Understanding Voltage Variations in Chip Multiprocessors using a Distributed Power-Delivery Network

Meeta S. Gupta*, Jarod L. Oatley*, Russ Joseph†, Gu-Yeon Wei* and David M. Brooks*

*Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA

{meeta, jloatley, wei, dbrooks}@eecs.harvard.edu

†Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL

rjoseph@ece.northwestern.edu

Abstract— **Recent efforts to address microprocessor power dissipation through aggressive supply voltage scaling and power management require that designers be increasingly cognizant of power supply variations. These variations, primarily due to fast changes in supply current, can be attributed to architectural gating events that reduce power dissipation. In order to study this problem, we propose a fine-grain, parameterizable model for power-delivery networks that allows system designers to study localized, on-chip supply fluctuations in high-performance microprocessors. Using this model, we analyze voltage variations in the context of next-generation chip-multiprocessor (CMP) architectures using both real applications and synthetic current traces. We find that the activity of distinct cores in CMPs present several new design challenges when considering power supply noise, and we describe potentially problematic activity sequences that are unique to CMP architectures.**

I. Introduction

Supply-voltage fluctuations have emerged as a serious cause for concern in high-performance processor design. These perturbations occur when the processor demands rapidly change current consumption over a relatively small time scale. Since the power-delivery subsystem can have substantial parasitic inductance, this current variation produces voltage ripple on the chip's supply lines. This is significant because if the supply voltage rises or drops below a specific tolerance range, the CPU may malfunction. This fundamental challenge is known as the *dI/dt problem* since the magnitude of these voltage ripples is affected by the instantaneous change of current with respect to time. Current fluctuations are primarily derived from dynamic resource utilization fluctuations, which are heavily influenced by architectural power-saving events such as clock- and power-supply gating and idle/sleep modes. Thus, analysis at the architecture-level is critical to allow designers to understand the impact of these techniques on power-supply voltage stability under a variety of power-delivery and package-modeling assumptions.

Previous architecture-level dI/dt studies ([1] and [2]) have used lumped models of the on-chip power-delivery network to capture the mid-frequency resonance. The major limitation of these architectural models is the global treatment of on-chip VDD/GND as single nodes, which fails to capture local on-die voltage variations across the chip. As the effects of supply variation play a more prominent role in performance and reliability, architects will have to pay closer attention to localized supply fluctuations due to package connections and the on-chip power-supply grid. In this paper, we describe an architecture-level, fine-grained, power-delivery model that captures localized voltage variations across the entire chip.

Current technology trends are moving towards chip multiprocessor (CMP) architectures like IBM's Cell processor [3] and Intel's Core Duo processor [4]. It is important to understand inter-core voltage variations for multiple cores on a CMP machine. Core utilization patterns and activity interactions between cores can lead to large inter-core voltage variations. In order to understand these inter-core variations, a fine-grained power-delivery network is needed to model these effects. Using a distributed power-delivery model of the on-chip power-supply grid, we explore the repercussions of different combinations of activity patterns.

The main contributions of our work are:

1) We provide a parameterizable, distributed, power-delivery model, which can be configured to closely match measured impedances found in the literature [5].
2) This paper investigates voltage variations across a CMP machine using both real and synthetic activity patterns.
3) We illustrate possible problematic activity sequences that are unique to CMP architectures.

The paper is organized as follows: Section II describes the modeling of a distributed power-delivery network. The different types of activities and their effects on voltage variations are studied in Section III. Section IV reviews prior research generally related to power delivery modeling. Finally, Section V concludes the paper.

II. Modeling the Power Delivery Network

This section presents a detailed yet flexible power-delivery model that captures the characteristic mid-frequency resonance, transients related to board and package interfaces, and localized on-chip voltage variations.

