Pre-RTL On-Chip Power Delivery Modeling and Analysis

A Dissertation

Presented to

the faculty of the School of Engineering and Applied Science

University of Virginia

In partial fulfillment

of the requirements for the degree

Doctor of Philosophy

Computer Engineering

by

Runjie Zhang

May 2015

Approvals

This dissertation is submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Computer Engineering

Runjie Zhang

Approved:

Kevin Skadron (Advisor)

Joanne Bechta Dugan (Chair)

Mircea R. Stan (Advisor)

Mary Lou Soffa

David Brooks

Accepted by the School of Engineering and Applied Science:

James H. Aylor (Dean)

May 2015

Abstract

The Power delivery network (PDN) is the electrical system that provides supply voltage to the transistors within a silicon chip. Due to the PDN's intrinsic resistance, capacitance, and inductance, the supply voltage can become noisy (drop or fluctuate) and cause timing errors, threatening program correctness. While CMOS technology scaling has resulted in exponentially greater transistor densities, threshold and supply voltages no longer decrease fast enough to prevent exponential growth in on-chip power density. Although the technology of three-dimensional integrated circuit (3D-IC) provides an alternative path toward the continued historical trend of device integration growth, it further increases the aggregated power density by stacking active silicon layers on top of each other.

As a result, power-delivery-related reliability issues are increasing, creating higher demands for the already scarce physical resources like controlled collapse chip connection (C4) pads and silicon area for the integrated on-chip decoupling capacitors. Unfortunately, there often exists a contention between power delivery needs and processor computation needs. Under these rising challenges, it becomes increasingly important to consider PDN design and optimization at early design stages, both to ensure an optimal design point selection in the complicated tradeoff space, and to prevent costly redesign due to power delivery issues.

In this dissertation, we build and validate a pre-RTL PDN model, called VoltSpot, to (1) identify the power-delivery difficulties for contemporary and near-future microprocessors; (2) understand the impact of the PDN's physical resource allocations on voltage noise, and explore the resulting tradeoff space considering the processor's performance and lifetime under electromigration stress; (3) design and evaluate both static and dynamic solutions to mitigate power delivery constraints; (4) assess the benefits and costs of novel power delivery schemes for 3D-IC.

Acknowledgments

I would like to thank my advisors Prof. Kevin Skadron and Prof. Mircea Stan for guiding me through my journey in graduate school. They taught me how to build and defend an argument, how to present work to peer researchers, and how to make improvements based upon feedback/criticisms. I have learnt much from them, not only about how to become an independent researcher, but also the spirit of professionalism and the art of balancing life and work. I owe a substantial amount of gratitude to their support and trust for allowing me to work remotely during my final semesters, so that I could put an end to my long-lasted two-body problem in much advance.

I wish to thank the members of my Ph.D. committee: Prof. Mary Lou Soffa, Prof. Joanne Bechta Dugan, and Prof. David Brooks for their valuable suggestions and comments. Prof. Brooks' earlier work on architecture-level power noise modeling and management not only helped me to paint the big picture in my mind, but also inspired an important branch of my research work. Through the past several years, Prof. Brett Meyer from McGill University has been continuously helping us with brainstorming new ideas and writing papers. As a UVa research scientist and a numerical expert, Dr. Ke Wang played a key role in the development of VoltSpot's circuit solver, and also brought many novel thoughts into our PDN research. A non-trivial portion of VoltSpot's source code was directly inherited from an existing thermal model HotSpot, which was developed by Dr. Karthik Sankaranarayanan, Dr. Wei Huang, and many others. Dr. Wei Huang also provided many valuable suggestions to us in the early-stage of VoltSpot development. My UVa colleague Kaushik Mazumdar and I collaborated closely in the 3D-IC-related projects. I learned a great deal of IC design and fabrication from our industrial collaborators, which include, but not limited to: Dr. Pradip Bose, Dr. Dale Becker, Dr. Sungjun Chun, Dr. Pritish Parida, Dr. Guoqing Chen, and

Dr. Cheng Zhuo. I want to express my sincere gratitude to all of them. This dissertation would not have been possible without their help.

I am grateful for having so many friends both inside and outside the CS/ECE departments. The members of Kevin's LAVA lab and Mircea's HPLP group have helped me in numerous ways. Especially, I thank Liang Wang, Shuai Che, Jack Wadden and Michael Boyer for their generous help in teaching me software development basics, debating research ideas, and sharing the location of free food. Also, I want to thank my fellow graduate students Yuchen Zhou, Yanqing Zhang, Wei Wang and Yan Huang for setting examples of excellent researchers and motivating me with peer-pressure.

I also want to express my gratitude to Prof. Xueqing Lou from Zhejiang University, China. As an undergraduate student, I took his computer organization class in 2007. That class and his vivid lectures made me discover my interest in computer architecture and decide to pursue a career in this area.

Finally, I want to thank my family for their unconditional love and support. As dedicated engineers themselves, my parents not only inspired my curiosity for the world, but also taught me everything I need to know about how to become a good person. Although my wife Chuhan and I spent the majority of the past six years living in separate cities, states, or even time zones, her wisdom and love could always travel through distance to share my joy and lead me through the most stressful, depressing, and even desperate times in my life. I dedicate this dissertation to all of them.

This work was supported by SRC grant 2009-HJ-2042, NSF grant no. CRI-0551630, NSF grant CNS-0916908 and MCDA-0903471, and DARPA MTO under contract no.HR0011-13-C-0022.

Contents

Acronym List **xviii**

1 Introduction **1**

1.1 Overview 1

1.2 Power Delivery Network and Supply Voltage Noise 3

1.3 PDN Resource Scarcity and Optimization 5

1.4 PDN Simulation and Modeling 7

1.5 Run-time Voltage Noise Mitigation Techniques 9

1.6 Electromigration and Long-term PDN Reliability Issues 10

1.7 3D-IC and Voltage-Stacked Power Delivery 12

1.8 Dissertation Organization 14

2 VoltSpot: A pre-RTL Power Delivery Network Model **15**

2.1 Overview 15

2.2 Related Work 16

2.3 Modeling Methodology 17

2.3.1 Input/Output 17

2.3.2 Structure 18

2.3.3 Numerical Engine 20

2.4 Improvements Beyond State-of-the-Art Models 20

2.4.1 Modeling Granularity 20

2.4.2 Explicit Multi-layer Model 22

2.5 Validation 23
2.6 Interfacing with Other Architecture Level Tools 25
2.7 Summary 27

3 Architecture Implications of Pads as a Scarce Resource 28
3.1 Overview 28
3.2 Background and Related Work 30
3.3 Simulation Setup 32
3.3.1 Multicore Scaling, Power Modeling and Floorplanning 32
3.3.2 Power Trace Sampling and Stressmark 33
3.3.3 PDN Parameters and Pad Location Optimization 34
3.4 Transient Voltage Noise: Scaling and Effects of C4 Pad Configuration 35
3.4.1 IR Drop vs. Transient Noise 35
3.4.2 Scaling Trends of Supply-Voltage Noise 35
3.4.3 Voltage Noise and Pad Configuration Effects 37
3.4.4 Sensitivity Study 39
3.5 Run-time Voltage Noise Mitigation 41
3.5.1 Dynamic Margin Adaptation 41
3.5.2 Recovering From Noise-Induced Errors 43
3.5.3 A Hybrid Technique 44
3.5.4 Trading Power Pads for Performance 45
3.5.5 Impact of I/O routing 46
3.6 Chip I/O Bandwidth and C4 EM Lifetime 47
3.6.1 Whole-Chip Electromigration MTTF Calculation 48
3.6.2 Tolerating Pad Failures with Runtime Mitigation 49
3.6.3 EM Lifetime Considering Pad Failure Tolerance 51
3.6.4 Impact of On-Chip Metal 52
3.7 Summary 52

4 A Statistical Analysis of EM-induced C4 Bumps Wearout 54

4.1 Overview . . . 54

4.2 Background and Related Work . . . 56

4.2.1 Single Bump EM Failure . . . 56

4.2.2 Consequences of Power-Supply Bump Failures . . . 57

4.2.3 Related Work on Statistical EM Analysis . . . 58

4.3 Statistical Simulation of Bump Failures . . . 59

4.3.1 Monte Carlo Simulation and Results Confidence Level . . . 59

4.3.2 Capturing the Impact of Current Re-distribution . . . 60

4.3.3 MCS Framework for Bump Failure Study . . . 61

4.3.4 Supporting Other Random Failures . . . 63

4.4 Simulation Methodology . . . 63

4.4.1 PDN Modeling . . . 63

4.4.2 Multicore Processor Power, Area and Floorplan Modeling . . . 64

4.5 Results and Discussions . . . 65

4.5.1 The Avalanche Effect: How Bad Is It? . . . 65

4.5.2 Achieving Target MTTF with Reduced Bump Count . . . 66

4.5.3 Impact of Multiple On-Chip Power Domains . . . 68

4.5.4 Using Graceful Performance Degradation Schemes to Extend Chip Lifetime 70

4.5.5 Transient Noise Evaluation . . . 72

4.6 Summary . . . 73

5 A Cross-Layer Exploration of Voltage-Stacked PDN in Many-Layer 3D-IC 75

5.1 Overview . . . 75

5.2 Background and Related Work . . . 77

5.2.1 Power Delivery Challenges and Solution in the Era of 3D-IC . . . 77

5.2.2 Voltage Regulation in V-S PDN . . . 78

5.2.3 System-level Evaluation of V-S . . . 79

5.3 Modeling Methodologies 80
5.3.1 SC Converter Modeling 80
5.3.2 PDN Modeling for 3D-IC 82
5.3.3 EM-induced System Lifetime Calculation 83
5.4 Simulation Setup 84
5.4.1 Many-core Processor Modeling 84
5.4.2 PDN Modeling and TSV Configurations 85
5.5 Results and Discussions 86
5.5.1 Heat Removal and Power Delivery in 3D-IC 86
5.5.2 EM-Induced TSV/C4 Pad Lifetime 88
5.5.3 Load-Imbalance-Induced Voltage Noise 89
5.5.4 System Power Efficiency 92
5.6 Summary 93

6 Transient Noise in Voltage-Stacked PDNs 95
6.1 Overview 95
6.2 Background and Related Work 96
6.2.1 Voltage Noise and Timing Margin in 3D-IC 96
6.2.2 Voltage Regulation in V-S PDN 97
6.2.3 System-level Supply Voltage Noise Modeling 97
6.3 V-S PDN Modeling Methodologies 99
6.3.1 A Transient Model for SC Converters 99
6.3.2 Validation 100
6.3.3 Whole-system Model 102
6.4 Simulation Setup 102
6.4.1 Many-core Processor Modeling 102
6.4.2 PDN Modeling 102
6.4.3 Workload Modeling 103

6.5 Results and Discussions 104
6.5.1 Cross-layer Noise Interference 104
6.5.2 Allocating On-Chip Capacitance: A Tradeoff Study 106
6.5.3 Impact of 3D Scaling 108
6.5.4 Impact of Package Impedance 109
6.6 Summary 110

7 Conclusions and Future Work **112**
7.1 Dissertation Summary 112
7.2 Future Directions 114
7.2.1 Pre-RTL Power Delivery Network Modeling 114
7.2.2 Statistical Analysis of PDN Wearout 115
7.2.3 3D-IC and Voltage Stacking 116
7.2.4 Other Directions 116

Bibliography **117**

List of Figures

1.1 An illustration of voltage noise and the evaluation metrics. 5

1.2 (a) Side view of a chip showing its connection to the package through C4 pads. (b) Top view of a C4 array showing an example allocation of power pads and I/O pads (real processors have many more C4 pads than illustrated here). 6

1.3 Illustrations of both static and run-time noise mitigation techniques. 10

1.4 The log-normal CDF of a single conductor's EM-induced failure time. The y-axis shows the conductor's failure probability, or $Prob[Lifetime] \leq t$. 11

1.5 Conventional and V-S power-delivery scheme for 3D-IC. 13

2.1 VoltSpot model structure 18

2.2 IBM power grid analysis benchmark: A snap shot. 23

2.3 Current/Voltage comparison between IBM data and VoltSpot results. 24

2.4 How VoltSpot interacts with other architecture level tools. 26

3.1 Voltage-emergency maps for different pad configurations for a 16nm, 16-core chip. A total of 1914 C4 locations are available. All three graphs have the same color scale; areas with warmer colors indicate more voltage violations. 30

3.2 Floorplan of the Penryn-like 16core processor. 32

3.3 Comparison of transient voltage noise and static IR drop over a 1K cycle window in the benchmark ferret. Periodic oscillation implies that *LC resonance* is the major cause. We build our noise-inducing stressmark based on this segment. 36

3.4 Voltage noise change across different pad configurations. Each MC needs 30 I/O pads. 38
3.5 Voltage noise change across different pad configurations under a different set of package parameters. 40
3.6 Speed up of recovery-based techniques with different timing margin settings. We evaluate our 16nm, 16core processor with 24MCs here. The baseline case enforces a 13% timing margin and thus guarantees timing is error-free. 43
3.7 Performance comparison between different noise mitigation techniques. Numbers after recover/hybrid in legend represent the cost (in cycles) of each recover, from error. The baseline case enforces a 13% constant timing margin. 45
3.8 Performance penalty of mitigating extra voltage noise caused by reduced power/ground pads. We use the hybrid technique with a conservative assumption of 50-cycle rollback-and-replay penalty per error. Each benchmark uses its own performance with 8MC case as baseline. 46
3.9 PDN pad failure's effect on noise mitigation performance overhead and expected EM lifetime. F indicates the number of failed pads. The baseline is an 8 MC chip with no failed pads. For *Recovery*, the performance overhead with 32 MC goes off-chart to 15% if 20 pads fail (25% if 60 pads fail). 50
3.10 Impact of on-chip metal occupancy on the EM lifetime of both C4 pads and on-chip wires. 52

4.1 The consequence of a power bump failure. 58
4.2 Floorplan of our example multi-core processor. 64
4.3 Normalized whole-system MTTF under different noise margin settings. The differences between the two sets of results are caused by the avalanche effect. 66

4.4 The required noise margin (in terms of absolute Vdd% IR drop) to achieve target MTTF for different bump allocations and multiple on-chip power domain settings. We assume that the target MTTF equals the EM-failure-free time of the single-domain chip with 70% bumps as power/ground. 68

4.5 Evaluation of two graceful performance degradation mechanisms. Margin adaptation technique gradually relax noise margin until a designated limit (different lines with 2-5% Vdd). Core desertion scheme abandons a core if its noise exceed the design target. All values are normalized to the amount of workdone by the baseline chip without any graceful degradation mechanisms enabled. 71

5.1 Stacked loads (three layers) with stacked SC converters (two), ideally providing Vdd voltage headroom to each load. Zoomed up single cell of 2:1 push-pull SC converter shown on the left. 79

5.2 Power efficiency and voltage drop model for SC converters. 80

5.3 Model validation results. 82

5.4 Conventional and Charge-recycled power-delivery scheme for 3D-IC. 83

5.5 Floorplan of our single-layer 16-core processor. 84

5.6 Temperature and voltage noise evaluation of our example many-core 3D processors. For voltage noise, we tested different power-supply TSV topologies described in Table 5.2. 87

5.7 EM-induced lifetime evaluation. All results are normalized to the lifetime of 2-layer V-S PDN 3D-IC. 88

5.8 Voltage noise evaluation of our 8-layer processor. For 3D-ICs without V-S, the worst-case IR drop happens when all layers are fully active. Therefore the assumption about workload-imbalance does not affect those evaluations. 90

5.9 A box-plot that shows the distributions of workload imbalance within and across different applications. 92

5.10 Power efficiency of 3D processors with V-S PDN. 93

6.1 A single cell of our 2:1 push-pull SC converter and its equivalent circuits in the two different clock phases. 98
6.2 An RC model for the interleaved SC converters and a whole-system view of a V-S PDN in a many-layer 3D-IC. 99
6.3 Validation results. 101
6.4 A plot of per-layer maximum noise amplitude over time. Only layer 3 has a noisy workload. 105
6.5 A 4-layer 3D stack's voltage noise amplitude under different PDN configurations and workload conditions. The x-axis numbers within each data cluster represent the percentage of die area allocated for explicit decap. The size of each SC converter equals 3% of an ARM core. 107
6.6 3D scaling's impact on worst-case voltage noise. 109
6.7 Package impedance's impact on 3D-IC's whole-stack maximum noise. The x-axis of each data group shows the normalized package impedance. 110

List of Tables

2.1 Impact of on-chip modeling granularity. At 25:1, the number of rows and columns in the on-chip grid roughly matches the number of top-level metal wires in the PDN stack. Among 1914 total C4 pads, 1254 pads were allocated to P/G. 21

2.2 Impact of modeling multiple metal layers. 22

2.3 Static and transient validation results against IBM benchmark 25

3.1 Characteristics of Penryn-like Multicore Processors 33

3.2 PDN Parameters 34

3.3 Voltage Noise Scaling Trend, Ideal (all pads allocated to power/ground), in fluidanimate benchmark 37

3.4 I/O pad breakdown for commercial processors. 37

3.5 Package parameters extracted from Intel Xeon 5500 series processor [29]. 39

3.6 Supply voltage change due to metal layer occupancy. The majority of this chapter assumes a 33.3% layer occupancy. The evaluation platform is our 16core, 8MC processor. 40

3.7 Dynamic Margin Adaptation and Scaling 43

3.8 Impact of package impedance. We increase the serial portion of package impedance as an emulation to the consequence of more complicated I/O routing in chip package. 47

3.9 C4 Pad EM Lifetime ScalingTrend 49

4.1 PDN Parameters used in this study. 64

4.2 Whole system's expected EM-failure-free time under different bump allocations. . 67

4.3 Evaluation of EM-induced system MTTF under different transient noise margin settings. A 1k cycle stressmark is used to evaluate chip transient voltage noise. All results are normalized to the EM-failure-free time of the baseline processor with 40% bumps assigned as power supply. 73

5.1 Major PDN modeling parameters 85

5.2 TSV configurations used in this study. 86

6.1 Primary PDN modeling parameters 103

6.2 Maximum voltage noise (%Vdd) per layer for different workloads on 3D-ICs with different PDN schemes. The "cross-layer mean" value averages all layers' maximum noise amplitude. 104

Acronym List

Definitions are collected here for easy reference. In general, the accepted definitions for terms are used, although some terms are used in a more restricted sense than their usual interpretation.

3D-IC Three-dimensional integrated circuits.

AC/DC Alternate current/Direct current.

ALU Arithmetic logic unit.

C4 pad/bump Controlled collapse chip connection. A method for interconnecting silicon dies to package.

CDF Cumulative distribution function, which defines the probability distribution of a random parameter.

CMOS Complementary metal-oxide-semiconductor.

CPM Critical path monitor. A structure proposed by Lefurgy et al. [51] to detect the critical path's available timing margin during run-time.

DRAM/SRAM Dynamic random-access memory/Static random-access memory.

ECC error-correcting code.

decap Decoupling capacitor.

DPLL Digital phase-lock loop.

DVFS Dynamic voltage and frequency scaling. A power management technique.

EM Electromigration. Refers to the gradual mass transportation in metal conductors induced by momentum transfer from electrons to atoms.

FBDIMM Fully buffered dual in-line memory module.

KoZ Keep-out zone for TSVs, allocated to prevent the thermal stress generated by TSVs from impacting the electrical performance of the nearby transistors.

LU decomposition A method to factor matrix as the product of a lower triangular matrix and an upper triangular matrix. 'LU' stands for 'lower upper'.

I/O Input/Output.

IR drop One type of voltage noise. Refers to the resistive drop across PDN wires.

ITRS International Technology Roadmap for Semiconductors.

IVR Integrated voltage regulators.

LdI/dt One type of transient voltage noise. Caused by the PDN's inductance and the change of circuit current consumption.

MC Memory controllers.

MCS Monte Carlo Simulation.

MIM capacitors Metal-insulator-metal capacitor.

MTTF Median time to failure.

MTTFF Median time to first failure. A metric we defined the show a group of conductors' expected lifetime until the first EM-induced failure.

NoC Network on chip.

PCB Printed circuit board.

P/G Power/Ground.

PDN Power delivery network.

RC/RLC Resistor and capacitor/Resistor, inductor, and capacitor.

RTL Register-transfer level. A design abstraction.

SC converter Switched-capacitor converter. A type of integrated voltage regulator.

SPICE Simulation Program with Integrated Circuit Emphasis. A general-purpose, open source analog electronic circuit simulator.

TDDB time-dependent gate oxide breakdown. A long-term failure mechanism for integrated circuit.

TOF Time of failure.

TSV Through-silicon-vias. TSVs vertically connect all layers in 3D-ICs and provide both inter-layer power delivery and I/O communication channels.

V-S Voltage-stacking. A power delivery arrangement that electrically connects circuit blocks in series.

VRM Voltage regulator modules.

Chapter 1

Introduction

1.1 Overview

In the past decades, the scaling of CMOS technology allowed the semiconductor industry to successfully keep an exponential growth rate in device integration. However, even though Moore's Law continues to bring exponentially greater transistor densities, threshold and supply voltages no longer decrease fast enough to prevent exponential growth in on-chip power density [58]. Although the technology of three-dimensional integrated circuits (3D-IC) provides an alternative path toward continued growth of device integration, it further increases the aggregated density of both current consumption and heat generation. Consequently, it is becoming increasingly challenging to deliver sufficient current to switching transistors as well as to remove the massive heat generated by silicon chips. Among these physical constraints, thermal issues have been recognized as a critical barrier to utilize transistors effectively [81]. However, power *delivery* may be an even more serious constraint, especially as research and development efforts focus on aggressive cooling technologies such as liquid and microchannel cooling ([78]), ameliorating if not eliminating the so called "temperature wall."

The power delivery network (PDN) is the electrical system that provides supply voltage to the underlying transistors within the silicon chip. A modern PDN usually consists of several levels of voltage regulator modules (VRM) and decoupling capacitors. Metal traces on both the printed circuit board (PCB) and in the chip package deliver supply current from off-board or off-chip VRMs

to controlled collapse chip connection (C4) pads, and on-chip metal layers further distribute current to the transistors [69].

Due to the power delivery network's intrinsic resistance, capacitance and inductance, the supply voltage will become noisy (drop or fluctuate) after traveling through the PDN. Since transistor delay is directly related to the voltage between its source and drain [74], any voltage variation outside the assumed design margin can cause a timing error, threatening program correctness. Furthermore, the PDN also suffers from long-term reliability issues like electromigration (EM). EM refers to gradual mass transport in metal conductors induced by momentum transfer from electrons to atoms. Under high current stress, EM can cause permanent damage to PDN wires, and degrade power delivery quality with increased voltage noise and voltage-noise-induced soft errors.

To guarantee power supply stability during silicon chip's entire lifetime, circuit designers and researchers have developed sophisticated methods to allocate and optimize a PDN's physical structure for both 2D and 3D chips (e.g. metal layer geometry [35, 36], decoupling capacitor distribution [67, 69], C4 pad location [99, 102, 106], and though-silicon-via, or TSV topology [23, 86]). These techniques help to reduce supply noise, and to date, noise can generally be handled with modest guardbands. However, power delivery difficulties arise and will get worse for several reasons. First, despite advances that have produced sophisticated on-chip PDN, the intrinsic resistance and inductance of the PDN circuit cannot be fully controlled with reasonable cost. This makes voltage fluctuation on the supply rails inevitable. Second, as current density increases with technology scaling and 3D scaling (i.e., stacking more layers of active silicon on top of each other), PDN reliability issues are exacerbated, creating higher demands for the already scarce physical resources like C4 pads and on-chip decaps. Finally, there is a contention between power delivery needs and processor computation needs. For example, C4 pads are the only connections between a silicon chip and the outside world thus both power supply and chip I/O signal links must utilize C4 pads exclusively. Under theses rising challenges, it becomes increasingly important to consider PDN design and optimization at early design stages, both to ensure an optimal design point selection in the complicated tradeoff space, and because a later redesign due to power delivery issues is costly.

This dissertation focuses on making cross-layer power delivery improvements in order to facil-

itate continued performance scaling. Our approach consists of the following major research tasks:

1. Design, implement, and validate a pre-RTL PDN model that is capable of both modeling PDN's physical structure in detail and simulating whole-application power supply noise. Integrate this model with other architecture-level performance, power, and floorplan tools.

2. Analyse the impact of different C4 configurations on power delivery quality. Quantitatively evaluate the resulting performance impacts with different run-time noise mitigation techniques.

3. Provide a statistical estimation for whole-chip's EM-induced lifetime that considers the failure possibility of all pads and the sequence of pad failures. Analyse the noise impact of multiple pad failures and explore the feasibility of tolerating failures with increased noise margin.

4. Extend research infrastructure to understand how 3D scaling affects voltage noise. Study a novel, charge-recycled power delivery scheme for 3D-ICs and compare its quality and costs against the traditional PDN structure.

1.2 Power Delivery Network and Supply Voltage Noise

A microprocessors' power delivery network is a vital system because it provides supply voltage and current to the underlying circuits. Because of the large scale and high current demands of contemporary silicon chips, modern PDNs are usually complicated and hierarchical systems. Starting from the motherboard's power source, a typical route that the supply current will travel through is: voltage regulation modules (via the metal traces on the PCB board), chip package (via the chip socket and package pins), on-chip PDN (via the package-level power planes and C4 pads), and finally the switching transistors (via the on-chip power supply metal stack). The power delivery system normally deploys several levels of voltage level converters to gradually and efficiently step down the supply voltage from the level provided to the motherboard (e.g., 12V [88]) to the transistors'

nominal supply voltage (e.g., 1V). Similarly, decoupling capacitors (decap) are also allocated along the power-delivery path as reservoirs of electric charge to stabilize the supply voltage.

As the main focus of this dissertation, 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). To suppress localized voltage noise, designers also allocate integrated capacitors as on-chip decap. Regardless of their hierarchy, on-chip PDNs were designed to keep the on-chip voltage as *spatially uniform* and *temporally steady* as possible.

In realistic scenarios, the power delivery network has a non-zero impedance due to its intrinsic resistance, inductance and capacitance. Consequently, the actual voltage level delivered to the power consumers (i.e., switching transistors) will inevitably deviate from the nominal value. This voltage-level deviation is referred to as voltage noise, of which there are three major sources: *IR drop*, *LdI/dt droop*, and *LC resonance*. The *IR drop* refers to the resistive drop across a metal conductor. Its magnitude is directly related to the PDN's resistance, and can be used as a noise metric in static PDN analysis. The *LdI/dt droop* is caused by the PDN wires' inductance. Because the voltage across an inductor is calculated as $V = LdI/dt$, the amplitude of LdI/dt droop is determined by both the PDN's inductance (L), and the change rate of the load current (dI/dt). Finally, the LC resonance occurs when the load current exhibits a periodic pattern with certain frequency that triggers an oscillation.

Figure 1.1 shows an example voltage trace and illustrates the commonly used noise metrics. Assuming a nominal supply voltage of 1V, the actual voltage seen by the circuit in this example fluctuates between 0.94 V and 0.97 V. At any moment, the difference between the actual voltage (blue line) and the nominal supply voltage (Vdd) is referred to as the *noise amplitude*, and it is usually normalized to Vdd and reported as a percentage value. Since each transistor's delay is directly related to the voltage between its source and drain [74], the delay of the circuit's critical paths will be affected by the voltage noise level. Because of this noise-induced delay variation, it is

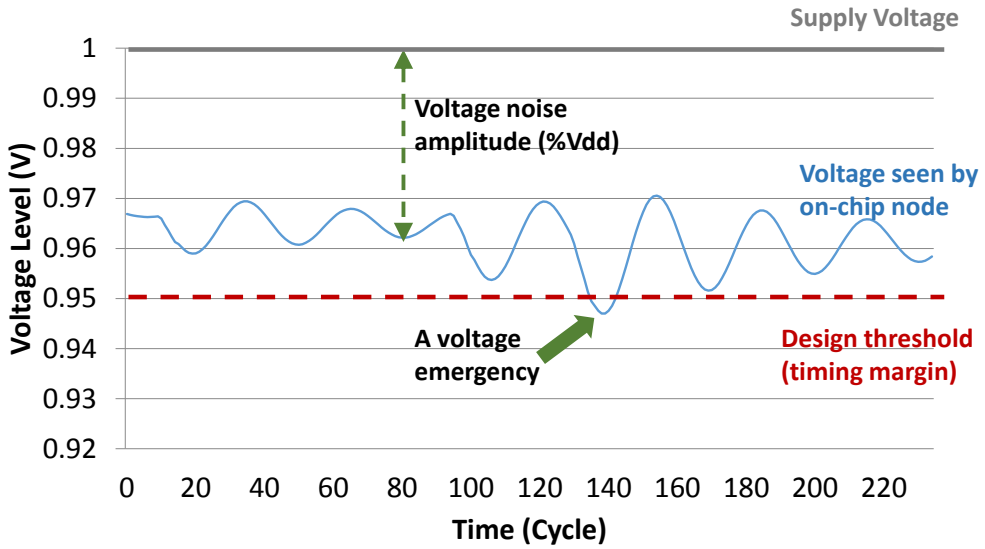

Figure 1.1: An illustration of voltage noise and the evaluation metrics.

a common practice for the circuit designers to reserve an extra voltage/timing margin on the critical paths to guard against large voltage noise [51]. Take timing margin for example: if we assume an x% noise amplitude will also increase the critical path's delay by x%, the circuit designers will need to allocation a timing margin of M% (i.e., slowdown the clock frequency by M%) to protect the circuit from a voltage noise of M% Vdd or smaller. If the voltage noise amplitude exceeds the designated margin (e.g., the short time period pointed by the green solid arrow in Figure 1.1), the delay of the critical path can exceed the clock period, which can cause a timing error and eventually lead to a soft error in the program's execution. These events of noise margin violations are also referred to as voltage emergencies. The count and frequency of such emergencies are also metrics we use in this dissertation to evaluate the noise quality of PDNs.

1.3 PDN Resource Scarcity and Optimization

Due to the silicon chip's high current consumption and high operating frequency, power supply network requires a significant amount of physical resources (e.g. metal wires, C4 pads, decoupling capacitors, etc.) to delivery stable voltage. Unfortunately, theses resources are expensive and phys-

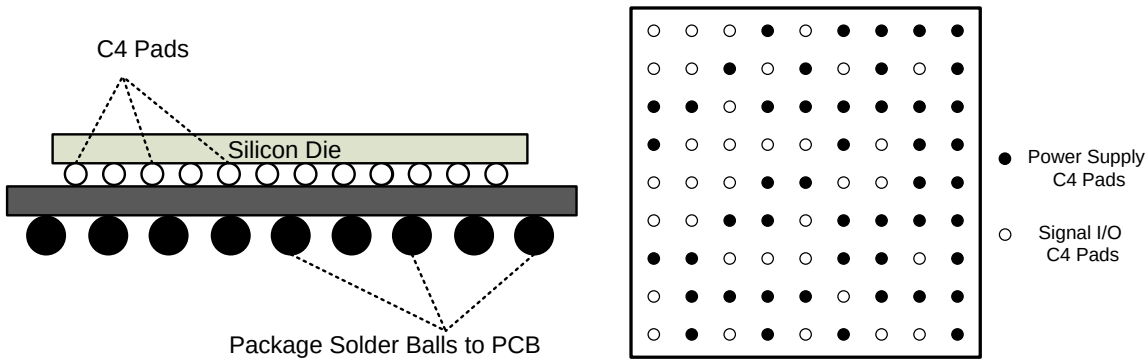

(a) Side view (b) Top view

Figure 1.2: (a) Side view of a chip showing its connection to the package through C4 pads. (b) Top view of a C4 array showing an example allocation of power pads and I/O pads (real processors have many more C4 pads than illustrated here).

ically constrained by fabrication technology limitations. For example, the C4 pads connect silicon die to its package electrically as well as physically, as illustrated in Figure 1.2. All the power/ground and I/O wires in the package are connected to these pads. However, the C4 pads do not scale with CMOS technology. In fact, according to the International Technology Roadmap for Semiconductors (ITRS) [32], pad density is predicted to remain constant for the foreseeable future, and pad composition and hence material properties (e.g. resistivity and maximum allowed current density) are improving slowly, if at all. According to Zhong et al., C4 pad location has a non-trivial impact on static supply voltage noise (IR drop) [106]. To optimize PDN quality under a given pad count, Sato et al. proposed an optimization algorithm for a pad ring structure [76]. For the more advanced flip-chip structure, Zhong et al. proposed an accelerated optimization algorithm based on simulated annealing (SA) [106] while Wang et al. proposed much faster solutions using on-chip voltage gradient to guide pad placement [90,91]. Wang's work also studied the impact of pad number on static voltage noise.

As another important component of PDN, the on-chip metal stack also has significant impact on supply noise. Mezhiba et al. characterized the electrical property of different metal layers within a on-chip PDN and proposed a compact multi-layer model for PDN simulation [60]. Jakushokas et al. derived a closed-form expression for on-chip inductance calculation [37] and proposed different approaches to optimize metal stack geometry [35, 36]. Decoupling capacitors are another example

of an expensive resource for PDN. On-chip decap helps significantly in stabilizing on-chip voltage fluctuation. However, most on-chip decap devices occupy silicon die exclusively, thus increasing the amount of decap will either reduce room for logic or on-chip memory, or increase total die area. To efficiently utilize the limited die area allowed for decap, both decap location optimization algorithms [47, 67, 68] and more efficient decap structures [11, 66] were explored by previous researchers. Similar to decap, the through-silicon-vias in 3D-ICs also occupy silicon area exclusively. Considering the fact that TSVs are the only channel for both cross-layer communication and power-supply connections, the allocation of TSVs introduces a tradeoff between the power delivery quality, the inter-layer communication bandwidth, and the area overhead. Prior work explored this tradeoff space by modeling TSVs with compact RLC elements [43, 95, 105], and optimizing the topology of TSVs [23, 86, 98].

The previous work not only reveals the scarcity of PDN resources, but also helps us to avoid drawing inappropriate conclusions due to sub-optimal PDN design. For example, in our pad-focused design-space-exploration, we adopt Wang's approach [91] to optimize location of power supply pads.

1.4 PDN Simulation and Modeling

Simulation is an essential step in the process of PDN design and verification. Due to its complexity and diversity, the power-delivery system is usually divided into several modules before being designed and verified separately. For example, in order to capture the complete frequency response of the metal planes in the chip package, package design teams usually use S-parameter models to simulate package PDN [83]. Meanwhile, the on-chip metal stack consists of parallel wires instead of 2D planes; therefore chip designers often use an extracted RC, or RLC netlist in their simulations [69]. Since the scope of this dissertation is the on-chip PDN, this section will primarily discuss the simulation and modeling of the on-chip metal stacks.

The three main steps of on-chip PDN simulation are: extracting metal grid parameters, estimating power consumptions, and solving equations. The goal of parameter extraction is to derive the

resistance, capacitance, and in some cases, inductance of the metal grid. Along with the topology of the metal wires, the extracted information describes the on-chip PDN and can be used directly to build matrices to solve the circuit [42]. Power consumption estimation can be performed at different granularities and different design stages. In the sign-off stage, where accuracy is the primary concern, detailed power values for individual transistors are calculated to minimize potential errors. For PDN simulations before placement and routing, power consumptions are usually estimated empirically based on each block's structure and anticipated activity. Although the empirical values are less accurate, they allow designers to estimate the power-delivery quality earlier in the design process. After extracting the circuit netlist and power values, designers create equations (usually in matrix form) and solve them numerically to get voltage and current results. Due to the complexity of the on-chip PDN, the dimension of the matrix can be several millions, or even billions. The problem of solving large linear systems efficiently has been a general research area of interest for many years. Various acceleration methods were proposed for PDN simulations ([7, 63, 100]). We will discuss our numerical method in Chapter 2.

Despite the efforts to speed up detailed PDN simulations, it still takes a significant amount of time (e.g., hours, or even days) to simulate the behavior the entire PDN over just a few nano seconds. In order to evaluate power-delivery noise at early design stages (e.g., pre-RTL), prior work proposed a lumped model that treats different PDN stages as single resistor-inductor (RL) or resistor-inductor-capacitor (RLC) pairs and lumps entire silicon die as one node. This methodology was widely adopted in previous PDN research and industrial practice [17, 40, 45, 71]. As processor frequency and power consumption scales up, with in die variation of voltage noise and current stress become significant [61]. To expose the spacial locality of on-chip PDN reliability issues, various distributed models were proposed to simulate PDN systems with different structures [19, 25, 48]. Regardless of their underlying physical design details, the basic methodology behind these distributed models is to model on-chip metal stack as a regular circuit mesh consisting R-only or RL pairs.

