DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY
Office for Scholarly Communication

Reducing Power Loss, Cost and Complexity of SoC Power Delivery Using Integrated 3-Level Voltage Regulators

Citation

Kim, Wonyoung. 2013. Reducing Power Loss, Cost and Complexity of SoC Power Delivery Using Integrated 3-Level Voltage Regulators. Doctoral dissertation, Harvard University.

Permanent link

http://nrs.harvard.edu/urn-3:HUL.InstRepos:10423839

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story .

Accessibility

©2013 - Wonyoung Kim

All rights reserved.

Thesis advisors

Gu-Yeon Wei and David Brooks

Author

Wonyoung Kim

Reducing Power Loss, Cost and Complexity of SoC Power Delivery using Integrated 3-Level Voltage Regulators

Abstract

Traditional methods of system-on-chip (SoC) power management based on dynamic voltage and frequency scaling (DVFS) is limited by 1) cores/IP blocks sharing a voltage domain provided by off-chip voltage regulators (VR) and 2) slow voltage scaling time ($<0.1\text{V}/\mu\text{s}$). This global, slow DVFS cannot track the increasingly heterogeneous, fluctuating performance requirements of individual microprocessor cores and SoC components. Furthermore, traditional off-chip VRs add significant area overhead and component cost on the board.

This thesis explores replacing a large portion of existing off-chip VRs with integrated voltage regulators (IVR) that can scale the voltage at a 50 mV/ns rate, which is 500 times faster than microsecond-scale voltage scaling with existing off-chip VRs. IVRs occupy 10 times smaller footprint than off-chip VRs, making it easy to duplicate them to provide per-core or per-IP-block voltage control. This thesis starts by summarizing the benefits of using IVRs to deliver power to SoCs. Based on a simulation study targeting a 1.6W, 4-core SoC, I show that greater than 20% energy savings is possible with fast, per-core DVFS enabled by IVRs. Next, I present two stand-alone

IVR test-chips converting 1.8V and 2.4V to 0.4-1.4V while delivering maximum 1W to the output. Both test-chips incorporate a 3-level VR topology, which is suitable for integration because the topology allows for much smaller inductors (1nH) than existing inductor-based buck VRs. I also discuss reasons behind lower-than-simulated efficiencies in the test-chips and ways to improve. Finally, I conclude with future process technologies that can boost IVR conversion efficiencies and power densities.

Contents

Title Page i
Abstract iii
Table of Contents v
List of Figures vii
List of Tables xi
Citations to Previously Published Work xii
Acknowledgments xiii

1 Challenges of Delivering Power to SoCs and the Need for Integrated Voltage Regulators 1
1.1 Challenge of delivering power to modern SoCs 2
1.2 Solution: Integrated Voltage Regulators 5

2 Basics of Voltage Regulators and Challenges of Integration 12
2.1 Basics of step-down voltage regulators 13
2.2 Challenges of Integrating Voltage Regulators 15
2.3 Evolution of IVRs 21

3 System-Level Energy Savings with Fast, Per-Core DVFS using Integrated Voltage Regulators 24
3.1 Prior Works on Fine-Grain DVFS 28
3.2 Potential of Fast and Per-Core DVFS Schemes 29
3.2.1 Simulation Framework 31
3.2.2 Offline DVFS Algorithm 32
3.2.3 Effects of Finer Temporal Resolution 34
3.2.4 Per-Core vs. Chip-Wide DVFS 36
3.3 Characteristics of On-Chip Regulators 40
3.3.1 Model and Simulation of Buck VR 40
3.3.2 Design trade-offs of IVRs 41
3.3.3 Regulator Efficiency 44
3.3.4 Load Transient Response 47

3.3.5 Voltage Scaling Time 49
3.3.6 On-Chip Regulators for Single and Multiple Power Domains . 50
3.4 Energy Savings for Per-Core and Chip-Wide DVFS using On-Chip Regulators 53
3.4.1 Comparison of Energy Savings 54
3.4.2 Power Domain Scalability 58

4 Fully-Integrated 3-Level Voltage Regulators 61
4.1 Buck, Switched-Capacitor and 3-Level IVRs 62
4.2 3-Level Voltage Converter 65
4.2.1 Design Parameters for 3-Level Converters 65
4.2.2 Comparison to Buck and SC Converters 69
4.3 3-Level Implementation: Open-Loop 74
4.3.1 Power FETs 76
4.3.2 Driver circuits 78
4.3.3 Passive elements 80
4.3.4 Feedback loop and shunt regulator 82
4.4 Measurement: Open-Loop 83
4.5 3-Level VR: Closed-Loop 94