Figure 1(a) presents our detailed model of the power-delivery network with a distributed on-chip power-supply grid. The off-chip network includes the motherboard, package, and off-chip decoupling capacitors and parasitic inductances, modeled via a ladder RLC network. Figure 1(b) illustrates the distributed on-chip grid model used in our analysis. The C4 bumps are modeled as parallel connections (via RL pairs) that connect the grid to the off-chip network, with each grid
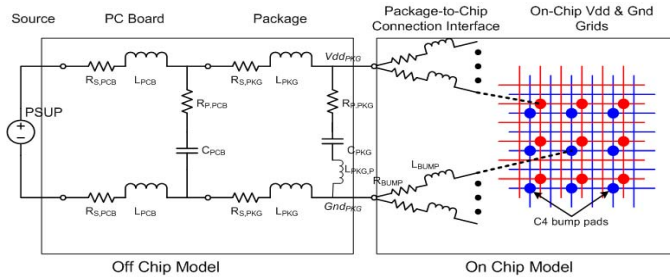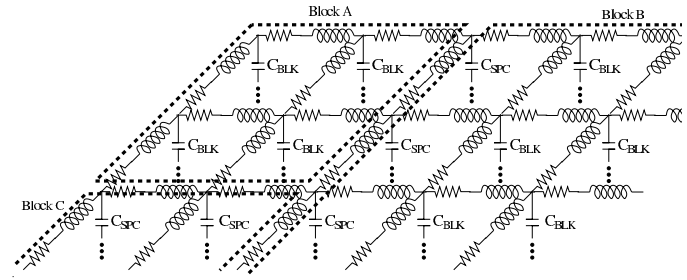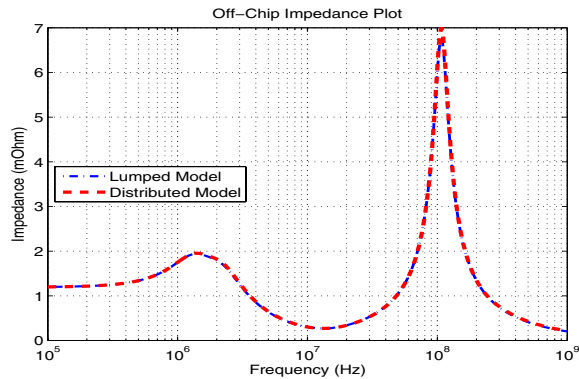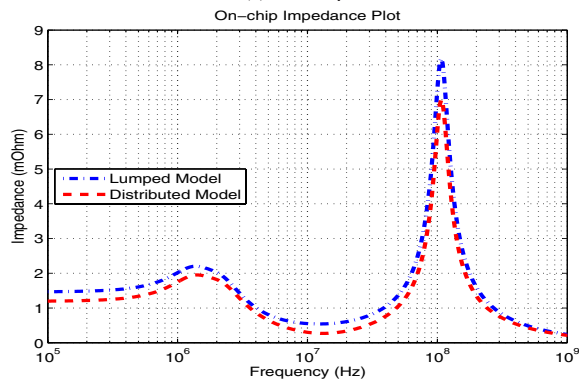
(a) Package model

Fig. 1. Power delivery model

(b) On-die grid model

(a) Off-chip

(b) OnChip

Fig. 2. Off-chip and on-die impedance plots

point having a bump connection. The on-chip grid itself is modeled as an RL network. The evenly distributed on-chip capacitance between the VDD and GND grids is modeled in two ways — C_{spc} represents the decoupling capacitance placed in the free space between functional units and C_{blk} represents the intrinsic parasitic capacitance of the functional units. In contrast, an on-chip lumped model would consist of a single RLC network connected across the package-to-chip interface. Table I provides the values of the resistances, inductances, and capacitances used for the PCB, package and on the die, for the lumped and distributed power-delivery models. These values were chosen to match the measured off-chip impedance of the Pentium 4 processor [5], [6]. Figure 2(a) plots the off-chip impedance for the lumped and distributed models, which closely match one another and are validated

with respect to the available Pentium 4 measurements [5]. The slight difference in the on-chip impedance, shown in Figure 2(b), can be attributed to the slightly higher bump resistances in the lumped model, which are required to match off-chip impedances. It is important to note these parameters can easily be modified to model different architectures and power-delivery networks.