Although previous work has explored on-chip PDNs' reliability issues with distributed models, existing approaches either focus on circuit level details that make whole-application simulation infeasible, or utilize coarse grained models that ignore or over simplify the impact of PDN physical

design details like C4 pad number and location. To the best of our knowledge, a fine-grained yet fast PDN model with architecture level integration is missing from the literature. Therefore an important part of this dissertation seeks to build and validate a PDN model (VoltSpot) that is fine-grained enough to model individual C4 pads, and integrate it with other architectural level tools.

1.5 Run-time Voltage Noise Mitigation Techniques

To protect silicon chips from unacceptable supply voltage noise level, circuit designers put conservative timing margins to guard against worst-case scenarios (as illustrated in Figure 1.3a). However, due to the gap between worse-case and average-case, designing a constant timing guardband and the coordinate supply voltage level for the worst-case noise would reduce efficiency during average-case execution [21]. Moreover, as the severity of supply voltage fluctuation increases with the scaling of current density, the guardband required for worst-case protection would also increase, which will further constraint the slow scaling of supply voltage and operating frequency in near-future technology nodes.

To reduce energy overhead from over-provisioning design margins, researchers proposed two streams of techniques. One school of thinking seeks to dynamically adjust circuits' timing and voltage margin. For example, Lefurgy et al. [51] proposed a technique to detect available timing margin with critical path monitors (CPM) at run-time and use per-core fast digital phase-locked loops (DPLL) to adjust core frequency as well as to guide supply voltage level. Figure 1.3b provides a simplified example of this adaptation scheme. Another methodology manages to reduce voltage noise or the penalty of noise-induced errors. For instance, noise avoidance techniques [17, 40, 70, 71] gather run-time on-chip voltage/current information through sensors or counters and try to avoid noise by throttling or inserting dummy instructions right before large voltage swing happens. Another genre of techniques propose to react after voltage emergencies. By monitoring dangerous noise level with voltage sensors [20] or detecting noise-induced errors with built-in error detection units (e.g. parity check, ECC, etc.) [12, 21], these techniques recover from noise emergencies (Figure 1.3c). Compared with avoidance techniques, recovery techniques are more robust to sensor

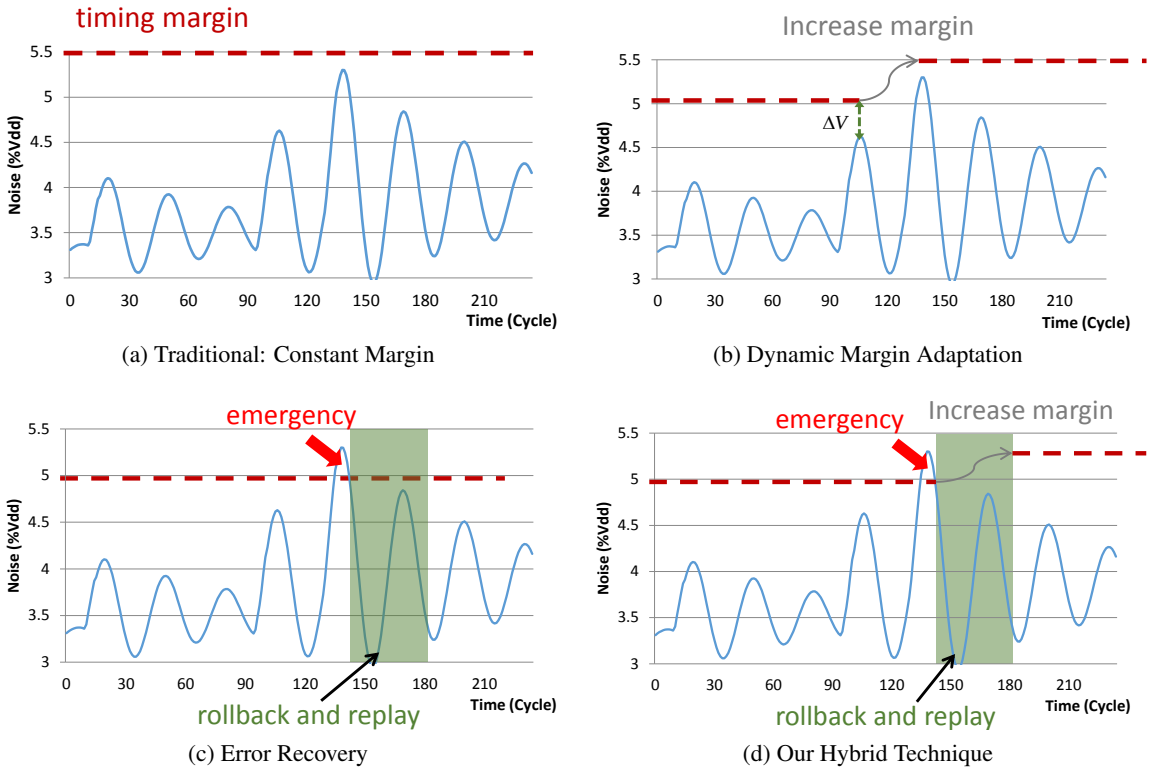

Figure 1.3: Illustrations of both static and run-time noise mitigation techniques.

delay and false alarms.

In this dissertation, we select a representative technique from each genre, as well as a hybrid technique that combines them together, to compare their performance with different PDN designs. Figure 1.3d illustrates the mechanism of our hybrid technique. Basically, it is capable of both recovering from noise emergencies, and dynamically adjusting timing margin. Chapter 3 will provide more details.

1.6 Electromigration and Long-term PDN Reliability Issues

Electromigration refers to the gradual mass transport in metal conductors induced by momentum transfer from electrons to atoms. Under the stress of high-density, uni-directional current, the metal atoms will be "pushed" by electrons toward the opposite direction of the current flow. Two undesirable consequences of EM are voids and hillocks in metal conductors, which will eventually break

the metal wire/pad (because of voids) or create shorts with neighboring wires (because of hillocks). Unlike noise-induced timing errors, which can happen in the time scale of several nano seconds, EM is a long-term effect that takes hours, days, or even years to accumulate its impact to a non-negligible level. A well-accepted model describes EM failure time with lognormal distribution [56] and the median time to failure (MTTF) is given by Black's equation [5]. Figure 1.4 presents both the log-normal distribution and the Black's equation, where J is current density, Q are material-specific constants, k is Boltzmann's constant, A is an empirical constant, and T is temperature in Kelvin. According to the MTTF equation, high current densities and temperatures speed up EM degradation and thus increase the likelihood of EM failures.

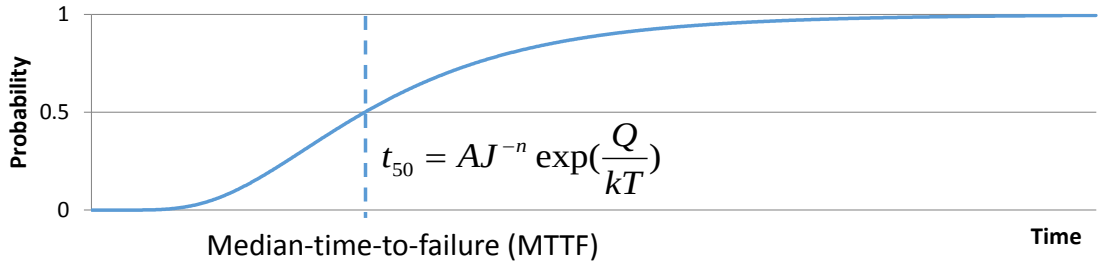

Figure 1.4: The log-normal CDF of a single conductor's EM-induced failure time. The y-axis shows the conductor's failure probability, or $Prob[Lifetime] \leq t$.

Even though modern PDNs deliver current with thousands of C4 pads and hundreds of thousands of on-chip metal wires, the sheer volume of current required by silicon chips along with limited cross-sectional area of metal conductors still create nontrivial EM threats. It is true that due to self-healing effects, pure bidirectional AC stressing current could extensively extend metal conductors' EM lifetime [55]. However, in PDN, especially power supply pads, bidirectional stressing current may not be pure AC waveforms with zero DC component. A more recent work [84] suggests that the metal lifetime under general AC current stress is determined by the DC component of the stressing current alone, as long as AC frequency exceeds a certain value (e.g. 10kHz). For this reason, we will use DC stress to analyze EM failures in PDN.

The failure mechanisms of a single copper wire [57, 89] or C4 bump [8, 97], as well as whole-system analysis of PDN grid [49, 87] or pad array [93] lifetime, have been extensively studied in the past. However, the architectural implications of PDN EM failure have not yet been discovered.

With our model VoltSpot, we study the mechanism of *multiple* pad failures and how it relates to pad allocation. Also, we seek to understand the performance impact of pad failures and explore potential solutions to mitigate the penalty of EM-induced failures at design- and/or run-time.

1.7 3D-IC and Voltage-Stacked Power Delivery

As the benefits of Dennard scaling (devices that are simultaneously smaller, faster and lower power) quickly vanish, it becomes increasingly challenging to keep the historical trend of device integration using the traditional planar fabrication approach. With recent advancements in semiconductor processing technologies, both homogeneous and heterogeneous systems (i.e. DRAM, SRAM, Flash, logic) were able to be designed and implemented in the same platform and stacked on top of each other [46]. This technology of three-dimensional integrated circuits not only shows a high potential to improve circuit performance [44, 85], but also provides a promising path to maintain the exponential growth in device integration.

However, in addition to the fabrication challenges, the tight integration of silicon layers also raises several fundamental technical difficulties. For instance, Todri et al. [86] showed that the upper-level silicon layers suffer from more severe voltage noise because they are further away from power supply C4 pads. On the contrary, the lower-level layers experience higher temperature because with convectional cooling system, the heat sink is attached to the top of the stack. As the number of physical layers in a 3D-IC stack is expected to increase in the future, the problems of delivering power to and removing heat from the 3D stack seem daunting. The main culprit is the fundamental mismatch between the volumetric (cubic) aspect of power consumption and dissipation in 3D-IC, and the fact that power *delivery* and affordable heat *removal* are limited to only the top or bottom 2D surface (quadratic). With the advance in the development of volumetric cooling technologies (e.g., micro-channel cooling [78]), power delivery may become an even more serious constraint.

To alleviate the power delivery constraints in the era of 3D-IC, prior work suggests using the idea of charge-recycling, or voltage-stacking (V-S) to build the power delivery network for 3D-

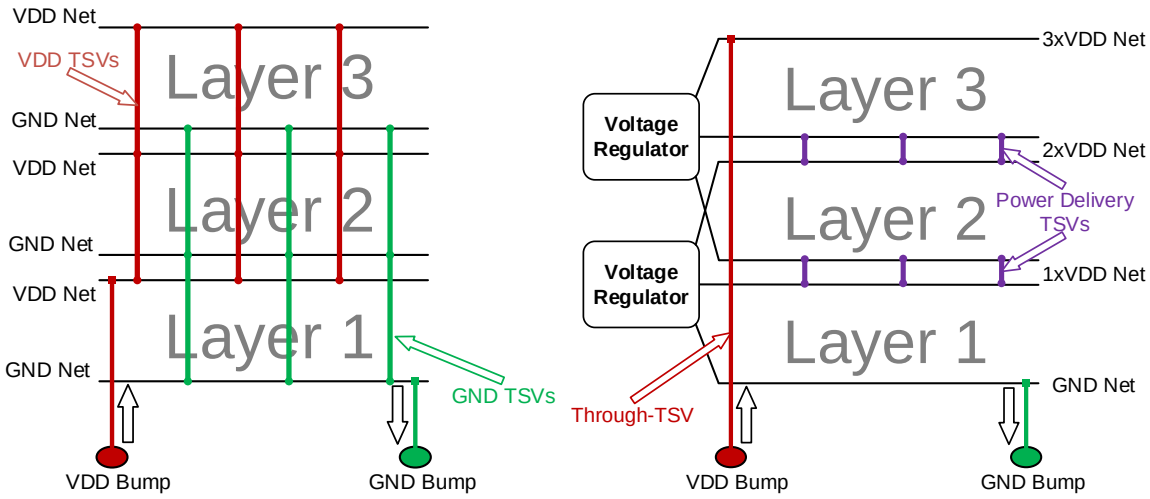

(a) Conventional PDN (b) V-S PDN

Figure 1.5: Conventional and V-S power-delivery scheme for 3D-IC.

IC [18, 34, 59]. Figure 1.5 illustrates both the conventional and the voltage-stacked PDN in an example 3-layer silicon stack. By linking all layers' Vdd and ground nets respectively with TSVs (Figure 1.5a), the conventional PDN parallelly connects the loads on different layers. Consequently, the off-chip current demand of a conventional-PDN-supported 3D-IC equals to the summation of all layers' current draw. On the other hand, the V-S scheme arranges the PDN in a way that the ground of one layer becomes the power supply connection for the next (Figure 1.5b). Within this serial connection, different layers shares the same current while their Vdd values are added. With the help of voltage-stacking's ability to "recycle" current between blocks, adding more layers to a 3D-IC only requires increasing the off-chip supply voltage while the current density within the PDN stays constant. For this reason, V-S provides a scalable solution to break the mismatch between 3D volume power dissipation and 2D surface power delivery.

However, V-S PDN incurs extra cost because in the general case when the currents of the various layers are not perfectly matched, explicit voltage regulation is required to constrain each layer's voltage noise below a designated threshold. While circuit solutions have been proposed for these explicit regulators [34,59], a cross-layer tradeoff study that examines the benefits of voltage-stacking's current reduction, the area overhead and power efficiency of explicit voltage regulation, and the supply voltage noise under different workload conditions, is missing from the literature. This disserta-

tion aims to build a PDN model for 3D-ICs that is capable of modeling both conventional PDNs and V-S PDNs. Based on this model, we perform detailed comparisons between the two power-delivery strategies and evaluate their noise, area overhead, and robustness against EM-wearout.

1.8 Dissertation Organization

The rest of this dissertation is organized as follows: Chapter 2 provides the design, implementation, and validation details of our pre-RTL PDN model, VoltSpot. Chapter 3 evaluates the trade-off between power-delivery quality and chip I/O communication bandwidth. We will show that with proper run-time noise mitigation techniques enabled, power/ground C4 pads can be aggressively replace with I/O pads with minimal performance impact from extra voltage noise. Chapter 4 builds a statistical simulation framework, analyzes the mechanism and consequences of multiple EM-induced C4 pad failures, and demonstrates that a mild increase in on-chip supply voltage noise guardband can tolerate pad failures and significantly increase system MTTF. By extending VoltSpot to support steady-state simulations for 3D-ICs, Chapter 5 shows that compared with conventional PDNs, V-S PDNs significantly improve EM-lifetime of C4 pads and TSV array. Chapter 6 evaluates and compares the transient voltage noise of both PDN schemes under different workload conditions and PDN design scenarios. We further prove that V-S is a practical, scalable, and affordable solution for 3D-ICs' power delivery challenge. Chapter 7 summarizes the lessons learned and discusses potential directions for future work.

Chapter 2

VoltSpot: A pre-RTL Power Delivery Network Model

2.1 Overview

The PDN of a modern processor usually consists of millions, or even billions of nodes, which require a significant amount of time for detailed simulation. In order to quickly explore the multi-dimensional space of PDN design and facilitate early-stage PDN optimization, we design, implement and validate *VoltSpot*, a pre-RTL PDN model for architecture-level PDN noise and reliability evaluation. VoltSpot utilizes a fine-grained grid model capable of capturing the relationship between PDN design details (e.g., C4 pad count and placement) and supply-voltage noise. Combined with other architecture-level tools, VoltSpot provides a versatile platform for investigating the effect of application- and time-dependent noise, evaluating design- and run-time noise mitigation techniques, and estimating vulnerability to lifetime-reliability problems such as electromigration. VoltSpot is implemented as a C library and is designed to be compatible with a variety of architectural tools. It is publicly available at `http://lava.cs.virginia.edu/VoltSpot`.

This chapter is organized as follows: Section 2.2 discusses the related work. Section 2.3 provides details about the key modeling methodologies. Section 2.4 elaborates our model's improvements beyond prior work. Section 2.5 validates VoltSpot against a power-grid analyze benchmark. Section 2.6 shows how VoltSpot interacts with other architecture-level tools. Section 2.7 summaries the chapter.

Most of the materials presented in this chapter are published in ISCA 2014 [103].

2.2 Related Work

Supply voltage noise comes from the intrinsic resistance, capacitance and inductance of the PDN. To reduce PDN impedance, which is directly related to voltage noise level, various levels of decoupling capacitors and voltage regulators was designated to compensate impedance over a wide frequency range. To capture the electrical characteristic of this hierarchical structure, a lumped model was proposed to simulate the low-frequency (PCB board), mid-frequency (package), and high-frequency (on-chip) response of PDN [69]. This model treats different PDN stages as single resistor-inductor (RL) or resistor-inductor-capacitor (RLC) pairs and it lumps the entire silicon die as one node. This methodology was widely adopted in previous PDN researches and industrial practices [17, 28, 40, 45, 71].

However, as processor frequency and power consumption scales up, with-in die variation of voltage noise and current stress become significant [61]. As a result, lumped models no longer suffice to provide sufficient insight into on-chip PDN. To expose the spacial locality of on-chip PDN reliability issues, various distributed models were proposed to simulate PDN systems with different structures [19, 25, 48]. Regardless of their underlying physical design details, the basic methodology behind these distributed models is to model on-chip metal stack as regular circuit mesh consisting R-only or RL pairs. This methodology has been adopted in the study of PDN design and optimization [47, 106] as well as architecture level voltage noise mitigation techniques [20, 24].

Although previous work has explored on-chip PDN's reliability issues with distributed models, existing approaches either focus on circuit level details that make whole application simulation infeasible, or utilize coarse grained models that ignore or over simplify the impact of PDN physical design details. In Chapter 3, we will show that certain PDN design details such as the number and location of C4 pads have a significant impact on the on-chip voltage noise, and therefore should be precisely modeled. To the best of our knowledge, a fine-grained yet fast PDN model with architecture level integration is missing from the literature. This chapter seeks to build and validate VoltSpot, and integrate it with other architectural level tools.

2.3 Modeling Methodology

2.3.1 Input/Output

As a PDN model, VoltSpot's main function is to profile a given power-grid design's voltage noise and current stress under certain workload conditions. Therefore, it requires three major inputs:

- Processor floorplan.
- Processor power-map.
- PDN configurations.

The floorplan input file describes the area, aspect ratio, and location of each block within the simulated processor. For the purpose of compatibility, it adopts the same text format from architecture-level thermal model HotSpot [81]. The power-map provides the power consumption of the processor. Like other pre-RTL tools, VoltSpot assumes that power density is uniform within each architectural block. This assumption can be relaxed by simply providing a higher-resolution floorplan and power map. Through the power-map input file, user can provide either the steady-state power consumption (i.e., one value for each block), or the dynamic power trace (i.e., multiple values per block) of the simulated processor. Finally, the configuration files specify the PDN's physical design details (e.g., metal stack geometry, decap density, C4 pad distribution, multiple voltage-domain setting, package impedance, TSV locations, etc.).

By default, VoltSpot outputs the per-cycle (in transient simulations), or a single (in steady-state simulations) maximum on-chip voltage noise value. Beyond that, it provides a rich set of built-in functions for the profiling and analysis of a wide variety of PDN statistics. For example, at the lowest level, VoltSpot can dump all nodes' voltage and/or all branches' current values at each simulated cycle. Give a timing margin, VoltSpot can also monitor each architecture block's maximum voltage noise over time and record whether or how often it experiences voltage emergencies. Users can easily extend these built-in functions to serve their own purposes.

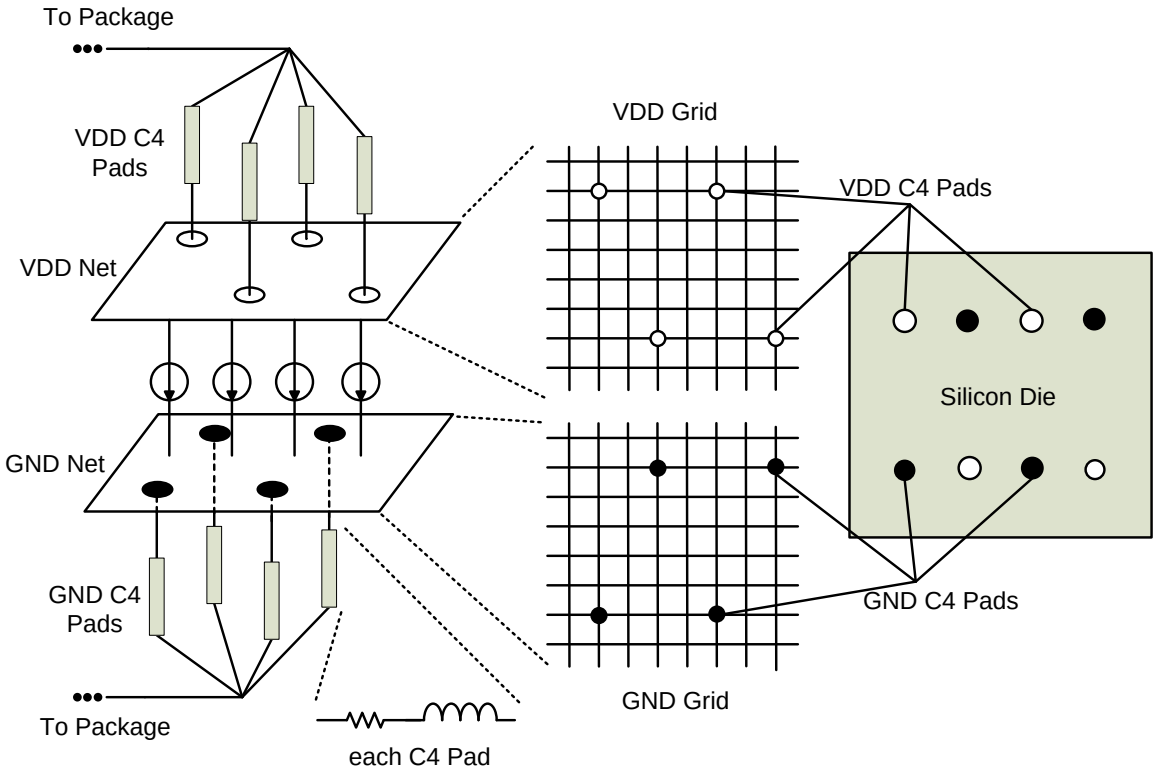

(a) On-chip model

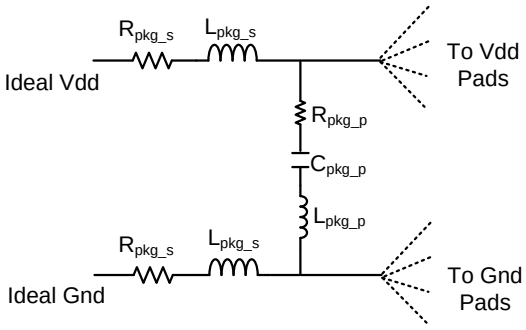

(b) Package model

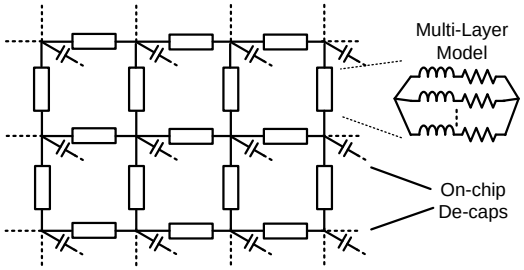

(c) Multi-layer on-chip grid model

Figure 2.1: VoltSpot model structure

2.3.2 Structure

Figure 2.1 illustrates the structure of VoltSpot. The major components are:

- On-chip Vdd/ground power supply metal stacks.
- C4 pads.

- Decoupling capacitors.
- Workloads.
- Chip package.

The typical regularity of the on-chip PDN's physical structure makes compact on-chip PDN modeling feasible. We adopt a well-accepted methodology [19, 24] that models the Vdd and ground nets as separate regular 2D circuit meshes. Section 2.4 will provide more details on the derivation of the resistor/inductor values. By collapsing the multi-layer PDN metal stack into a 2D virtual gird, VoltSpot effectively ignores the resistance and inductance of all the vias, which vertically connect the different metal layers. Compared with on-chip wires, vias have much lower impedance due to their size. We note that this abstraction significantly reduces the problem size and thus enables application-level noise simulation.

C4 pads are modeled as individual resistor-inductor branches attached to separate nodes in the on-chip virtual grid. With a high modeling granularity, VoltSpot can capture both number and locations of the C4 pads precisely. On-chip decoupling capacitors are modeled as distributed capacitors connecting the Vdd and ground grids. We currently assume that they are uniformly distributed within the die area, and it is straightforward to extend VoltSpot to support non-uniformly allocated on-chip decap. Ideal current sources model the load (i.e., the power of the switching transistors and associated leakage), and the current values are calculated as $I = \frac{Power}{SupplyVoltage}$. In transient simulations, we perform linear interpolation between adjacent power samples in order to avoid pessimistic LdI/dt noise estimations caused by the non-realistic step-functions in the power trace.

Since our main focus is the on-chip PDN, we model off-chip components such as the package with lumped RLC elements and assume the PCB provides an ideal power supply. VoltSpot therefore captures both mid- (package) and high-frequency (on-chip) responses of the PDN [69]. We note that Figure 2.1 focuses on 2D-IC. VoltSpot also supports 3D-ICs' PDN simulation, and we will discuss the related extensions in Chapter 5 and Chapter 6.

2.3.3 Numerical Engine

VoltSpot models PDNs with up to tens of thousands of RLC components. To solve such large-scale circuits efficiently and accurately at each time step, we choose the implicit trapezoidal numerical method, an A-stable method with 2^{nd}-order accuracy [9]. This method is the default ordinary differential equation solver in SPICE and it is also widely used in circuit and PDN simulations [73]. Since the solver's error monotonically increases with simulation time step, we set our time step to 50ps to keep the numerical error of node voltage below 10^{-5}V. To solve the fundamental equation of $AX = B$, we use an open-source sparse matrix solver, SuperLU [54], and optimize memory efficiency with multiple minimum-degree reorderings, significantly reducing fill-ins in sparse LU decomposition.

2.4 Improvements Beyond State-of-the-Art Models

VoltSpot makes two key improvements over alternative models in the literature:

- Transient PDN grid modeling at the granularity of pad pitch or smaller
- Multi-layer-metal PDN modeling using multiple, parallel RL branches.

This section discusses these improvements in detail.

2.4.1 Modeling Granularity

Due to their coarse modeling granularity, previous pre-RTL PDN models are incapable of either modeling pads in detail, or precisely capturing localized noise. Our study also show that, in some cases, even the finest grained on-chip grid from previous models (e.g., 12x12 [19]) underestimates the localized voltage noise amplitude by 20% and voltage emergency count by 3x. Alternatively, VoltSpot makes the on-chip grid size a function of C4 pad array size. For example, for a flip-chip design with 2,500 C4 bumps (e.g., distributed as a 50x50 array), the minimum modeling grid size will also be 50x50. It is left to the user's discretion whether to use a finer-grained model (at the price of reduced simulation speed) to attain higher precision. To calculate the resistance of on-chip

grid, VoltSpot uses the metal resistance equation $R = \rho * l/A$, where ρ is the resistivity per unit area of the metal, A is cross-sectional area, and l is wire length. For on-chip inductance calculation, VoltSpot adopts the equation from [36]:

$$L_{eff} = \frac{\mu_0 l}{N\pi}\left[ln\left(\frac{w+s}{w+t}\right) + \frac{3}{2} + ln\left(\frac{2}{\pi}\right)\right] \tag{2.1}$$

where N, μ_0, l, w, t, s are the number of power and ground pairs, permeability of the vacuum, length, width and thickness of a single wire, and the spacing between wires. In most cases, the number of on-chip metal wires (with in one layer) is larger than the dimension of VoltSpot's on-chip grid. To translate the impedance of single wires into the RL values of the branches in VoltSpot's on-chip virtual grid, we first calculate the total resistance and inductance of each on-chip PDN metal layer and then derive grid RL so that the total layer resistance and inductance in the modeled grid match the physical PDN:

$$Z_x = \frac{N_{rows}}{N_{metal\_rows}} * \frac{Z_{wire}}{N_{cols} - 1} \tag{2.2}$$

where $Z_x, Z_{wire}, N_{rows}, N_{cols}, N_{metal\_rows}$ are the impedance of horizontal branches in the on-chip grid, the impedance of a single wire, number of rows and columns in the on-chip grid, and the number of metal wires within that layer (assuming that the layer consists all horizontal wires). Z_y can be calculated similarly.

Grid Node Count : Pad Site Count	1:1	4:1	9:1	16:1	25:1
Number of Voltage Emergencies (5% Vdd Threshold)	9629	9285	9380	9616	9763
Number of Voltage Emergencies (8% Vdd Threshold)	804	713	706	720	712
Average Noise (% Vdd)	7.63	7.53	7.51	7.53	7.54
Maximum Noise (% Vdd)	11.83	11.65	11.59	11.59	11.58

Table 2.1: Impact of on-chip modeling granularity. At 25:1, the number of rows and columns in the on-chip grid roughly matches the number of top-level metal wires in the PDN stack. Among 1914 total C4 pads, 1254 pads were allocated to P/G.

Table 2.1 shows the relationship between modeling granularity and captured voltage noise (both violation count and noise amplitude). Using a 16nm, 16-core processor as evaluation platform and fluidanimate benchmark as workload, we observe that increasing PDN modeling granularity beyond

an average of four grid nodes per C4 pad site (i.e., using a 100x100 grid to model a chip with a 50x50 C4 array) improves precision by less than 5% (in terms of violation count). This is due to the limitation of the granularity of our power model, which assumes a uniform power distribution within architecture blocks. For these reasons, we set the grid-node-to-pad-ratio to 4:1. We also note that VoltSpot can handle finer-grained power input as long as user provide the floorplan and power map accordingly.

2.4.2 Explicit Multi-layer Model

To further improve accuracy, VoltSpot explicitly models the multi-layer structure of physical PDN. Previous work suggests that the electrical properties of on-chip metal layers heavily depend on their geometry (width, pitch, etc.) [60]. For example, the upper metal layers (which are usually wide and thick) have relatively high inductance and low resistance while the lower metal layers (which are usually narrow and thin) have higher resistance but lower inductance. As a result, high frequency current tends to go through lower layers while low frequency current favors upper layers. Therefore a single RL circuit thus cannot accurately describe the entire stack. We replace the single RL pair in on-chip grid with multiple parallel RL branches (as shown in Figure 2.1c) and calculate RL values based on the different metal layers' geometry.

# of Layers Modeled	2	4	6
Number of Voltage Emergencies (5% VDD Threshold)	28311	9285	8860
Number of Voltage Emergencies (8% VDD Threshold)	4929	713	681
Average Noise (% Vdd)	9.50	7.53	7.48
Maximum Noise (% Vdd)	13.63	11.65	11.63

Table 2.2: Impact of modeling multiple metal layers.

Table 2.2 illustrates both the importance and diminishing returns of explicitly modeling multiple metal layers. Using a metal stack with a similar geometry to [92], we sweep the number of modeled metal layers and evaluate voltage noise with the same platform processor and benchmark we used in Table 2.1. Our results show that the single RL pair model using values extracted from the top metal layers overestimates PDN inductance and reports a voltage noise amplitude 26% larger than

a multi-pair-branch model that considers four layers of PDN metal. Also, modeling six metal layers instead of four only improves precision by less than 1% (in terms of noise amplitude). Therefore we choose to model four layers of metal in the rest of our analysis.

2.5 Validation

We validated VoltSpot using an IBM PDN analysis benchmark suite [64]. The suite consists of detailed PDN structural information for six chips with different die sizes, silicon designs and metal layer counts. The PDN structure is given in SPICE format and includes metal wires' geometric properties and resistance. Other information, such as C4 pads placement, on-chip decap distribution and workload pattern, can also be extracted from the SPICE file. Besides the PDN structure, this benchmark suite also provides both steady-state and transient SPICE simulation results for all or selected on-chip nodes. As in VoltSpot, the loads in these benchmarks are modeled as ideal current sources.

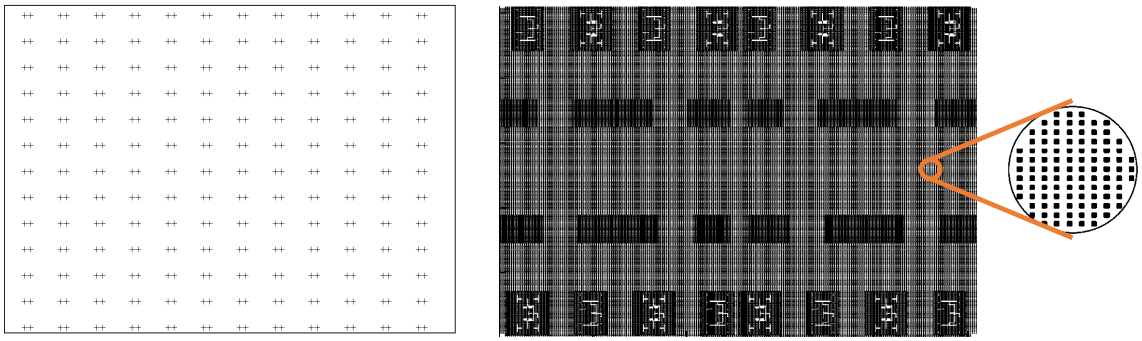

(a) VDD pad location, benchmark #4 (b) Current source position, benchmark #6

Figure 2.2: IBM power grid analysis benchmark: A snap shot.

The focus of VoltSpot is capturing within die current and voltage variation. We therefore evaluate VoltSpot's accuracy by comparing both simulated static C4 currents and transient on-chip voltage droops with those derived from the reference SPICE netlists accompanying the IBM benchmarks. Figure 2.3a visualizes the steady-state results comparison 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 2.3a 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. Figure 2.3b illustrates the dynamic voltage droop of an on-chip node in PG2. For each benchmark, IBM suite provides the dynamic voltage trace for 20 nodes that scattered across the silicon die. Among all the benchmarks, VoltSpot gives the largest error with PG2 and among all the 20 nodes in PG2, we select the one with the largest error to plot Figure 2.3b. Even with the worst-case node, VoltSpot still captures both the amplitude and the temporal variation of voltage droop with small deviation.

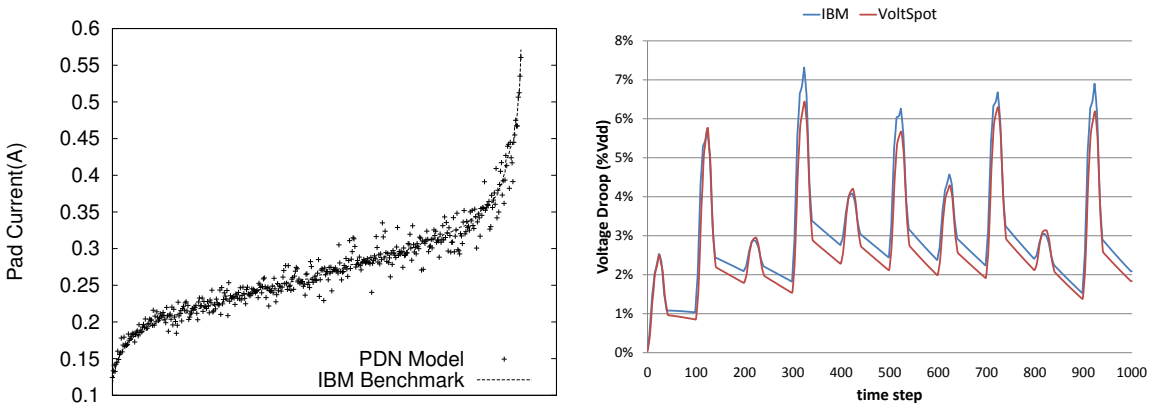

(a) Steady-state pad current comparison for PG3

(b) Transient voltage droop comparison for PG2

Figure 2.3: Current/Voltage comparison between IBM data and VoltSpot results.

