous observations hold for different PDN designs. Figure 6 and Figure 7 present results for varying metal pitch and metal width. We did not change pad pitch in order to keep the number of total pads consistent with ITRS projections. The number of power pads after the optimization is the main metric here, because within acceptable IR-drop values, what eventually affects performance is the available I/O bandwidth.

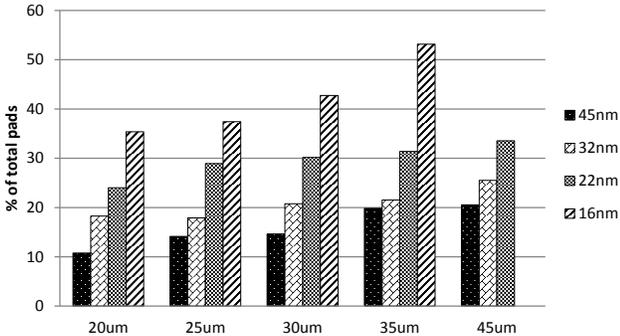


Figure 6: PDN pad requirement's sensitivity to metal pitch. Each bar represents the percentage of total pads required by PDN to achieve 5% IR drop or less. 16nm does not have data for 45um pitch because at that pitch the PDN cannot reduce the max IR drop below 5% even if all C4 pads serve as power pad.

Either decreasing metal pitch or expanding metal width can increase the number of pads available for I/O because they both add more metal to the PDN and thus help reduce IR drop by lowering resistance. However, adding more metal for power delivery means that the cost of the chip will rise and/or signal routing will become more difficult. Changing these physical parameters will not fundamentally alter the basic I/O bandwidth scaling trend—as technology scales forward, it will be critical that bandwidth, routing, IR drop and chip cost are carefully balanced.

5.5 Temperature vs. IR Drop

Both IR drop and temperature are physical design constraints that closely relate to chip power density. A robust system should be designed with both factors in mind. We

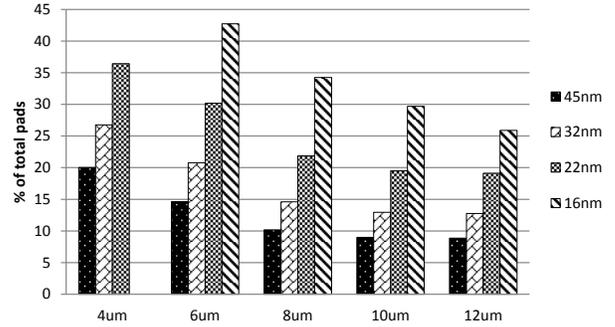


Figure 7: PDN pad requirement's sensitivity to metal width. 16nm does not have data for 4um width because with that width, the PDN cannot reduce the max IR drop below 5% even if all C4 pads serve as power pad.

found similarities between the architecture-level PDN model and compact temperature models like HotSpot [14] and integrated our model with HotSpot. Figure 8 combines chip max temperature with max IR drop. The temperature results are based on both an air cooling system and a liquid cooling system. Our results to date indicate similar trends for both power delivery and thermal limits. With air cooling system, both temperature and on-chip IR drop will cause reliability issues starting from 16nm. Switching to liquid cooling solutions will be helpful to bring down chip temperature but IR drop will still stay as scaling bottleneck. The platform we built provides an infrastructure for future studies.

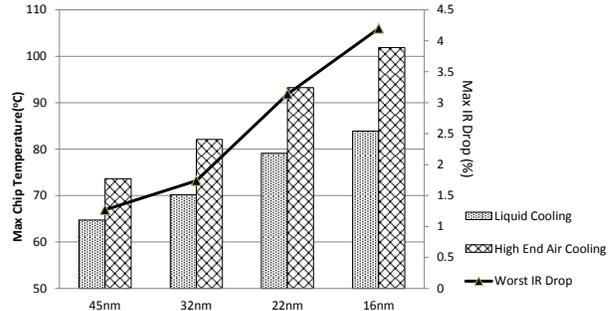


Figure 8: A comparison between chip max temperature and worst IR drop across different technologies

6. RELATED WORK

In the past, researchers have extensively studied PDN's physical structure, modeling methodology and solving and optimization algorithms. However, most previous studies focused only on the circuit level. At the architecture level, Gupta et al.[7] proposed a transient model for on-chip voltage fluctuation study and Healy et al.[8] proposed both noise-aware floorplanning and a mechanism for run-time inductive-noise control. However, neither work considers the location and number of C4 pads in their model and thus neither is capable of I/O-bandwidth tradeoff studies. To the best of our knowledge, we are the first to provide a parameterizable steady-state PDN model that incorporates the number and location of C4 pads.