Resistance	Value	Inductance	Value	Capacitance	Value
$R_{pcb,s}$	0.094 mohm	L_{pcb}	21 picoH	C_{pcb}	240 μF
$R_{pcb,p}$	0.1666 mohm				
$R_{pkg,s}$	1 mohm	L_{pkg}	120 picoH	C_{pkg}	26 μF
$R_{pkg,p}$	0.5415 mohm	$L_{pkg,p}$	5.61 picoH		
$R_{bump,lumped}$	0.3 mohm	$L_{bump,lumped}$	0.5 pH		
$R_{bump,grid}$	40 mohm	$L_{bump,grid}$	72 pH		
$R_{ondie,lumped}$	0.1 mohm			$C_{decoupl}$	335 nF
R_{grid}	50 mohm	L_{grid}	5.6 fH	C_{blk}	0.12nF
				C_{spc}	1.5nF

TABLE I

PARAMETERS FOR THE POWER DELIVERY MODEL

Voltage regulator modules (VRM) typically have response frequencies in the sub-MHz range, which is much lower than the challenging higher frequencies associated with the entire power-delivery network. For simplicity, the power supply is modeled as a fixed voltage source, which is scaled with respect to the average current draw to deliver 1V at the bump nodes, mimicking the feedback loop associated with the VRM.

Our architectural simulation framework consists of a four-core setup, shown in Figure 3, with each core divided into five microarchitectural blocks: FPU (floating point unit), OOO (which combines the rename, regfile, resultbus and window units on a core), INT (Integer ALU), Fetch (which combines the Icache and branch predictor) and Data (representing the Data cache and Load-Store Queue). Each block's power, derived from architectural simulations [7], is distributed evenly across the grid points according to their respective areas. To have a reasonably accurate model with low simulation overhead, we use a 12x12 grid, with each core having 36 grid points. A fast circuit solver, based on *preconditioned Krylov subspace iterative methods* [8], utilizes a SPICE netlist of the entire power-delivery network and per block current profiles to simulate on-die voltages.

The power consumption of a CMP typically varies per core due to variations in the application profiles for each core, as
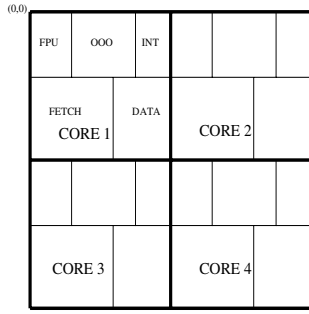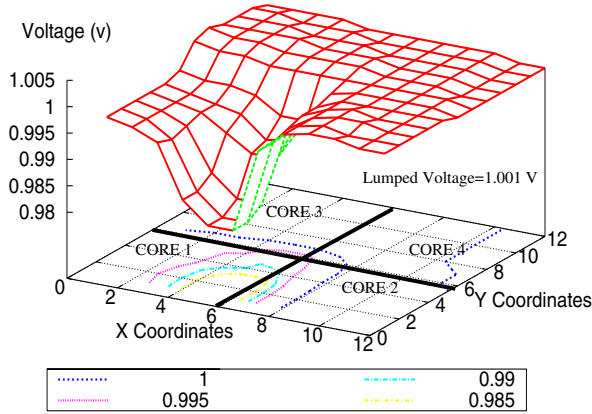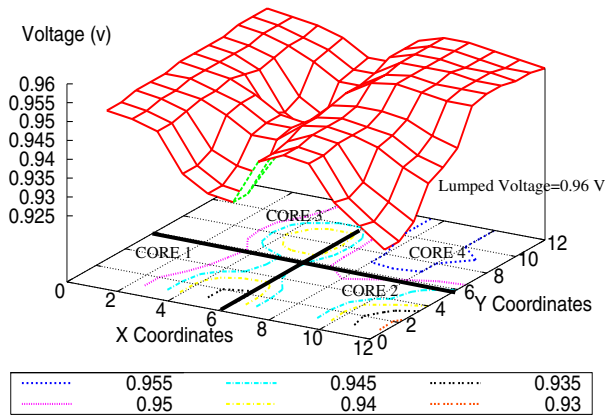
Fig. 3.   A four-core chip floorplan

(a) Core 1 running *bzip*

(b) Core 1-3 running *bzip*

Fig. 4.   Voltage variation across the chip for a snapshot of *bzip*

well as the active/idle state of each core. Figure 4 presents two different types of scenarios: Figure 4(a) shows only Core 1 running *bzip*, with the remaining 3 cores idle. We can see significant voltage variations between Core 1 and the rest of the chip. In the second example, shown in Figure 4(b), Cores 1,2 and 3 are running *bzip* and only Core 4 is idle. Again, significant voltage variations are observed across the chip. In contrast, a lumped power-delivery model would only provide single voltage values of 1.01v and 0.96v, respectively; failing to capture the voltage variations across the chip. Hence, we see the necessity of using a distributed on-chip power-delivery model. The next section focuses on understanding

these variations in the context of CMP workload scenarios.