For each benchmark (Bench), Table 2.3 shows the number of circuit nodes (# Nodes), metal layers (# of Layers) and power supply pads (# of Pads) and the validations results. The table also shows whether the resistance of vias is ignored in the SPICE model. Even though the variation of pad current within the same chip could be as large as 5x (observed in PG3), VoltSpot still accurately captures all pads' current with an average error of 5.2% (Pad Current Error). In transient validation results, Voltage Error Average shows the average node voltage mismatch across all simulated time steps and all given on-chip nodes, while Voltage Error Max Droop compares the max droops observed during the entire transient simulation. Both metrics give low error even for the bench-

marks that include detailed via information. The last metric, Voltage Error Correlation, shows the coefficient of determination (R^2) between VoltSpot results and IBM SPICE results. In statistics, R^2 indicates how well data points fits a statistical model. We borrow this concept to evaluate how well VoltSpot results agree with SPICE simulation results. Since $R^2 = 1$ means a perfect match, our results, which all exceed 0.96, indicate a strong temporal tracking of VoltSpot results. These results demonstrate that VoltSpot provides high-quality estimation for on-chip voltage fluctuation and that vias can be safely omitted from the model.

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

Table 2.3: Static and transient validation results against IBM benchmark

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

2.6 Interfacing with Other Architecture Level Tools

VoltSpot uses the same specification of the blocks and their power values as HotSpot [81], which allows straightforward integration with other architecture level tools. Figure 2.4 illustrates a pre-RTL PDN analysis tool-chain that combines VoltSpot with McPAT [53] (an power and area model),

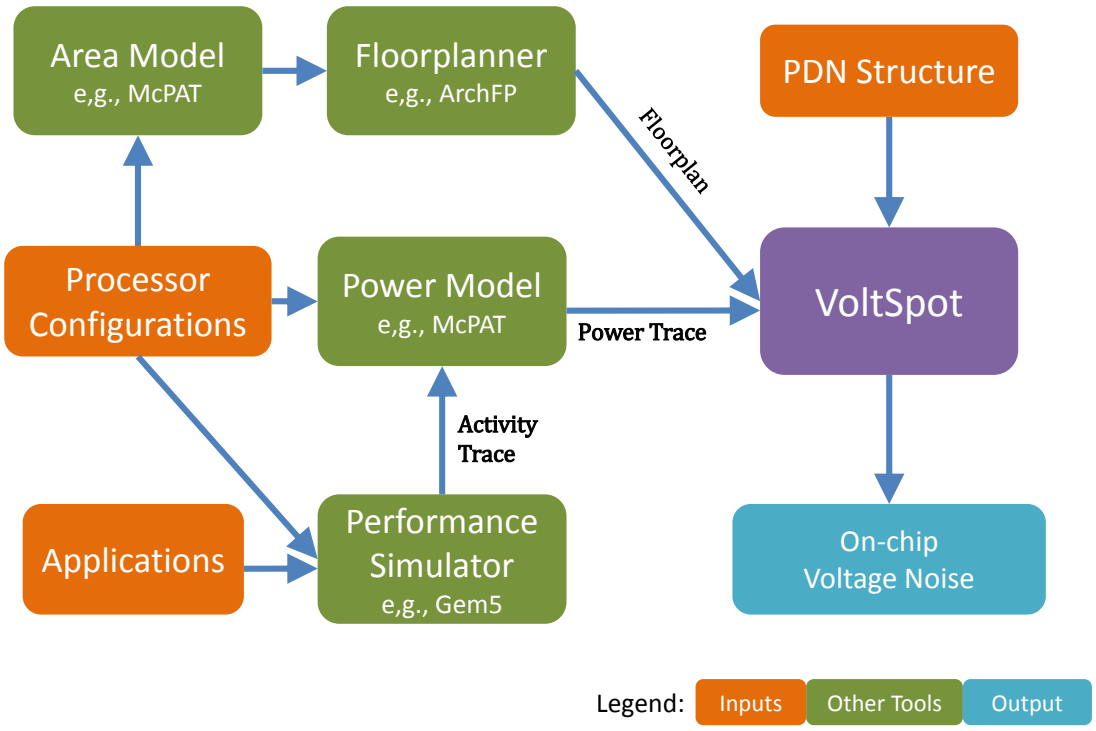

Figure 2.4: How VoltSpot interacts with other architecture level tools.

Gem5 [4] (a multi-core performance simulator), and ArchFP [13] (a pre-RTL floorplanner). The typical steps for architecture-level PDN simulations are:

1. Select a set of applications;
2. Configure the processor's architecture and its PDN structure;
3. Generate all architecture blocks' activity trace using the performance simulator;
4. Convert activity traces into power traces using the power model;
5. Calculate each block's area and generate the processors' floorplan using the floorplanner;
6. Run VoltSpot with the power trace, floorplan, and PDN configuration as inputs.

With the power and flexibility of these tools combined, users can study both supply voltage noise and EM-induced lifetime for a wide range of processor architectures, PDN designs, and applications.

2.7 Summary

Power delivery quality is becoming a limiting factor in multicore processor design. Consequently, it becomes increasingly important to evaluate whole-system PDN noise and the effects of design- and run-time mitigation techniques, to estimate the PDN's robustness against EM-induced wearout, and to jointly optimize the chip architecture and the power-delivery system. Such explorations are not possible with physical level modeling due to complexity, and higher-level models (e.g., lumped models) are too inaccurate.

In this chapter, we introduce our pre-RTL, fine-grained PDN modeling tool (VoltSpot), validate it against a power grid analysis benchmark, and combine it with other architecture-level tools for pre-RTL PDN analysis. VoltSpot makes it possible to accurately simulate application-specific noise, analyze electromigration wearout, and evaluate mitigation techniques. It in turn gives architects the opportunity to make critical, early design trade-offs (e.g., power supply pads vs. I/O pads). VoltSpot is implemented as a C library and is publicly available as an open-source tool.

Chapter 3

Architecture Implications of Pads as a Scarce Resource

3.1 Overview

Power delivery difficulties arise and will get worse for several reasons. First, despite advances that have produced sophisticated on-chip PDNs, the intrinsic resistance and inductance of the PDN circuit cannot be fully controlled with reasonable cost. This makes voltage fluctuation on the supply rails inevitable, threatening correct processor functionality with noise-induced short-term timing errors. Second, the PDN also suffers from long-term reliability threats such as electromigration (EM). EM results in permanent failures and directly affects chip lifetime and voltage stability. Finally, there is a contention between power delivery needs and processor computation needs: the only connections between a silicon chip and the outside world are the C4 pads, required by both power supply and chip I/O signal links. The demand for I/O is expected to grow exponentially with on-chip processor counts, while the demand for power supply pads also increases as current consumption scales up. Unfortunately, C4 density is expected to remain flat in the foreseeable future [32].

We therefore hypothesize that C4 pads are a scarce resource valuable to both circuit designers and chip architects, and provisioning of C4 pads should be exposed for architectural exploration. Using the pre-RTL analysis tool chain described in Chapter 2, we find that despite their integral role in the PDN, power-ground (P/G) pads can be aggressively reduced to their electromigration limit with minimal performance impact due to voltage noise—but *only* with a suitable noise-mitigation

strategy. The reason is that even though the number of voltage-noise events increases significantly, both due to technology scaling and further exacerbated by fewer P/G pads, the change in noise magnitude is small, and can be addressed with modest changes in noise mitigation and guardbanding.

The major contributions of this chapter are:

- We explore the impact of C4 pad configuration on power-supply noise with the assistance of cycle-accurate performance simulation and power modeling. Our results indicate that replacing some power pads with I/O pads only marginally increases the *amplitude* of supply noise, even though the *number* of noise events increases dramatically.

- We compare different state-of-the-art run-time noise mitigation techniques and observe that a hybrid mechanism that combines dynamic margin adaptation and noise-induced error recovery is the most robust to technology scaling and worst-case noise-inducing power viruses. With this hybrid technique, noise introduced by reducing the number of P/G pads can be mitigated with negligible overhead (1.5% slowdown). The key insight is that a very small increase in the default timing guardband eliminates problems due to the much greater frequency of small/medium noise events.

- We also study the effect EM-induced PDN pad failure has on transient noise. With reduced P/G pad count, earlier and more frequent PDN pad failures are expected. However, when appropriate noise mitigation strategies are available, performance degrades gracefully. EM ultimately limits the extent to which P/G pads can be reduced in favor of increasing I/O pads, because reducing the number of P/G pads increases the remaining pads' current.

These findings in turn enable a reduction in the number of P/G pads, with negligible performance overhead, in favor of a substantial increase in I/O bandwidth. This work is published in ISCA 2014 [103].

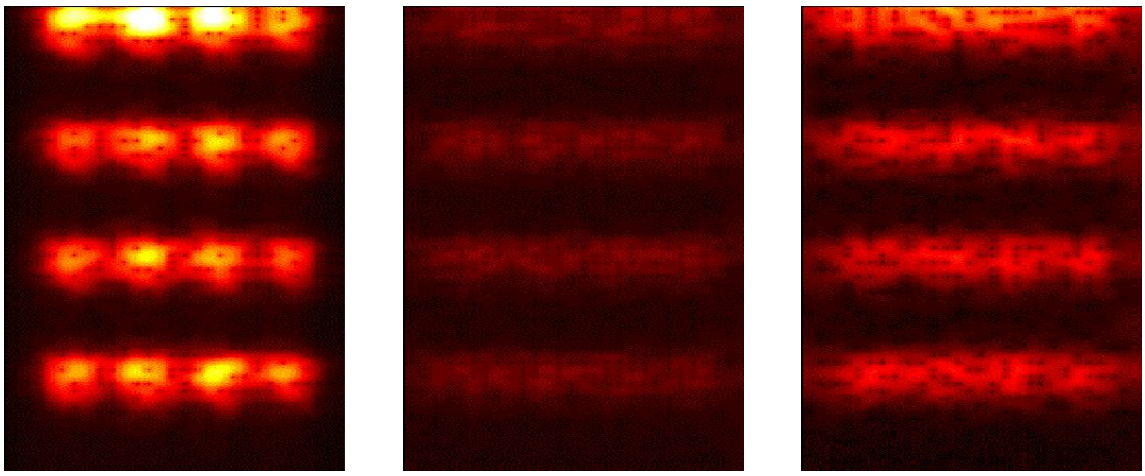

(a) 960 power supply pads with low quality placement

(b) 960 power supply pads with optimized placement

(c) 540 power supply pads with optimized placement

Figure 3.1: Voltage-emergency maps for different pad configurations for a 16nm, 16-core chip. A total of 1914 C4 locations are available. All three graphs have the same color scale; areas with warmer colors indicate more voltage violations.

3.2 Background and Related Work

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

C4 pads play a crucial role in power delivery. Due to growing power density, chips manufactured in future technology nodes (e.g.,16 nm) will require 150-200 Amp or more. However, the maximum feasible C4 pad density does not scale with CMOS technology [32], and pad composition and hence material properties (e.g., resistivity and maximum allowed current density) are improving slowly, if at all. If I/O pad count increases to keep pace with core count, there will be increasing competition between I/O and power-ground needs for scarce C4 pad sites.

As we discussed in Chapter 2, existing pre-RTL, architecture-level PDN studies often make simple assumptions about C4 pads, and none have considered the power/ground vs. I/O tradeoff.

They either use lumped PDN models that collapse all pads into a single resistor-inductor pair [17, 20, 71], or use coarse-grained models that are not able to accurately capture the effect of pad count and location [19,96]. Figure 3.1 illustrates the type of effect that pad count and location can have on on-chip voltage noise. For this figure, we simulate a 16-core processor with our fine-grained model running a PDN-stressing workload for 100K cycles. We count the number of cycles in which a voltage emergency (defined for this figure as a cycle with an average voltage droop larger than 5% Vdd) occurs on the chip. Figure 3.1a and 3.1b illustrate configurations with the same number of power supply pads, but only the latter case has optimized pad locations. Due to the sub-optimal allocation of pads, Figure 3.1a experiences 6x more emergency cycles compared with Figure 3.1b. Figure 3.1c illustrates a configuration with optimized pad locations, but 40% fewer pad count than in Figure 3.1b. Although the optimized locations in Figure 3.1c prevent extreme voltage-noise hotspots, the system still experiences up to 3x more emergency cycles than Figure 3.1c. Clearly, both pad count and pad locations have a large effect on on-chip voltage noise.

Other researchers have studied the effect of C4 optimization during pre-RTL design. For example, Wang et al. [90] evaluated the impact of pad placement on chip IR drop and proposed an optimization algorithm to minimize global IR drop. Zhang et al. [102] also used IR drop as a figure of merit and examined the impact of C4 pad count on technology scaling. Although the PDN models used by these works are fine-grained enough to precisely model C4 pads, they cannot simulate a PDN's transient behavior.

Runtime noise mitigation techniques have been proposed to either avoid excessive voltage noise [17, 40, 71], or recover from noise-induced errors [20]. To reduce the energy overhead of guarding against worst case, several proposals [2, 21, 51] dynamically adjust circuit timing margins to save energy during average-case execution while guaranteeing functionality in the worst case. We will evaluate both error recovery and margin adaptation methods and show the necessity of combining these into a hybrid strategy.

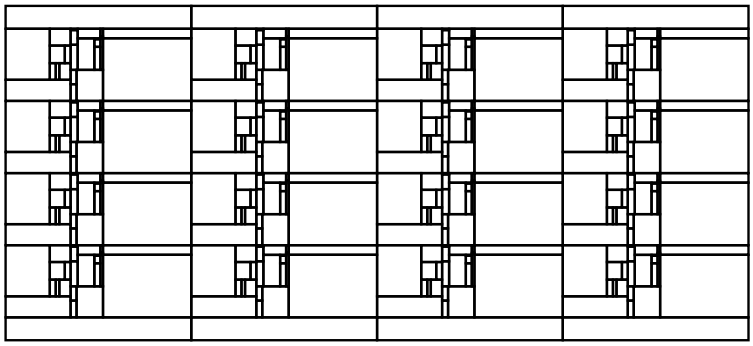

Figure 3.2: Floorplan of the Penryn-like 16core processor.

3.3 Simulation Setup

3.3.1 Multicore Scaling, Power Modeling and Floorplanning

To study the effect of technology scaling trends on PDN noise in the near future and explore the resulting trade-offs in architectural design choices due to power-delivery voltage-noise limitations, we create a series of multicore processor configurations scaled down to 16nm. For a reasonably modern configuration focused on single-thread performance, we chose a 3.7GHz 45nm Intel Penryn-like processor [16] as the baseline. It has two 32-bit 4-way out-of-order cores. Each core contains a 32kB L1 instruction cache and a 32kB L1 data cache. Unified L2 caches private to each core are each 3MB. For each technology node, we hold the processor architecture constant but assume that the number of cores (and therefore the number of L2s) doubles. For scalability and consistency, we assume a mesh network-on-chip (NoC). Figure 3.2 shows the floorplan of our 16nm, 16-core processor.

To get chip-wide and application-specific power consumption and area for all technology nodes, we use McPAT [53], an architecture-level power model, and integrate it with Gem5 [4], a multi-core performance simulator. Table 3.1 shows the area and peak power (including leakage power) results for our designs. We use ArchFP [13] to generate our floorplans.

Besides tolerating voltage noise, design margins also deal with application-induced thermal variation, process variation, aging, and test inaccuracy [51]. In this work, we only address the portion of margin needed for voltage noise and omit the margin for other sources of variations.

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

Table 3.1: Characteristics of Penryn-like Multicore Processors

3.3.2 Power Trace Sampling and Stressmark

To accelerate simulation, we borrow the idea of statistical sampling [22, 94]. The idea behind statistical sampling is to simulate in detail only short samples from much longer applications and estimate the behavior of the whole application based on the sampled segments. In performance studies, the accuracy of sampling heavily depends on how structures are "warmed up" before each sample, since the state of many architectural blocks (e.g., cache, branch predictor) accumulates over a long history and significantly affects performance [22]. Fortunately, warming up a PDN simulation is much simpler: the only factor that might affect results accuracy is the electrical charge in decoupling capacitors, which have short time-constants. Our study shows that 1000 cycles of warm-up at 3.7GHz is sufficient.

According to [94], a total of 2 million instructions (including instructions for warm up) sampled at equal intervals is sufficient to estimate a multi-billion-cycle application's IPC with $\pm 3\%$ error and 99.7% confidence. For this reason, we take 1000 samples at equal intervals from the end-to-end execution of simmedium inputs for 11 benchmarks in the Parsec 2.0 benchmark suite [3]. Two benchmarks (facesim and canneal) were incompatible with our performance simulation infrastructure and were omitted. Each sample contains 2000 cycles of per-cycle power information, of which the first 1000 cycles in each sample are used for warm-up.

We take two further steps to ensure that we simulate worst-case behavior. First, the sampled power traces are taken from 2-core simulations; for technology nodes with more than two cores, we replicate the 2-core power trace to 4, 8 or 16 cores. In this way, transient current fluctuation representative of that in Parsec benchmarks occurs simultaneously in each pair of cores, increasing

the stress on the PDN. Second, we construct a stressmark to simulate a voltage noise virus by selecting the most noisy (in terms of noise amplitude) power trace among all samples and replicating it 1000 times. We illustrate the stressmark's noise pattern in Figure 3.3.

3.3.3 PDN Parameters and Pad Location Optimization

Table 3.2 lists the major physical PDN parameters we use with VoltSpot. For on-chip metal, we use copper and adopt a metal layer structure similar to an Intel 45nm metal stack [92]. Table 3.2 reports the width (W), pitch (P), and thickness (T) of the global, intermediate, and local layers. For on-chip decoupling capacitors, we use deep trench capacitors, which have the highest capacitance density [66]. The die area allocated to on-chip decap is a design parameter and will be discussed in Section 3.5. We assume SnPb is the primary material for C4 pads. Typical pad diameter and resistivity are derived from [93]. Pad spacing was selected so that our pad density matches ITRS [32] projections. Package resistance, inductance and capacitance come from [28]. We performed sensitivity studies on package and pad parameters (e.g., SnAg pads) and observe that the effect of pad allocation (Section 3.4) or wear-out (Section 3.6) on voltage noise is insensitive to these variables: the impedance of the PDN's pad layer mostly depends on pad configuration.

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

Table 3.2: PDN Parameters

As mentioned before in Section 3.2, the number and locations of C4 pads have a significant effect on power delivery quality. To avoid sub-optimal pad allocation, we adopt the simulated-

annealing algorithm described in [90] and extend it to jointly optimize both Vdd and ground pad locations.

3.4 Transient Voltage Noise: Scaling and Effects of C4 Pad Configuration

3.4.1 IR Drop vs. Transient Noise

Power supply noise comes from three major sources: static IR drop, LdI/dt droop, and *LC resonance*. IR drop is the consequence of PDN resistance at any given time irrespective of processor behavior sequence. Instead, dynamic noise is triggered by certain chip behaviors such as a sudden change in power consumption, which will create large LdI/dt noise, or reoccurring power consumption patterns at or near an *LC resonance* frequency. Figure 3.3 compares an on-chip node's resistive drop (a function of instantaneous current and resistance only) and transient noise (from all three noise sources) over 1000 clock cycles. The plotted transient voltage noise is the result of transient PDN simulation which represents the actual voltage fluctuation at that on-chip node. The IR drop line was calculated as if the power consumption at each cycle were static and it represents the contribution of IR drop to overall voltage noise. We observe that IR drop only constitutes a small fraction of the overall noise, suggesting that considering IR drop alone (as have all prior studies of C4 pads) is insufficient in PDN design and optimization. For this reason, for the remainder of the chapter we will use aggregate transient voltage droop as the key metric for evaluating the quality of a PDN; IR drop will not be distinguished as a separate issue.

3.4.2 Scaling Trends of Supply-Voltage Noise

As manufacturing process technologies scale, current density grows as power density increases and supply voltage decreases. We compared maximum voltage droop and the number of voltage noise events across process technology nodes. Here we define a transient voltage droop greater than a certain threshold as a voltage-droop *violation*. We assumed a fixed PDN metal stack structure across technologies. This is because in our study, we focus on the upper layers of metal, where

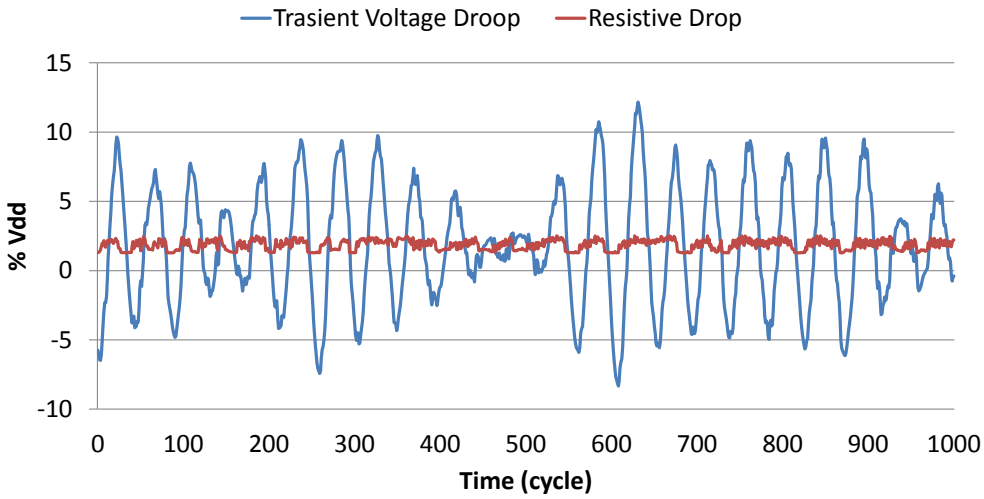

Figure 3.3: Comparison of transient voltage noise and static IR drop over a 1K cycle window in the benchmark ferret. Periodic oscillation implies that *LC resonance* is the major cause. We build our noise-inducing stressmark based on this segment.

wires are bulky and thus less affected by technology. Our sensitivity studies show that the impact of metal width on voltage noise amplitude is small ($\pm$50% metal-width changes max noise-amplitude by less than 0.5% Vdd. Section 3.4.4 provides more detail). Other parameters such as decap density and C4 pad density are also kept constant for fair comparisons between technologies. For this scaling limit study, we allocate all available C4 pads to power, to study the upper bound of PDN quality. The impact of C4 pad placement will be discussed in detail in Section 3.4.3. We simulate fluidanimate, one of the most noisy applications in the suite. The results of our experiment are summarized in Table 3.3.

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

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

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

Table 3.3: Voltage Noise Scaling Trend, Ideal (all pads allocated to power/ground), in fluidanimate benchmark

because a static guardband must countenance worst-case behavior.

3.4.3 Voltage Noise and Pad Configuration Effects

Our primary interest in this work is the contention between power supply and chip I/O; it is therefore important to quantify the effect of C4 pad allocation on power supply noise. Since power noise increases rapidly with technology scaling, we henceforth focus on the 16nm node and use a 16-core Penryn-like processor as the platform for our analysis. To better illustrate the trade-off between I/O pad count and processor off-chip memory bandwidth, we assume all on-chip memory controllers (MC) are single-channel and convert the number of available I/O pads into the number of supported MCs (different MC channels cannot share C4 pads [33]); more MCs means fewer Vdd and GND pads, and vice versa.

	Intel i7-900 Extreme [30]	Intel Xeon E7-8800 [31]	IBM Power7 [80]
Memory Controller	121x3	27x4	76x4
Inter-Chip Links	85x1	84x4	80x5
Misc	85	78	64
Total	533	522	768

Table 3.4: I/O pad breakdown for commercial processors.

Table 3.4 provides detailed pad requirements for three commercial processors. Based on these pad requirement breakdowns, we assume that our 16-core chip includes four inter-chip links (85 pads each) and a total of 85 miscellaneous pads (including clock, dynamic voltage scaling control, sensing, debug, testing, etc.). We assume all MCs are FBDIMM interfaces (as in the Xeon E7-8800) to represent the trend toward narrower, more serial interfaces; FBDIMM requires far fewer

pads (about 30 per MC channel) than DDR3 (80-120 pads per channel). The total number of C4 pads for our 16nm chip is 1914; all pads not used for I/O are dedicated to power supply.

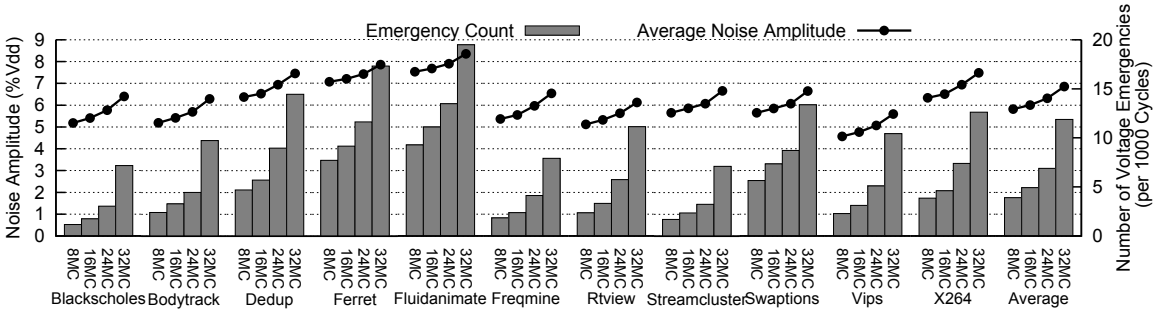

Figure 3.4: Voltage noise change across different pad configurations. Each MC needs 30 I/O pads.

Figure 3.4 shows the relationship between supply voltage noise and I/O configuration. We evaluate voltage noise with two different metrics: noise violation count and noise amplitude. The violation rate bars assumes a 5% voltage-droop violation threshold and report the number of noise violations (averaged across all 1000 samples). The lines in Figure 3.4 illustrate the maximum observed voltage noise (averaged across all samples). Independent of the application, increasing the number of MCs worsens voltage noise in terms of both violation event rate and noise amplitude, though noise amplitude increases only marginally (up to 1.5% Vdd).

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

3.4.4 Sensitivity Study

3.4.4.1 Impact of Package Parameters

Package Resistance (R_pkg_s) ($m\Omega$)	0.1
Package Inductance (L_pkg_s) (pH)	1
Package Resistance (R_pkg_p) ($m\Omega$)	1.5
Package Inductance (L_pkg_p) (pH)	1
Package Capacitance (C_pkg_p) (μF)	250

Table 3.5: Package parameters extracted from Intel Xeon 5500 series processor [29].

VoltSpot utilizes lumped RLCs to model package components and our previous analyses assume a set of fixed package parameters. To explore the sensitivity of our results against chip package, we extracted a different set of parameters from the package designed for Intel Xeon 5500 series processors [29] and configured VoltSpot accordingly. Table 3.5 lists the new parameters. After simulating Parsec benchmark with different pad configurations, we examined the relationship between pad allocation and transient voltage noise (Figure 3.5). Compared with the previously used package, the new package has much larger capacitance which helps to reduce the overall impedance of the entire PDN. As a result, both average and maximum (not shown in Figure 3.5) voltage noise amplitude are lower than the previous case. On the contrary, switching to the new package significantly increases the number of voltage emergencies with %5 Vdd as threshold, indicating a more severe supply voltage fluctuation (with smaller amplitude). Even though the new set of package parameters have contradictory impact on noise amplitude and emergency count, our major observation in Section 3.4.3 that increasing the number of MCs only marginally increases the noise amplitude still holds true: from 8MC to 32MC, average noise amplitude increases by only 1.5% Vdd. We therefore conclude that the impact of pad allocation on noise amplitude is insensitive to package parameters.

3.4.4.2 Impact of On-Chip Metal

The on-chip metal stack plays a vital role in distributing current from C4 pads to transistors as well as maintaining the spatial uniformity of supply voltage. Wires for both signal routing and

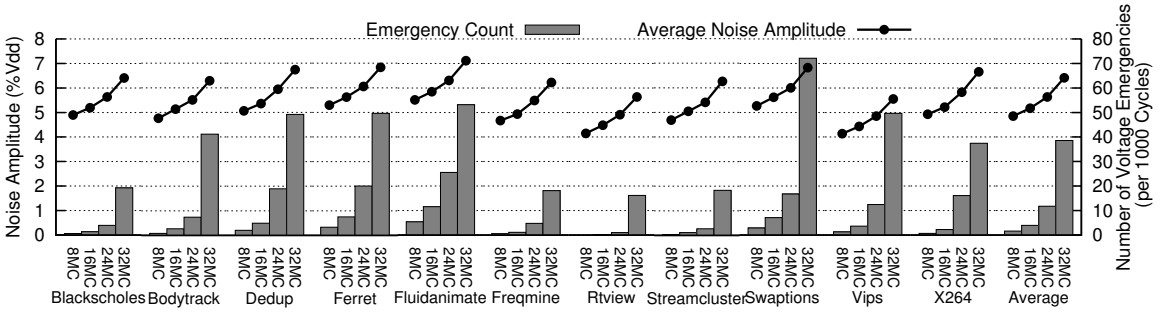

Figure 3.5: Voltage noise change across different pad configurations under a different set of package parameters.

power delivery share the metal stack and therefore although increasing the width of PDN wires helps to decrease on-chip resistance, it also increases the complexity of on-chip signal routing. As an exploration to our results' sensitivity against metal allocation, we simulated fluidanimate with our baseline processor and tested a range of metal occupancy. Here, the occupancy of a metal layer is defined as the width of PDN wires divided by the metal pitch of that layer. Note that since VoltSpot focuses on top level layers, we assume that all wires are allocated for power supply. This assumption can be easily relaxed by increasing the effective pitch (according the ratio between signal wire count and power supply wire count within that layer) used in the calculation of on-chip grid RL values. The simulation results (listed in Table 3.6) shows that although increasing the metal layer occupancy (i.e., the width of power supply wires) reduces both noise amplitude and violation count, the absolute noise reduction is fairly insignificant (doubling occupancy from one third to two thirds only reduces average noise amplitude by 3%). This indicates that the impedance of on-chip metal wires only consists a small portion of the entire PDN's impedance thus reducing stack impedance alone would not significantly suppress on-chip voltage noise.

Metal Layer Occupancy (%)	16.7	33.3	50	66.6
Violations (8% Threshold)	4792	3253	2826	2599
Average Noise (% Vdd)	7.99	7.53	7.37	7.28
Maximum Noise (% Vdd)	11.94	11.66	11.60	11.56

Table 3.6: Supply voltage change due to metal layer occupancy. The majority of this chapter assumes a 33.3% layer occupancy. The evaluation platform is our 16core, 8MC processor.

3.5 Run-time Voltage Noise Mitigation

VoltSpot makes it possible to evaluate run-time noise mitigation strategies in the context of architectural design decisions. In this section, we evaluate the performance overhead of several run-time mitigation techniques, and the effect of the number of integrated memory controllers. Increasing the number of memory controllers improves performance by reducing memory access latency. However, pads for I/O must be taken from those budgeted for power supply, increasing the demands on (and performance overhead of) noise mitigation. For fair comparisons between (i) dynamic margin adaptation and (ii) error recovery techniques, we present all results in terms of speedup. We assume a constant supply voltage and adjust the timing margin only. Since there is a roughly linear relationship between supply voltage droop and circuit delay within a reasonable range [74], we assume a voltage droop of $X\%$ Vdd increases circuit delay by $X\%$.

3.5.1 Dynamic Margin Adaptation

As an alternative to wasteful, fixed margins, Lefurgy et al. [51] propose a technique to detect available timing margin with critical path monitors (CPM) at run-time and use fast digital phase lock loops (DPLL), capable of reducing frequency by 7% within 5ns, to increase timing margin until the integral frequency control loop catches up and raises voltage as needed. Note that this emergency response to voltage droop is a one-shot control, immediately dropping frequency by the maximum (7%) needed to preserve correct operation in the presence of the worst possible voltage droop. However, voltage could continue to drop rapidly before the DPLL can adjust, so the processor must always underclock (relative to the current voltage) by enough to preserve correct behavior given the worst-case voltage slew rate in 5ns. The magnitude of this guardband has not been made public.

In this work, we only focus on performance. We adapt the above methodology by assuming that the voltage remains fixed at the maximum, and only clock speed is reduced in response to voltage droop.[1] We assume ideal voltage sensing in each core, and per-core DPLLs to respond to per-core voltage-droop behavior. We observe that most applications exhibit phases of low and high voltage

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

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

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

We determine the necessary safety margin (S) as a function of technology node (using a brute-force search). Both S and the worst-case margin grow significantly from 45nm to 16nm, as shown in Table 3.7. We observe that the required safety margin increases by almost 2%, significantly reducing the performance benefits of margin adaptation: on average, the average portion of the 13% worst-case margin that can be removed for performance improvement shrinks from 27% to 9%. The problem is that margin adaptation must be very conservative to guard against potential (but rare) worst-case voltage droops. Since margin adaptation only removes margin during low-noise program phases, we choose fluidanimate instead of our stressmark (Section 3.3.2) for this analysis, otherwise the margin controller could not reduce margin at all due to the stressmark's constantly noisy behavior.

In these simulations, we hold the PDN design (power supply pads density, on-chip decap density, metal structure, etc.) constant across technologies. It is possible to reduce voltage droop slew rate and margin adaptation safety margins by adding more on-chip decap as technology scales. However, our design space exploration study shows that, to keep the 16 nm chip's performance

overhead on a par with that of 45 nm chip, at least 15% more die area must be allocated to decap, a cost equivalent to two cores.

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

Table 3.7: Dynamic Margin Adaptation and Scaling

3.5.2 Recovering From Noise-Induced Errors

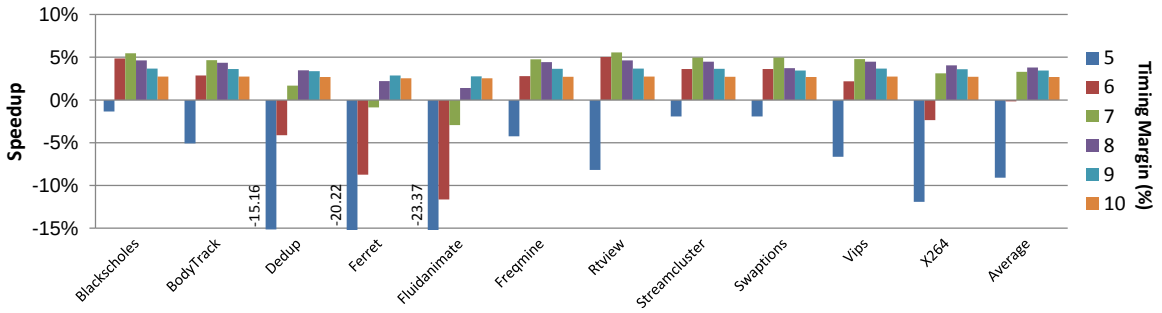

Figure 3.6: Speed up of recovery-based techniques with different timing margin settings. We evaluate our 16nm, 16core processor with 24MCs here. The baseline case enforces a 13% timing margin and thus guarantees timing is error-free.

An alternative to dynamic margin adaptation is to roll back and recover when a timing error is detected [20]. Such techniques address a key weakness of dynamic margin adaptation: false positives where noise reduces timing margin but would not ultimately cause an error. To analyze the overhead of noise recovery, we first simulate benchmarks to completion and collect noise amplitude data. Then, we perform post-processing to determine, given an allowed voltage droop and recovery overhead in clock cycles, the total performance overhead in cycles. Figure 3.6 shows the performance effect of different timing margin settings. Using a design with constant margin of 13% as our baseline (no recovery needed), we evaluated recovery-based methods' performance with different benchmarks on our 16 nm, 16-core chip with 24 memory controllers.

We observe that as we remove timing margin, the processor runs faster but also experiences more errors. As a result, the recovery penalty associated with removing too much margin over-

whelms the benefit of increased clock frequency. In extreme cases (e.g., fluidanimate with only 5% margin), aggressive margin settings introduce so many errors that it significantly hurts processor performance. We assume here that each error recovery requires 30 cycles (rollback 10 cycles and replay at half frequency [20]) and observe that on average, 8% timing margin gives the best performance.