5 Technologies on the Horizon 108

Bibliography 112

List of Figures

1.1 Power delivery in mobile and server systems . . . 3
1.2 Illustration showing how shared voltage domains and slow DVFS lead to wasted energy. . . . 4
1.3 Large amounts of decaps in the board, package and SoC die decrease the speed of voltage scaling . . . 4
1.4 Teardown images of iPhone 4S and iPhone 5 show that off-chip VRs and required passive elements occupy a significant area on the board 4
1.5 Power delivery using IVRs in a server system . . . 6
1.6 IVR enables nanosecond timescale DVFS compared to microsecond DVFS with existing off-chip VRs . . . 7
1.7 IVR enables fast voltage scaling by reducing the amount of decap to charge/discharge . . . 7
1.8 Efficiency and footprint data for TI voltage regulator products plotted using TI's WEBENCH, a simulator provided by TI for its voltage regulator products [13] . . . 8
1.9 IVRs can reduce I^2R loss on the board by delivering power at a high voltage and low current and converting to a lower voltage at the point of load. . . . 10

2.1 Buck converter schematics . . . 14
2.2 Dimensions, DC resistance and inductance of various chip inductors from Coilcraft [4]. Images of inductors is roughly to scale. . . . 16
2.3 IVRs can be integrated on the package-level (a,b) or in the logic IC die (c). . . . 17
2.4 There has been increasing interest on IVRs by industry and research communities. . . . 21

3.1 Three power-supply configurations for a 4-core CMP. . . . 27
3.2 DVFS transition times with an IVR . . . 30
3.3 Benefits of fine-grained DVFS scheme for *mcf* and *fft*. . . . 35
3.4 Per-Core DVFS for multi-threaded applications. . . . 37

3.5 Snapshot of *ocean* with per-core and chip-wide DVFS. 38
3.6 Snapshot of *fft* with per-core and chip-wide DVFS. 38
3.7 Per-core DVFS for multi-programming scenarios. 39
3.8 Power delivery network using (a) only off-chip and (b) both off-chip and on-chip VRs. 42
3.9 Conversion loss, voltage variation, and voltage scaling time of a VR with different parameters. 44
3.10 VR efficiency and power vs. output voltage for different activity factors. 45
3.11 Voltage fluctuation of off-chip and on-chip VRs during step and sine wave load current transient 46
3.12 Example of reducing voltage fluctuations by selectively disabling clock gating. 50
3.13 Snapshot of output voltage, frequency, and load current traces with DVFS. 51
3.14 Total energy overhead with different regulator settings for *facerec* . . 52
3.15 Detailed breakdown of energy consumption for the processor and VR for single power domain (global) and multiple domains (per-core) DVFS. 55
3.16 Relative energy consumption of on-chip VR configurations compared to a off-chip VR with DVFS. 56
3.17 Loss, inductor size, and area of on-chip VRs for different numbers of power domains. 58

4.1 Power FET and output filters of (a) buck, (b) switched-capacitor, and (c) 3-level VRs 65
4.2 Simulated conversion efficiencies of 3-level VRs with fixed and optimal design parameters. Table shows the range of design parameters used in simulations. 66
4.3 Design parameters that maximize efficiencies across duty cycle, output voltage and load current ranges. 67
4.4 Simulated peak-to-peak inductor current ripple ($\Delta I_{L,PP}$) of 3-level and buck VRs in continuous conduction mode (CCM). 68
4.5 Simulated conversion efficiencies of buck VRs across inductance values ($L/R = 2.5nH/\Omega$). 70
4.6 Simulated conversion efficiencies of 3-level and buck VRs across inductor qualities. 71
4.7 Simulated conversion efficiencies of 3-level and switched-capacitor VRs across inductor qualities. 72
4.8 Simulated conversion efficiencies of 3-level and SC VRs with and without bottom-plate parasitic capacitance. 73
4.9 Block diagram of 3-level converter with slow digital feedback control and fast shunt regulation. Finer duty cycle control is necessary to avoid limit-cycling. 75