To better understand variations both within and across different cores, it is necessary to have a distributed on-chip power delivery model. In the rest of the paper, we focus on the distributed model for the CMP processor.

III.   ANALYSIS OF VOLTAGE VARIATIONS

Voltage variations within a CMP architecture are a strong function of different workloads and current profiles associated with each core. In this section, we classify the different kinds of load current profiles and understand their effects on voltage variations within each core and across the chip.

A. Classification of Activity Patterns

In order to facilitate a thorough analysis of using a distributed power-delivery network model in CMP architectures, we begin by classifying current consumption profiles based on a suite of SPEC benchmarks. Figure 5 illustrates snapshots of interesting current profiles for four of the SPEC benchmarks —*equake, apsi, bzip*, and *mcf* for a single core. The current for the SPEC benchmarks were measured using an architectural power model based on Wattch [7]. Based on the observed characteristics, we broadly classify current consumption profiles into three categories:

1) Step Currents: This type of current profile commonly occurs when a core suddenly changes state. For example, a sudden increase/decrease in activity after long stalls due to various events like cache misses/branch mispredicts. This can also occur when the firmware enables sleep/active transitions that power up/down cores.

2) Pulse Currents: These are sudden and short duration increase/decrease in activity of the core which can again be caused due to long stalls. Figures 5(a) and 5(b) shows two examples of isolated pulses, with varying pulse widths.

3) Resonating Currents: Periodic behavior is largely associated with recurring activity patterns generally attributed to loops in an application. In particular, a periodic sequence of current pulses occurring at or near the resonant frequency of the power-delivery network are of most interest. These resonating currents are shown in Figures 5(c) and 5(d), occurring for *bzip* and *mcf*, respectively.

Given the observed application profiles we can simplify the analysis by substituting in synthetic current profiles in order to interrogate the power-delivery network for a wide range of problematic scenarios. In this paper, we focus on the effects of step currents and sequences of pulse currents on the power-delivery network leading to voltage variations. Current pulses of long enough duration can be classified as step currents. The worst case analysis can be achieved by using two states for each core: *Max-power* and *Min-power*. A *max power* state refers to when the core is drawing maximum power, which corresponds to 10W/core in our simulations. The *min power* state refers to the core consuming minimum power from the system, which corresponds to 4W/core. In our remaining
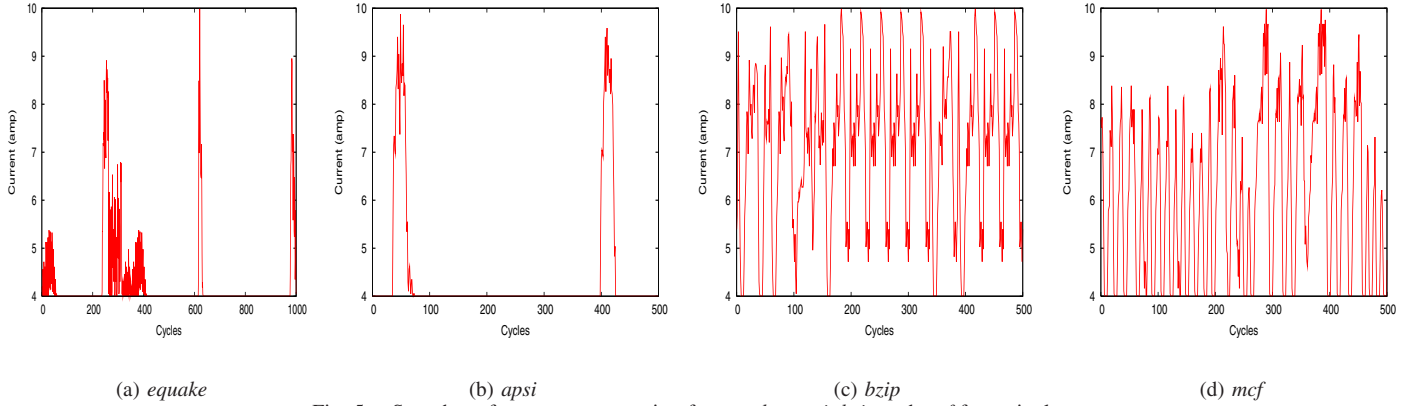
(a) *equake* (b) *apsi* (c) *bzip* (d) *mcf*