Gammie et al.[3] suggest a hierarchical PDN structure for mobile application processors. Having both global power

plane and local power plane provides the benefit of supporting fine-grained power management mechanisms like power gating. Although VoltSpot assumes a single-level PDN, it is still capable of studying power grids with different granularity because of its configurability. VoltSpot's simple design, coupled with its ability to model individual blocks, makes it straightforward to extend to support more sophisticated PDN structures, and this is an important direction for future work.

7. CONCLUSIONS AND FUTURE WORK

Power delivery limits are becoming a problem in the design of microprocessors. In this paper, we introduce VoltSpot, an architecture-level power delivery network model, and validate the model against an IBM power grid benchmark suite. We study both electromigration in C4 pads and the worst-case on-chip IR drop. Our results, based on a series of scaled multicore processors, indicate that IR drop will at least triple from 45nm to 16nm and will pose a more severe constraint on future designs than pad electromigration. Furthermore, using a first-order optimization algorithm, we estimate the available I/O bandwidth under a 5% IR-drop constraint and find that, starting from 16nm, microprocessors will be unable to keep per-core bandwidth constant due to growing demand for power pads. VoltSpot is designed to be a portable library for use with a variety of performance and power models such as McPAT and HotSpot, and hence provides an infrastructure for a variety of future research opportunities.

VoltSpot suggests a number of direction for future work. We plan to investigate PDN stress under different workloads and different architectures. This requires integration between our model and architecture level performance simulators. Such infrastructure will also enable us to study run-time techniques, for example scheduling or throttling, to mitigate IR drop in high-current workload phases. We also plan to incorporate the transient aspects of power delivery modelling into our model and study effects such as Ldi/dt . Another feature we plan to add to VoltSpot is the support of different PDN organizations. This will be particularly helpful for study of fine-grained power management. Moreover, we are also interested in evaluating different pad number/location optimization algorithms in the context of an IR-drop aware floorplanner.

Acknowledgments

This work was supported in part by NSF grant no. CRI-0551630.

8. REFERENCES

- [1] *Intel Pentium 4 Processor in the 423 pin package / Intel 850 Chipset Platform*. Intel, 2002.
- [2] T. Chen and C. Chung-PingChen. Efficient large-scale power grid analysis based on preconditioned krylov-subspace iterative methods. *DAC*, June 2001.
- [3] G. Gammie, A. Wang, H. Mair, R. Lagerquist, M. Chau, P. Royannez, S. Gururajarao, and U Ko. Smartreflex power and performance management technologies for 90 nm, 65 nm, and 45 nm mobile application processors. *Proceedings of the IEEE*, 98(2):144–159, Feb. 2010.
- [4] S. Gee, L. Nguyen, J. Huang, and K. Tu. Mean time to failure in wafer level-csp packages with snpb and snagcu solder bumps. In *IWLPC*, pages 159–167, 2005.
- [5] 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*, pages 14–17, Nov 2007.
- [6] G. Faust, B. H. Meyer, and K. Skadron. Rapid prototyping of CMP floorplans. Technical Report CS-2012-02, University of Virginia, Mar 2012.
- [7] M. S. Gupta, J. L. Oatley, R. Joseph, G. Wei, and D. M. Brooks. Understanding voltage variations in chip multiprocessors using a distributed power-delivery network. *DATE*, 2007.
- [8] M. B. Healy, F. Mohamood, H. S. Lee, and SK. Lim. Integrated microarchitectural floorplanning and run-time controller for inductive noise mitigation. *TODAES*, 16(4):46:1–46:25, Oct 2011.
- [9] ITRS, 2011. <http://www.itrs.net>.
- [10] A.M. Joshi, L. Eeckhout, L.K. John, and C. Isen. Automated microprocessor stressmark generation. In *HPCA*, pages 229–239, Feb. 2008.
- [11] S. Li, JH. 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*, December 2009.
- [12] S.R. Nassif. Power grid analysis benchmarks. In *ASPDAC*, pages 376–381, March 2008.
- [13] M. Shao, Y. Gao, LP. Yuan, and M.D.R. Wong. IR drop and ground bounce awareness timing model. In *IEEE Computer Society Annual Symposium on VLSI*, pages 226–231, May 2005.
- [14] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, D. Tarjan, and K. Sankaranarayanan. Temperature-aware microarchitecture. In *ISCA*, June 2003.
- [15] N. H.E. Weste and D. M. Harris. *CMOS VLSI Design A Circuit and Systems Perspective*. Addison-Wesley, 4th edition, 2011.
- [16] Y. T. Yeh, C. K. Chou, Y. C. Hsu, Chih Chen, and K. N. Tu. Threshold current density of electromigration in eutectic snpb solder. *Applied Physics Letters*, 86(20), May 2005.