4.10 Schematic of the proposed 3-level power converter. Signal timing diagrams illustrate different operating modes. 77
4.11 Schematic and waveforms that drive power FETs when duty cycle is over 50%. 79
4.12 High-level architecture of the 3-level converter test-chip prototype. . . 84
4.13 Die micrograph of the converter with dimensions of main blocks. Flying capacitors are placed underneath the inductors to reduce area overhead. The table shows converter specifications. 85
4.14 Measured snapshot of fast dynamic voltage scaling of the converter operating in open-loop. Voltage scales from 1.4V to 0.4V and vice versa within 20ns. 86
4.15 Measured efficiency of converter operating in open-loop. 87
4.16 Measured conversion efficiency with optimal switching frequencies and number of phases. 88
4.17 Open-loop measurement of peak-to-peak output voltage ripple of the 3-level converter with DC load current. Ripple changes across duty cycles, switching frequencies and number of phases. 89
4.18 Histogram of voltage noise measured in open-loop with and without shunt regulator for connected and disconnected power domains of two sectors. 91
4.19 Comparison of open-loop measurement of on-die voltage noise without shunt regulator, with reactive shunt, and with predictive shunt. . . . 92
4.20 High-level diagram of the feedback control in the second version 3-level VR test-chip. 95
4.21 Illustration of how the nonlinear control works. Both PFETs turn on whenever V_{OUT} drops below V_{LOW} (a), while both NFETs turn on whenever V_{OUT} spikes above V_{HIGH} (b). 96
4.22 Die photo of the second version 3-level VR test-chip. Similar to the first version, flying capacitors are placed under the inductors to save die area. 97
4.23 Snapshot of layout showing that two connections betweeen power FETs — N_{BOT}-N_{MID} and P_{BOT}-P_{MID} — have high resistance, which significantly degrades conversion efficiency. This is due to a mistake of using too narrow and long paths to connect different power FETs. 99
4.24 Measured efficiencies are lower than expected due to parasitic resistance caused by wires connecting C_{FLY} (RCfly), wires connecting nFET and pFET power switches (RNP) and bondwires (Rbondwire). Simulated efficiencies including three parasitics match well with measured efficiencies. Higher efficiencies are possible with better inductors with higher Q. 101
4.25 Measured output voltage across duty cycles at 0A load current in open-loop. 103

4.26 Measured voltage traces show that using $N_{DUTYDVFS}$ in closed-loop operation allows the voltage to scale faster. Load current ranges between 0.33A (at 0.6V output voltage) and 0.38A (1.2V output voltage). . . 104
4.27 Measurement shows that voltage scaling is slower in closed-loop than in open-loop. Both operate with 0.33-0.38A load current. 105
4.28 Measured voltage fluctuation with load current steps of various magnitudes (labeled in each subplot) when converter operates in open- and closed-loop. 106
4.29 Measured voltage traces show that nonlinear control reduces voltage droops/spikes during load current steps in closed-loop operation. . . . 106
4.30 Measurement results show that magnitude of voltage fluctuation changes as frequencies of load current steps change. Regulator operates at 200MHz switching frequency. Magnitude of load current step is 0.51A and the current transition occurs within 50ps based on simulation. . . 107
4.31 Measurements with same settings as Figure 4.30 except that regulator operates at 100MHz instead of 200MHz. 107

5.1 Technologies that will impact IVR designs include better transistors [56], thick metal layers and integrated magnetics [34], dense capacitors [103], 2.5D silicon interposers [17] and 3D die stacking [57]. 109

List of Tables

1.1 Comparison of existing off-chip voltage regulators offered by TI [13, 12] and IVR published by IBM [28]. 9

3.1 Processor configuration and system parameters for SESC. 31
3.2 Benchmark Suite. 32
3.3 Characteristics of the on-chip VR (all percentage (%) numbers are relative to the processor energy with DVFS). 54

4.1 Specifications of on-chip spiral inductors modeled using ASITIC [11] and MOS capacitors. 81
4.2 Breakdown of conversion loss of the 3-level converter for three design points. 90
4.3 Comparison with prior IVR designs. 94

Citations to Previously Published Work

Portions of this dissertation have appeared in the following publications:

"A Fully-Integrated 3-Level DC/DC Converter for Nanosecond-Scale DVFS"
IEEE Journal of Solid-State Circuits (JSSC), Jan. 2012
Wonyoung Kim, David Brooks, Gu-Yeon Wei

"A Fully-Integrated 3-Level DC/DC Converter for Nanosecond-Scale DVS with Fast Shunt Regulation "
IEEE International Solid-State Circuits Conference (ISSCC), Feb. 2011
Wonyoung Kim, David Brooks, Gu-Yeon Wei

"System Level Analysis of Fast, Per-Core DVFS using On-Chip Switching Regulators"
IEEE International Symposium on High-Performance Computer Architecture (HPCA), Feb. 2008
Wonyoung Kim, Meeta S. Gupta, Gu-Yeon Wei, David Brooks

"Enabling On-Chip Switching Regulators for Multi-Core Processors using Current Staggering"
Workshop on Architectural Support for Gigascale Integration (ASGI at ISCA), Jun. 2007
Wonyoung Kim, Meeta S. Gupta, Gu-Yeon Wei, David Brooks

Acknowledgments

I am who I am thanks to those surrounding me. I used to joke that whenever I read someone's thesis, I would read the acknowledgement section first. That was partly true and not entirely a joke because I thought while the body of the thesis shows what kind of research the person did, the acknowledgement section shows the human side of how the person spent his/her PhD years. It's not just about thanking people, but also about looking back at your graduate school years and cherishing invaluable memories. So here I am, writing this one day before the thesis submission deadline, thinking about all the good things that happened to me in the last six and a half years and the great people that made those days especially memorable (although I still don't know if Gu approved my thesis).