3.5.3 A Hybrid Technique

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

Figure 3.7 depicts a performance comparison between the techniques discussed above. The chip evaluated here is also a 16-core, 24 MC processor. "Ideal" bars represent the performance gain achieved by an oracle margin controller that always maintains the minimum required margin without causing any timing errors. As a sensitivity study, we explored three different rollback penalties for the recovery technique. Using the analysis in Figure 3.6, we select the optimal timing margin setting for each rollback penalty assumption. As we expected, the margin-adaptation-only technique has lower speedup compared with the recovery-based technique. We also observe that the recovery-based techniques' performance is minimally sensitive to rollback penalty. This is because, with proper margin setting, errors happen rarely, thus spending more cycles on each error would not significantly degrade overall performance. In contrast, the performance of the hybrid technique is much more sensitive to error recovery overhead, because it relies on the occurrence of errors

to trigger margin adjustment, and therefore experiences more errors than a well-tuned recovery technique. The average results from the Parsec benchmark suite show that the hybrid technique only barely outperforms the recovery-only technique when recovery cost is low.

While evaluating mitigation techniques' performance with our stressmark, we observe that recovery-only experiences significant slowdown (right most bars in Figure 3.7; note that the stressmark's performance was excluded from Parsec average calculation). This is because, in order to achieve optimal performance with normal applications, we relax timing margin toward the average case. However, since the stressmark constantly excites PDN resonance, it will experience frequent errors (12 per 1000 cycles) under this tight margin setting, and thus makes the recovery technique suffer from frequent recovery penalties. In contrast, the hybrid technique can quickly adapt the margin to the required level at the beginning of the stressmark and avoid all remaining errors. In summary, although recovery-only performs better with typical workloads even when the recovery cost is high, the hybrid technique is more robust to worst-case behavior. If guarding against worst-case noise is the priority, designers should choose hybrid over recovery-only techniques.

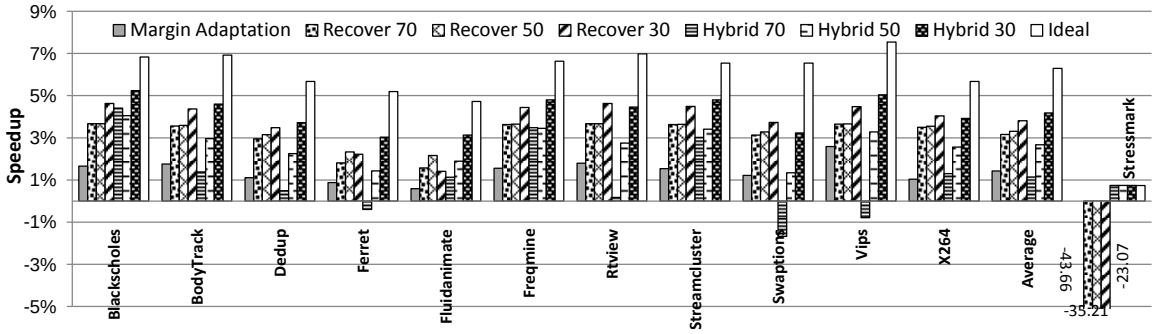

Figure 3.7: Performance comparison between different noise mitigation techniques. Numbers after recover/hybrid in legend represent the cost (in cycles) of each recover, from error. The baseline case enforces a 13% constant timing margin.

3.5.4 Trading Power Pads for Performance

Our primary interest is the tradeoff between power supply pads and I/O pads. Section 3.4.3 points out that trading power/ground pads for I/O connections will degrade power delivery quality. Based on the preceding discussions, we choose the hybrid noise mitigation technique and assume a pes-

simistic per-error recovery cost of 50 cycles. Figure 3.8 shows the pad-induced noise mitigation overhead for different applications. Each application uses its own performance with the 8 MC case as the baseline. While increasing MC count proportionally improves performance, it also exacerbates supply voltage noise. Our results indicate that the performance penalty to mitigate the extra noise is fairly low, even if we aggressively increase chip MC count from 8 to 32, reducing power/ground pad allocations from 1254 to 534. The reason is that even though the number of voltage-noise events increases significantly, the change in amplitude of most of these events is small and modest changes in noise mitigation and guardbanding can address these more frequent noise events.

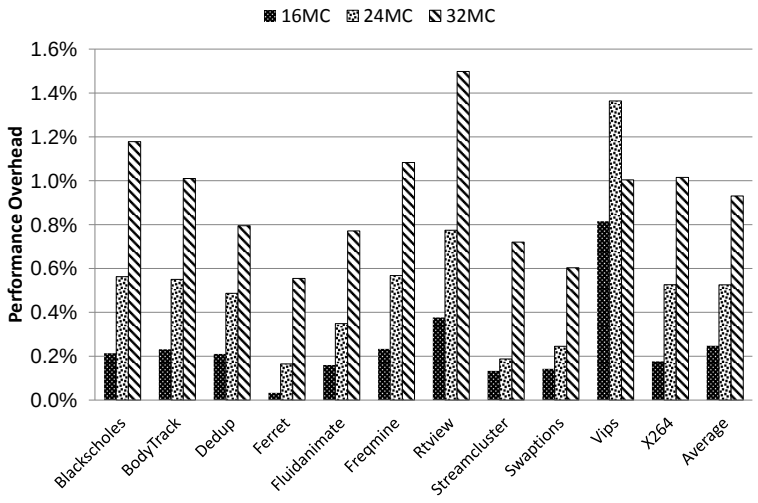

Figure 3.8: Performance penalty of mitigating extra voltage noise caused by reduced power/ground pads. We use the hybrid technique with a conservative assumption of 50-cycle rollback-and-replay penalty per error. Each benchmark uses its own performance with 8MC case as baseline.

3.5.5 Impact of I/O routing

One side-effect of increasing the number of I/O channels is more complicated package I/O routing: lateral I/O wires could "cut" through package metal layers and split power delivery planes into separate islands. As a result, the impedance of package PDN will increase. We performed a first-order analysis of this effect by increasing the impedance of the package's serial portion (R_{pkg_s} and L_{pkg_s} in Figure 2.1b). This adjustment emulates the change of the impedance between C4 pads and

the power supply from PCB board. Table 3.8 shows the impact of this impedance change on power supply noise. Using our 16-core processor with 32 MC as evaluation platform, we increased both R_{pkg_s} and L_{pkg_s} by up to 100% and observe that the amplitude of on-chip voltage droop varies by less than 0.2% Vdd. This is because that although larger R_{pkg_s} increases the PDN's impedance at low frequency, it also helps to reduce its impedance at resonant frequencies due to damping. With a proper run-time technique, this mild noise amplitude increase could be handled with low overhead. With the extension of distributed package modeling, VoltSpot could be used to perform more detailed analysis of I/O routing's impact on PDN impedance—an interesting area for future work.

% Increase of Package Impedance	0	25	50	75	100
Violations (8% Threshold)	1736	1840	1935	2045	2138
Average Noise (% Vdd)	8.36	8.39	8.42	8.44	8.47
Maximum Noise (% Vdd)	11.75	11.83	11.92	11.87	11.76

Table 3.8: Impact of package impedance. We increase the serial portion of package impedance as an emulation to the consequence of more complicated I/O routing in chip package.

3.6 Chip I/O Bandwidth and C4 EM Lifetime

Electromigration failures are the result of continuous, high-density current flow. It has been shown that a metal conductor's lifetime under pure AC current is much longer than DC lifetime [55]. However, in real circuits, especially power supply pads, bidirectional current may not be pure AC waveforms. More recent work [84] suggests that the lifetime of metal under high-frequency AC current stress is determined by the DC component of the stressing current alone. We therefore focus on DC stress only and use the processor's peak power consumption for worst-case analysis. To be more specific, we use 85% (a ratio suggested by [102]) of theoretical peak power reported by McPAT as a stressmark power input. Although I/O pads are also vulnerable to EM, changing the number of power/ground pads would not affect per-I/O pad current. For this reason, we do not evaluate EM's effect on I/O pads.

3.6.1 Whole-Chip Electromigration MTTF Calculation

For a single C4 pad, EM-induced failure times follow a log-normal distribution with probability density function (PDF):

$$f(t) = \frac{1}{\sqrt{2\pi}\sigma t} exp\left(-\frac{(ln(t/t_{50}))^2}{2\sigma^2}\right), \tag{3.1}$$

where σ is the standard deviation (determined experimentally to be 0.5 [56]) The median time to failure is determined by Black's equation [5] and adjusted to consider current crowding and Joule heating [8]:

$$t_{50} = A(cJ)^{-n} exp\left(\frac{Q}{k(T+\Delta T)}\right) \tag{3.2}$$

where J is current density, n and Q are material-specific constants (for SnPb solder bump, $n = 1.8, Q = 0.8eV$ [38]), k is Boltzmann's constant, $c = 10$, $\Delta T = 40°C$ [8], A is an empirical constant, and T is temperature in Kelvin. Given the MTTF, we can calculate the probability of failure, $F(t)$, for a single pad after any time t.

Only using the failure probability of a single pad, even the pad with highest current density, to estimate whole chip's robustness against EM is insufficient because all pads are threatened by EM. Any pad could potentially fail before any others. To quantitatively evaluate the chip's EM lifetime, we derive a new cumulative distribution function that describes the distribution of time t when the first PDN pad failure occurs:

$$P(t) = 1 - \prod_i (1 - F_i(t)) \tag{3.3}$$

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

Table 3.9 shows the scaling trend of average on-chip current density, single-pad worst current and both MTTF and MTTFF for C4 pads. MTTF/MTTFF results are normalized to the 45nm MTTFF value. When we examine the whole chip's robustness against EM, MTTFF is much worse than the worst single element's expected lifetime. For example, if every single pad in a 45nm chip

were designed to have a 10 year MTTF under the worst case, the median time to first PDN pad failure in the entire chip would be around 3.4 years. As technology scales, a power delivery system designed to work 10 years at 45nm would only be EM-failure-free for about 2.4 years at 16nm. All our calculations assume a temperature of $100°C$ to represent the worst-case scenario.

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

Table 3.9: C4 Pad EM Lifetime ScalingTrend

3.6.2 Tolerating Pad Failures with Runtime Mitigation

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

To evaluate PDN pad failures' effect on voltage noise, we remove power pads from the previously studied 16nm, 16-core chip and simulate the EM-damaged chips with the benchmark fluidanimate. Since EM-induced failure is a stochastic event, any power supply pad could fail at any time. Thus the number of possible pad configurations is huge (e.g., for a chip with 1200 power supply pads, there are more than 10^{43} ways to have 20 failures). It is impossible to traverse all failure combinations, so as a "practical worst case" estimation, we choose to fail the pads with the highest current density. There are three reasons behind this methodology. First, a single pad's MTTF is inversely related to current density, so pads with the highest current density tend to fail first. Second, pads carrying high current are closer to blocks with high power consumption and those blocks (e.g., the ALU) are more likely to produce large voltage noise. Therefore, removing supporting pads for those blocks introduces more voltage fluctuation and thus gives a better estimation for the

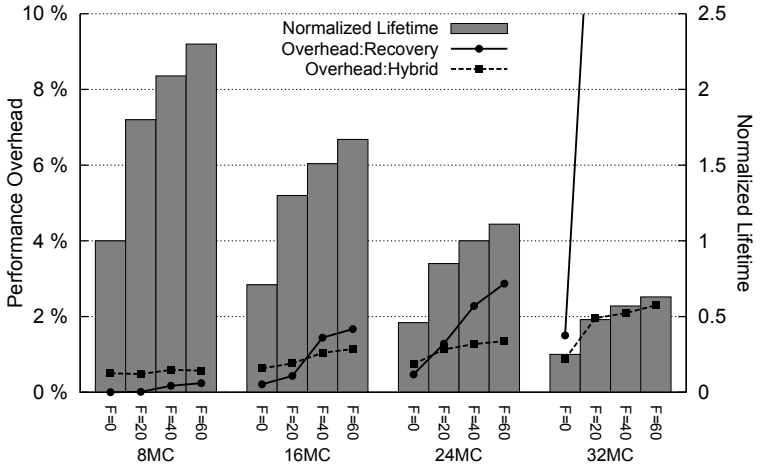

Figure 3.9: PDN pad failure's effect on noise mitigation performance overhead and expected EM lifetime. F indicates the number of failed pads. The baseline is an 8 MC chip with no failed pads. For *Recovery*, the performance overhead with 32 MC goes off-chart to 15% if 20 pads fail (25% if 60 pads fail).

worst consequences of PDN pad failure. Third, although PDN pad failures happen sequentially, and any PDN pad failure could change on-chip current distribution and thus change the distribution of current among pads, such phenomena will not reduce the failure risk of initially high-current pads: EM is an effect that accumulates over time. For these reasons, we use single-pass approximation to choose failed pads.

The lines in Figure 3.9 show the noise mitigation performance overhead of different pad-failure-tolerances F and different I/O configurations. We examine both recovery-only and hybrid techniques with a conservative recovery cost of 50 cycles per error. The baseline is the performance of recovery-only with an 8 MC chip and no PDN pad failures. For the same reason noted in Section 3.4.3, noise amplitude only increases mildly despite the significant increase in voltage violation rate that results from a limited number of PDN pad failures. As a result, the recovery technique, which enforces a constant timing margin, suffers from more frequent rollbacks as pads fail. The hybrid technique, however, can avoid this penalty by dynamically adjusting the timing margin. Even with a relatively high recovery cost, the hybrid mechanism more effectively tolerates PDN pad failures, especially in wide-I/O chips, and in particular as technology scales.

Regardless of noise mitigation technique selection, we observe that for chips with abundant

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

3.6.3 EM Lifetime Considering Pad Failure Tolerance

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

The bars in Figure 3.9 show the normalized expected lifetime of different pad-failure tolerance levels F across different chip I/O configurations. If we do not allow any PDN pad failures ($F = 0$), increasing the memory controller count from 8 to 24 reduces EM lifetime by half with modest noise-mitigation performance overhead. However, by tolerating a small number of PDN pad failures, the whole chip's expected lifetime is extended significantly. For example, if 40 pads are allowed to fail, increasing processor's MC count from 8 to 24 would not hurt the expected system lifetime. According to our performance results, the penalty of tolerating 40 PDN pad failures with a 24 MC chip is negligible (1%). However, there are limits to how far this approach can extend lifetime. Going beyond 24 MCs and giving up more power pads to accommodate 32 MCs will place too much pressure on the rest of the power pads and even PDN pad failure tolerance cannot extend chip lifetime enough to match the baseline case. Thus we conclude that the power and I/O pad tradeoff is ultimately limited by C4 EM lifetime–to 24 MCs in the scenarios we evaluate.

3.6.4 Impact of On-Chip Metal

Since the damage of electromigration is directly related to the density of stressing current, increasing the width of on-chip metal wires could help to alleviate the EM pressure on PDN metal wires. Using the technique described in Section 3.6.1, we evaluated the MTTFF value for both on-chip wires and C4 pads. Figure 3.10 shows the impact of on-chip metal occupancy on both C4 pads' and on-chip wires' expected lifetime. Even though increasing metal occupancy significantly increase the lifetime for on-chip wires, the change in C4 pads' MTTFF value is negligible. This is because wide metal reduces the current density in on-chip wires with increased cross-sectional area. For C4 pads, although increased metal width reduces the effective resistance between pads, it only marginally decreases max pad current and does not affect average current at all. We therefore conclude that the EM lifetime of C4 pads is insensitive to metal occupancy.

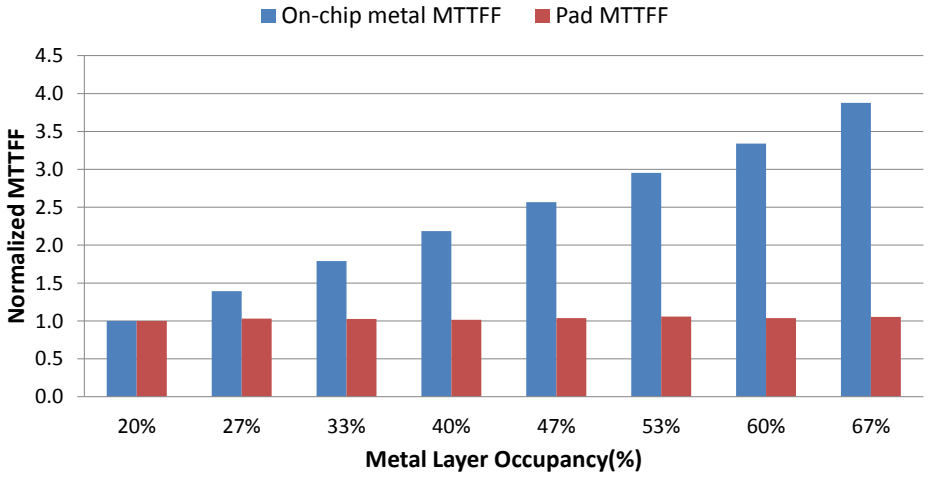

Figure 3.10: Impact of on-chip metal occupancy on the EM lifetime of both C4 pads and on-chip wires.

3.7 Summary

Power delivery quality is becoming a limiting factor in multicore processor design. In this chapter, we combine VoltSpot with other architecture-level tools to quantitatively evaluate voltage droop and electromigration (EM) lifetime changes subject to power/ground and I/O C4 pad tradeoffs. Our

results suggest that with noise-aware pad allocation, replacing power-supply pads with I/O pads only mildly increases supply-voltage noise amplitude. By evaluating the performance of different run-time voltage noise mitigation schemes, we conclude that the penalty of mitigating extra noise caused by reducing power-supply pads is negligible, as long as run-time noise control is carefully designed. We also discover that, with the help of dynamic noise mitigation, chips can tolerate a small number of C4 pad EM failures with low performance overhead. Combined with a detailed EM lifetime study, we show that with a performance overhead of just 1%, chip I/O bandwidth can be tripled without sacrificing EM lifetime.

Chapter 4

A Statistical Analysis of EM-induced C4 Bumps Wearout

4.1 Overview

The road toward building reliable silicon chips with higher performance and lower energy consumption is full of challenges from various physical constraints. Electromigration has long been recognized as an important lifetime reliability issue. Among the various circuit components of a silicon chip, Controlled Collapse Chip Connection bumps are particularly vulnerable to EM damage because the materials they are made of (e.g., SnPb or SnAgCu) have orders of magnitude lower current density tolerance compared with the material used in on-chip wires (e.g., copper) [97]. Moreover, C4 bumps suffer from the effect of Joule heating and current crowding [8], which both accelerate EM damage.

Due to the sheer volume of power supply bumps that today's high-performance microprocessors usually have (e.g., a few thousand [32]), the whole system's expected EM-failure-free time is much shorter than a single bump's mean-time-to-failure, because any bump could fail after even a relatively short period of time [69]. As a result, guaranteeing EM-failure-free operation during the targeted system lifetime requires allocating a large amount of C4 bumps to the power delivery network to amortize per-bump current and improve single-bump MTTF. As non-ideal technology scaling brings exponentially increasing transistor density without being able to reduce the supply voltage accordingly, power and current consumption density grow rapidly. As we demonstrated in Chapter 3, near future silicon chips will require more power-supply C4 bumps to maintain EM

lifetime.

Unfortunately, higher power bump count raises the cost of the chip package, with increased design and fabrication complexity. For example, the difficulty of silicon chip assembly, especially getting sufficient underfill material to flow smoothly through the C4 bump array, is directly related to the density of populated C4 bumps [104]. Besides, more power bumps create more footprints (i.e., contacting areas for electrical connection) in the chip package, which increases the complexity of I/O routing. To make matters worse, higher power bump requirements also reduce the available bandwidth for off-chip I/O communication, which eventually degrade processor performance. This is because the bump sites, which are projected to have a constant density in the future [32], are shared exclusively by both power supply and I/O channel. It is therefore becoming critically important to improve whole-chip robustness against EM-induced bump wearout in order to reduce the requirement for power supply bumps.

Once a power supply bump fails, the PDN's effective impedance seen by its nearby circuit will increase, because now the supply current has to come from neighbouring power bumps through the lateral on-chip PDN wires. In Chapter 3, we observe that because the on-chip wires are designed to have very low impedance, the voltage noise increment caused by a moderate number of bump failures is usually very small. This observation suggests that allocating a mild extra on-chip voltage noise margin (i.e., over-volting and/or under-clocking) could help tolerate bump failures, improve whole-chip's EM-robustness, and reduce the power bump requirement. However, as power bumps start to fail, the current they used to carry will be re-distributed to their neighbouring bumps and accelerate EM wearout. This may create an "avalanche" effect where a cluster of bumps break shortly after the first bump failure happens. Our simulation results show that this failure-acceleration effect does exist and designers should carefully consider its impact.

In this chapter, we design and implement a Monte Carlo Simulation (MCS) framework to analyze the mechanism and consequences of multiple, EM-induced, random power-C4 bump failures. By integrating this framework with VoltSpot (Chapter 2), we are able to capture the impact of current re-distribution and study the whole system's MTTF under given power bump configuration and voltage noise margin. Our framework helps to navigate designers through the complex trade-

off involving C4 bump allocations, EM-induced MTTF, voltage noise margin, and multiple on-chip power domains. This chapter's major contributions are:

- We develop an MCS framework for multiple EM-induced random bump failure simulation. It helps designers provision bump allocation under different design constraints, to reduce packaging cost and support more off-chip I/O channels in current and near-future technology nodes.
- We observe that the EM-induced bump-wearout power delivery quality degradation is a relatively slow process. Therefore, a small extra voltage noise margin can tolerate a significant increase in permissible bump failures and extend system lifetime. Simulation results indicate that with an extra noise margin of 0.5% Vdd IR drop, designers can reduce power bump count by 43% without shortening system MTTF.
- By comparing results against a simplified model that ignores the impact of current redistribution, we prove the existence of bump failures' avalanche effect and show that ignoring its impact will over-estimate system MTTF by up to 80% and under-estimate the required noise margin for EM-wearout by over 50%.
- We evaluated the impact of having multiple on-chip voltage domains and observe that splitting power grid into multiple islands will degrade bump array's robustness against EM-wearout. Under a fixed MTTF target and voltage noise limit, finer grained power domain setting requires up to 15% more power bumps.
- We design and evaluate two run-time techniques that effectively extend system's lifetime beyond designed MTTF with graceful performance degradation.

4.2 Background and Related Work

4.2.1 Single Bump EM Failure

EM refers to the phenomenon of gradual mass transport in metal conductors induced by momentum transfer from electrons to atoms. It creates an open or short-circuit in the PDN and permanently

degrades the quality of power delivery. The EM-induced failure mechanism for a single C4 bump has been extensively studied in the past. The cumulative distribution function (CDF) of a single bump's failure probability over time can be well described by a lognormal distribution [56]:

$$F(t) = \int_0^t \frac{1}{\sqrt{2\pi}\sigma t} e^{-(\ln t - \ln t_{50})^2/2\sigma^2} \tag{4.1}$$

where σ is the lognormal distribution's standard deviation (determined experimentally to be 0.5 [56]). The mean-time-to-failure is determined by Black's equation. For C4 bumps, the MTTF equation must be adjusted to consider the impact of current crowding and Joule heating [8]:

$$t_{50} = A(cJ)^{-n} exp\left(\frac{Q}{k(T+\Delta T)}\right) \tag{4.2}$$

where J is current density, n and Q are material-specific constants (for C4 solder bump, $n = 1.8, Q = 0.8eV$ [6]), k is Boltzmann's constant, $c = 10$, $\Delta T = 40°C$ [8], A is an empirical constant, and T is temperature in Kelvin. With equation (4.1) and (4.2), we can calculate the probability of failure for a single bump after any time t. We note that although EM wearout can also cause a short-circuit in the on-chip metal stack, C4 bumps are less likely to suffer from such a phenomenon, since the spacing between C4 bumps is usually much larger than on-chip metal. Thus in this chapter, we assume that all EM-induced bumps failures cause an open circuit. Also, because I/O bumps carry bidirectional current, they are characterized by longer times to the EM-induced failure [27]. Furthermore, I/O bump failures are independent of each other and do not affect power delivery quality. For these reasons, we only study the failure of power bumps in this chapter.

4.2.2 Consequences of Power-Supply Bump Failures

Figure 4.1 provides a simplified illustration of bump failure on a 1-dimensional PDN. Once a power supply C4 bump fails, it can no longer deliver current to the silicon chip. Transistors near the failed bump therefore have to draw current from neighbouring power bumps through the lateral on-chip PDN wires. This phenomenon has two consequences. First, the increased current density in those neighbouring bumps (the two remaining ones in Figure 4.1b) will exacerbate EM wearout, causing

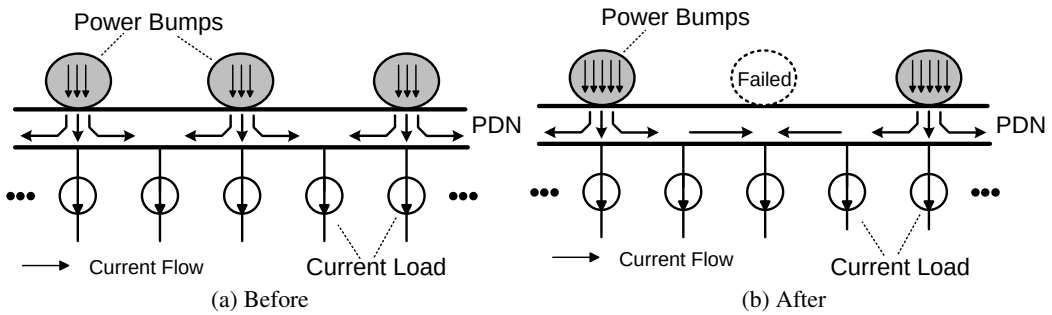

Figure 4.1: The consequence of a power bump failure.

an avalanche effect where clustered bumps fail more quickly after the first failure within that cluster. Second, the effective PDN impedance seen by the transistors near the failed bump (e.g., the current source in the middle of Figure 4.1b), will increase, which intensifies the supply voltage noise in those regions. The increased noise will result in frequent timing errors and consequently, system failure. We note that the severity of the avalanche effect directly affects the rate of power delivery quality degradation (i.e., the increment of on-chip voltage noise over time) and the expected chip lifetime. If bump failures are scattered across the entire chip, the per-failure-induced increment of voltage noise will be smaller, and therefore the chip will survive longer before on-chip voltage noise exceeds the design threshold. If all failed bumps are clustered together, on-chip voltage noise will rapidly exceed the design threshold after a few bumps fail.

4.2.3 Related Work on Statistical EM Analysis

In Chapter 3, we found that tolerating multiple EM-induced bump failures with extra noise guard-band could significantly extend system MTTF. However, our first-order analysis made several simplified assumptions about the mechanism of EM-induced C4 failures. For example, because Chapter 3 ignores the impact of bump-failure-induced current redistribution, the potential damage of the avalanche effect could not be captured. Also, we previously assumed that the bumps with the largest current always fail first, which neglects the stochastic nature of bump failures. In order to evaluate the impact of failure-induced bump-current redistribution, and to accurately estimate the system's expected lifetime, we improve our previous work with a statistical analysis approach and

more accurate assumptions.

Li et al. [52] proposed a statistical model for multi-via EM lifetime estimation. Fawaz [14] designed a MCS framework to estimate the on-chip power grid's expected lifetime under given a IR drop target. Although the failure times of a single via or copper line follow the same probability distribution (i.e., lognormal) as a C4 bump, the prior work is not adequate for the study of multiple bump failures, because it either ignores the failure-induced current redistribution [14], or is only capable of modeling a small number of vulnerable elements (e.g., four vias in [52]). To the best of our knowledge, a whole-system analysis of multiple C4 bump failures that considers the impact of the avalanche effect is missing from the literature.

4.3 Statistical Simulation of Bump Failures

4.3.1 Monte Carlo Simulation and Results Confidence Level

In this chapter, our primary interest is system's robustness against EM-wearout. This is usually measured with system mean-time-to-failure, where the time-of-failure (TOF) is defined as when the maximum on-chip voltage noise exceeds a pre-determined threshold. Because the failure of each bump is a probabilistic event, multiple bump failures become a stochastic process. Considering the total number of power-supply bumps, the permutation space of possible bump failure sequences will be enormous even if we only allow a small portion of bumps to fail (we use permutation instead of combination because the order of bump failures matters due to the current re-distribution effect). For example, for a silicon chip with 1500 total C4 bumps, there are more than 10^{126} ways to fail 40 bumps. For this reason, an analytical approach would be infeasible for the study of multiple bump failures.

Fortunately, since we are interested in the mean value of a distribution (system lifetime), Monte Carlo Simulation (MCS) can be utilized to get reliable results within reasonable time. The idea of MCS is to take random samples from the population-of-interest and use the *sample* arithmetic mean to estimate the *true* mean. The quality of this estimation heavily depends on the number of samples taken. Prior work [39] suggests that when sampling from an unknown distribution, the following

equation can be used as the stopping criteria:

$$N \geq \left(\frac{z_{\alpha/2} s_N}{|\bar{x}_N| \varepsilon / (1+\varepsilon)} \right)^2, (for N \geq 30) \tag{4.3}$$

where N is the minimum number of samples required, $\bar{x}_N$ is sample mean, ε is the estimation's relative deviation from the population's true mean μ. $z_{\alpha/2}$ is the $(1-\alpha/2)$-percentile of a random variable Z, where Z has standard normal distribution. s_N is sample standard deviation, which can be calculated by:

$$s_N = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} (x_i - \bar{x}_N)^2} \tag{4.4}$$

In our work, we set $\varepsilon = 0.005$ and $z_{\alpha/2} = 2.32\,(\alpha = 0.02)$ therefore our MTTF results have a confidence interval of $\pm 0.5\%$ around the true mean at a confidence level of 98%. In our study, this confidence level can be achieved with $N = (800, 1500)$.

4.3.2 Capturing the Impact of Current Re-distribution

As we discussed in Section4.2.2, the severity of the avalanche effect caused by failure-induced current re-distribution directly affects system's robustness against EM. To model the EM-wearout acceleration phenomenon caused by failure-induced current re-distribution, we need to adjust bumps' failure probability distribution functions whenever their current changes. This CDF adjustment has been discussed in previous work [52]. It involves two major steps: stressing current translation and conditional probability calculation. The rule of current translation is given by equation:

$$\left(\frac{I_{prev}}{I_{new}} \right)^n = \frac{{t_{prev}}'}{t_{prev}} \tag{4.5}$$

where n is the same as in equation 4.2, I_{prev} and I_{new} are previous and new current, t_{prev} is the stressed time under previous current and ${t_{prev}}'$ is the translated stress time. Under this equation, the probability of failure under current I_{prev} over time t_{prev} would be the same as if been stressed under I_{new} over time ${t_{prev}}'$. With the translated stressed time, we can calculate the new CDF with

conditional probability:

$$F_{new}{}'(t) = \frac{F_{new}(t + t_{prev}{}') - F_{new}(t_{prev}{}')}{1 - F_{new}(t_{prev}{}')} \tag{4.6}$$

For any time period t after the moment of current change, this equation calculates the probability of bump failure under current I_{new}, given the condition that the bump hadn't fail after the first time period of $t_{prev}{}'$. F_{new} is the unmodified CDF under stressing current I_{new}.

4.3.3 MCS Framework for Bump Failure Study

```
1  Initialize an empty set S;
2  while stopping criteria (equation 4.3) is not met do
3      foreach bump do
4          calculate current density and generate CDF;
5          generate TOF;
6      end
7      while max on-chip noise < threshold do
8          bump_to_fail = the alive bump with earliest TOF;
9          current_time = bump_to_fail.TOF;
10         remove bump_to_fail from PDN;
11         re-evaluate on-chip voltage noise;
12         foreach alive bump do
13             re-calculate current density;
14             adjust failure probability CDF;
15             generate TOF;
16         end
17     end
18     add current_time as a new sample to S;
19 end
20 MTTF = mean(S);
```

Algorithm 1: Our MCS framework for failure study.

Algorithm 1 illustrates the details of the failure model used in our MCS framework. Within a trial of bump failure simulation, we first calculate the current density of each bump and generate per-bump failure time CDF (line 4) using equation 4.1 and 4.2. Then we randomly generate the TOF of each bump using the inverse transform sampling method [10]. To be more specific, we first generate a random number p for each bump from a uniform distribution between 0 and 1. The bump's TOF then equals to $F^{-1}(p)$ where F^{-1} is the inverse function of the bump's CDF. By selecting the bump

with earliest TOF and removing it from the PDN (line 8-10), we can simulate one random bump failure event. Every time a bump fails, all remaining bumps' CDF will be adjusted (line 13-14) and inverse-transform-sampled (line 15) to capture the impact of current re-distribution.

To monitor the degradation of power delivery quality over time, we re-evaluate on-chip noise level whenever a failed bump is removed (line 11). As soon as the noise level exceeds the design threshold, we terminate the trial and record the last failed bump's TOF as the TOF of the whole system (line 18). Each iteration of the outer while loop (line 2-19) simulates a trial of bump failures and generates one system TOF sample. The MCS will stop as soon as the stopping criteria (equation 4.3) is satisfied and the system's MTTF equals to the arithmetic average of all samples (line 20).

```
1  Initialize an empty set S;
2  while stopping criteria (equation 4.3) is not met do
3      foreach bump do
4          calculate current density and CDF;
5          generate TOF;
6      end
7      while max on-chip noise < threshold do
8          bump_to_fail = the alive bump with earliest TOF;
9          current_time = bump_to_fail.TOF;
10         remove bump_to_fail from PDN;
11         re-evaluate on-chip voltage noise;
12     end
13     add current_time as a new sample to S;
14 end
15 MTTF = mean(S);
```

Algorithm 2: Simplified bump failure model

To evaluate the severity of the avalanche effect, we also implement a simplified baseline model that only calculates bump current, CDF and TOF once in every trial and does not adjust bump current or CDF after bump failures (Algorithm 2). This way, the simplified model ignores the impact of current re-distribution and the results' differences between these two models directly measures the severity of the avalanche effect. We note that the model used in Chapter 3 would not be a proper baseline for this study because it assumes that bumps with the highest current will always fail first, which completely neglects the stochastic nature of bump failures and thus unable to provide system MTTF.

4.3.4 Supporting Other Random Failures

Although Algorithm 2 and 1 are designed to simulate multiple EM-induced C4 bump failures, they can be easily extended to study other vulnerable components under EM stress or other reliability threats. For example, by replacing C4 bumps with other vulnerable units (e.g., on-chip metal wires under EM stress or transistors under the threat of time-dependent gate oxide breakdown, or TDDB) and plug in the corresponding failure probability CDF, we can evaluate the consequence of multiple failures and the whole-system's MTTF. If the modeled failure events are not independent to each other, the system re-evaluation and CDF re-calculation mechanism in Algorithm 1 (line 11-16) can be used to capture the interaction between random failures and their system-level impacts. Moreover, our framework has built-in interfaces with architecture-level power, thermal and voltage noise models therefore it directly supports the exploration of different failure mechanisms' dependency on current, temperature and voltage.

4.4 Simulation Methodology

4.4.1 PDN Modeling

The MCS framework described in 4.3 uses per-bump current to calculate/adjust failure probability CDF. It also depends on the evaluation of on-chip voltage noise to determine the whole system's TOF. To get the PDN's current and voltage profile, we integrate our MCS framework with VoltSpot. The parameters used in our study are listed in Table 4.1. We note that although existing circuit level PDN models (e.g., the SPICE model in [64]) are also capable of simulating PDN current and noise, they usually involve millions of nodes, which significantly increase simulation time. As a result, these models can not provide enough MCS trials to derive high quality results within a reasonable time.

In order to avoid pessimistic results caused by sub-optimal power bump placement, we optimized the location of all C4 bumps in all of our test cases. The optimization algorithm is adopted from prior work [91].

On-Chip Metal Resistivity (ρ)	1.68e-8
Global PDN Layers Width/Pitch/Thickness (μm)	10 / 30 / 3.5
Intermediate PDN Layers W/P/T (μm)	400 / 810 / 720
Local PDN Layers W/P/T (μm)	120 / 240 / 216
C4 Bump Diameter/Pitch (μm)	100 / 285
C4 Bump Resistance ($m\Omega$)	10
Package Resistance ($m\Omega$)	0.015

Table 4.1: PDN Parameters used in this study.