Fig. 5. Snapshot of current consumption for *equake, apsi, bzip* and *mcf* for a single core

Fig. 6. Effect of *powering on* cores

Fig. 7. Effect of staggering the cores on the voltage drop
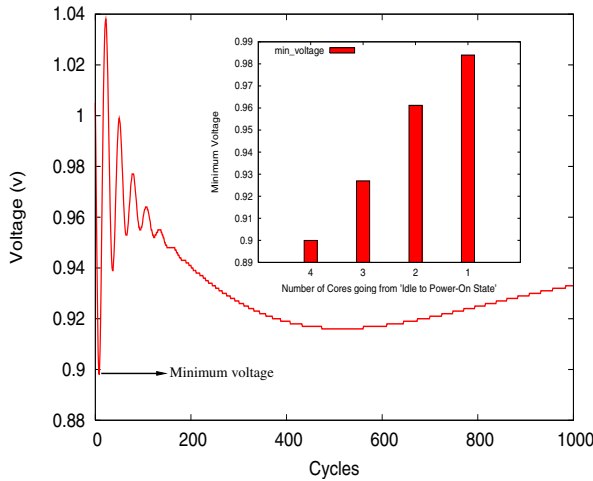
analysis we model steps and pulses with these max/min power levels to mimic powering up/down cores or activities observed in the SPEC benchmarks.

B. Voltage Variations given Step Currents

Current steps can induce large voltage fluctuation around the nominal voltage. A drop in voltage is the more alarming scenario as this can cause timing violations. Figure 6 shows the voltage variation for a node on the chip when all four cores are powered on at the same time. Given that a step is comprised of signals across a wide range of frequencies, the initial drop in voltage and the subsequent ringing can be attributed to the high frequency resonance (100MHz) in the power-delivery network. The voltage dip that occurs at 500 cycles can be attributed to the low frequency resonance. The voltage eventually stabilizes to the nominal voltage of the system (1V). Figure 6 (inset) plots the minimum voltage with respect to the number of simultaneously engaged cores. As expected, the worst drop is observed when all the cores are switched on simultaneously.

To avoid this worst case condition, a staggering mechanism can be used to gradually ramp the current profile with assistance from the firmware. The inter-core delay for switching on the cores is called the *stagger interval*. Figure 7 (inset)
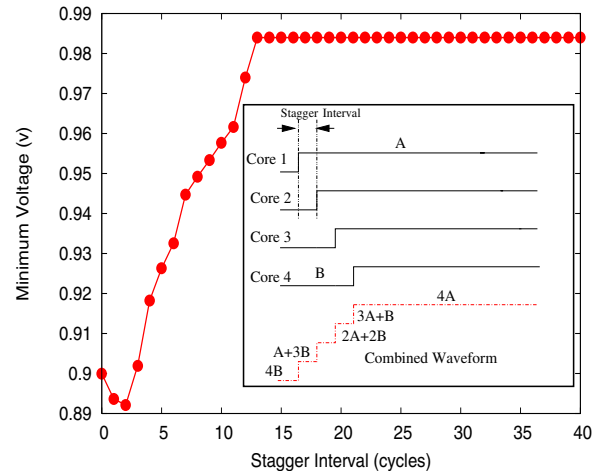
illustrates one such staggering mechanism. The combined waveform reflects the overall current consumed by the chip. Figure 7 shows that increasing stagger intervals can reduce voltage fluctuations. As stagger intervals increase beyond three clock cycles, the worst case minimum voltage across the chip improves and eventually stabilizes as the stagger interval extends beyond ten clock cycles. At this point each core behaves independently and is equivalent to a single core switching on (Figure 6 (inset)).