I was really fortunate to have Gu and David as my co-advisors. When I first arrived at Harvard, I was a 21 year old, just-out-of-undergrad student who had little clue about how to do good research, how to write good papers, and how to present my ideas clearly to other people. Gu and David basically taught me a lot of things that I am proud of today. I remember when we were preparing for my first talk to be given at HPCA 2008. I gave a practice talk and apparently it was so terrible that just giving me feedback on the spot was not going to fix it. Gu asked me to video-tape myself, write down what I said and send that to him so that he could fix the script and I could memorize it. Even after I memorized the whole thing, we practiced 6-7 more times to make the whole presentation sound natural, as if I didn't memorize anything. I remember David coming to Ben and my hotel room to have me rehearse the talk one last time. Eventually, I gave a satisfying talk at the conference and was really excited to see well-known academics ask questions about my research.

Some people might say this seems pretty tough, but I am really thankful for those experiences and I think those trainings made me who I am today. I am not sure if I would have the patience to sit through 6-7 practice talks for anyone else, but Gu and David was willing to do that for the students. I also have very fond memories of having long - sometimes up to 3 hour long - meetings with Gu, David and Meeta when we were working on the HPCA paper. Gu and David's conversations were not just limited to our paper, but would go off on tangents discussing various topics in the semiconductor industry and other random research ideas not necessarily related to our paper. I would just sit there and try to absorb anything I could hear and process since everything was so new to me. Those were really great learning experiences for me as a G1/G2 PhD student and I think it heavily affected my research moving forward. It was also a great feeling to be able to hear and realize how much passion my advisors have on what they work on. I still chat with David on gchat on all sorts of topics to the point that Jiye started complaining about it =)

I also want to thank my committee members. The first time I met Paul was when he served in my quals committee. I had heard great things about him before, but I realized during the quals that he was genuinely interested in learning new things. He would ask questions not to criticize me, but because he was really interested in the answers. Tanay Karnik was kind enough to fly all the way from Hillsboro to serve in my PhD committee. I had seen his name in numerous IVR papers from Intel and always wanted to talk to him. So I just walked up to him at ISSCC to ask several questions and he was kind of enough to have a discussions with me then and also afterwards via email.

I can't even imagine what my PhD life would have been like without our crazy group members. Crazy here is a compliment. I feel excited, entertained and motivated by crazy people. All of our group people are not only great researchers (or research machines injecting coffee and spitting out papers), but are also people who are interested in a wide variety of things and love to have lively discussions. The VLSI group people when I first got here - Hayun, Andrew, Ankur, Ruwan, Amber and Mark - were all kind enough to teach me all the nitty-gritty things I should know to start designing chips. Once I started working with Meeta around March 2007, she taught me so many things, although she might not have realized that. I think that was when I started going over to MD307 often to hang out with Alex, who would be often focused in day-trading, and Ben, who would always be very well organized and dressed formally. I really enjoyed going to ISCA and HPCA with Alex and Ben. It was my first time going to a conference and they introduced me to a lot of people and taught me what to do in conferences. I have lots of good memories having endless conversations with Vijay about random stuff. Vijay might like to think they were philosophical conversations, but I recall most of them were not =P I learned a lot from his work ethics and perfectionist tendencies when writing papers. I was always worried whenever group members left that life would be more boring without those leaving. It turned out there were always new members who were as interesting as those who left. Tao, Saekyu, Silvia, Mario in MD311 and Svilen, Kevin, Simone, Brandon, Mike, Bob, Sophia, Amanda in 307. It even feels almost meaningless to list people from two groups separately since we always have lunch together and somehow find different topics to talk about everyday. I'll never forget the fun conversations

over lunch and feeling sleepy after filling myself up with truck food. I'll definitely miss Svilen's daily greetings of "What's going on Wonyoung", Kevin's socks that used to lie around in 307, Kevin and Mike's absurd debates, learning about China from Tao and Silvia and countless memories of crunching before tape-out and paper deadlines. I can bet with pretty high confidence that our group has one of the most exciting group dinners. I hope we'll have a chance to gather altogether in the future.