4.4.2 Multicore Processor Power, Area and Floorplan Modeling

To study the severity of EM in the near future high-performance processors, we build a multicore processor based on a 45nm Intel Penryn processor [16] and scaled it down to 16nm. It has 16 32-bit 4-way out-of-order cores. Each core contains a 32kB L1 instruction cache and a 32kB L1 data cache. Unified L2 caches private to each core are each 3MB. The chip area was calculated with McPAT [53], an architecture-level power model. Application-specific power consumption was derived by integrating McPAT with a performance simulator Gem5 [4]. We use ArchFP [13] to generate our floorplans. Figure 4.2 shows the floorplan of our 16nm, 16 core processor. It has an area of 159.4 mm^2 and a total number of 1914 C4 bump sites. With a supply voltage of 0.7V, the 3.7GHz processor's peak power consumption is 151.7 W. We assume that the silicon chip has an uniform temperature of 100°C, which provides a lower bound of system MTTF and therefore an upper bound of the minimum required noise margin.

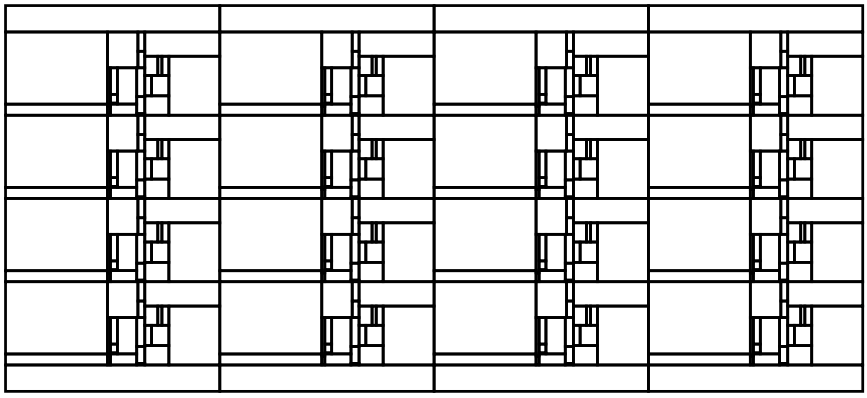

Figure 4.2: Floorplan of our example multi-core processor.

In the past, Tao et al. [84] discovered that the lifetime of metal under high-frequency AC current stress is determined by the DC component of the stressing current alone. For this reason, we analyze bump EM stress in steady-state only. To be more specific, we first simulate the Parsec 2.0 benchmark suite [3] with our power and performance simulators and then extract the average power consumption of the entire suite. With this power map as an input to steady-state VoltSpot simulation, we can capture the whole system's EM wearout under the processor's average behavior.

4.5 Results and Discussions

4.5.1 The Avalanche Effect: How Bad Is It?

In this chapter, the avalanche effect of bump failures refers to the phenomenon that once a power bump fails, the current it use to carry got re-routed to its neighbouring bumps and causes a chain effect of accelerated wearout and failures in the neighbouring bumps. To analyze the severity of the avalanche effect, we perform MCS on our 16-core processor with both detailed and simplified failure models described in Section 4.3. Figure 4.3 illustrates the whole chip's MTTF with different extra noise margins reserved for EM wearout. The baseline case, which is a conservative design with no extra margin and 40% of total bumps assigned to power supply, has a maximum on-chip IR drop of 1.24% Vdd. All results are normalized to the MTTF of this baseline design. As we tolerate larger IR drop (x-axis represents the delta increase of noise margin), we significantly extend system MTTF by allowing more bumps to fail. Since the only difference between the two models is whether they consider the impact of current re-distribution, the differences between their results are directly caused by the avalanche effect. Therefore, Figure 4.3 indicates that the avalanche effect does exist and in extreme cases (e.g., with 5% Vdd extra noise margin), ignoring it will over-estimate system MTTF by 80%.

We also evaluated the computational complexity of the two models. In general, the detailed model requires more simulation time, because failure time CDFs need to be updated for every bump after each bump failure. However, the simplified model requires more trials for each MC experiment, because bump failures are not correlated in space and more bump failures are needed

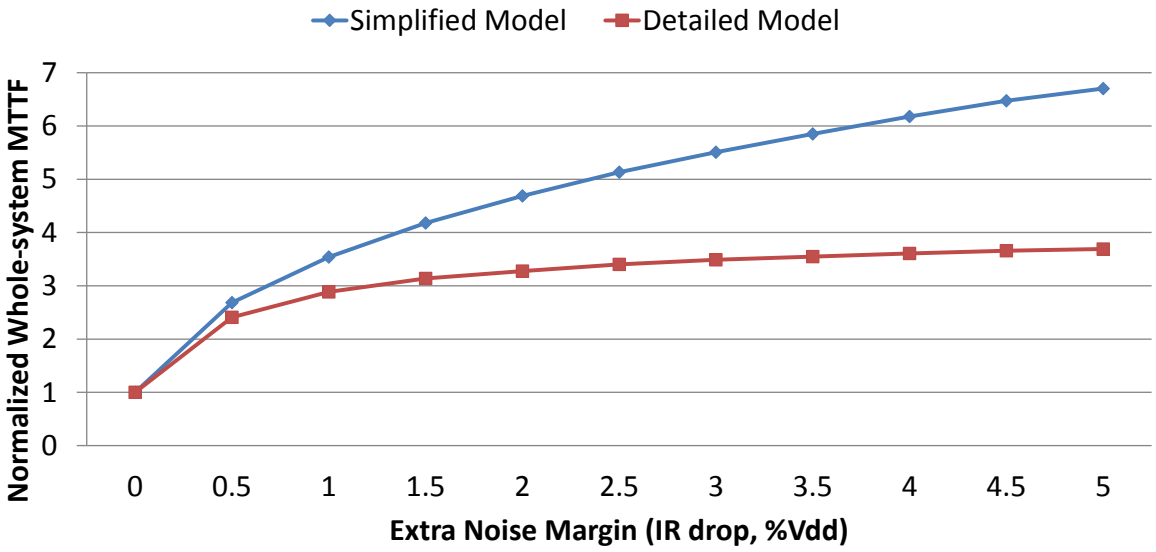

Figure 4.3: Normalized whole-system MTTF under different noise margin settings. The differences between the two sets of results are caused by the avalanche effect.

to reach chip failure, while the detailed model requires fewer trials as extra noise margin increases, because there are fewer paths to failure. As a result, the simulation overhead of the detailed model ranges from 80% (at 0.5% extra noise margin) to break-even (at 4.5% extra margin).

An interesting observation is that the severity of the avalanche effect will increase as we further relax noise margin to allow more bumps to fail. This is because as EM-stress gradually fails a cluster of power bumps, the amount of current re-distributed to the remaining bumps will accumulate and exacerbate the acceleration of EM-wearout. If we only allow a small portion of bumps to fail (i.e., the left-most data points of Figure 4.3), the impact of current-redistribution is fairly small and a mild increase in noise margin can tolerate those bump failures and significantly increase system's MTTF.

4.5.2 Achieving Target MTTF with Reduced Bump Count

The number of C4 bumps allocated as power supply is an important design parameter because it directly affects both the complexity of chip package design and the available physical width of off-chip I/O communication channels. Although both design considerations favor fewer power supply bumps, reducing power bump count will increase bump current and therefore degrade system's EM-

robustness. Table 4.2 shows our baseline chip's expected lifetime without any EM-induced bump failures. We evaluated different bump allocation schemes and observe that reducing power-bump count significantly shortens EM-failure-free time.

% of Bumps Allocated as Power Supply	70%	60%	50%	40%	30%
Normalized System EM-failure-free Time	1.00	0.90	0.72	0.54	0.36

Table 4.2: Whole system's expected EM-failure-free time under different bump allocations.

Fortunately, the observations in Section 4.5.1 indicate that by assigning a mild extra noise margin to tolerate bump failures, the MTTF of a chip with fewer power bumps could be extended to match the MTTF of a chip with more bumps. In other words, at the cost of increased noise margin, power bump count can be reduced without degrading system's MTTF.

With our MCS framework, we analyzed the tradeoff between chip power bump count and the required on-chip voltage noise margin under a fixed system MTTF target. Using a target MTTF that equals the EM-failure-free time of a baseline chip with 70% of total bumps allocated as power/ground, Figure 4.4a shows the required noise margin for different bump allocation schemes (In this subsection, we only discuss the single domain case, which is the blue line with diamond-shaped marker at the bottom). We note that the noise margin values here represent absolute Vdd% IR drop and they include the guard band to tolerate both bump failures and the power delivery quality degradation due to reduced bump count (e.g., if we reduce power bump ratio from 70% to 30%, the worst on-chip IR drop will increase from 0.91% to 1.59% even without any bump failures). We observe that a mild increase in noise margin is sufficient even if we significantly reduce power bump count. For example, with an extra noise margin of 0.53% Vdd (from 0.91% to 1.44%), designers can reduce power bump count by 43% (from 70% to 40%) without shortening whole-system's MTTF. This increases the physical bandwidth for off-chip I/O communication by up to 2x (I/O bump count from 30% to 60%) while only incurring less than 1% slowdown (addressable by under-clocking to increase timing margin).

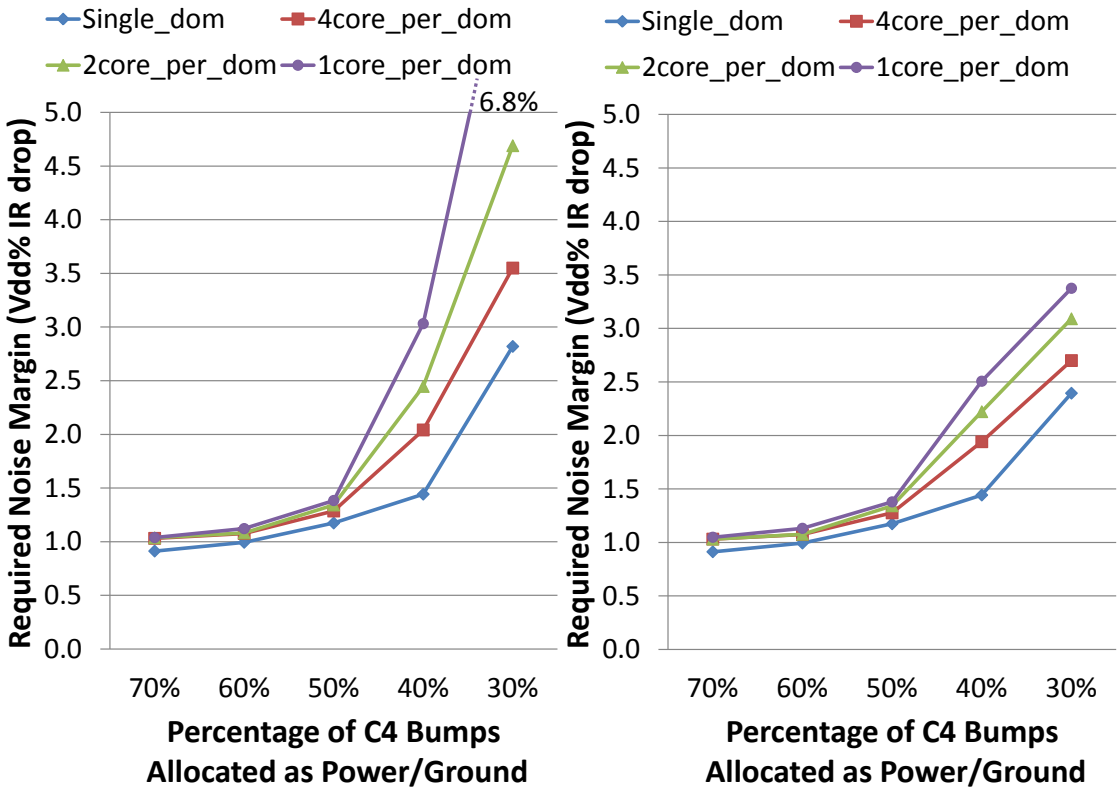

Figure 4.4: The required noise margin (in terms of absolute Vdd% IR drop) to achieve target MTTF for different bump allocations and multiple on-chip power domain settings. We assume that the target MTTF equals the EM-failure-free time of the single-domain chip with 70% bumps as power/ground.

4.5.3 Impact of Multiple On-Chip Power Domains

Contemporary processors usually have multiple on-chip power islands instead of supporting the entire silicon chip with a single power domain. Splitting the on-chip power delivery grid into multiple isolated domains provides finer-grained spatial support for dynamic voltage and frequency scaling (DVFS), which improves the energy efficiency of multi-core processors. Also, isolating high-frequency digital blocks (e.g., CPU cores) from mixed-signal blocks (e.g., PHY unit in memory controllers) helps to improve the integrity of analog signals. However, separating the power grid into mutually insulated domains inevitably "cuts off" a portion of lateral current flow in the on-chip PDN. This not only degrades power delivery quality due to increased PDN impedance, but

also undermines C4 bumps' robustness against EM-wearout. This is because for the bumps near domain boundaries, the amount of redistributed current due to bump failures within one domain will be higher since the bumps in neighbouring domains could no longer help to share the current load.

To study the impact of power domains on system EM robustness, we extended VoltSpot by splitting the virtual on-chip Vdd grid into different islands according to the chip floorplan and the power domain specification for each block. The lateral circuit branches in the virtual grid that cross domain boundaries are removed so that different domains are mutually insulated from each other. With this extension, we performed statistical analysis with our MCS framework and Figure 4.4a shows the required noise margin to achieve the target system MTTF under different power domain settings. Using the same MTTF target and bump allocations as Section4.5.2, we tested three cases with 4, 2, and 1 core(s) per domain, which give us 4, 8 and 16 different domains respectively. Simulation results indicate that having multiple power domains does exacerbate EM-wearout, and the required noise margin to guarantee targeted MTTF will increase significantly as we split the power grid into more domains or reduce power bump count. If the chip design has strict limits for both MTTF and on-chip noise margin, supporting more power domains will require more power bumps. For example, with a fixed noise margin of 2% Vdd, a single-domain chip needs only about 35% of total bumps to supply power while a one-core-per-domain chip has to increase this ratio to around 55%. These results indicate that although more on-chip power islands improves system energy efficiency with finer-grained support for voltage/frequency tunning, it also comes with the price of a higher demand for power bumps.

To further explore the impact of the avalanche effect, we performed the same set of analysis with the simplified model (Figure 4.4b). By ignoring the current re-distribution phenomenon, the simplified failure model produces optimistic estimations of EM-wearout and therefore ends up with a much smaller noise margin requirement. In the worst case scenario with 1 core per domain and 30% bumps as power supply, the simplified model underestimates the required noise margin by over 50%, which leads to a design that will fail long before the target MTTF. We therefore conclude that it is critically important to evaluate the impact of the avalanche effect in the design tradeoff study

that involves power bump count, noise margin, EM-induced system MTTF and multiple on-chip power domains. Failure to do so can lead to incorrect pre- and post-RTL design decisions.

4.5.4 Using Graceful Performance Degradation Schemes to Extend Chip Lifetime

The EM-wearout of a silicon chip's C4 bump array is a gradual process that slowly increases on-chip voltage noise. Due to the non-uniform current distribution between different bumps and the presence of the avalanche effect, different on-chip regions will suffer from different levels of EM-induced power delivery quality degradation. These observations suggest the possibilities of using graceful performance degradation techniques to extend silicon chip's lifetime beyond the original MTTF and thus get more work done with lower performance. For example, with the ability to further relax noise margin (e.g., slow down clock frequency) after worst-on-chip IR drop reaches the design-target, the silicon chip could tolerate more bump failures and operate longer with reduced performance. Another possible scheme is to abandon a core as soon as its IR drop exceeds the design-target. This way, the processor loses throughput but the remaining cores could operate longer with unchanged voltage and frequency level.

We implemented both margin adaptation and core desertion techniques in our MCS framework. For the margin adaptation scheme, we assume a fixed voltage level and only adjusting global clock frequency to accommodate extra voltage noise. Since there is a linear dependency between the delay of transistors and supply voltage noise level, we assume that a 1% extra IR drop requires slowing down clock frequency by 1%. After the on-chip noise exceeds the original design target, the margin adaptation controller will gradually slowdown the entire processor according to the actual voltage noise level. It is worth mentioning that in practical designs, designers could set a hard limit on noise margin beyond which the processor or core will be considered as failed. This help to maintain a lower performance bound and guarantee silicon chip's functionality under worst case voltage noise. We include this hard limit as a design parameter in our framework. For the core-desertion technique, the controller will power-gate a core as soon as its noise level exceeds design-target and the abandoned core will be never used again. We note that both schemes assume ideal voltage sensing that captures worst per-core IR drop. The discussion of noise detection's role

exceeds the scope of this chapter but is an interesting direction for future work.

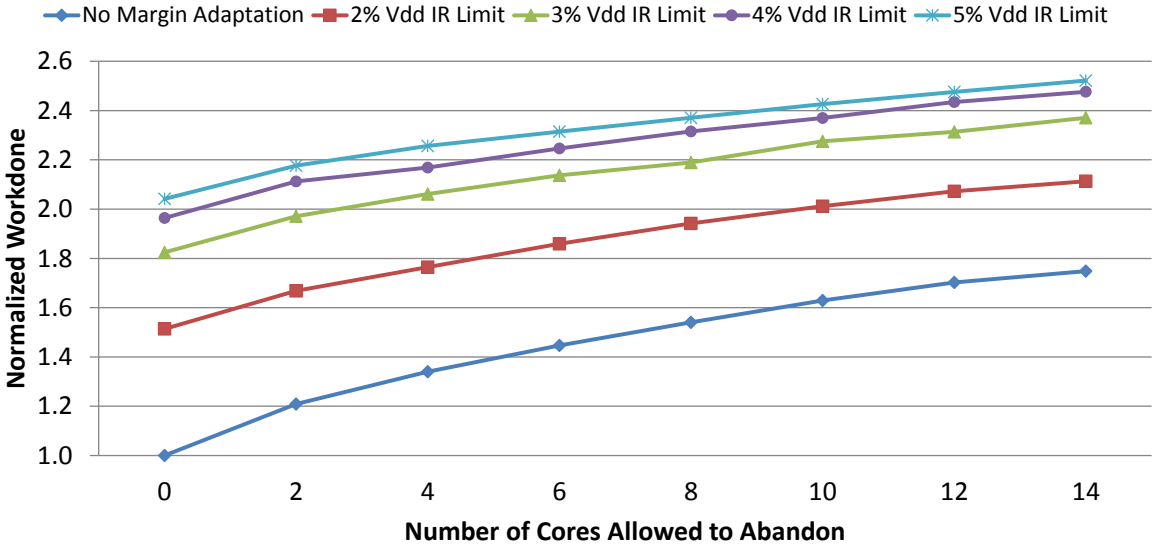

Figure 4.5: Evaluation of two graceful performance degradation mechanisms. Margin adaptation technique gradually relax noise margin until a designated limit (different lines with 2-5% Vdd). Core desertion scheme abandons a core if its noise exceed the design target. All values are normalized to the amount of workdone by the baseline chip without any graceful degradation mechanisms enabled.

To study the effectiveness of these two schemes, we pick a design point from Figure 4.4a and apply the graceful degradation techniques with different settings. The baseline design has 40% bumps allocated to power supply and uses a single on-chip power domain. It requires an IR drop margin of 1.44% Vdd to meet the MTTF target. Figure 4.5 shows the evaluation results with the metric of aggregated workdone, which is calculated as $\sum clock\_frequency * number\_of\_functional\_cores * time\_duration$. All results are normalized to the workdone of our baseline design without any graceful performance degradation scheme. By testing the core-desertion technique that abandons up to 14 cores (out of 16) and the margin adaptation technique with different hard noise limit, we found that the aggregated workdone can be significantly increased (up to 2.5x) with graceful degradation schemes.

An interesting observation is that if we only apply one technique, the benefit of slightly increasing noise margin would be equivalent to abandoning many cores. For example, by gradually relaxing noise margin from 1.44% to 2% Vdd, the chip could get 51% more work done. To com-

plete the same amount of extra work, core-desertion technique must abandon at least 7 cores. The reason behind margin adaptation technique's high-efficiency is similar to what we've observed in Section 4.5.1, where a slight noise margin increase could tolerate more bump failures and significantly improve chip lifetime. On the contrary, core-desertion technique is not as effective because the silicon chip usually suffers from multiple EM-damaged regions therefore shutting down only a few cores could not extend the lifetime of the entire chip. In conclusion, graceful performance degradation techniques could keep a processor running after The designed MTTF and thus significantly increase the amount of work done by the processor.

4.5.5 Transient Noise Evaluation

The majority of this study uses steady-state IR drop as the metric to evaluate power delivery quality. Although this methodology has been widely used by researchers and engineers [69], it is also important to understand the impact of bump failures on transient noise. Unfortunately, even though VoltSpot supports transient noise simulation, evaluating PDN transient behavior will be many orders of magnitudes slower than steady-state simulation. For example, simulating 2M cycles (which is the minimum requirement to get an average noise behavior of one application [103]) would take over 200 hours, while one steady-state simulation takes less than 0.1 second. Considering the fact that each trial in our MCS has to evaluate on-chip noise once after every bump failure and it requires up to 1000 trials to get high-quality MTTF results, it become impossible to rely on whole-application transient noise evaluation in our analysis.

In order to validate whether our IR-drop-based observations are consistent with transient noise evaluation, we significantly reduce simulation time by only evaluating the worst-case transient noise. Using a 1k-cycle stressmark that triggers the largest transient voltage drop among the entire Parsec 2.0 benchmark suite, we are able to perform our MCS within an acceptable time period (e.g., 2 weeks). To be more specific, we replace the steady-state noise evaluation in Algorithm 1 (line 11) with transient simulation of the stressmark and use the maximum on-chip transient voltage drop during the entire stressmark to determine the failure of the chip (line 7). Table 4.3 shows the MTTF results under transient noise evaluation. Similar to the steady-state results, a small extra

margin reserved for EM wearout (e.g., 0.5% Vdd) allows the system to tolerate bump failures and significantly extend MTTF (e.g., 2x). Also, the fact that larger noise margin gives diminishing return on MTTF indicates that the avalanche effect also applies to transient behavior. These results give us confidence that our observations derived from IR-drop evaluation applies to the cases where transient noise is a major concern. A more detailed transient study is left for future work.

Extra Noise Margin % Vdd	0.5	1	1.5	2	2.5	3
Normalized MTTF	1.99	2.19	2.34	2.43	2.52	2.56

Table 4.3: Evaluation of EM-induced system MTTF under different transient noise margin settings. A 1k cycle stressmark is used to evaluate chip transient voltage noise. All results are normalized to the EM-failure-free time of the baseline processor with 40% bumps assigned as power supply.

4.6 Summary

In this chapter, we develop a Monte Carlo Simulation framework for the study of multiple EM-induced power-supply C4 bump failures. Our results indicate that the degradation of power delivery quality, in terms of maximum on-chip voltage noise, is a relatively slow process and a small extra noise guardband enables the system to tolerate the consequence of multiple bump failures and significantly extends system MTTF. As a result, target lifetime can be achieved with significantly reduced power supply bump count (e.g., by 43%) and a slightly increased on-chip noise margin (e.g., 0.5% Vdd IR drop). Also, we show that an avalanche-like wearout acceleration effect exists and ignoring it will over-estimate system MTTF by up to 80% and under-estimate the required noise margin for EM-wearout by over 50%. Moreover, we explored the impact of splitting on-chip power grid into multiple domains and observe that the lack of lateral PDN current flow across domain boundaries degrades bump array's EM-robustness. Consequently, chips with finer-grained domains will require more power supply bumps to maintain the target MTTF.

The design of C4 bumps under EM-induced reliability constraint is a multi-dimensional design space that consists of bump allocations, system MTTF, voltage noise margin, multiple on-chip

power domains, and etc. The MCS framework we develop here can guide a designer through this complex trade-off space and help designers to: 1. Better provision bump allocation in different design scenarios to reduce packaging cost and support more off-chip I/O channels in near-future technology nodes; 2. Better reserve timing margin for EM-wearout to achieve targeted MTTF; 3. Better design on-chip power islands to balance system energy-efficiency, I/O bandwidth and cost. Our framework can be easily extended to study the impact of on-chip temperature gradient. It also supports the design and evaluation of run-time wear leveling and/or graceful performance degradation techniques.

Chapter 5

A Cross-Layer Exploration of Voltage-Stacked PDN in Many-Layer 3D-IC

5.1 Overview

Because the benefits of Dennard scaling (devices that are simultaneously smaller, faster and lower power) are quickly vanishing, three-dimensional integrated circuits (3D-IC) are becoming an essential path to maintain exponential growth in device integration. However, 3D-IC raises several fundamental technical difficulties in addition to the fabrication challenges. Because the number of physical layers in a 3D-IC stack is expected to increase in the future, the problems of delivering power to the 3D stack seem daunting. The main culprit is the fundamental mismatch between the volumetric (cubic) aspect of power consumption and dissipation in 3D-IC, and the fact that power *delivery* is limited to only the bottom 2D surface (quadratic).

To alleviate the power delivery constraints in the era of 3D-IC, previous research proposals [18, 34, 59] suggest using the idea of voltage-stacking (V-S) to build the power delivery network for 3D-IC. V-S simply refers to the power delivery arrangement of two or more circuit blocks such that the ground of one block becomes the power supply connection for the next: the blocks are connected as a series stack for power delivery, with all of them sharing the same current while their Vdd values are added. With the help of voltage-stacking's ability to "recycle" current between blocks, adding more layers to a 3D-IC only requires increasing the off-chip supply voltage while

the current density within the PDN stays constant. For this reason, V-S provides a scalable solution to break the mismatch between 3D volume power dissipation and 2D surface power delivery.

To make the envisioned 3D-IC V-S practical, explicit voltage regulation is required for the general case when the currents of the various layers are not perfectly matched. While circuit solutions have been proposed for these explicit regulators [34, 59], a cross-layer tradeoff study that examines the benefits of voltage-stacking's current reduction, the area overhead and power efficiency of explicit voltage regulation, and the supply voltage noise under different workload conditions, is missing from the literature. For example, it is intuitive that compared with regular PDNs, V-S PDNs are more robust to electromigration wearout due to the reduced current density in through-silicon-via and Controlled Collapse Chip Connection pads. However, it is not clear how 3D-IC scaling (i.e., more layers) affects TSV/C4 array's EM lifetime, or whether designers can improve regular PDN's EM-robustness to match V-S PDNs' lifetime by simply allocating more power supply TSVs and pads.

In this chapter, we first extract noise/area/power information from a circuit implementation of a charge-recycled voltage regulator. Based on circuit-level insights, we extend VoltSpot into a system-level 3D PDN model to explore the impact of V-S on 3D-IC's PDN design. With the ability to profile each TSV and C4 pad's current and on-chip IR drop, 3D VoltSpot enables a detailed comparison of power delivery quality between regular PDNs and V-S PDNs in the context of 3D-IC. Our results show that by significantly improving PDN's EM-induced lifetime with high power-efficiency and minimal extra voltage noise, V-S provides a scalable and practical solution to the power delivery challenge in the era of many-layer 3D-IC. Our major contributions are:

- A system-level PDN model for 3D-ICs that supports the study of EM-induced reliability and supply voltage noise for both regular and voltage-stacked PDN. With a fine-grained modeling granularity and the ability to capture on-chip voltage regulators' power efficiency and output voltage drop, our model can assist system designers with evaluating the benefits and costs of design scenarios with different number of regulators and different TSV/C4 pad allocations.

- A detailed analysis of voltage stacking's impact on power-supply C4 pad and TSV array's

EM-induced lifetime. Our analysis indicates that although stacking more silicon layers quickly degrades regular PDNs' EM-lifetime, V-S PDNs' EM robustness is much less sensitive to silicon-layer count. For an 8-layer 3D processor, V-S improves the EM-induced lifetime of C4 pad and TSV array by up to 5x.

- Demonstrating the importance of workload imbalance (i.e., the power consumption difference between two adjacent layers in a 3D-IC) as a V-S design consideration and quantifying its impact on supply voltage noise, power efficiency and area overhead. Simulation results show that with the same total area overhead and an average workload imbalance ratio extracted from full applications, a V-S PDN's IR drop is only marginally larger (i.e., 0.75% Vdd) than a regular PDN.

This work is published in DAC15 [101].

5.2 Background and Related Work

5.2.1 Power Delivery Challenges and Solution in the Era of 3D-IC

As the number of layer increases, the mismatch between 3D volume power dissipation and 2D surface power delivery will inevitably increase the current density in both power supply C4 pads and TSVs. Higher current density not only exacerbates supply voltage noise (e.g., IR drop, Ldi/dt and LC resonance), but also accelerates EM-induced wearout. This indicates that the power delivery constraints will become more severe in the era of 3D-IC.

Voltage-stacking connects the PDN of two or more circuit blocks in series such that the ground of one block becomes the power supply for the next. When voltage-stacked, the current consumed by one block will be reused (charge-recycled) by the next block: all the blocks will share the same current while their Vdd values are added. While the scalability of V-S is limited in the regular 2D technology due to the fact that all blocks have to share the common substrate, 3D-IC offers a straightforward platform for supporting multiple blocks with V-S. This is because 3D-IC provides physical separations between layers thus allocating one block (e.g., a CPU core) per layer will

significantly reduce the complexity of V-S PDN design. Compared with the conventional power delivery scheme for an N-layer 3D-IC, V-S reduces the off-chip and cross-layer current density by up to N times through recycling charges between layers. This not only reduces the resistive noise (i.e., IR drop) across the PDN, but also significantly improves PDN's EM-induced reliability. With the ability to support an arbitrary number of layers with constant current, voltage-stacking provides a scalable solution for breaking the 3D-IC power delivery wall.

5.2.2 Voltage Regulation in V-S PDN

A major design challenge of V-S arises from the fact that V-S will try to compensate for any current-consumption mismatch between the stacked loads by re-distributing the intermediate voltage. This effect of workload imbalance gives rise to voltage noise, which can disrupt the functionality of the stacked circuits. Explicit regulators have been proposed to handle this accumulation of charge imbalance at the intermediate nodes. Unlike conventional regulation schemes, where the regulators provide 100% of the current required by the loads, V-S requires differential converters that only handles the current-mismatch between the layers, and thereby converters with smaller passives can attain higher efficiency than conventional regulation. These differential converters have a "push-pull" ability that can either source or sink charges depending on the behavior of the loads.

To regulate workload imbalance, pioneering work proposed a push-pull linear regulator [72] for V-S PDN. Although linear regulators have low area overheads, they suffer from poor power efficiency due to their resistive nature, especially when the current imbalance is large. More recent work proposed a push-pull switched-capacitor (SC) to recycle the charge imbalance between the stacked loads [59]. Because of their energy-storage capability, these SC converters provide higher power-efficiency at the cost of a larger silicon area dedicated to the capacitors. Besides linear regulators and SC converters, off-chip bidirectional buck-boost converters have also found usage as differential converters for other types of applications such as power management in photovoltaic cells [79]. Since various surveys and comparisons of switching regulators in the literature [82] show that, with the rapid improvement of capacitive technology, switched-capacitors are going to surpass inductive converters, we focus on SC converters in this chapter and leave the study of inductive

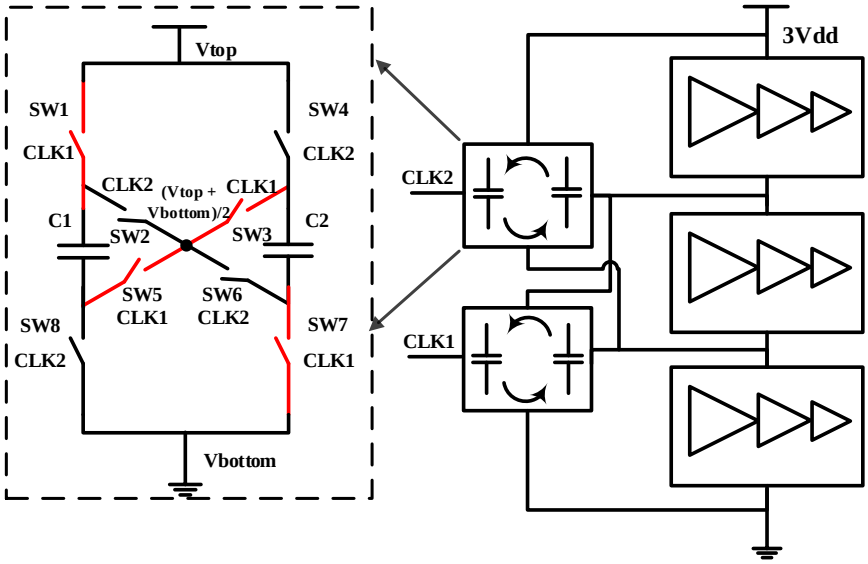

Figure 5.1: Stacked loads (three layers) with stacked SC converters (two), ideally providing Vdd voltage headroom to each load. Zoomed up single cell of 2:1 push-pull SC converter shown on the left.

converters for future work. Figure 5.1 shows the structure of the SC converters we implement/model in this chapter. It involves two fly-capacitors (C1 and C2) interchanging their positions periodically, thereby shuttling excess charge between the stacked loads to "source" or "sink" them as the loads demand. To support many-layer 3D-IC, we extend this converter for two stacked loads [59] into a scalable, multi-output ladder SC.

5.2.3 System-level Evaluation of V-S

Although researchers have previously identified V-S as a promising solution to alleviate the power delivery constraints in 3D-IC [75], the impact of V-S on PDN current density and the resulting implications to PDN reliability has not been closely investigated. In this chapter, we build a system-level PDN model and simulate an example many-core 3D processor to directly compare the EM-induced lifetime of regular and V-S PDNs. Another important aspect of V-S design is the voltage noise at the intermediate nodes. Zhou et al. proposed a whole-system PDN model to study SC converters' impact on power delivery noise [107]. However, they only studied the traditional 2D-IC case without voltage stacking. To the best of our knowledge, a system-level noise evaluation

for SC-converter-supported V-S PDNs is missing from the literature. Adopting a design methodology similar to [107], we combine a resistive model of SC converters [77] with a 3D extension of VoltSpot to evaluate V-S PDN's noise level. Since the supply noise in V-S PDN is strongly correlated with the workload-imbalance between the adjacent layers [59], we also examine a large range of workload-imbalance and quantify its impact on voltage noise, system power efficiency, and PDN area overhead.

5.3 Modeling Methodologies

5.3.1 SC Converter Modeling

A cross-layer design exploration of the benefits and overheads of V-S in 3D-IC requires incorporating circuit-level insights with architecture-level study. To accurately capture the power efficiency, output voltage drop and area overhead of SC converters, we implement a 2:1 push-pull SC converter (as shown in Figure 5.1) in a commercial 28nm CMOS technology. It has integrated fly capacitors (8nF total), an optimum switching frequency of 50MHz, and 4-way interleaving. Each SC converter can provide up to 100mA current to the load. Using the Cadence ADE environment and Spectre simulator, we simulate this converter and use the results to derive a compact model for system-level exploration.

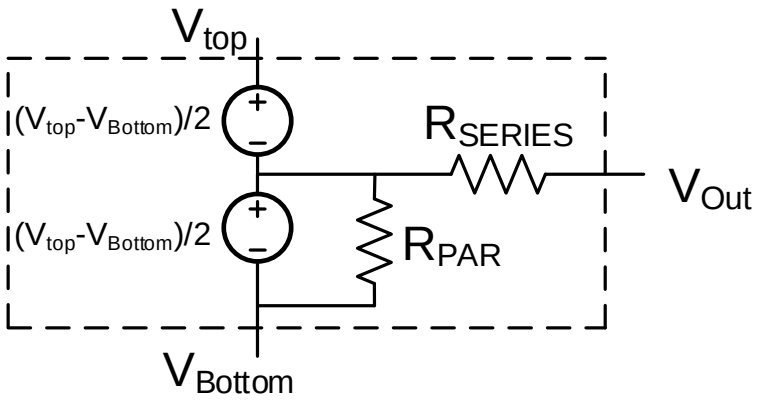

Figure 5.2: Power efficiency and voltage drop model for SC converters.