C. Voltage Variations given Periodic Current Pulses

Resonating currents are periodic current pulses occurring with frequencies within the resonant band of the power-delivery network. Figure 8 plots the peak voltage swing observed across the chip when the current consumption of all four cores simultaneously switch between max and min power at different frequencies with 50% duty cycle. As anticipated by the impedance plot of the power-delivery network, worst case voltage swings occur in the vicinity of 100MHz. Previous studies [1], [9] for single core machines have highlighted the detrimental effects of resonating currents on supply voltage stability. In this section, we explore the effect of resonating currents in CMP machines.

Given resonating currents, the resulting voltage ripple initially grows and then settles to a periodic waveform around the
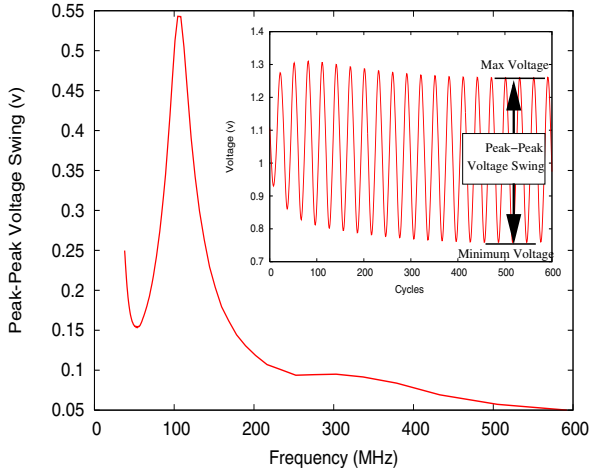
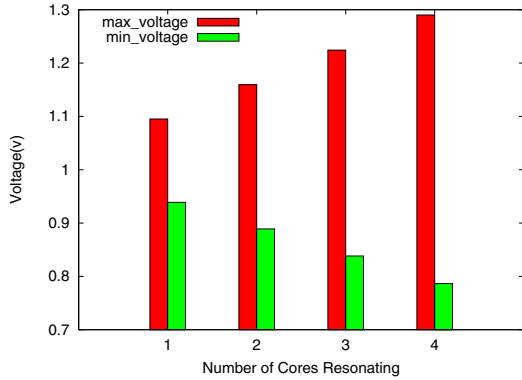Fig. 8. Periodic currents of different frequencies

Fig. 9. Effect of number of resonating cores on the peak voltage swing

nominal voltage (as shown in Figure 8 (inset)). In steady state, small current pulses can induce large peak-to-peak swings, becoming the focus of our analysis. Resonance can be further classified into: *Locally Resonant*, where each core individually has periodic current pulses at the resonant frequency; and *Globally Resonant*, where the aggregate current, globally seen across the die, has or appears to have current pulses at the resonant frequency of the power-delivery network's impedance. We further investigate the combination and interaction of these two types of resonating currents:

1) Locally and Globally Resonant: This is a scenario where each core has resonating current and the combined (or average) current pulses across all of the cores is also at the same resonant frequency. Figure 9 plots worst-case minimum and maximum voltages seen across the chip as the number of active cores increase. As expected, swings grow as the number of resonating cores increases due to the higher aggregate current amplitudes. The theoretical worst-case condition occurs when current pulses across all of the cores are aligned in phase.

2) Locally Resonant but Globally Non-Resonant : Locally, the cores are resonating but due to phase differences the combined view seen by the system is not a resonating wave. For conditions where the resonating currents across the four cores are phase-shifted with respect to one another, currents between the cores can interact to cancel out some of the effects
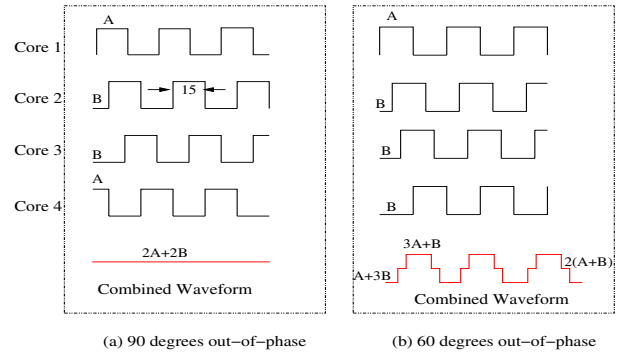