I would also like to thank Glenn for his tremendous help throughout the years. I think Glenn might be the single most important person in our group since he boosts the productivity of everyone! Arriving at Harvard with little idea on how a server works, he was very patient in teaching me all the nitty-gritty things I should know to make the best use of our high-performance linux machines. One unfortunate thing that might bother Glenn is that I still open multiple emacs windows...

I have long-lasting memories with friends outside of lab. Taeg Sang hyung, Myung Jin nuna, Will and HwanChul. At some point we were spending so much time together that it was hard to believe we will be apart someday. Everyone is now scattered across different states and countries, but I'm sure we'll have a chance to get together to do some stupid things (we'll need Will for that haha). Hong Ha, Jung Ook, Yejin, all born in 84, were really good friends. Hye Young nuna, In Keun hyung - it was unfortunate we met so late into my PhD years, but it was always a joy to discuss Korean and US politics with them. I also have fun memories with my high-school friends, Heesang and Leebong, both being energetic and passionate people that I love to mingle with.

I cannot thank my family enough for all the love they've given me. Although far

away in Korea, my parents and my older brother were always supportive of whatever I wanted to do with my life (although it's true I've never set out to do something really crazy). They were always there to give me love and advice full of wisdom. I miss them very much and often wish I could've visited Korea more often during my PhD. Thankfully, Jiye was here for me at Boston. This journey would have been so much more boring, dull and less colorful had Jiye not been there the whole time with me. Memories are best when shared with someone you love. Since coming to Boston as 21-year old students with no clue about research or living in the US, Jiye and I have gone through so many experiences together and I am thankful to be able to share so much memories with her. After 3 years of married life and 7 years since we started dating in December 2005, she still surprises me with her thoughtfulness and understanding towards other people and herself.

Thank you so much for everything everyone. I am who I am thanks to all of you.

Chapter 1

Challenges of Delivering Power to SoCs and the Need for Integrated Voltage Regulators

Contents

1.1 Challenge of delivering power to modern SoCs 2
1.2 Solution: Integrated Voltage Regulators 5

1.1 Challenge of delivering power to modern SoCs

The rise of mobile computing places ever-increasing demands on high performance and low power for future microprocessor designs, not only for the mobile devices but also for the back-end servers needed to support their proliferation. In light of these demands, chip architects have moved towards tightly integrated system-on-chips (SoC) that incorporate multiple cores and heterogeneous components (e.g., memory controllers, hardware accelerators, etc.) into a single chip. Such complex SoC systems require sophisticated power delivery schemes to manage power efficiently.

Figure 1.1 shows a high-level diagram of how power is typically delivered from a high-voltage source to an SoC that operates at around 1V in mobile and server systems. In mobile handsets and laptops, power comes from batteries operating at around 3.7V and 5-15V, respectively. Server systems deliver power at higher voltages such as, for example, 480VDC in Facebook's datacenters [5] and 110VAC from wall plugs for smaller servers, and convert them down to 12V at the motherboard where the SoC sits. The conversion from a higher voltage to 12V is not drawn in Figure 1.1. Off-chip voltage regulators (VRs or often called DC-DC converters) convert the high 3.7V or 12V down to a voltage range that the SoC can operate under, which is 0.7-1.1V in this case. This form of power delivery has the following problems.

1. **Wastes power due to shared voltage domains across multiple cores and IP blocks:** Since one VR can deliver one voltage, the number of required VRs and associated board components is proportional to the number of SoC voltage domains. Multiple cores typically share a single voltage partly because it is difficult to duplicate bulky off-chip VRs due to board area overhead and