Figure 5.2 shows the efficiency and noise model for the SC converters. We adopt an analytical

methodology introduced in [77]. Based on the switch topology, the charge multiplier vectors ($a_{c,i}$ and $a_{r,i}$) are derived to calculate the slow (R_{SSL}) and fast switching (R_{FSL}) asymptotic limits of SC converter output impedance. The optimized R_{SSL} and R_{FSL} are given as:

$$R_{SSL} = \frac{1}{C_{tot} f_{SW}} \left(\sum_{i}^{n} |a_{c,i}| \right)^2 \tag{5.1}$$

$$R_{FSL} = \frac{1}{G_{tot} D_{cyc}} \left(\sum_{i}^{n} |a_{r,i}| \right)^2 \tag{5.2}$$

where C_{tot} is the fly capacitance, G_{tot} is the total switch conductance, f_{SW} is the switching frequency and D_{cyc} the duty cycle (assumed 50%) of SC Converter switching cycle. The R_{SERIES} in Figure 5.2 captures the switching and conductance losses while R_{PAR} captures the various parasitic losses of switch parasitic capacitance, bottom-plate capacitance and gate-drive loss. This model also captures the resistive voltage drop of the SC converters through R_{SERIES} , which can be calculated as: $R_{SERIES} = \sqrt{{R_{SSL}}^2 + {R_{FSL}}^2}$. For the SC converter we implemented, $R_{SERIES} = 0.6\Omega$.

As shown in Figure 5.1, the voltage-headroom (i.e., the potential difference between V_{Top} and V_{Bottom}) of the SC converters in many-layer 3D-ICs is dependent on the adjoining layers' workload imbalance. In order to incorporate this dependency in our cross-layer study, we make both V_{Top} and V_{Bottom} as inputs to our SC converter model and calculate the ideal output voltage (i.e., without the IR drop on R_{SERIES}) as $(V_{Top} + V_{Bottom})/2$.

To verify the accuracy of this model, we compare the estimated power-efficiency and output voltage drop against circuit simulation results of a SC converter for a 2-layer 3D-IC under fixed capacitance and different load current. We test two different frequency modulation strategies. The closed-loop scheme modulates SC converter's switching frequency dynamically with the load current while the open-loop control scheme keeps the frequency constant at all time. Figure 5.3 shows that our model accurately captures power-efficiency and output voltage drop for both control policies. According to Figure 5.3, closed-loop converters have higher power-efficiency. However, because it requires the implementation of feedback loops, the closed-loop policy is more complex to model. For simplicity, we use open-loop SC converters in our 3D V-S PDN analysis and leave

the evaluation of closed-loop control for future work.

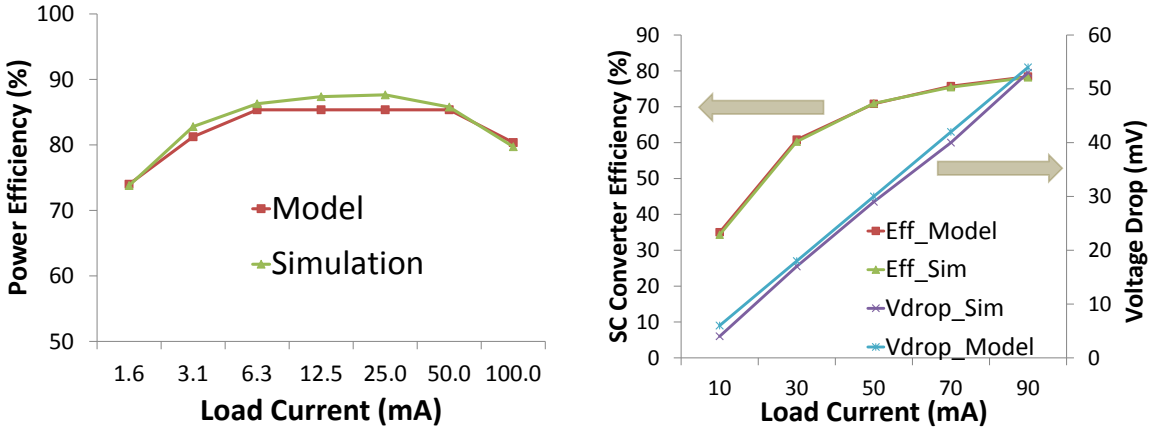

(a) Close-loop Control

(b) Open-loop Control

Figure 5.3: Model validation results.

We implement our SC converters with MIM capacitors and the resulting area of each converter is $0.472mm^2$. Considering the fact that the fly-caps contribute to the majority of SC-converters' area and MIM capacitors have low density, we also calculate the converters' area overhead with other high-density integrated capacitors. For example, if implemented with ferreolectric [82] or trench capacitors [66], the area of each converter would be $0.102mm^2$ or $0.082mm^2$.

5.3.2 PDN Modeling for 3D-IC

In order to quickly explore the multi-dimensional space of 3D-IC's PDN design and evaluate the cost and benefits of different design scenarios, we extend VoltSpot to support 3D-IC. Figure 5.4 illustrates our extensions.

To model the traditional PDN for 3D-IC, we simply add more layers of silicon on top of each other and connect all layers' Vdd nets and ground nets with TSVs (Figure 5.4a). To model V-S PDN, we connect all layers' Vdd nets and ground nets in series with regular TSVs and provide the off-chip supply voltage (i.e., the single layers' Vdd multiplied by the number of layers) to the top layer via through-TSVs (Figure 5.4b). TSVs are modeled as resistors. The resistive model for SC converters has been described in Section 5.3.1. We uniformly distribute them within each core.

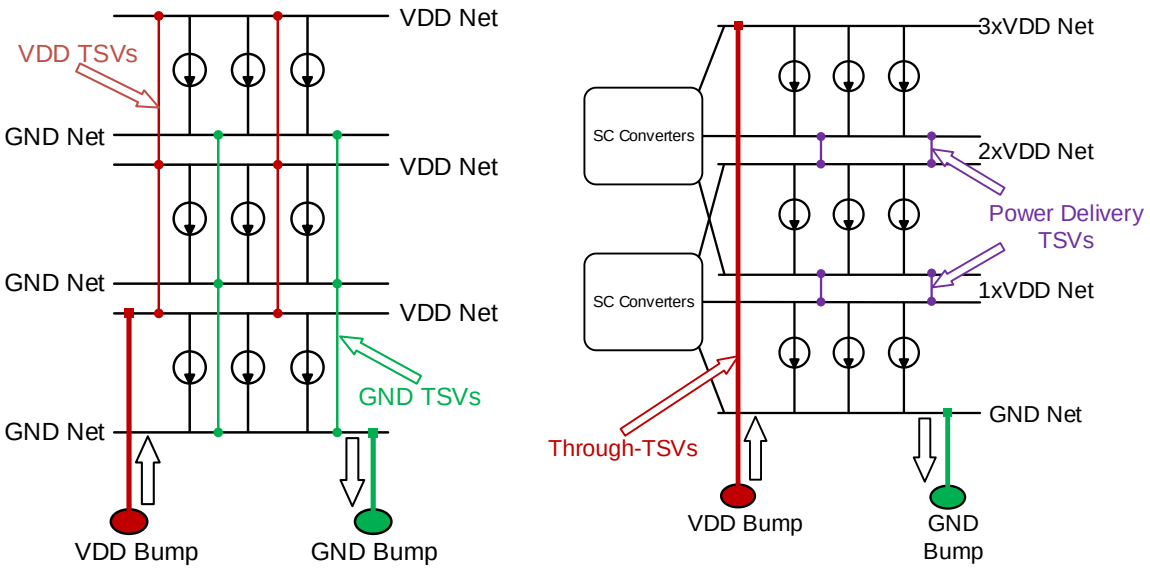

(a) Conventional PDN (b) V-S PDN

Figure 5.4: Conventional and Charge-recycled power-delivery scheme for 3D-IC.

With the fine-grained pre-RTL modeling capability of VoltSpot, this extended 3D-IC PDN model provides a detailed current profile for both the C4 pad and TSV arrays. It also captures on-chip IR drop for both regular PDN and V-S PDN under given workload behaviors. This model provides a key link to the tool chain that allows designers to explore the complex tradeoff space that involves power delivery architecture, C4 pad/TSV allocation, voltage regulation scheme, PDN noise/reliability, and workload characteristics.

5.3.3 EM-induced System Lifetime Calculation

As we discussed in previous chapters, a metal conductor's EM-induced lifetime follows a lognormal distribution and the mean-time-to-failure can be estimated with Black's equation [5]. For a group of conductors (e.g., C4 pad or TSV array), all elements are subject to EM-induced wearout. Equation 5.3 calculates the cumulative distribution function (CDF) of multiple metal conductors' EM failure probability $P(t)$ based on the failure probably of each conductor $F_i(t)$.

$$P(t) = 1 - \prod_i (1 - F_i(t)) \tag{5.3}$$

Using the detailed per-pad/TSV current information generated by our PDN model, we first determine $F_i(t)$ for each pad/TSV's CDF. After that, we calculate $P(t)$ and use the time value which makes $P(t) = 0.5$ as a lifetime estimation that represents the whole pad/TSV array's expected lifetime until the first EM-induced failure. We will use this metric (expected EM-damage-free-lifetime) in the remainder of this chapter to evaluate PDN's robustness against EM stress.

5.4 Simulation Setup

5.4.1 Many-core Processor Modeling

In order to establish realistic 3D-IC design scenarios for our cross-layer exploration, we select a 40nm, dual-core ARM Cortex A9 IP implementation running at 1.0 GHz [1] and replicate it 8 times to build a single-layer, 16-core processor. The reason for selecting ARM processors is that they are power-efficient and therefore can be used to build many-layer 3D-ICs without relying on aggressive, volumetric cooling solutions. We use McPAT [53], an architecture-level power/area/timing model to derive the area and power consumption of the single-layer processor. The processor floorplan was generated by ArchFP [13]. With a supply voltage of 1V, this 1GHz single-layer processor has a peak power consumption of 7.6 W and an area of 44.12 mm^2. Figure 5.5 shows the processor's floorplan generated by ArchFP [13].

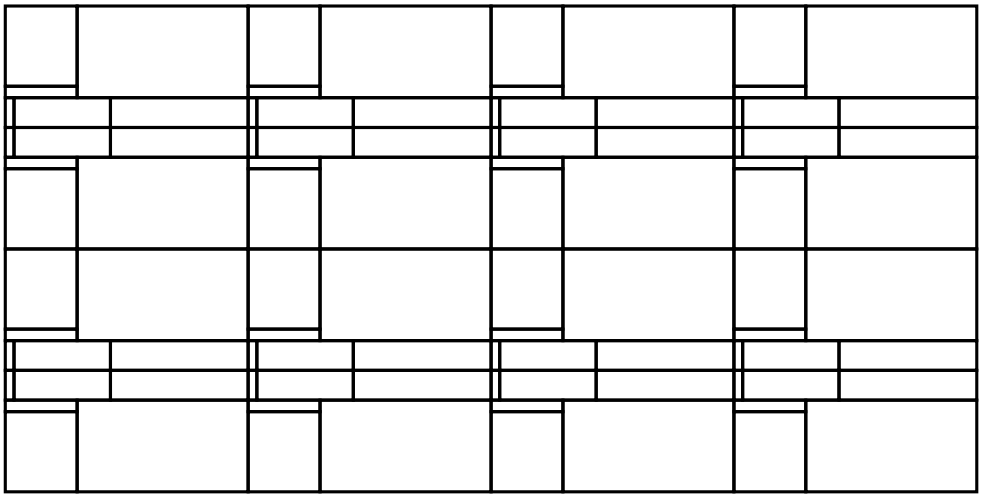

Figure 5.5: Floorplan of our single-layer 16-core processor.

Although many-layer, especially many-logic-layer 3D-ICs pose various fabrication challenges [75], the possibility of manufacturing 3D stacks economically has been exemplified by existing commercial products (e.g., the Micron hybrid memory cube with 4-8 layers [62]). To study the voltage noise in both short-term and long-term future 3D-ICs, and to evaluate how 3D scaling affects PDN design tradeoff, we build a series of example 3D systems with 2 to 8 layers stacked together. With the help of a pre-RTL thermal model, HotSpot [81], we find that we can build 3D-ICs with up to 8 layers of our example 16-core processor while maintaining the hotspot temperature below 100 Celsius (which is a typical upper limit [81]) with a conventional air-cooling solution. With 16 ARM cores per layer, the peak power consumption of these 3D processors ranges from 30.4W to 60.8W.

5.4.2 PDN Modeling and TSV Configurations

One of the major extensions we made to VoltSpot in this chapter is adding an explicit model for TSVs. The diameter, pitch and resistance values of TSVs come from prior work [43]. As suggested by Pathak et al. [65], the thermal stress generated by TSVs could potentially impact the electrical performance of the nearby transistors. Therefore each TSV requires a keep-out zone (KoZ) to space away other active devices. We use the size of this KoZ to calculate the TSV array's total area occupancy. We assume that all TSVs have equal size and resistance, and they are uniformly distributed within each silicon layer. Other PDN modeling parameters are adopted from Chapter 3 and listed in Table 5.1.

C4 Pad Pitch (μm)	200
C4 Pad Resistance ($m\Omega$)	10
Minimum TSV Pitch (μm)	10
TSV Diameter (μm)	5
Single TSV's Resistance ($m\Omega$)	44.539
TSV Keep-Out Zone's Side Length (μm)	9.88
On-chip PDN's Pitch,Width,Thickness (μm)	810,400,720

Table 5.1: Major PDN modeling parameters

The number of TSVs allocated for PDN is a design parameter for system designers. More TSVs

provide more vertical current delivery channels, therefore reducing both average TSV current and the effective inter-layer PDN resistance. At the cost of higher area overhead due to the KoZs, increasing the number of power-supply TSVs not only reduces voltage noise, but also improves the TSV array's reliability against EM-induced wearout. To explore the tradeoff between power delivery quality and TSVs' area overhead, we examine three TSV topologies in our study that represent a conservative (Dense), an aggressive (Few) and an average (Sparse) design scenario. Table 5.2 gives more details about each configuration's TSV count and area overhead.

	Effective Pitch(um)	Number of TSVs per Core	Total Area Overhead
Dense TSV	20	6650	24.2%
Sparse TSV	40	1675	6.1%
Few TSV	240	110	0.4%

Table 5.2: TSV configurations used in this study.

Besides modeling the PDN of 3D-ICs, we also evaluate 3D-ICs' impact on chip temperature using a pre-RTL thermal model HotSpot [81]. HotSpot uses compact RC elements to model the heat transfer within a silicon chip and between chip and the cooling system. It is capable of estimating the temperature of different layers in a 3D-IC using conventional, passive cooling solution (e.g., air or liquid cooling using a heat sink). More advanced volumetric cooling solutions like micro-channel cooling is not supported yet. However, since those technologies will further reduce the constraints of heat removal, utilizing them would only make power delivery an even more severe problem, which further motivates our charge-recycled power delivery scheme. The evaluation of advanced cooling technologies is left for future work.

5.5 Results and Discussions

5.5.1 Heat Removal and Power Delivery in 3D-IC

The technology of 3D-ICs bring challenges to both power delivery and heat removal. Using our system-level simulation infrastructure, we examine both temperature and on-chip IR drop of our

example many-core 3D processors. Figure 5.6 shows the temperature results with both air cooling and a more aggressive liquid cooling technique, and the voltage noise evaluations of different PDN strategies and different power-supply TSV topologies. Since the basic building block of our 3D processor is a power-efficient ARM core, the power density of each layer is relatively low ($0.17W/mm^2$). Therefore the traditional convection-based air cooling is sufficient to keep the maximum temperature of a silicon stack with up to eight layers under 100 Celsius degrees (which is normally considered the upper limit [81]). In this chapter, we adopt an air cooling solution for its low cost and thus evaluate 3D processors with up to eight layers.

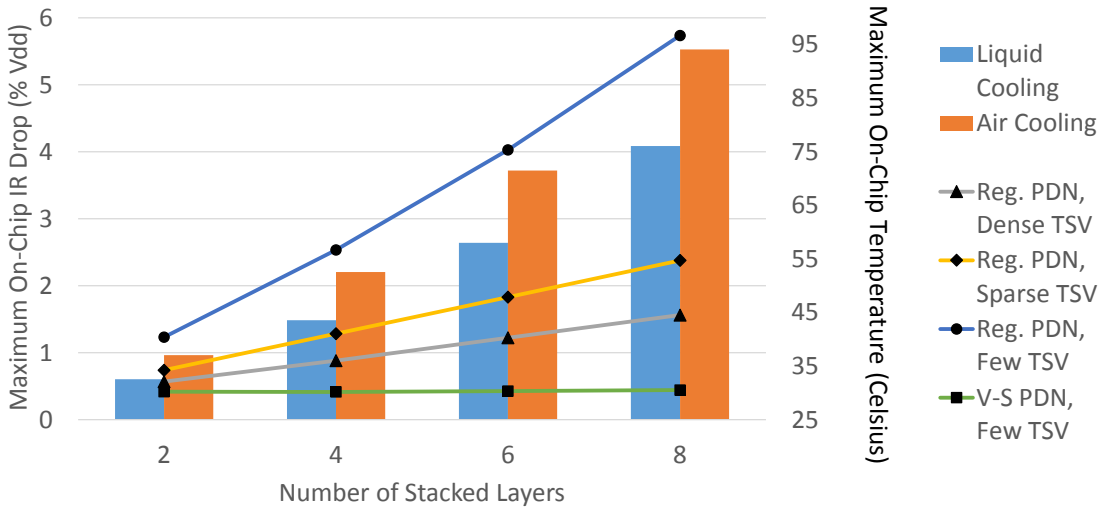

Figure 5.6: Temperature and voltage noise evaluation of our example many-core 3D processors. For voltage noise, we tested different power-supply TSV topologies described in Table 5.2.

The lines in Figure 5.6 illustrate the 3D-ICs' maximum IR drop with both V-S PDN and traditional PDN. Although utilizing the "Dense TSV" topology without V-S only incurs an IR drop of 1.56% Vdd in a 8-layer 3D processor, this noise level is still more than three times higher than the noise of a V-S PDN with the "Few TSV" topology. More importantly, V-S PDN's voltage noise shows a perfect scalability such that adding more layers does not increase IR drop at all. These results prove that V-S provides a scalable and fundamental solution for 3D-IC's power delivery challenge. We note that the IR drop for the V-S case is calculated based on the scenario where all layers are fully active. Although such behavior is the worst case for traditional PDN in terms of voltage noise, it is in fact the best case for V-S PDN. This is because when all cores are consuming

the same power, the voltage regulators do not provide any current and therefore no voltage drop will incur at the regulators' outputs. We will further investigate the general cases where the workloads are not balanced in Section 5.5.3.

5.5.2 EM-Induced TSV/C4 Pad Lifetime

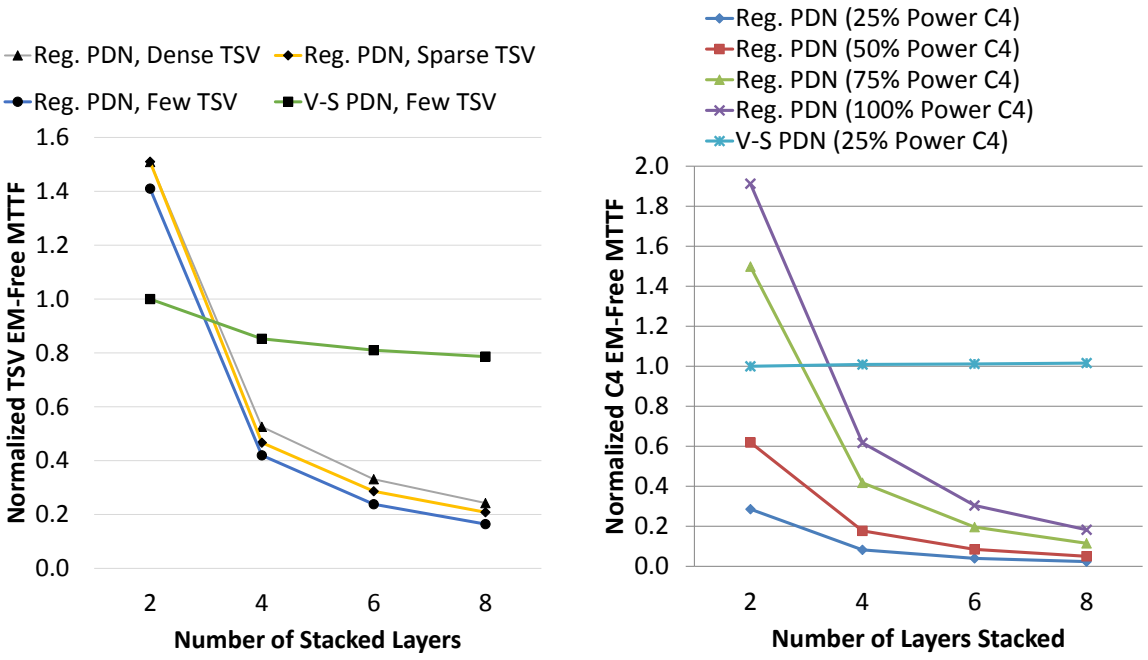

(a) Power-Supply TSVs

(b) Power-Supply C4 Pads

Figure 5.7: EM-induced lifetime evaluation. All results are normalized to the lifetime of 2-layer V-S PDN 3D-IC.

Using the methodology described in Section 5.3.3, we evaluated the expected EM-damage-free lifetime for both regular and V-S PDN's TSV (Figure 5.7a) and C4 pad (Figure 5.7b) arrays. And show that the V-S provides significant benefits. As we stack more layers, the increasing current density significantly reduces the lifetime of the regular PDN's TSV array by up to 84%. At the same time, the V-S PDN's MTTF only slightly degrades. This is because while the current density of TSVs in the V-S PDN is independent of layer count, adding more layers still requires more TSVs to support them, which increases the risk of TSV failures. We also observe that the V-S PDN's TSV array has a shorter lifetime compared to the regular PDN when the number of stacked layers

is small (e.g., 2 layers). This is because in the V-S PDN, we connect each Vdd C4 pad with only one TSV, to provide supply voltage/current directly to the top-layer. Since the number of Vdd pads (32 per-core in this case) is smaller than the number of Vdd TSVs in a regular PDN (55 per core in the "Few TSV" case), the Vdd TSVs in the V-S PDN have higher average current, and this limits the whole-system MTTF. Regardless of this side effect, the EM-induced lifetime of V-S PDNs in 3D-ICs with more layers still surpasses that of the regular PDN by more than 3x.

Similarly, the regular PDN's C4 MTTF quickly degrades with 3D-IC scaling. For the V-S PDN, stacking more layers neither increases the number of total pads, nor raises the total off-chip current demand of the V-S PDN, and therefore its C4 array's EM-damage-free lifetime is independent of layer count. For the 8-layer 3D processor, the gap in the C4 array's MTTF between the V-S PDN and the regular PDN can be up to 5x. This indicates that, because V-S extends the pad array's EM lifetime, it reduces the requirement for power supply pads and allows more pads to be used for I/O. Please note that since the C4 array's EM robustness is insensitive to the TSV topology, we use a fixed topology ("Sparse TSV") in all the evaluations.

Another interesting observation is that, for the regular PDN, adding more TSVs or C4 pads only marginally increases MTTF. Even with aggressive allocations (e.g., "Dense TSV" topology or even allocating 100% of pads as power supply), the regular PDN's MTTF is still far inferior to that of the V-S PDN. We therefore conclude that for many-layer 3D-ICs, it is not feasible to improve the regular PDN's EM-robustness to the same extent as with the V-S PDN by simply allocating more power-supply TSVs and C4 pads.

5.5.3 Load-Imbalance-Induced Voltage Noise

Integrated-voltage-regulation is necessary in V-S PDN, because when the current consumptions of two adjacent layers do not match, the voltage regulators need to either provide or sink the difference. This introduces extra voltage noise due to the regulators' output voltage drop (Figure 5.1) and the lateral impedance of the on-chip PDN. While larger workload-imbalance increases noise with higher current demand for the SC converters, having more regulators distributed across the silicon die reduces IR drop by amortizing the per-converter current load and reducing the average load-to-

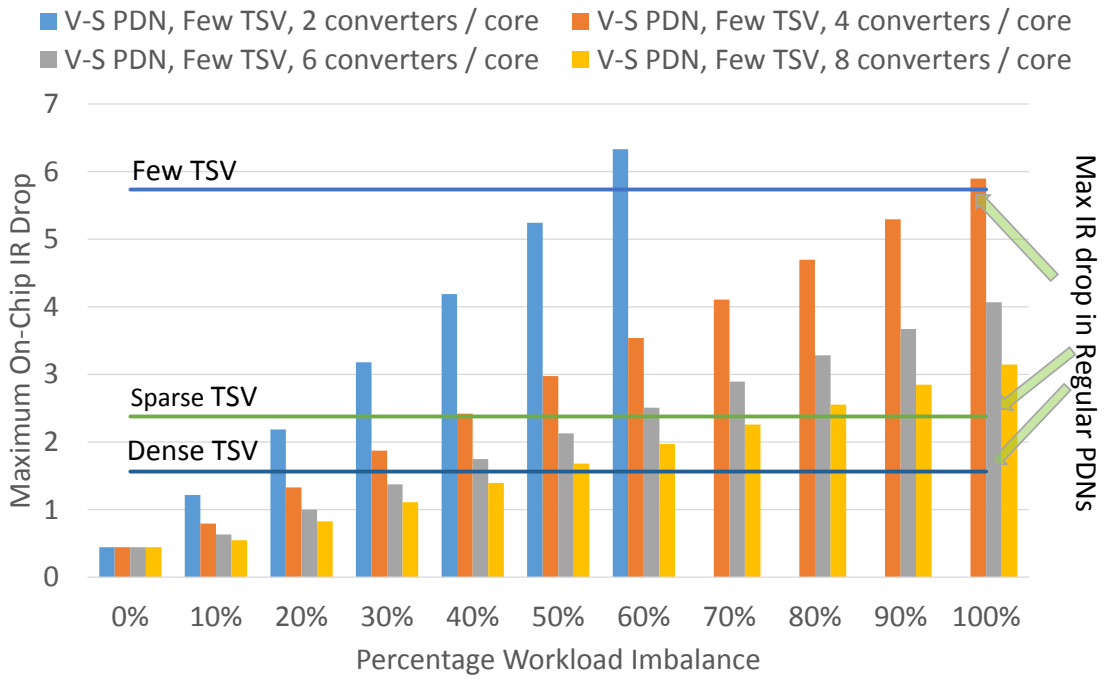

Figure 5.8: Voltage noise evaluation of our 8-layer processor. For 3D-ICs without V-S, the worst-case IR drop happens when all layers are fully active. Therefore the assumption about workload-imbalance does not affect those evaluations.

regulator distance. Figure 5.8 shows the noise levels of PDNs for a 8-layer 3D-IC under different regulator configurations and workload behavior conditions. We assume that the power consumption of the silicon layers has an interleaved "high-low" pattern, where the high-power layers are always fully active and the low-power layers consume X% lower dynamic power (e.g., 100% imbalance means that the low-power layers are idle and only consume leakage power). This pattern serves as a good benchmark, because it requires the converters on all layers to source/sink the same amount of current, therefore imposing the most stress on the PDN. We use it to study the worst-case noise of V-S PDNs. We note that since our SC converters are designed to have a maximum load of 100 mA, Figure 5.8 skips all data points that violate this limit.

The lines in Figure 5.8 illustrate the maximum on-chip IR drop of regulator PDNs with different TSV configurations. Regular PDNs rely on TSVs to provide all current to all layers, and therefore the worst-case IR drop always happens when all layers are fully active. For this reason, regular PDNs' maximum IR drop results are irrelevant to the imbalance of workloads. Since adding one

SC converter to an ARM core incurs around 3% area overhead (assuming the converters are implemented with high-density capacitors discussed in Section 5.3.1), a V-S PDN with 8 converters per core and "Few TSV" topology occupies the same area as a regular PDN with "Dense TSV" topology. If we compare the voltage noise of theses two cases, we find that the V-S PDN has lower IR drop when the workload-imbalance ratio is lower than 50%. When larger imbalance exists, V-S PDN's IR drop surpasses regular PDN by up to 1.58% Vdd.

To give an example of workload-imbalance in full applications, we simulate the Parsec 2.0 benchmark suite [3] with performance simulator Gem5 [4] and adopt the methodology of statistical sampling from prior work [94]. We simulate one thousand 2k-cycle samples from each application and calculate their average power consumption with McPAT. Figure 5.9 shows the distribution of each application's power consumption. The top/bottom bars in Figure 5.9 represent the max/min values of each distribution. The edges of the boxes are the 25th and 75th percentiles, and the central marks are the medians. We observe that although the samples from different applications have large differences in power consumption, the samples from the same application show much smaller variance in demand. For example, while the maximum workload imbalance among the samples from all applications is more than 90%, the best-case application (blackscholes) shows a maximum imbalance of 10% across all its samples. On average, the applications have a maximum-imbalance ratio of 65%, which makes the V-S PDN's IR drop only 0.75% larger than the regular PDN. These results also indicates that by scheduling different instances of the same application, or different threads from the same instance onto the cores in the same core-stack, we can reduce the workload-imbalance and a V-S PDN's noise. [1]

In summary, the voltage noise in V-S PDN is dependent on the imbalance of different layers' power consumptions. With large workload-imbalance, V-S PDN experiences more severe IR drop than regular PDN. Furthermore, unlike regular PDN, V-S PDN's voltage noise is insensitive to the number of layers. With the advance of cooling technologies that allows designers to stack more layers, V-S PDN will out-perform regular PDN in terms of noise even in the presence of large

[1] Prior work have shown that run-time task scheduling can further reduce workload-imbalance [41]. Also, GPU and SIMD architectures are designed to exploit data-level parallelism, which implies more balanced behavior. These topics are beyond this chapter's scope, but provides interesting directions for future work.

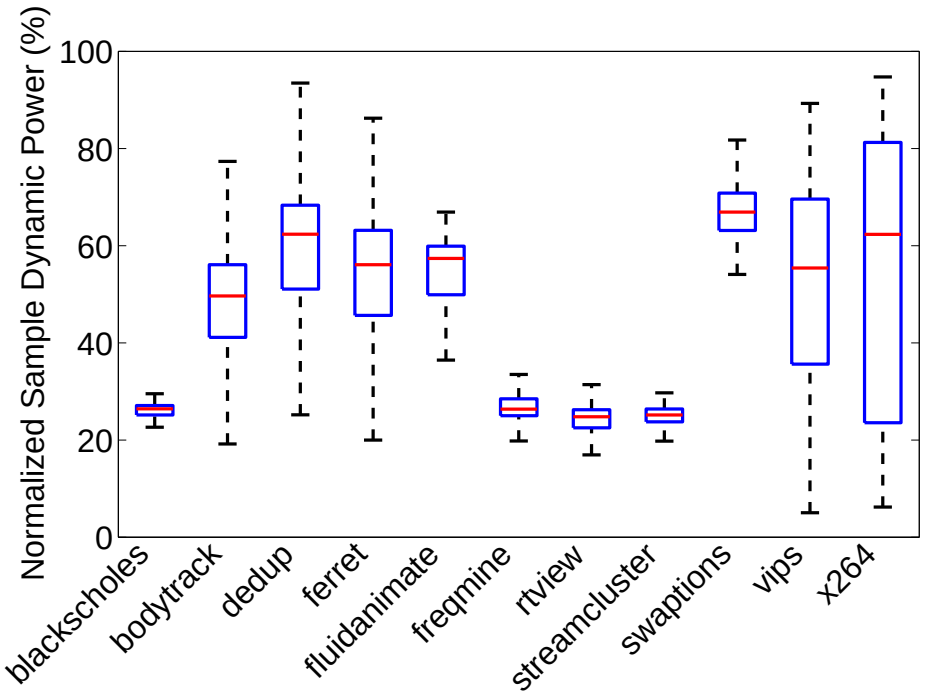

Figure 5.9: A box-plot that shows the distributions of workload imbalance within and across different applications.

workload imbalance.

5.5.4 System Power Efficiency

Figure 5.10 shows the power efficiency (i.e., the total power consumed by the processors divided by the total power drawn from the off-chip power source) results for 3D-ICs with V-S PDN. As the amount of workload-imbalance increases, the SC converters need to compensate by delivering more power. Consequently, the power overhead of voltage regulation increases. When we compare V-S PDNs with different numbers of SC converters, we observe that increasing the number of converters reduces power efficiency. As we discussed in Section 5.3.1, this is because our open-loop converters do not modulate their switching frequency at run-time, and therefore each converter's efficiency reduces as more converters are allocated to share the current load. Closed-loop control is an area for future work. Considering the fact that placing more converters can reduce on-chip IR drop, the allocation of SC converters in V-S PDN becomes a tradeoff between on-chip voltage noise and system-level power efficiency. Our models can help designers to choose the optimal design point

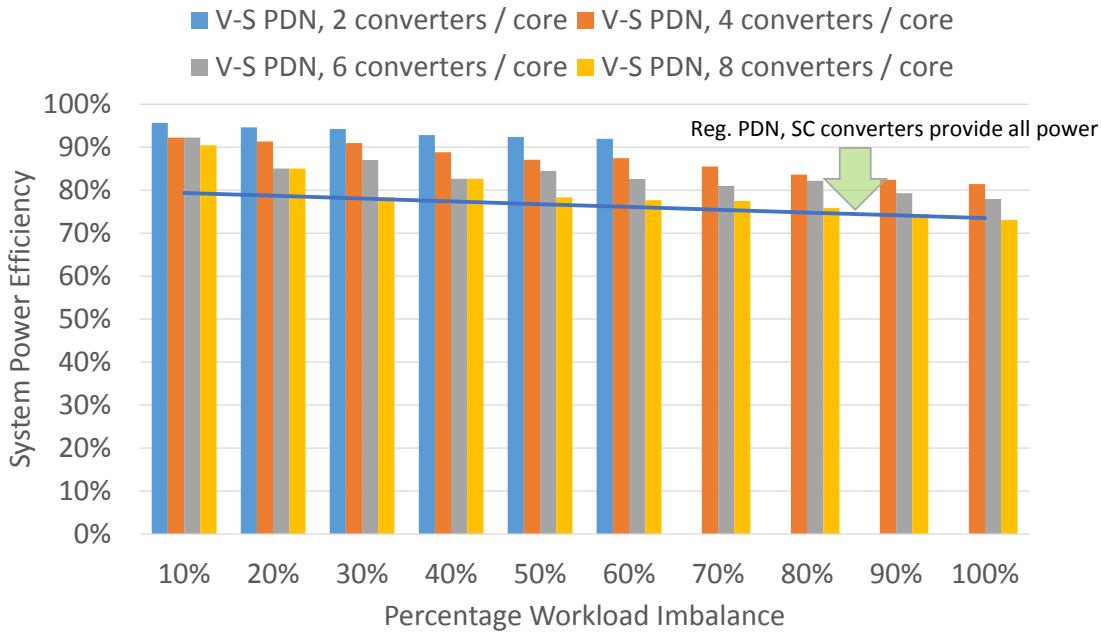

Figure 5.10: Power efficiency of 3D processors with V-S PDN.

based on their specific design objectives.

Figure 5.10 also shows the power efficiency of using SC converters in 3D processors with regular PDN. Unlike V-S PDN, where the voltage regulators only need to compensate for the differential power consumption between layers, SC converters in regular PDN have to provide current to all layers. As a result, V-S PDNs have higher power efficiency.

5.6 Summary

3D-IC provides an essential mechanism for the industry to stay on the historical scaling trend of device integration while raises power delivery challenges with reduced EM-lifetime and increased voltage noise. In this chapter, we build a system-level PDN model for 3D-ICs to study a charge-recycled, voltage-stacking PDN structure and compare it with the regular, non-voltage-stacked PDNs in the context of 3D-IC. Our EM-robustness analysis on both C4 pad and TSV arrays indicates that V-S PDN's EM-induced MTTF significantly surpasses (e.g., 5x longer) the regular PDN's lifetime. By implementing and validating a resistive model for V-S PDN's voltage regulator (i.e., SC converters), we analyze the on-chip voltage noise and observe that with the same total

area overhead, the V-S PDN has lower IR drop than the regular PDN when the workload-imbalance ratio is below 50%. Under the average workload imbalance ratio extracted from full applications (65%), a V-S PDN's IR drop is no greater than 0.75% Vdd beyond the noise level of a regular PDN. Combined with the observation that both EM-lifetime and IR drop of V-S PDNs are insensitive to many-layer 3D-ICs' layer count, our study demonstrates that V-S provides a scalable and practical solution to the power delivery challenge in the era of many-layer 3D-IC.