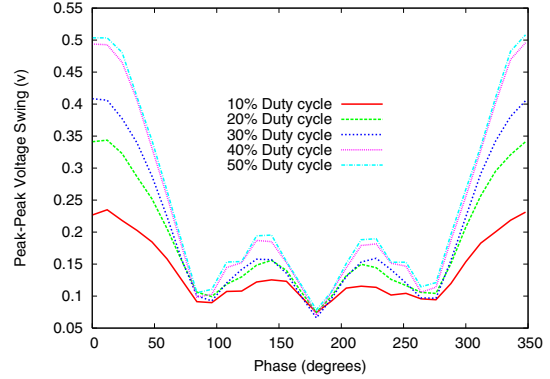Fig. 10. Examples of cores resonating out-of-phase

Fig. 11. Effect of phase difference on the peak voltage swings

of the locally resonating currents at the global scale.

When 50% duty cycle current pulses are $90°$ out of phase, with one another, as shown in Figure 10(a), the currents combine to appear as constant current with fixed amplitude at the global scale. It is important to note that due to the distributed power-supply grid model with non-zero impedance between cores, localized fluctuations exist, but interaction between the cores would cancel out resonant behavior that was seen when all of the phases were aligned. On the other hand, a lumped model would overly underestimate the potential problem given that it lacks the localized view of resonating. Figure 10(b) presents the case where resonating current pulses are each offset by $60°$. In this case, the combined currents have periodicity at the resonant frequency, but the stepwise waveform leads to smaller voltage fluctuations. Figure 11 summaries the effect of varying the phase shift between resonant currents across the four cores, and a range of duty cycles, on the resulting peak-to-peak voltage swing magnitudes seen across the CMP. As seen before the worst-case condition is when all current pulses are aligned in phase (0 or 360). And generally, larger duty cycle means higher overall current draw and, hence, larger voltage swings. Interestingly, in this four core CMP example, interactions between cores lead to the most canceling when current pulses are phase-shifted by multiples of $90°$. Given this dependence on the number of cores, a 16 core CMP may exhibit similar dips for phase differences occurring in multiples of $22.5°$.

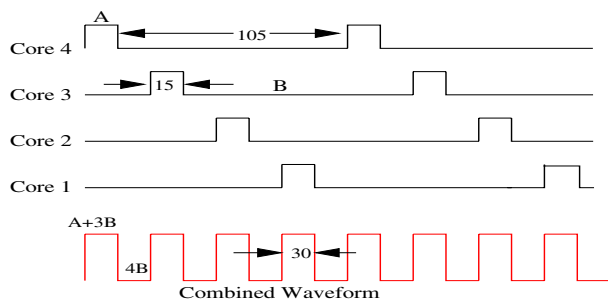*3) Locally Non-Resonant but Globally Resonant :* While the previous two conditions were examples of resonating

Fig. 12. Example of a locally non-resonant- globally resonant input

(a) 25 Mhz, cores 90 degrees out–of–phase   (b) 100 Mhz, synchronized cores
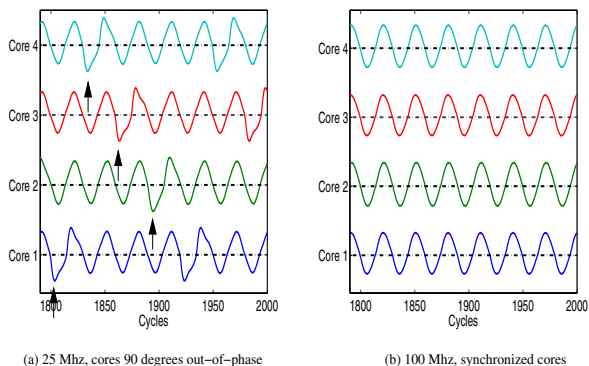
Fig. 13. Snapshot of voltages for the four cores

currents occurring in local cores, we now consider the opposite scenario. Each local core does not consume currents that pulse at the resonant frequency, but, as shown in Figure 12, the combined waveforms looks like resonating current. Moreover, given the tightly coupled power-supply grid with low impedance connections between the cores, Figure 13(a) shows that resonant voltage behavior is seen across each of the cores. In fact, there is little difference to the condition where the combined current waveform is evenly distributed across the four cores, whose resulting voltage waveforms are plotted in Figure 13(b). The only difference is the higher local ripples that occur according to the local current pulses. Hence, simply avoiding current pulses occurring at the resonant frequency alone at the core level may not prevent resonant behavior at the global scale across the entire CMP. This example further emphasizes the need to understand and model inter-core interactions at various levels of the system and design process, from application-derived current profiles to the low-level power-supply grid network.