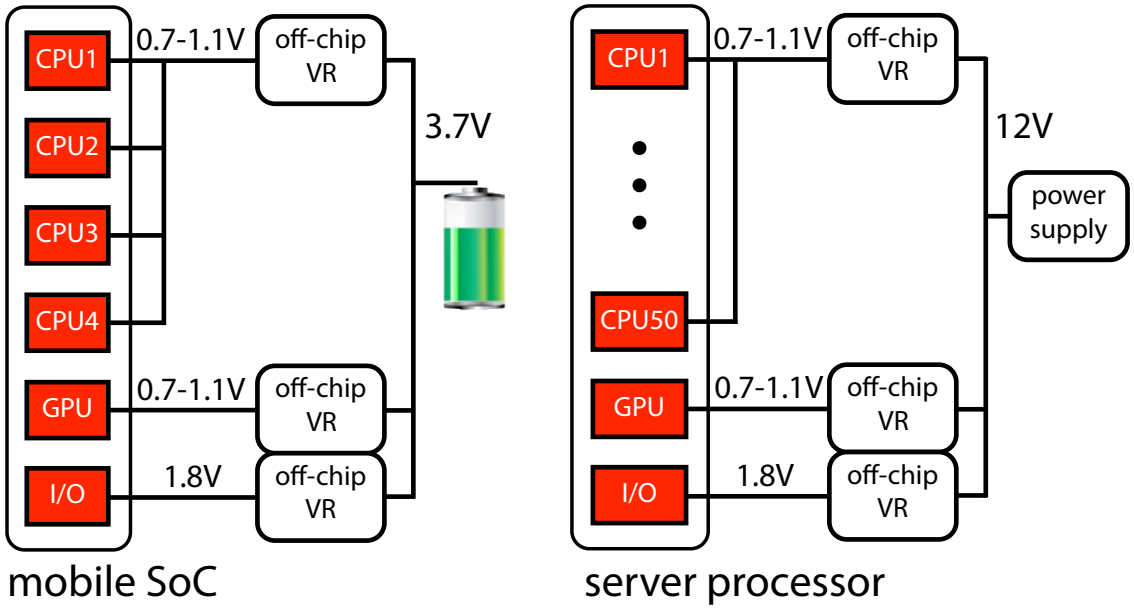

Figure 1.1: Power delivery in mobile and server systems

the challenge of routing large numbers of voltage rails on the board. Since performance demands can vary widely across cores [51], a shared voltage cannot track the different demands, which leads to wasted energy (Figure 1.2).

2. **Wastes power due to slow DVFS:** Existing off-chip VRs can scale the voltage at microsecond timescales [32], which is not fast enough to track fast-changing CPU demands [84]. The mismatch between the voltage and CPU demand results in wasted energy (Figure 1.2). Voltage scaling is slow with off-chip VRs because of large amounts of decouling capacitors (decap) on the board, package and SoC die (Figure 1.3). There are parasitic inductance on the path connecting the off-chip VR and SoC die. Since the parasitic inductance can cause large voltage fluctuation, designers place decaps on the board, package and SoC dies to suppress voltage noise. Whenever the off-chip VR changes

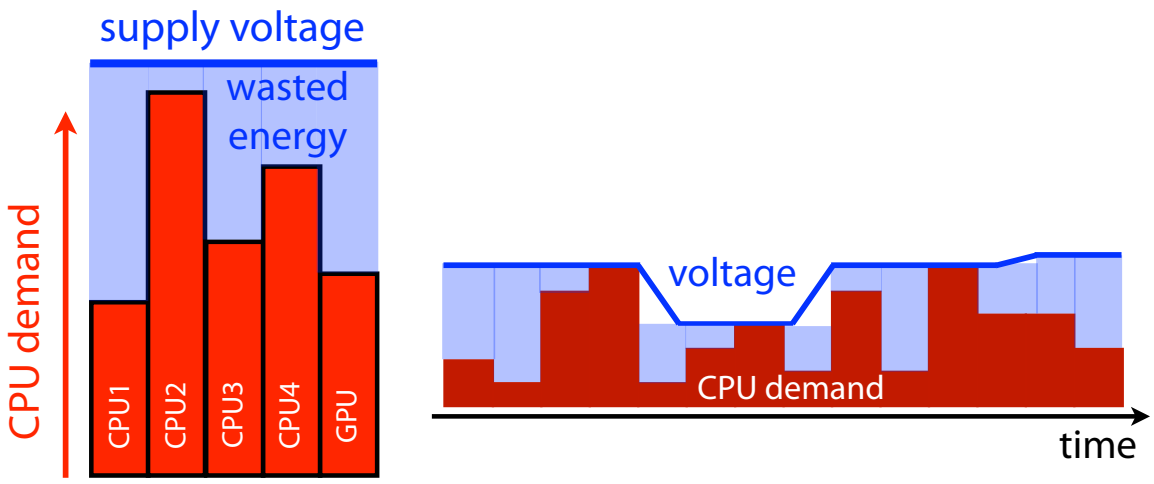

Figure 1.2: Illustration showing how shared voltage domains and slow DVFS lead to wasted energy.