This work is done in collaboration with Kaushik Mazumdar, a Ph.D. student in the Electrical and Computer Engineering department at University of Virginia. Kaushik's major contributions include: 1) the circuit-level implementation of the SC converter; 2) the structure and parameters of our resistive IVR model; 3) the circuit simulation results for validation. My contributions are: 1) extending VoltSpot to support 3D-IC; 2) integrate the SC converter model with VoltSpot; 3) whole-system PDN simulation and noise, EM-robustness calculation; 4) workload imbalance characterization.

Chapter 6

Transient Noise in Voltage-Stacked PDNs

6.1 Overview

In Chapter 5, we observe that with reduced current density in C4 bumps and through-silicon-vias, V-S significantly improves 3D-IC's robustness against EM-induced PDN wearout. However, V-S PDNs are not guaranteed to have lower supply voltage noise compared with the traditional power delivery scheme: when the power consumption in the various layers are significantly imbalanced, the voltages at the intermediate nodes in the V-S stack deviate from the nominal value. Although Chapter 5 explored V-S PDN's steady-state noise, it is necessary to study 3D-IC's transient noise before comparing the overall noise quality between V-S PDNs and traditional PDNs. This is because the transient simulation not only exposes the impact of all types of voltage noise (i.e., IR drop, LdI/dt, and LC resonance), but also allows us to explore the impact of PDN components like on-chip decap and package, which cannot be captured in steady-state simulations.

In this chapter, we first design and validate a transient model for the voltage regulators in V-S PDNs. We then integrate it with VoltSpot, producing the first platform to enable whole-system, transient simulation for many-layer 3D-ICs' V-S PDN. Using an example low-power, ARM-based manycore 3D processor, we then compare the supply noise between voltage-stacked and traditional PDNs, and explore the impact of (a) cross-layer noise, (b) on-chip decoupling capacitance, and (c) package impedance. This chapter's major contributions are:

- We design and integrate a compact model for SC converters with VoltSpot. To the best of

our knowledge, we are the first to develop a model that simulates transient voltage noise in 3D-IC systems using V-S PDNs.

- We demonstrates that unlike the traditional PDN in which voltage fluctuations can propagate from one layer to another, the V-S PDN provides stronger isolation for the cross-layer noise interference. As a result, V-S PDNs can better utilize run-time noise mitigation techniques such as dynamic margin adaptation [51] to improve system efficiency.

- Under different workload-behavior assumptions, we explore the impact of both on-chip capacitance allocation and the trend of 3D scaling on 3D-ICs' transient voltage noise. Using a layer-wise amortized metric that averages all layers' maximum noise, we observe that with the same on-chip area overhead for integrated capacitors, the V-S PDN provides up to 60% lower worst-case noise amplitude than the traditional PDN.

- By simulating packages with a wide range of impedance, we observe that the V-S PDN's noise level is insensitive to the quality of chip package. We therefore conclude that the V-S scheme can reduced 3D-IC's packaging cost.

Consequently, we conclude that V-S achieves lower noise and lower cost compared with traditional 3D PDNs.

6.2 Background and Related Work

6.2.1 Voltage Noise and Timing Margin in 3D-IC

As we discussed in Chapter 3, transistor delay is directly proportional to the source-to-drain potential differences [74]. Therefore, it is a common design practice to assign a timing margin to critical paths to avoid noise-induced timing errors. Besides allocated at design-time to guard against the worst-case scenario, the timing margin can also be dynamically adjusted to improve system efficiency. For example, Lefurgy et al. [51] proposed a technique that detects available timing margin at run-time with critical path monitors. Using digital phase-lock-loops, their scheme can rapidly

change clock frequency to save energy during average-case execution (i.e., reduce margin) while guaranteeing functionality in the worst case (i.e., increase margin).

As more layers of active device layers are stacked together, the aggregate current demand increases, and the amplitude of the voltage noise grows proportionally with the layer count if the PDN impedance is kept constant [23]. To maintain a traditional PDN's robustness against voltage noise, 3D-IC designers will have to keep increasing timing margin (which degrades system performance with lower clock frequency), and/or reducing PDN impedance (which increases PDN cost with extra area overhead for on-chip decoupling capacitance or higher packaging complexity). Unfortunately, neither of these two approaches are scalable to many-layer 3D-ICs.

6.2.2 Voltage Regulation in V-S PDN

V-S introduces extra voltage noise caused by the workload imbalance between device layers. This is because when layers are connected in series, the ratio of their effective resistances (which are inversely proportional to their power consumptions) directly affects the voltage levels of the intermediate nodes. Consequently, layers with higher power will experience greater voltage drops. To regulate this noise, prior work proposed using explicit regulators with V-S PDN [59, 72]. Considering the rapid improvement of capacitive technology, we focus on switching-capacitor (SC) converters in this chapter due to their regulation efficiency [82].

Figure 6.1 shows the detailed circuit structure of the V-S SC converter we adopt from the literature [59]. Each converter cell consists of two fly-capacitors (C1 and C2) and eight switches (SW1-8). By periodically interchanging the positions of the fly-caps (i.e., phase CLK1 and CLK2 in Figure 6.1), the SC converter can either "source" or "sink" the charge difference between the stacked loads to regulate the voltage at its output.

6.2.3 System-level Supply Voltage Noise Modeling

In the past, researchers constructed system-level models to examine the supply voltage noise in both 2D ([19, 103]) and 3D ([23, 101]) chips. While prior work has demonstrated that stacking more layers of active silicon using the traditional PDN structure will monotonically increase on-chip

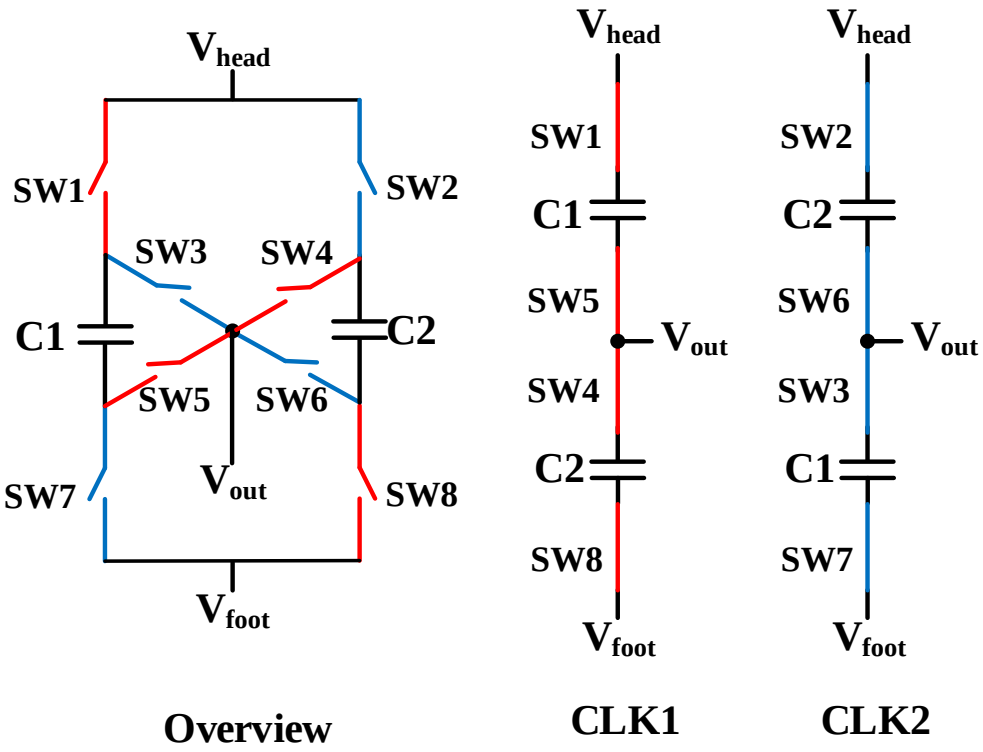

Figure 6.1: A single cell of our 2:1 push-pull SC converter and its equivalent circuits in the two different clock phases.

noise [23], it is still not clear whether, or in which scenarios the V-S scheme provides better power delivery quality (in terms of transient noise) for 3D-ICs. To answer this question, we build a whole-system evaluation platform for V-S PDNs by designing a compact RC model for SC converters and integrating it with VoltSpot.

The topic of SC converter modeling has been discussed in the past. However, prior work either focused on the traditional 2D-IC case without voltage stacking ([107]), or only studied the static noise (i.e., IR drop) of SC converters (our own work in [101]). To the best of our knowledge, ours is the first work to model transient voltage noise in SC-converter-supported V-S PDNs and compare V-S PDNs with traditional PDNs.

6.3 V-S PDN Modeling Methodologies

Despite that fact that contemporary, planar-IC processors' PDNs are already large systems with up to several billion nodes, 3D integration and voltage stacking further increase the PDN's complexity with more device layers and new components like TSVs and voltage regulators. Consequently, circuit-level simulations will be extremely computational-intensive and incapable of supporting whole-system design-space exploration studies. To enable a system-level study of V-S PDN's voltage noise, we design and validate a transient model for the SC converters and integrate it with VoltSpot. This section discusses our modeling methodology and the validation results.

6.3.1 A Transient Model for SC Converters

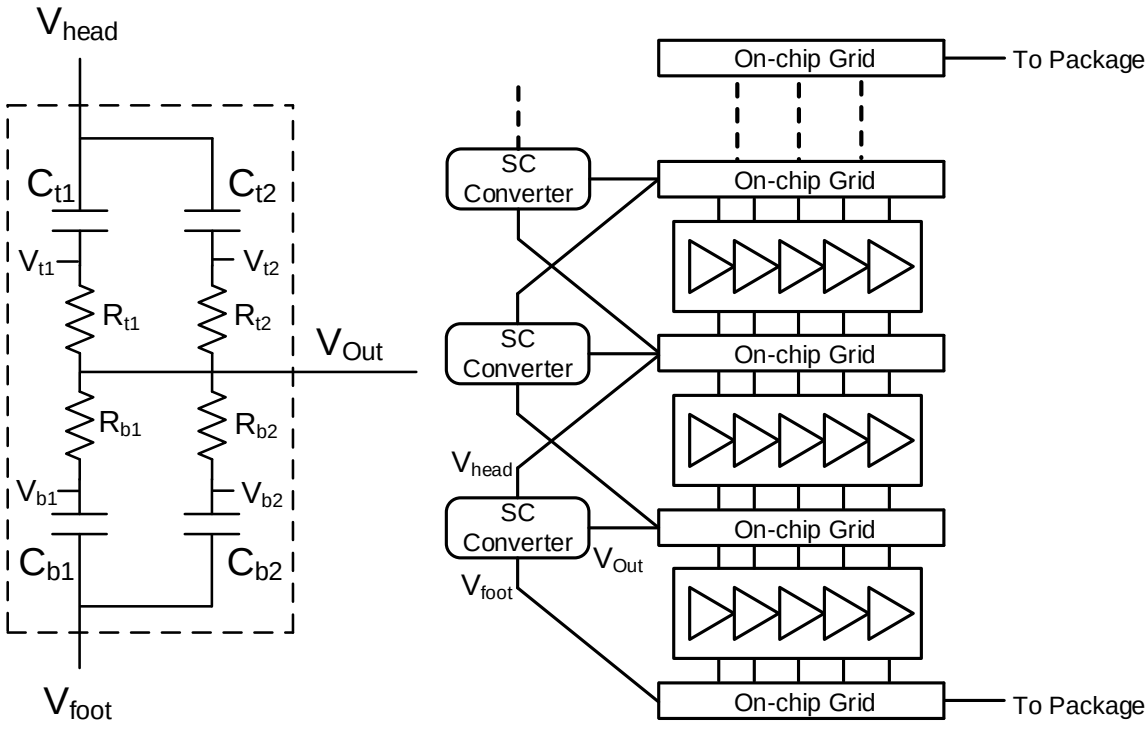

(a) Interleaved SC Converter (b) V-S PDN for Many-layer 3D-IC

Figure 6.2: An RC model for the interleaved SC converters and a whole-system view of a V-S PDN in a many-layer 3D-IC.

Figure 6.2a shows the compact RC circuit we use to model the interleaved SC converters. Each pair of top and bottom RC branches represent a cell of the converter that is controlled by a separate

clock signal. At each clock edge, we exchange the position of the top and bottom fly-caps to model the switching activities of the converter cell. That is, we calculate $V_{head} - V_{t1}{}' = V_{b1} - V_{foot}$, where $V_{t1}{}'$ is the voltage value after the clock edge while V_{b1} is the value before. Note that although we exchange the positions (i.e., electric charge) of the fly-caps at each clock edge, the resistance of each top and bottom branch is kept unmodified. This is because each time we "flip" the position of the fly-caps, we also change the set of switches to conduct the current (Figure 6.1). Fortunately, the switches are designed in a symmetric way such that both the top and bottom RC branch in the two different clock phases have the same equivalent resistance [59]. Therefore, we can collapse the eight switches into two resistors (R_t represents SW1&5 and SW2&6, R_b represent SW4&8 and SW3&7) and reduce the model's complexity. From circuit simulations, we extract that $R_t = 4.208\Omega$, and $R_b = 4.68\Omega$ ($R_t \neq R_b$ because NMOS and PMOS have different channel resistances).

As common design technique to smooth the output voltage ripple, designers divide the single-cell converters into multiple sub-cells and interleave their switching clocks [59]. To model this interleaved structure, we simply instantiate a pair of top/bottom RC branches for each sub-cell, scale the capacitance values according to the number of total sub-cells, and shift the phase of each sub-cell's control clock. Figure 6.2a illustrates an example model for a two-way interleaved SC converter. Similar to [59], we assume that all the sub-cells have identical structure, and therefore, the same RC values.

6.3.2 Validation

Similar to Chapter 5, we implement a 4-way interleaved, 2:1 push-pull SC converter in a commercial 28nm CMOS technology to validate our modeling methodology. It has an optimum switching frequency of 50MHz and a total capacitance of 8nF. Each SC converter can source/sink up to 100mA current to/from the load at a nominal voltage of 1V. Using the Cadence ADE environment and the Spectre simulator, we simulate this converter in a two-layer, voltage-stacked system (i.e., $V_{head} = 2V, V_{foot} = 0V$) and compare results against the output of our RC model.

Figure 6.3a shows the DC results comparison under constant workload conditions. Since the SC converter's output voltage is directly related to its output current, we attach an ideal current

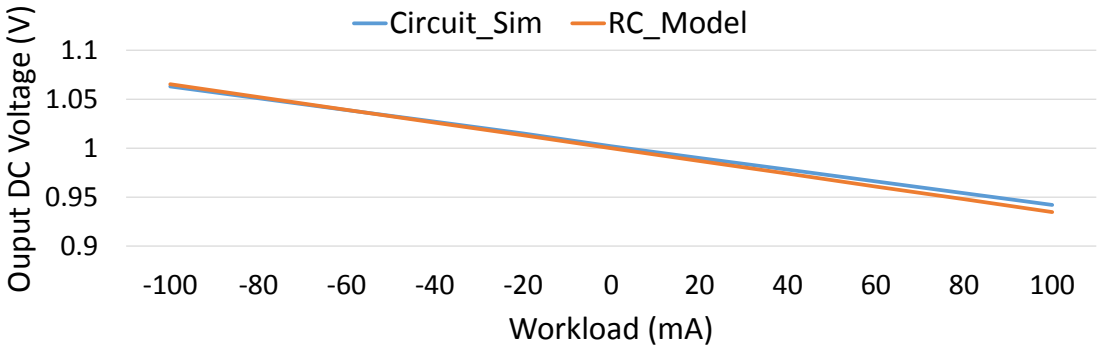

(a) Output DC Voltage

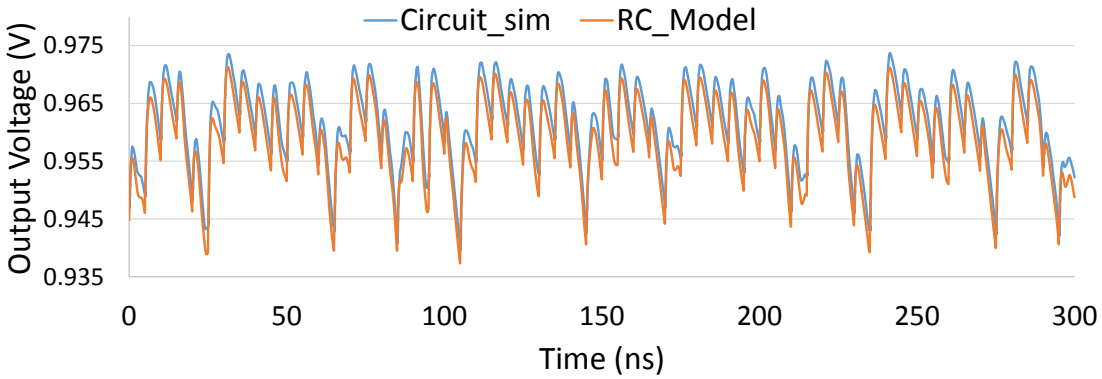

(b) Output Transient Voltage Trace

Figure 6.3: Validation results.

source directly to the V_{out} port and sweep the test cases from maximum sourcing (positive 100mA) to maximum sinking (negative 100mA). Under a constant workload, the output voltage shows a periodic rippling behavior caused by the converters' switching activities. Validation results show that our model's maximum DC error is 75mV, or 0.75% Vdd.

We also use a time-varying load current to validate our model. Figure 6.3b shows the output voltage trace over 300 ns. The load current is sampled from Parsec 2.0 benchmark raytrace [3]; it induces an average current of 66.3mA in an ARM Cortex A9 core. Over the entire simulated time window, the output voltage trace of our model matches well with circuit simulation in term of DC component, AC amplitude, and slew rate. Overall, our model can capture the SC converter's transient output voltage with less than 72mV error at all times.

6.3.3 Whole-system Model

To study the interaction between the SC converters and the on-chip PDN grid, and to evaluate V-S PDNs' overall noise quality, we combine our SC converter model with VoltSpot. Section 6.4.2 will discuss the parameters we use and the modifications we made to VoltSpot in detail. Figure 6.2b shows the structure of the whole-system model we build for many-layer V-S PDNs. For each SC converter, we connect its three ports (i.e., V_{head}, V_{out}, and V_{foot}) to three consecutive layers in the voltage-stacked power grids. We note that ideally, $V_{out} = (V_{head} + V_{foot})/2$, which indicates that any change in either V_{head} or V_{foot} will also affect the regulator's output voltage. Our model directly captures this inter-layer voltage dependency.

6.4 Simulation Setup

6.4.1 Many-core Processor Modeling

In this chapter, we utilize the same set of example power-efficient 3D processors used in Chapter 5 to study transient noise. Please refer to Section 5.4 for more detailed descriptions.

6.4.2 PDN Modeling

Besides integrating our SC converter model with VoltSpot, we also extend it to support transient simulations for 3D-IC. Here, the major modification is an explicit resistor-inductor model for the TSVs. We adopt TSV parameters from prior work [43]. Similar to prior work [26], we ignore TSV capacitance in this chapter because it is usually orders-of-magnitude smaller than the on-chip and package decoupling capacitance. Other PDN modeling parameters are also adopted from prior work [103] and listed in Table 6.1.

By default, VoltSpot utilizes ideal current sources to model the load (i.e., switching transistors). In order to model the voltage-stacked PDN organization, we replace the current sources with time-varying resistors. This is a necessary modification, because V-S PDN connects multiple layers of load in series, and using a resistive load model eliminates potential current source cutsets (if it exists, the solution is not unique) in the modeling circuit. The load resistance is calculated as

Minimum C4 Pad Pitch (μm)	150
Single Pad Resistance/Inductance ($m\Omega/pH$)	10 / 7.2
Minimum TSV Pitch (μm)	10
Single TSV Resistance/Inductance ($m\Omega/pH$)	44.5 / 36.3
On-chip PDN's Pitch,Width,Thickness (μm)	810,400,720
Package Capacitance (μF)	7.3
Package Resistance/Inductance ($m\Omega/pH$)	0.054 / 10.8

Table 6.1: Primary PDN modeling parameters

$R = Vdd^2/Power$. This modification increases the model's computational complexity with more frequent LU-decomposition operations. This is because, unlike the previous VoltSpot where the modeling circuit is time-invariant (only the current excitation changes), this new model changes the load resistors over time to match the power consumption. To explore a broader design space within an affordable simulation time (e.g., 1 hour to simulate 1k cycles), we adopt the modeling methodology from Huang et al. [26] and only simulate a "slice" of the entire 3D stack. Since each layer of our example 3D processor is a homogeneous 16-core ARM chip, we utilize the symmetry and simulate a reduced system of 2 cores per layer.

6.4.3 Workload Modeling

Using an integrated tool flow that combines McPAT with performance simulator Gem5 [4], we simulate the Parsec 2.0 benchmark suite [3] and extract dynamic power consumption traces to build realistic test cases for our noise study. Due to the simulation-speed constraint, we simulate 2k-cycle-long samples of power traces instead of whole-applications (which normally execute for billions of cycles). We collect a pool of workload samples by taking a large number of snapshots (i.e., 1000) from each benchmark. We then profile each sample's average power consumption and maximum noise amplitude when running alone (on a 2D-IC). This sample pool, along with the profile information, helps us to construct different multi-layer load patterns to study 3D-IC's voltage noise. Section 6.5 will give more details about the specific load combinations we use in different explorations.

6.5 Results and Discussions

6.5.1 Cross-layer Noise Interference

The technology of 3D stacking significantly improves silicon chips' level of integration because through-silicon-vias provide high-density connections between different layers. However, while signal TSVs enable low-latency cross-layer communications, power-supply TSVs bring active silicon closer to each other. As a result, large supply-voltage noise triggered by one layer can propagate to other layers and undermines the whole stack's robustness against noise-induced timing error.

To study whether or how different layers' voltage variations affect each other in traditional and V-S PDN, we pick one noisy workload and three less noisy ones from our sample pool and assign them to our 4-layer example 3D processor. The first row in Table 6.2 shows each workload sample's maximum noise amplitude when running alone on a single-layer chip. Figure 6.4 shows each layer's maximum voltage drop (%Vdd) over time.

In the traditional PDN (Figure 6.4a), voltage noise in all layers is clearly highly correlated, a consequence of the layers' high-density, parallel interconnection. Supply voltage fluctuations in one layer affect the entire 3D stack through the vertical connections (i.e., TSVs). Conversely, the V-S PDN connects layers in series and regulates voltage levels with SC converters. Consequently, it breaks the inter-layer noise correlation (Figure 6.4b). Table 6.2 shows each layer's maximum noise amplitude over the entire simulated time window. Compared with a 2D PDN, the traditional 3D PDN significantly reduces Task3's noise because the decoupling capacitors (decap) on adjacent layers help to stabilize local voltage variation. However, other layers' voltage noise is also affected by Task3. The V-S PDN isolates Task3's noise so that other layers have lower noise.

	Task1 Layer1	Task2 Layer2	Task3 Layer3	Task4 Layer4	Cross-layer Mean
Single-layer	4.0	3.0	10.9	2.8	N.A.
Traditional	3.7	4.2	4.2	4.3	4.1
V-S	2.8	1.9	3.6	2.3	2.7

Table 6.2: Maximum voltage noise (%Vdd) per layer for different workloads on 3D-ICs with different PDN schemes. The "cross-layer mean" value averages all layers' maximum noise amplitude.

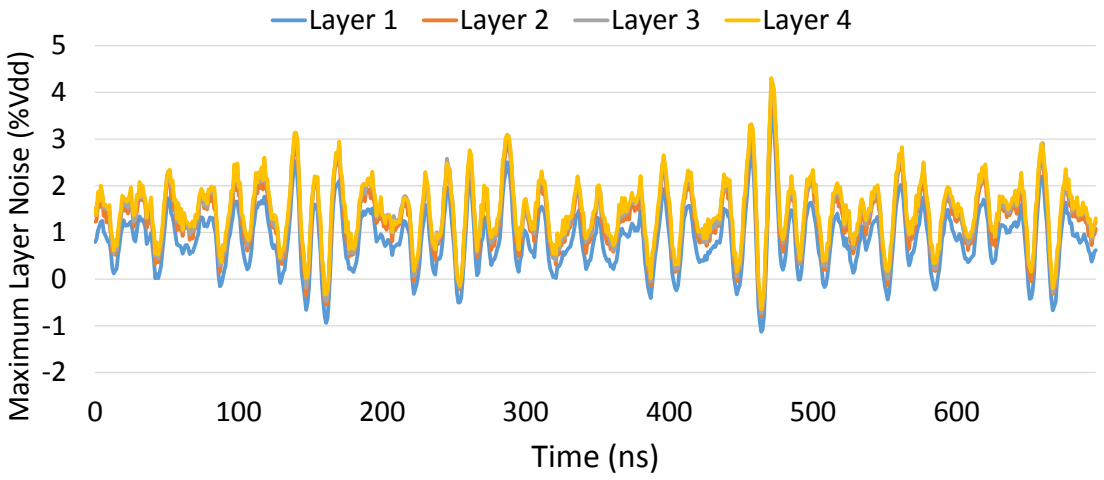

(a) Tradional PDN

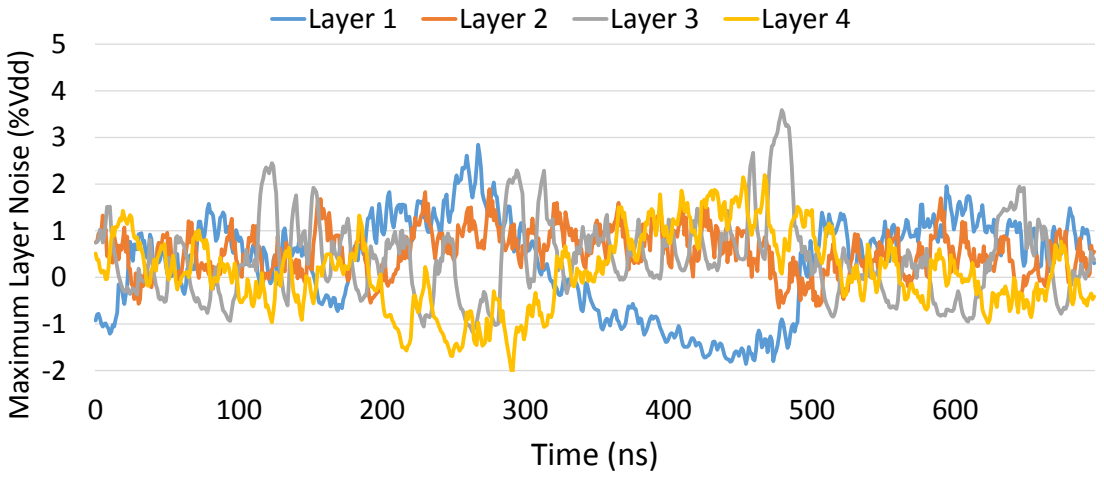

(b) V-S PDN

Figure 6.4: A plot of per-layer maximum noise amplitude over time. Only layer 3 has a noisy workload.

With dynamic margin adaptation (see Section 6.2.1, Chapter 3, and also reference [51]), each layer can adjust its timing margin according to its own maximum noise amplitude. Consequently, less noisy layers can run faster. Given the linear relationship between noise amplitude and transistor delay, we assume that $x\%$ Vdd noise also requires an $x\%$ decrease in clock frequency. The last column in Table 6.2 shows the arithmetic mean of all four layers' maximum noise amplitude. This cross-layer mean metric shows the whole-stack's average slowdown when we use per-layer margin adaptation. By isolating the cross-layer noise interaction, V-S PDN can improve system performance with less slowdown.

6.5.2 Allocating On-Chip Capacitance: A Tradeoff Study

The on-chip integrated capacitors can serve as either explicit decap for both traditional and V-S PDNs, or as fly-caps for V-S PDNs' SC converters. Because of their high area overhead, the total amount of on-chip capacitance is usually limited. It is therefore important to understand the tradeoff between the allocation of explicit decap and SC converters in the V-S PDN before we compare the overall area overhead and voltage noise quality between the two schemes.

6.5.2.1 Workload selection

In order to understand 3D-ICs' voltage noise level under a wide range of workload conditions, we construct different scenarios to stress both traditional and V-S PDNs. Starting from our sample pool, we first sort all workloads by average power consumption and then select the top, medium, and bottom one-percentile samples as candidate-groups, categorized as high (H), medium (M), and low (L). Using these candidate groups, we build the following three classes of multi-layer workloads. The first class (All_H, All_M, and All_L) assigns different samples from the same group to different layers in the 3D-IC. The second class (H/M and H/L) selects samples from any two candidate groups and assigns them to the 3D stack in an interleaved fashion. This pattern is particularly stressful for V-S PDNs because it forces all layers' SC converters to provide the same large amount of current, and the SC converters' output voltage drop is directly proportional to the load. In fact, the interleaved high-low (H/L) combination is the worst-case scenario for V-S PDNs. The last group (H_lkstp, M_lkstp, and L_lkstp) constructs a "lock-step" execution pattern by replicating the same workload to the entire stack. With all layers' power consumption changing simultaneously, this group can excite large LdI/dt and LC resonance voltage noise in the PDN. As an estimation to the worst-case scenario, we select the workloads with the highest single-layer noise within each H, M, and L candidate group.

6.5.2.2 Tradeoff study

Using our example 4-layer 3D processor, we simulated both V-S PDNs and traditional PDNs with different on-chip capacitance allocations. Figure 6.5a shows the cross-layer-mean noise amplitude

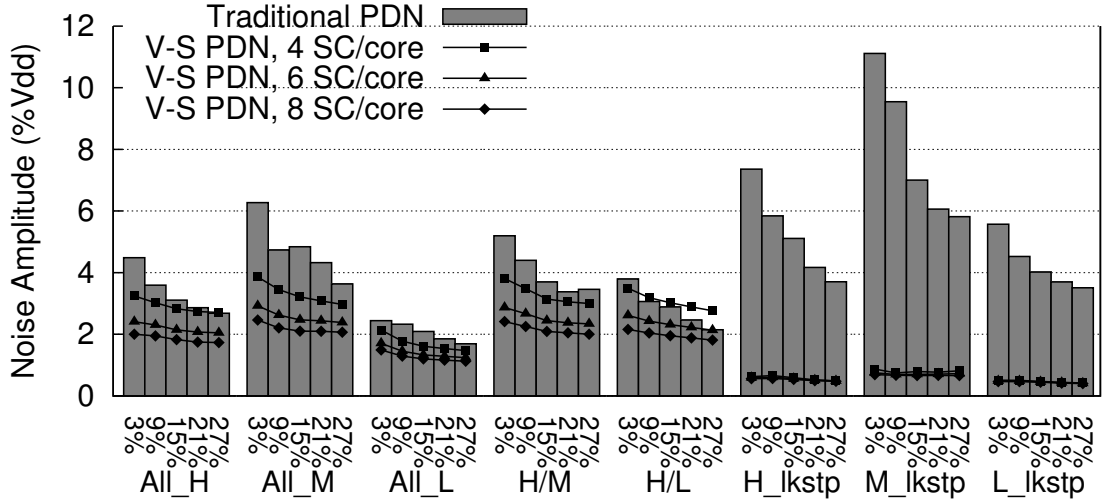

(a) Cross-layer mean

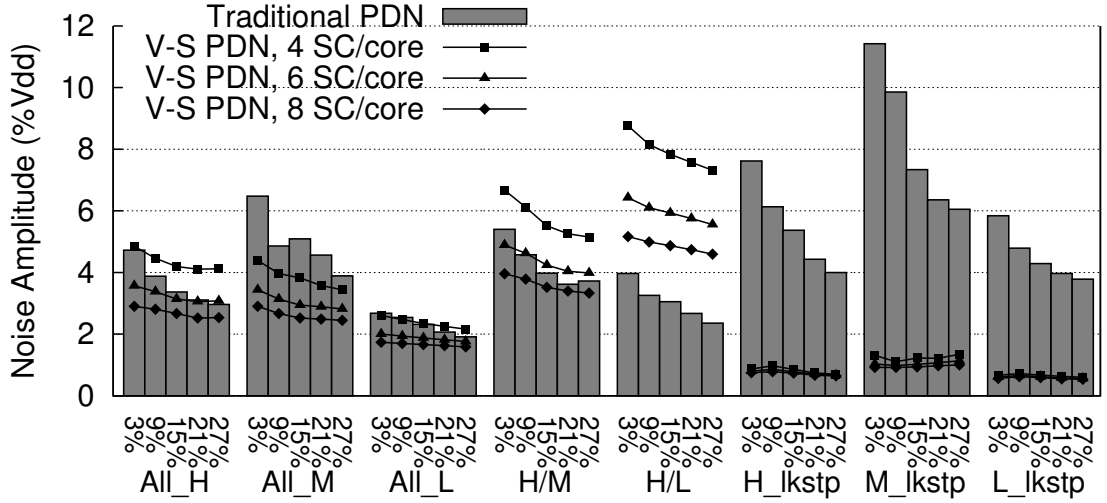

(b) Whole-stack maximum

Figure 6.5: A 4-layer 3D stack's voltage noise amplitude under different PDN configurations and workload conditions. The x-axis numbers within each data cluster represent the percentage of die area allocated for explicit decap. The size of each SC converter equals 3% of an ARM core.

For both PDN schemes, we sweep the percentage of die area allocated for explicit decap (x-axis within each data group). For V-S PDNs, we assign different number of SC converters to each core (lines with different markers). We note that all SC converters have the same amount of capacitance and switching frequency. Using an advanced, high-density technology (e.g., trench capacitors [66]), each SC converter occupies $0.082mm^2$, which is 3% of an ARM core. Therefore, the V-S PDN's on-chip capacitance area equals $decap\_area + number\_SC\_percore * 3\%$.

According to Figure 6.5a, the V-S PDNs' overall noise is not as sensitive to the amount of explicit decap as traditional PDNs, especially in the lock-step scenarios, where the traditional PDN suffers from LC resonance. This is because the SC converters not only help to smooth local LdI/dt noise with the built-in fly-capacitors, they also isolate the on-chip PDN from the package RLC loop, so that the package LC resonance is greatly suppressed. Consequently, designers can significantly reduce the amount of explicit decap in V-S PDNs. If we compare two PDN designs with the same amount of on-chip area allocated for overall capacitance (i.e., a V-S PDN with 4 per-core converters and 3% decap allocation, and a traditional PDN with 15% decap allocation), we observe that under their respective cross-layer means, the V-S PDN's noise is significantly lower than the traditional PDN's. This means that if per-layer runtime margin adaptation is used, the performance loss will be significantly lower for V-S.

Figure 6.5b shows the maximum noise amplitude observed in any layer for all test cases. The observation that the V-S PDN's cross-layer mean noise (Figure 6.5a) is significantly lower than its global maximum noise (Figure 6.5b) further proves the superior cross-layer noise isolation of V-S. This suggests that if a static worst-case noise margin is used, the V-S PDN will be worse. V-S PDN performance is only better when we utilize the per-layer dynamic margin adaptation.

6.5.3 Impact of 3D Scaling

To explore the effect of 3D scaling (i.e., stacking more layers) on both the V-S and traditional PDN's noise, we simulate our example 3D processors with two to eight layers, using the eight workload combinations. To make fair comparisons, we pick the design points described in Section6.5.2 that allocate a 15% on-chip area for capacitors in both PDN schemes.