IV. RELATED WORK

Previous architectural studies analyzing power-delivery systems have utilized models very similar to the simple lumped model described in Section II [10], [11]. These approaches capture the transient behavior of the system via an impulse response and simulation is performed via convolution. Joseph et al. [1] and Powell and Vijaykumar [2] use a single lump model which captures the mid-frequency resonance. In contrast, our work provides a fine-grained view of the localized supply droops across the chip.

Previous work studying dI/dt issues in microprocessors have

mainly focused on throttling approaches to mitigate voltage swings in single-core microprocessors [1], [2], [9], [10]. This work focuses on the inductive noise problem in the context of CMP architectures and primarily considers issues that are specific to core-to-core interactions in these machines.

V. CONCLUSIONS

As the industry trends towards aggressive power management and voltage scaling in future multi-core designs, it is increasingly important for architects to understand the potential for voltage fluctuations within this new paradigm. This paper presents a distributed power-delivery model that is designed to analyze local on-chip voltage variations to allow architects to understand the impact of inter-core interactions. We analyze this system across a range of current loads using SPEC benchmarks and synthetic current traces. We find that powering on all cores simultaneously can lead to a significant voltage drop in the system and that staggering this activity can be beneficial. Resonating current pulses can cause significant voltage swings, but if cores resonate out-of-phase, swings can be reduced. We also find that in some cases current behavior that would not be resonant within a local core, can become resonant when combined with the activity of other cores.

This paper is an initial attempt to understand the voltage variations in a CMP system. A more detailed model of the CMP architecture with different kinds of applications would lead to more insights into dI/dt effects on CMPs and possible solutions. Future research should consider more gating styles including Vdd-gating; understanding the impact of isolated per-core power domains; and studying more multi-threaded workload scenarios.

REFERENCES

[1] R. Joseph, D. Brooks, and M. Martonosi, "Control Techniques to Eliminate Voltage Emergencies in High Performance Processors," in *Int'l Symposium on High-Performance Computer Architecture*, 2003.
[2] M. D. Powell and T. N. Vijaykumar, "Pipeline muffling and a priori current ramping: architectural techniques to reduce high-frequency inductive noise," in *Int'l Symposium on Low Power Electronics and Design*, 2003.
[3] J. A. Kahle *et al.*, "Introduction to the Cell Processor," *IBM Journal of Research and Development*, vol. 49, no. 4, 2005.
[4] A. Mendelson *et al.*, "CMP Implementation in Systems Based on the Intel Core Duo Processor," *Intel Tech. Journal*, vol. 10, no. 2, May 2006.
[5] K. Aygun, M. J. Hill, K. Eilert, R. Radhakrishnan, and A. Levin, "Power Delivery for High-Performance Microprocessors," *Intel Technology Journal*, vol. 9, no. 4, Nov. 2005.
[6] Intel, "Intel Pentium 4 Processor in the 423 Pin/Package /Intel 850 Chipset Platform," February 2002.
[7] D. Brooks, V. Tiwari, and M. Martonosi, "Wattch: a Framework for Architectural-level Power Analysis and Optimizations," in *27th Annual International Symposium on Computer Architecture*, 2000.
[8] T.-H. Chen and C. C.-P. Chen, "Efficient Large-Scale Power Grid Analysis Based on Preconditioned Krylov-Subspace Iterative Methods," in *38th conference on Design automation*, 2001.
[9] M. Powell and T. Vijaykumar, "Exploiting Resonant Behavior to Reduce Inductive Noise," in *Int'l Symp. on Computer Architecture*, Jun 2004.
[10] E. Grochowski, D. Ayers, and V. Tiwari, "Microarchitectural Simulation and Control of di/dt-induced Power Supply Voltage Variation," in *Int'l Symposium on High-Performance Computer Architecture*, 2002.
[11] W. El-Essawy and D. Albonesi, " Mitigating Inductive Noise in SMT Processors," in *International Symposium on Low Power Electronics and Design*, August 2004.