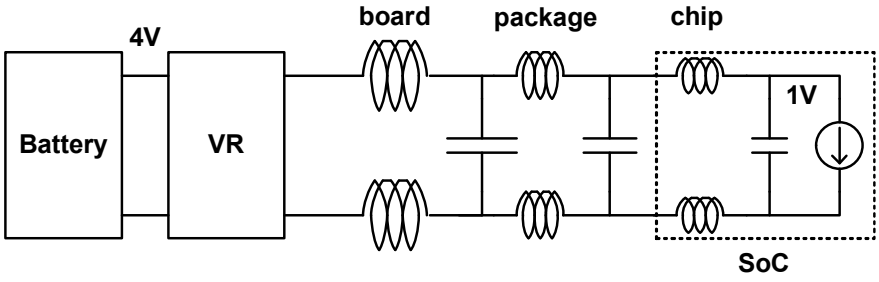

Figure 1.3: Large amounts of decaps in the board, package and SoC die decrease the speed of voltage scaling

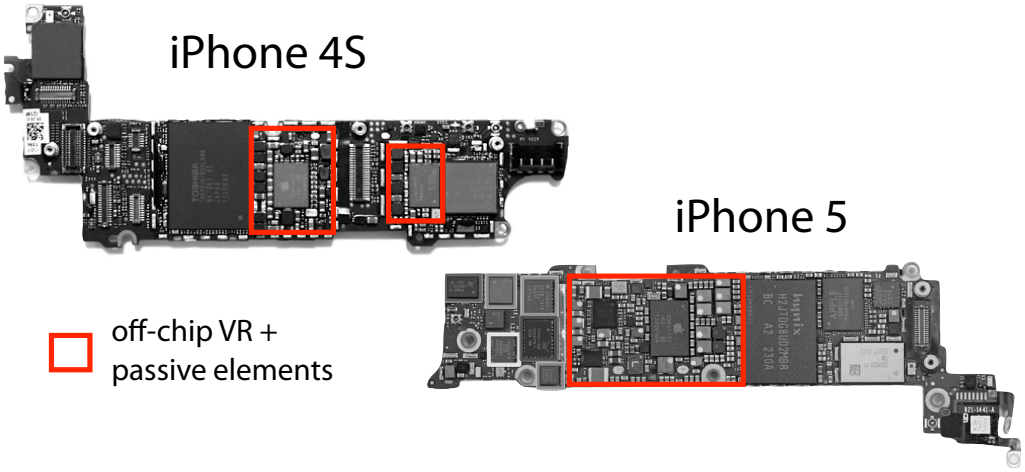

Figure 1.4: Teardown images of iPhone 4S and iPhone 5 show that off-chip VRs and required passive elements occupy a significant area on the board

the voltage, it has to charge/discharge all of the decaps, which makes voltage transition slow.

3. **Occupies large board area:** Reducing board area is important especially for portable electronics because a smaller board leaves more room to fit a larger battery in a constrained space, which enables longer battery life. Figure 1.4, a teardown image of iPhone 4S and iPhone 5 [8], shows that off-chip VRs and required board-level inductors and capacitors occupy a significant area. It also shows that off-chip VR area has not decreased over phone generations.

4. **Costly due to multiple board-level components:** Existing off-chip VRs usually consist of three board-level components – power switches, inductors and capacitors. Some VRs use a separate feedback controller chip, but others integrate them in the same die as the power switches. As the number of SoC voltage domains increase, the number of board components required for off-chip VRs increase proportionally, increasing cost and complexity of board design. As logic ICs become more complex, there can be up to 10 voltage domains [14], which requires roughly 30 board components for off-chip VRs. A VR solution using fewer board components has the potential to reduce component cost and simplify power delivery on the board.

1.2 Solution: Integrated Voltage Regulators

What if we can design a VR that drastically reduces the number and size of required board-level passives? What if the entire VR solution — including power