Figure 6.6 plots all test cases' maximum noise amplitude (both whole-stack max and cross-layer mean) across all workload conditions. In general, stacking more layers together increases voltage noise in both types of PDNs. If a constant noise margin is applied to all layers at design-time, this margin has to accommodate the worst-case whole-stack maximum noise. Consequently, V-S structure requires smaller margin in 3D-ICs with 2 layers or more than 6 layers. With a per-layer dynamic margin adaptation technique enabled, the whole-stack's average margin will be no

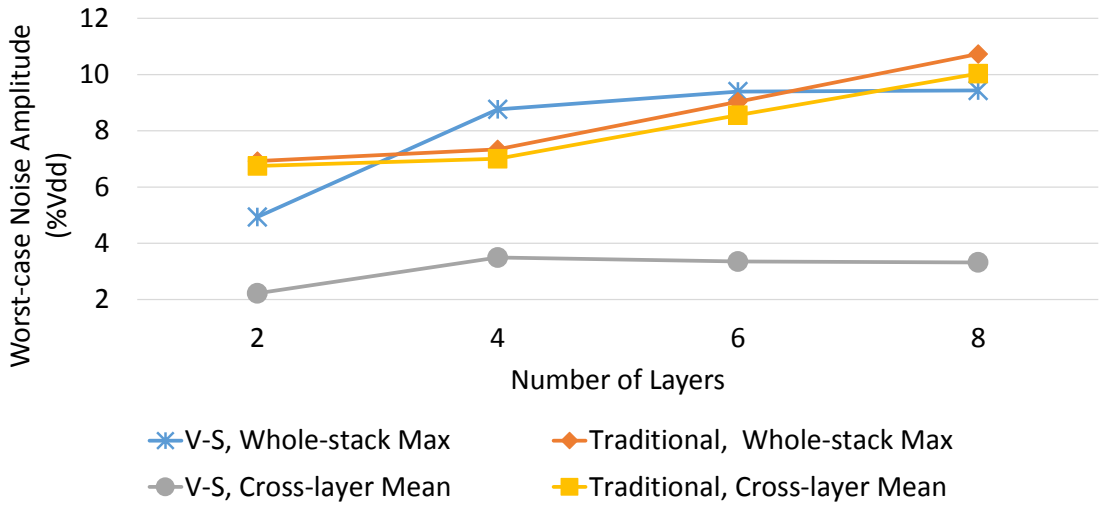

Figure 6.6: 3D scaling's impact on worst-case voltage noise.

larger than the worst-case cross-layer mean value. As a result, V-S PDNs always requires smaller timing margin, regardless of layer count. In the 8-layer 3D-IC, V-S PDN's noise is 60% lower than traditional PDN.

One interesting observation is that a 2-layer V-S PDN's whole-stack maximum noise is significantly lower than the maximum noise of V-S PDNs with more layers. This is because in a 2-layer V-S PDN, the output voltage variations of the SC converters only affect one supply net (either foot-bounce or head-droop) of any layer while the other net is directly connected to the off-chip voltage source via C4 pads. As silicon layers are added, foot-bounce and head-droop can be added to the same layer, which significantly increases noise.

6.5.4 Impact of Package Impedance

Chip package impedance has a significant impact on the supply-voltage noise [69]. Although package designs with lower impedance can provide more current with lower noise, they usually have higher cost due to their increased complexity (e.g., more layers of power planes to reduce the package resistance and inductance, more package decap, etc.). To explore the impact of chip package quality, we simulate a series of package designs by applying different scaling factors to the resistance, inductance, and capacitance values of our PDN model's lumped package model. For ex-

ample, to get a package with 200% impedance, we double the baseline package model's RL values (listed in Table 6.1) and reduce the package capacitance by half. We note that this scaling factor does not change the package RLC loop's resonance frequency.

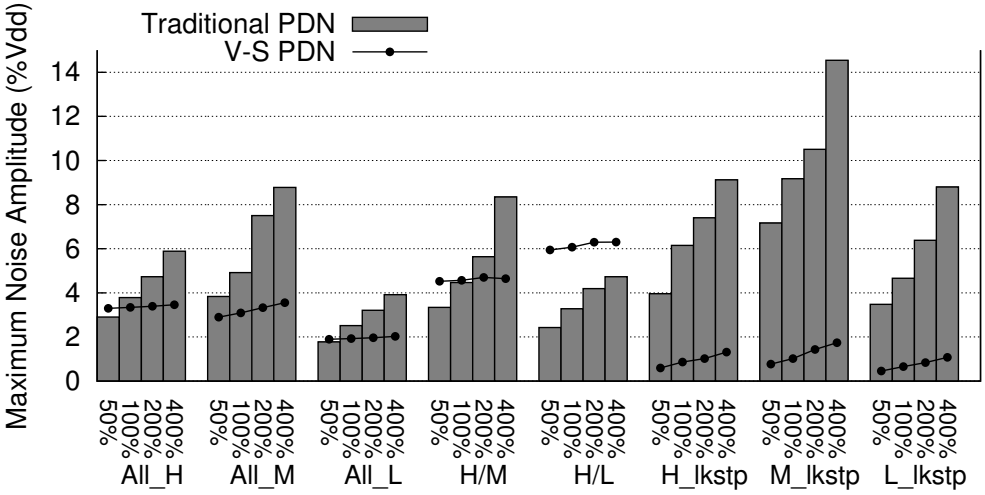

Figure 6.7: Package impedance's impact on 3D-IC's whole-stack maximum noise. The x-axis of each data group shows the normalized package impedance.

Figure 6.7 illustrates how package impedance affects both V-S and traditional PDNs' noise in a 4-layer 3D processor. Compared with the traditional PDNs, the maximum noise in V-S PDNs is much less sensitive to the package quality. For example, a 300% impedance increase only raises the V-S PDN's worst-case noise by 0.23% Vdd. Since the V-S PDN reduces off-chip current significantly, package impedance contributes much less noise overall. By relaxing the constraint on package impedance, the V-S PDN is expected to reduce the cost of 3D-IC's packaging.

6.6 Summary

In this chapter, we build a whole-system transient PDN model to: 1. Examine voltage-stacked 3D-ICs' transient noise under different workload conditions; 2. Compare voltage noise between V-S PDN and traditional PDN in the context of 3D scaling; 3. Explore the impact of various PDN design parameters. Our simulation results show that, compared with a traditional PDN, the V-S PDN provides stronger isolation for the cross-layer noise interference, but suffers higher noise in

the particular case of highly imbalanced workloads. This is mitigated if dynamic, per-layer margin adaptation is used to respond to severe noise. If so, V-S PDN can better reduce timing margin and improve system performance. Without incurring extra on-chip area overhead for the integrated capacitors, the V-S PDN's cross-layer-mean noise amplitude under the worst-case scenario is up to 60% lower than the traditional PDN. Furthermore, we observe that the V-S PDN allows lower packaging cost for 3D-ICs. Overall, we demonstrate that the V-S PDN provides a low-noise, low-cost, and scalable solution to the challenges of 3D-ICs' power delivery.

My colleague Kaushik Mazumdar (from U.Va.'s ECE department) and I worked together on this research project. His contributions include: 1) the circuit-level implementation of the SC converter; 2) the structure and parameters of our transient SC converter model (joint effort); 3) the circuit simulation results for validation. My contributions are: 1) extending VoltSpot to support transient simulations for 3D-IC; 2) the structure and parameters of our transient SC converter model (joint effort); 3) integrate the SC converter model with VoltSpot; 4) workload characterization/selection, whole-system PDN simulation, and transient noise evaluation.

Chapter 7

Conclusions and Future Work

7.1 Dissertation Summary

In this dissertation, we first discuss the power-delivery difficulties for contemporary and near-future processors, and address the importance of early-stage PDN design and evaluations. Then we build an architecture level PDN model and use it to study the impact of supply voltage noise and EM-wearout on system design.

Pre-RTL PDN Modeling We design and implement a pre-RTL PDN model VoltSpot that is capable of both modeling PDNs' physical structure in detail and simulating whole-application power supply noise. Compared with previous pre-RTL PDN models, VoltSpot makes two key improvements: (1) transient PDN grid modeling at the granularity of pad pitch or smaller, and (2) multi-layer-metal PDN modeling using multiple, parallel RL branches. VoltSpot has been validated against an IBM PDN analysis benchmark suite [64]. Validation results show that for representative silicon chips, VoltSpot has high accuracy in modeling both steady-state on-chip current distribution (up to 3.3% error in pad current) and transient voltage fluctuation (errors in capturing max voltage drop go up to 0.54% Vdd). Furthermore, we integrate VoltSpot with architecture-level performance (Gem5 [4]), power (McPAT [53]), and floorplan (ArchFP [13]) tools for the purpose of exploring the impact of PDN resource allocation on application-level voltage noise behavior.

Scarcity of C4 Pads By simulating PDN configurations with different number of pads dedicated to power supply, we found a strong correlation between power pad count and supply voltage

noise emergency rate. However, although removing power pads increases emergency rate drastically, the increment of noise amplitude is relatively mild as long as the locations of power pads are optimized. To evaluate the performance impacts of mitigating the extra voltage noise caused by C4 pad configuration change, we compared different state-of-the-art run-time noise mitigation techniques. Our simulation results indicate that a hybrid mechanism that combines dynamic margin adaptation [51] and noise-induced error recovery [20] is the most robust to technology scaling and worst-case noise-inducing power viruses. With this hybrid technique, noise introduced by replacing P/G pads with I/O pads (e.g., triples I/O bandwidth) can be mitigated with negligible overhead (e.g., 1.5% slowdown). The key insight is that a very small increase in the default timing guardband eliminates problems due to the much greater frequency of small/medium noise events.

Statistical analysis of EM-induced C4 failures To study C4 pads' reliability subject to EM damage and the performance impact of pad failures, we first derive per-pad DC current from VoltSpot and use Monte Carlo simulation to estimate whole-chip expected lifetime under a given supply-noise tolerance. Our results indicate that power pad failures tend to happen early. For example, in a silicon chip with 1900 power supply pads, the expected lifetime to the first EM-induced pad failure is only one-third of the expected lifetime of the single pad that carries the highest current among the entire chip. With the MCS framework we built around VoltSpot, we quantitatively examine both the mechanism and consequences of power pad failures. Our results indicate that the degradation of power delivery quality due to power pads' EM wearout is a relatively slow process and a small extra noise guardband (either allocated constantly or adjusted dynamically) enables the system to tolerate the consequence of multiple pad failures and significantly extends system MTTF. As a result, target lifetime can be achieved with significantly reduced power supply pad count (e.g., by 43%) and a slightly increased on-chip noise margin (e.g., 0.5% Vdd IR drop).

Power Delivery in 3D-IC Starting with a circuit implementation of a switched-capacitor converter, we build and validate a compact model for this charge-recycled voltage regulator and integrate it with VoltSpot to study the costs and benefits of utilizing voltage-stacking in 3D-ICs. Our EM-robustness analysis on both C4 pad and TSV arrays indicates that V-S PDN's EM-induced MTTF significantly surpasses (e.g., 5x longer) the conventional PDN's expected lifetime. Transient

simulation results show that compared with the traditional PDN scheme, V-S provides stronger isolation for cross-layer noise interference, which in turn grants higher performance benefits for run-time noise mitigation techniques, such as dynamic margin adaptation. With the same die area overhead for integrated capacitors, V-S PDNs provide up to 60% lower transient noise under the most noise-incurring workload behaviors. Furthermore, we show that V-S PDNs significantly reduce the packaging cost, because their noise is almost insensitive to the package impedance (e.g., a 300% impedance increase only raises worst-case noise by less than 0.3% Vdd). Our study demonstrates that V-S provides a practical, scalable, and affordable solution for many-layer 3D-IC's power delivery challenge.

7.2 Future Directions

7.2.1 Pre-RTL Power Delivery Network Modeling

While this dissertation has closely examined the power delivery challenges in both 2D and 3D silicon chips, our evaluations are mostly focused on the on-chip PDN, and the chip package is modeled as a lumped RLC loop. One interesting direction for future work is to extend VoltSpot with a distributed package model that captures the impact of package-PDN design details (e.g., number of power planes, number and locations of package pins, distribution of package decap, etc.). With an improved off-chip modeling granularity, VoltSpot will be able to support the study of package design tradeoffs. For example, because the number of affordable metal layers in the package is usually much smaller than the layer count of the on-chip metal stack, package I/O routing generally requires a significant amount of metal resource [50]. This metal-layer scarcity creates yet another resource contention between power delivery and chip I/O, which can be resolved or balanced at early design stages with the extended version of VoltSpot.

Even though VoltSpot provides a first-order support for multiple on-chip voltage domains, it does not model the potential hierarchical structure where several localized power grids attach to a global network through header transistors or IVRs. Since such organizations are critical for power management techniques like power-gating and DVFS, another future step is to include this hier-

archical formation in VoltSpot. Such extensions enable detailed analysis of fine-grained power-gating/DVFS and their implications on C4 pad/decap requirement.

7.2.2 Statistical Analysis of PDN Wearout

In our statistical simulations, we assumed that the consequence of EM-induced failure is creating a short circuit instantaneously at the time of failure. This assumption can be relaxed to consider the phenomenon that before breaking the metal conductor, EM will gradually increase its resistance [15]. A more detailed failure model can improve the accuracy of MTTF estimation, and more importantly, prevent designers from over provisioning the scarce resources for PDN. Furthermore, our MCS framework can also be extended to study other failure mechanisms such as thermal migration and thermal-stress, or other components of the PDN such as TSVs in 3D-IC. These extensions will be helpful for identifying bottlenecks in near-future processors' PDNs and to design and evaluate solutions to improve PDN long-term reliability.

Another direction for future work is the design and evaluation of run-time scheduling techniques that balance the level of wearout in different chip regions. Due to the stochastic nature of EM-induced conductor failures and the phenomenon of current-redistribution-induced PDN wearout acceleration, different regions of the PDN might experience different level of damage. For example, once a power C4 pad fails, the "avalanche" effect accelerates the EM wearout in its nearby region. Depending on the severity of the avalanche effect, the C4 pads in this region could start breaking consecutively and fail the entire system before the rest of the pad array receives significant EM damage. One solution to prevent this potential catastrophic failure is to reduce a region's current stress soon after it starts to suffer from EM-induced open-circuits. Unfortunately, the exact control sequence of this technique cannot be determined at design-time, because any bump could fail and initiate a cascade of failure events. Therefore, it would be interesting to study how to detect EM-wearout at runtime and how to adjust on-chip power consumption (e.g., through scheduling or frequency scaling) to avoid catastrophic failure and extend system MTTF.

7.2.3 3D-IC and Voltage Stacking

Chapter 5 and Chapter 6 studied the benefits and costs of using an SC-converter-supported V-S power-delivery scheme in 3D-ICs. Although an SC converter provides high power efficiency and low-cost integration, it requires that all layers must have the same nominal supply voltage. This requirement becomes a limitation for the design and implementation of heterogeneous 3D systems, where components with different supply voltages can be stacked together. Fortunately, with the help of inductive voltage regulators, voltage-stacked PDN can support heterogeneous systems [79]. An interesting future direction is to build a compact model for inductive IVRs and integrate it with 3D VoltSpot to study the design tradeoff of inductive-IVR-supported V-S PDNs.

7.2.4 Other Directions

In general, VoltSpot directly supports, or can be extended to facilitate the study of a wide variety of power-delivery-related issues. As we illustrated in our previous work [90, 91], VoltSpot's fine-grained on-chip modeling capability enables the evaluation of pre-RTL C4 placement optimization algorithms. Besides C4 pads, VoltSpot also makes it possible to quickly evaluate the impact of other PDN design details, such as the distribution of on-chip decap, the configuration of IVRs, and the topology of TSVs (in 3D-ICs). Combined with the ability to examine processors with different architecture configurations and floorplan arrangements, our pre-RTL tool chain provides a versatile platform for the development and evaluation of early-stage PDN optimization algorithms. It would be interesting to develop optimization schemes that either focus on single PDN component, or jointly optimize multiple design parameters. Another possible future direction is to expand the optimization algorithm's cost function to include not only supply voltage noise, but also EM-induced lifetime.

Bibliography

[1] ARM. http://www.arm.com/products/processors/cortex-a/cortex-a9.php.

[2] Anys Bacha and Radu Teodorescu. Dynamic reduction of voltage margins by leveraging on-chip ECC in Itanium II processors. In *International Symposium on Computer Architecture (ISCA)*, 2013.

[3] Christian Bienia. *Benchmarking Modern Multiprocessors*. PhD thesis, Princeton University, Jan 2011.

[4] Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R Hower, Tushar Krishna, Somayeh Sardashti, et al. The gem5 simulator. *ACM SIGARCH Computer Architecture News*, 39(2):1–7, 2011.

[5] J.R. Black. Electromigration: A brief survey and some recent results. *IEEE Transactions on Electron Devices*, 16(4):338–347, 1969.

[6] S Brandenburg and S Yeh. Surface mount international conference and exposition. In *SMI*, volume 98, page 337, 1998.

[7] Tsung-Hao Chen and Charlie Chung-Ping Chen. Efficient large-scale power grid analysis based on preconditioned krylov-subspace iterative methods. In *Design Automation Conference (DAC)*, 2001.

[8] WJ Choi, ECC Yeh, and KN Tu. Mean-time-to-failure study of flip chip solder joints on Cu/Ni (V)/Al thin-film under-bump-metallization. *Journal of applied physics*, 94(9):5665–5671, 2003.

[9] In Jae Chung. *Modeling and Hybrid Simulation of On-chip Power Delivery Network Using an Unconditionally Stable Electromagnetic Field Solver*. ProQuest, 2007.

[10] Luc Devroye. *Non-Uniform Random Variate Generation*. Springer, 1986.

[11] Philip W Diodato. Embedded DRAM: more than just a memory. *IEEE Communications Magazine*, 38(7):118–126, 2000.

[12] Dan Ernst, Nam Sung Kim, Shidhartha Das, Sanjay Pant, Rajeev Rao, Toan Pham, Conrad Ziesler, David Blaauw, Todd Austin, Krisztian Flautner, et al. Razor: A low-power pipeline based on circuit-level timing speculation. In *International Symposium on Microarchitecture (MICRO)*, pages 7–18, 2003.

[13] Gregory Faust, Runjie Zhang, Kevin Skadron, Mircea Stan, and Brett Meyer. ArchFP: Rapid prototyping of pre-RTL floorplans. In *International Conference on Very Large Scale Integration (VSLI-SoC)*, 2012.

[14] Mohammad Fawaz. Electromigration reliability analysis of power delivery networks in integrated circuits. Master's thesis, U. of Toronto, 2013.

[15] M Gall, C Capasso, D Jawarani, R Hernandez, H Kawasaki, and PS Ho. Statistical analysis of early failures in electromigration. *Journal of Applied Physics*, 90(2):732–740, 2001.

[16] Varghese George, Sanjeev Jahagirdar, Chao Tong, K Smits, Satish Damaraju, Scott Siers, Ves Naydenov, Tanveer Khondker, Sanjib Sarkar, and Puneet Singh. Penryn: 45-nm next generation Intel core 2 processor. In *Asian Solid-State Circuits Conference (ASSCC)*, 2007.

[17] Ed Grochowski, David Ayers, and Vivek Tiwari. Microarchitectural simulation and control of dI/dt-induced power supply voltage variation. In *International Symposium on High-Performance Computer Architecture (HPCA)*, 2002.

[18] Jie Gu and Chris H Kim. Multi-story power delivery for supply noise reduction and low voltage operation. In *International Symposium on Low Power Electronics and Design (ISLPED)*, 2005.

[19] Meeta S Gupta, Jarod L Oatley, Russ Joseph, Gu-Yeon Wei, and David M Brooks. Understanding voltage variations in chip multiprocessors using a distributed power-delivery network. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2007.

[20] Meeta Sharma Gupta, Krishna K Rangan, Michael D Smith, Gu-Yeon Wei, and David Brooks. DeCoR: A delayed commit and rollback mechanism for handling inductive noise in processors. In *International Symposium on High Performance Computer Architecture (HPCA)*, 2008.

[21] Meeta Sharma Gupta, Jude A Rivers, Pradip Bose, Gu-Yeon Wei, and David Brooks. Tribeca: design for PVT variations with local recovery and fine-grained adaptation. In *International Symposium on Microarchitecture (MICRO)*, 2009.

[22] John W Haskins Jr and Kevin Skadron. Accelerated warmup for sampled microarchitecture simulation. *ACM Transactions on Architecture and Code Optimization (TACO)*, 2(1):78–108, 2005.

[23] Michael B Healy and Sung Kyu Lim. Distributed TSV topology for 3-D power-supply networks. *IEEE Transactions on VLSI*, 20(11), 2012.

[24] Michael B Healy, Fayez Mohamood, Hsien-Hsin S Lee, and Sung Kyu Lim. Integrated microarchitectural floorplanning and run-time controller for inductive noise mitigation. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 16(4), Oct 2011.

[25] Dennis J. Herrell and Benjamin Beker. Modeling of power distribution systems for high-performance microprocessors. *IEEE Transactions on Advanced Packaging*, 22(3):240–248, August 1999.

[26] Gang Huang, Muhannad Bakir, Azad Naeemi, Howard Chen, and James D Meindl. Power delivery for 3D chip stacks: Physical modeling and design implication. In *Conference on Electrical Performance of Electronic Packaging (EPEP)*, 2007.

[27] Xin Huang, Tan Yu, Valeriy Sukharev, and Sheldon Tan. Physics-based electromigration assessment for power grid networks. In *Design Automation Conference (DAC)*, 2014.

[28] Intel. Intel Pentium 4 Processor in the 423 pin package / Intel 850 Chipset Platform, 2002.

[29] Intel. Voltage regulator module (VRM) and enterprise voltage regulator-down (EVRD) 11.1. *Design Guidelines*, March 2009.

[30] Intel. Intel Core i7-900 Desktop Processor Extreme Edition Series and Intel Core i7-900 Desktop Processor Series on 32-nm Process, 2011.

[31] Intel. Intel Xeon Processor E7-8800/4800/2800 Product Families, 2011.

[32] ITRS, 2012. http://www.itrs.net.

[33] Bruce Jacob, Spencer Ng, and David Wang. *Memory systems: cache, DRAM, disk*. Morgan Kaufmann, 2010.

[34] Pulkit Jain, Tae-Hyoung Kim, John Keane, and Chris H Kim. A multi-story power delivery technique for 3D integrated circuits. In *International Symposium on Low Power Electronics and Design (ISLPED)*, 2008.

[35] Renatas Jakushokas and Eby G Friedman. Line width optimization for interdigitated power/ground networks. In *Symposium on Great Lakes Symposium on VLSI (GLSVLSI)*, 2010.

[36] Renatas Jakushokas and Eby G Friedman. Multi-layer interdigitated power distribution networks. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 19(5):774–786, 2011.

[37] Renatas Jakushokas, Mikhail Popovich, Andrey V Mezhiba, Selçuk Köse, and Eby G Friedman. Inductance model of interdigitated power and ground distribution networks. In *Power Distribution Networks with On-Chip Decoupling Capacitors*, pages 323–335. Springer, 2011.

[38] JEDEC. Failure Mechanisms and Models for Semiconductor Devices. *JEDEC Publication JEP122G*, 2011.

[39] Richard Arnold Johnson, Irwin Miller, and John E Freund. *Probability and statistics for engineers*. Prentice-Hall, 2011.

[40] Russ Joseph, David Brooks, and Margaret Martonosi. Control techniques to eliminate voltage emergencies in high performance processors. In *International Symposium on High Performance Computer Architecture (HPCA)*, 2003.

[41] Svilen Kanev. *Motivating Software-Driven Current Balancing in Flexible Voltage-Stacked Multicore Processors*. PhD thesis, Harvard University Cambridge, Massachusetts, 2012.

[42] William H Kao, Ch-Yuan Lo, Mark Basel, and Raminderpal Singh. Parasitic extraction: current state of the art and future trends. *Proceedings of the IEEE*, 89(5):729–739, 2001.

[43] Guruprasad Katti, Michele Stucchi, Kristin De Meyer, and Wim Dehaene. Electrical modeling and characterization of through silicon via for three-dimensional ICs. *IEEE Transactions on Electron Devices*, 57(1), 2010.

[44] Dae Hyun Kim and Sung Kyu Lim. Through-silicon-via-aware delay and power prediction model for buffered interconnects in 3D ICs. In *International Workshop on System Level Interconnect Prediction*, 2010.

[45] Youngtaek Kim and Lizy Kurian John. Automated dI/dt stressmark generation for microprocessor power delivery networks. In *International Symposium on Low Power Electronics and Design (ISLPED)*, 2011.

[46] John U Knickerbocker, PS Andry, B Dang, RR Horton, CS Patel, RJ Polastre, K Sakuma, ES Sprogis, CK Tsang, BC Webb, et al. 3D silicon integration. In *Electronic Components and Technology Conference (ECTC)*, 2008.

[47] Selçuk Köse and Eby G Friedman. Simultaneous co-design of distributed on-chip power supplies and decoupling capacitors. In *International SOC Conference (SOCC)*, 2010.

[48] Keunmyung Lee and Alan Barber. Modeling and analysis of multichip module power supply planes. *IEEE Transactions on Components, Packaging, and Manufacturing Technology*, 18(4):628–639, November 1995.

[49] Ki-Don Lee and Paul S Ho. Statistical study for electromigration reliability in dual-damascene Cu interconnects. *IEEE Transactions on Device and Materials Reliability*, 4(2):237–245, 2004.

[50] Ren-Jie Lee and Hung-Ming Chen. Fast flip-chip pin-out designation respin for package-board codesign. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 17(8):1087–1098, 2009.

[51] Charles R Lefurgy, Alan J Drake, Michael S Floyd, Malcolm S Allen-Ware, Bishop Brock, Jose A Tierno, and John B Carter. Active management of timing guardband to save energy in POWER7. In *International Symposium on Microarchitecture (MICRO)*, 2011.

[52] Dian Li, Zhong Guan, Malgorzata Marek-Sadowska, and Sani R Nassif. Multi-via electromigration lifetime model. *International Conference on Simulation of Semiconductor Processes and Devices (SISPAD)*, 2012.

[53] Sheng Li, Jung Ho Ahn, Richard D Strong, Jay B Brockman, Dean M Tullsen, and Norman P Jouppi. McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures. In *International Symposium on Microarchitecture (MICRO)*, 2009.

[54] Xiaoye S. Li. An Overview of SuperLU: Algorithms, Implementation, and User Interface. *ACM Transactions on Mathematical Software (TOMS)*, 31(3):302–325, 2005.

[55] B-K Liew, NW Cheung, and Chenming Hu. Projecting interconnect electromigration lifetime for arbitrary current waveforms. *IEEE Transactions on Electron Devices*, 37(5):1343–1351, 1990.

[56] JR Lloyd. On the log-normal distribution of electromigration lifetimes. *Journal of Applied Physics*, 50(7), 1979.

[57] JR Lloyd and JJ Clement. Electromigration in copper conductors. *Thin solid films*, 262(1), 1995.

[58] Chris A Mack. Fifty years of Moore's law. *IEEE Transactions on Semiconductor Manufacturing*, 24(2):202–207, 2011.

[59] Kaushik Mazumdar and Mircea Stan. Breaking the power delivery wall using voltage stacking. In *Symposium on Great Lakes Symposium on VLSI (GLSVLSI)*, 2012.

[60] Andrey V Mezhiba and Eby G Friedman. Electrical characteristics of multi-layer power distribution grids. In *International Symposium on Circuits and Systems (ISCAS)*, 2003.

[61] Andrey V Mezhiba and Eby G Friedman. Scaling trends of on-chip power distribution noise. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 12(4):386–394, 2004.

[62] Micron. http://www.micron.com/products/hybrid-memory-cube.

[63] Sani R Nassif and Joseph N Kozhaya. Fast power grid simulation. In *Design Automation Conference (DAC)*, 2000.

[64] S.R. Nassif. Power grid analysis benchmarks. In *Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2008.

[65] Mohit Pathak, Young Joon Lee, Thomas Moon, and Sung Kyu Lim. Through-silicon-via management during 3D physical design: When to add and how many? In *International Conference on Computer-Aided Design (ICCAD)*, 2010.

[66] Chengwen Pei, Roger Booth, Herbert Ho, Naoyoshi Kusaba, Xi Li, M-J Brodsky, Paul Parries, Hulling Shang, Rama Divakaruni, and Subramanian Iyer. A novel, low-cost deep trench decoupling capacitor for high-performance, low-power bulk CMOS applications. In *International Conference on Solid-State and Integrated-Circuit Technology (ICSICT)*, 2008.

[67] Mikhail Popovich, Eby G Friedman, Radu M Secareanu, and Olin L Hartin. Efficient distributed on-chip decoupling capacitors for nanoscale ICs. *IEEE Transactions on Very Large Scale Integration Systems*, 16(12):1717–1721, 2008.

[68] Mikhail Popovich, Eby G Friedman, Michael Sotman, Avinoam Kolodny, and Radu M Secareanu. Maximum effective distance of on-chip decoupling capacitors in power distribution grids. In *Symposium on Great Lakes Symposium on VLSI (GLSVLSI)*, 2006.

[69] Mikhhail Popovich, Andrey V Mezhiba, and Eby G Friedman. *Power Distribution Networks with On-chip Decoupling Capacitors*. Springer, 2008.

[70] Michael D Powell and TN Vijaykumar. Pipeline muffling and a priori current ramping: architectural techniques to reduce high-frequency inductive noise. In *International Symposium on Low Power Electronics and Design (ISLPED)*, 2003.

[71] Michael D Powell and TN Vijaykumar. Exploiting resonant behavior to reduce inductive noise. In *International Symposium on Computer Architecture (ISCA)*, 2004.

[72] Saravanan Rajapandian, Zheng Xu, and Kenneth L Shepard. Implicit DC-DC downconversion through charge-recycling. *IEEE Journal of Solid-State Circuits (JSSC)*, 40(4), 2005.

[73] Anand Ramachandran. *Methodologies for Transient Simulation of Hybrid Electromagnetic/Circuit Systems with Multiple Time Scales*. PhD thesis, University of Illinois at Urbana-Champaign, 2009.

[74] Martin Saint-Laurent and Madhavan Swaminathan. Impact of power-supply noise on timing in high-frequency microprocessors. *IEEE Transactions on Advanced Packaging*, 27(1), 2004.

[75] Sachin S Sapatnekar. Addressing thermal and power delivery bottlenecks in 3D circuits. In *Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2009.

[76] Takashi Sato, Masanori Hashimoto, and Hidetoshi Onodera. Successive pad assignment algorithm to optimize number and location of power supply pad using incremental matrix inversion. In *Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2005.

[77] Michael D Seeman. A design methodology for switched-capacitor DC-DC converters. Technical report, DTIC Document, 2009.

[78] Deepak Sekar, Calvin King, Bing Dang, Todd Spencer, Hiren Thacker, Paul Joseph, Muhannad Bakir, and James Meindl. A 3D-IC technology with integrated microchannel cooling. In *International Interconnect Technology Conference (IITC)*, 2008.

[79] Pradeep S Shenoy, Katherine A Kim, Brian B Johnson, and Philip T Krein. Differential power processing for increased energy production and reliability of photovoltaic systems. *IEEE Transactions on Power Electronics*, 28(6), 2013.

[80] Balaram Sinharoy, R Kalla, WJ Starke, HQ Le, R Cargnoni, JA Van Norstrand, BJ Ronchetti, J Stuecheli, J Leenstra, GL Guthrie, et al. IBM POWER7 multicore server processor. *IBM Journal of Research and Development*, 55(3):1–1, 2011.

[81] Kevin Skadron, Mircea R Stan, Wei Huang, Sivakumar Velusamy, Karthik Sankaranarayanan, and David Tarjan. Temperature-aware microarchitecture. In *International Symposium on Computer Architecture (ISCA)*, 2003.

[82] Michiel Steyaert, Tom Van Breussegem, Hans Meyvaert, Piet Callemeyn, and Mike Wens. DC-DC converters: From discrete towards fully integrated CMOS. In *European Solid-State Device Research Conference (ESSDERC)*, 2011.

[83] Madhavan Swaminathan and Ege Engin. *Power integrity modeling and design for semiconductors and systems*. Pearson Education, 2007.

[84] Jiang Tao, Boon Khim Liew, Jone F Chen, Nathan W Cheung, and Chenming Hu. Electromigration under time-varying current stress. *Microelectronics Reliability*, 38(3), 1998.

[85] Thorlindur Thorolfsson, Kiran Gonsalves, and Paul D Franzon. Design automation for a 3DIC FFT processor for synthetic aperture radar: a case study. In *Design Automation Conference (DAC)*, 2009.

[86] Aida Todri, Sandip Kundu, Patrick Girard, Alberto Bosio, Luigi Dilillo, and Arnaud Virazel. A study of tapered 3-D TSVs for power and thermal integrity. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 21(2):306–319, 2013.

[87] Aida Todri and Malgorzata Marek-Sadowska. Electromigration study of power-gated grids. In *International Symposium on Low Power Electronics and Design (ISLPED)*, 2009.

[88] Gabriel Torres, 2010. http://www.hardwaresecrets.com/printpage/Everything-You-Need-to-Know-About-The-Motherboard-Voltage-Regulator-Circuit/616.

[89] KN Tu. Recent advances on electromigration in very-large-scale-integration of interconnects. *Journal of Applied Physics*, 94(9):5451–5473, 2003.

[90] Ke Wang, Brett H. Meyer, Runjie Zhang, Kevin Skadron, and Mircea R. Stan. Walking pads: Fast power-supply pad-placement optimization. In *Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2014.

[91] Ke Wang, Brett H Meyer, Runjie Zhang, Micrea Stan, and Kevin Skadron. Walking pads: Managing C4 placement for transient voltage noise minimization. In *Design Automation Conference (DAC)*, 2014.

[92] N. H.E. Weste and D. M. Harris. *CMOS VLSI Design A Circuit and Systems Perspective*. Addison-Wesley, 4th edition, 2011.

[93] S.L. Wright, R. Polastre, H. Gan, L.P. Buchwalter, R. Horton, P.S. Andry, E. Sprogis, C. Patel, C. Tsang, J. Knickerbocker, J.R. Lloyd, A. Sharma, and M.S. Sri-Jayantha. Characterization of micro-bump C4 interconnects for Si-carrier SOP applications. In *Electronic Components and Technology Conference (ECTC)*, 2006.

[94] Roland E Wunderlich, Thomas F Wenisch, Babak Falsafi, and James C Hoe. SMARTS: Accelerating microarchitecture simulation via rigorous statistical sampling. In *International Symposium on Computer Architecture (ISCA)*, 2003.

[95] Zheng Xu, Xiaoxiong Gu, Michael Scheuermann, Kenneth Rose, Buckwell C Webb, John U Knickerbocker, and Jian-Qiang Lu. Modeling of power delivery into 3D chips on silicon interposer. In *Electronic Components and Technology Conference (ECTC)*, 2012.

[96] Guihai Yan, Xiaoyao Liang, Yinhe Han, and Xiaowei Li. Leveraging the core-level complementary effects of PVT variations to reduce timing emergencies in multi-core processors. In *International Symposium on Computer Architecture (ISCA)*, 2010.

[97] 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), 2005.

[98] Hao Yu, Joanna Ho, and Lei He. Allocating power ground vias in 3D ICs for simultaneous power and thermal integrity. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 14(3):41, 2009.

[99] Ting Yu and Martin DF Wong. A novel and efficient method for power pad placement optimization. In *International Symposium on Quality Electronic Design (ISQED)*, 2013.

[100] Zhiyu Zeng, Zhuo Feng, Peng Li, and Vivek Sarin. Locality-driven parallel static analysis for power delivery networks. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 16(3):28, 2011.

[101] Runjie Zhang, Kaushik Mazumdar, Brett Meyer, Ke Wang, Kevin Skadron, and Micrea Stan. A cross-layer design exploration of charge-recycled power-delivery in many-layer 3D-IC. In *Design Automation Conference (DAC)*, 2015.

[102] Runjie Zhang, Brett H. Meyer, Wei Huang, Kevin Skadron, and Mircea R. Stan. Some limits of power delivery in the multicore era. In *Proceedings of the Workshop on Energy Efficient Design (WEED)*, 2012.

[103] Runjie Zhang, Ke Wang, Brett H. Meyer, Mircea R. Stan, and Kevin Skadron. Architecture implications of pads as a scarce resource. In *International Symposium on Computer Architecture (ISCA)*, 2014.

[104] Zhuqing Zhang and CP Wong. Recent advances in flip-chip underfill: materials, process, and reliability. *IEEE Transactions on Advanced Packaging*, 27(3), 2004.

[105] Xin Zhao, Yang Wan, Michael Scheuermann, and Sung Kyu Lim. Transient modeling of TSV-wire electromigration and lifetime analysis of power distribution network for 3D ICs. In *International Conference on Computer-Aided Design (ICCAD)*, 2013.

[106] Yu Zhong and Martin DF Wong. Fast placement optimization of power supply pads. In *Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2007.

[107] Pingqiang Zhou, Dong Jiao, Chris H Kim, and Sachin S Sapatnekar. Exploration of on-chip switched-capacitor DC-DC converter for multicore processors using a distributed power delivery network. In *Custom Integrated Circuits Conference (CICC)*, 2011.