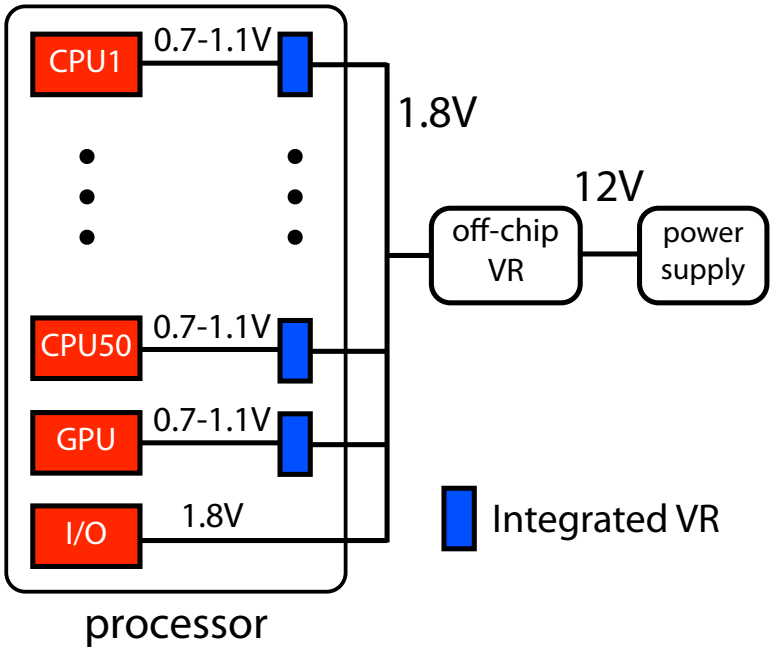

Figure 1.5: Power delivery using IVRs in a server system

switches and passives — could be small enough to be integrated in the SoC die or package? There has been a rising interest in building integrated VRs (IVR) occupying much smaller footprint and using fewer discrete components than off-chip VRs [42, 88, 91]. To tackle the aforementioned problems of existing off-chip VRs, this thesis builds upon prior works and studies system-level benefits of IVRs and proposes ways to build more efficient IVRs.

Figure 1.5 shows an example of how IVRs can change power delivery in a server system. An off-chip VR converts 12V to an intermediate voltage, which is 1.8V in this example, and multiple IVRs integrated in the processor die or package deliver different voltages to each core/IP-block depending on their processing demands. Following are potential benefits of this power delivery scheme using IVRs.

1. **1000 times faster voltage scaling than off-chip VRs:**. Figure 1.6 compares

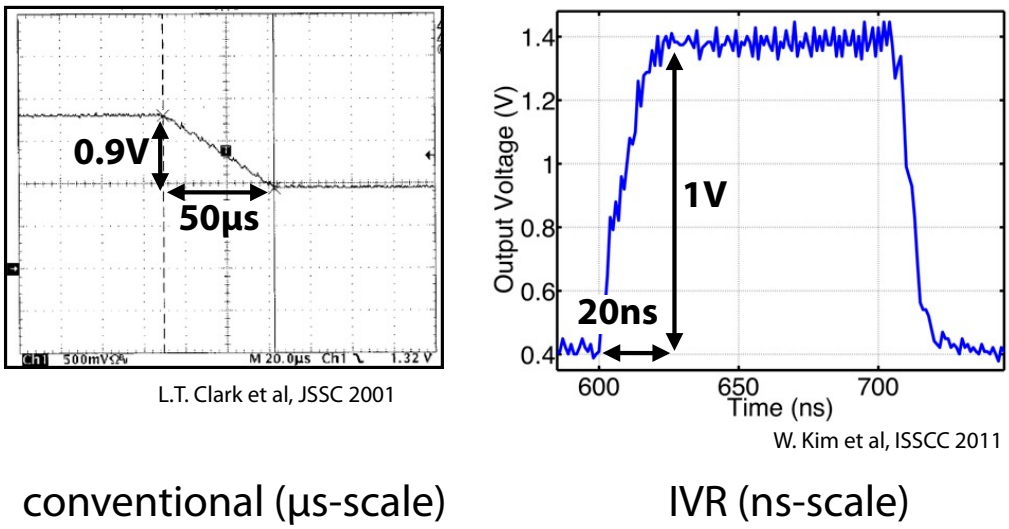

Figure 1.6: IVR enables nanosecond timescale DVFS compared to microsecond DVFS with existing off-chip VRs

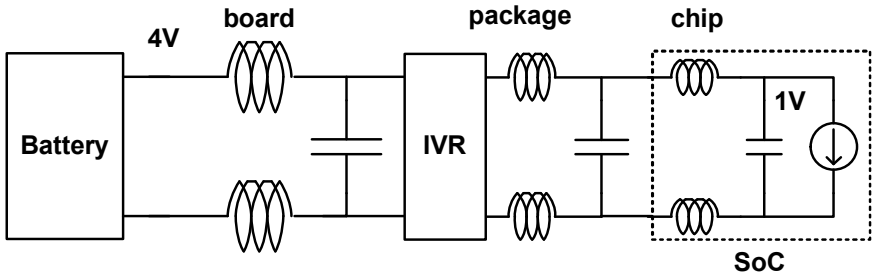

Figure 1.7: IVR enables fast voltage scaling by reducing the amount of decap to charge/discharge

voltage traces of a typical off-chip VR [32] and measured results of an IVR test-chip [53]. The IVR can scale the voltage across 1V within 20ns, which is more than 1000 times faster than the microsecond time-scale in the off-chip VR case. Nanosecond timescale voltage scaling is possible because IVRs are placed close to the processor, either on the same die or package, and hence need to charge/discharge less capacitance than conventional off-chip VRs (Figure 1.7).

2. **10 times smaller footprint than the smallest off-chip VRs commer-**

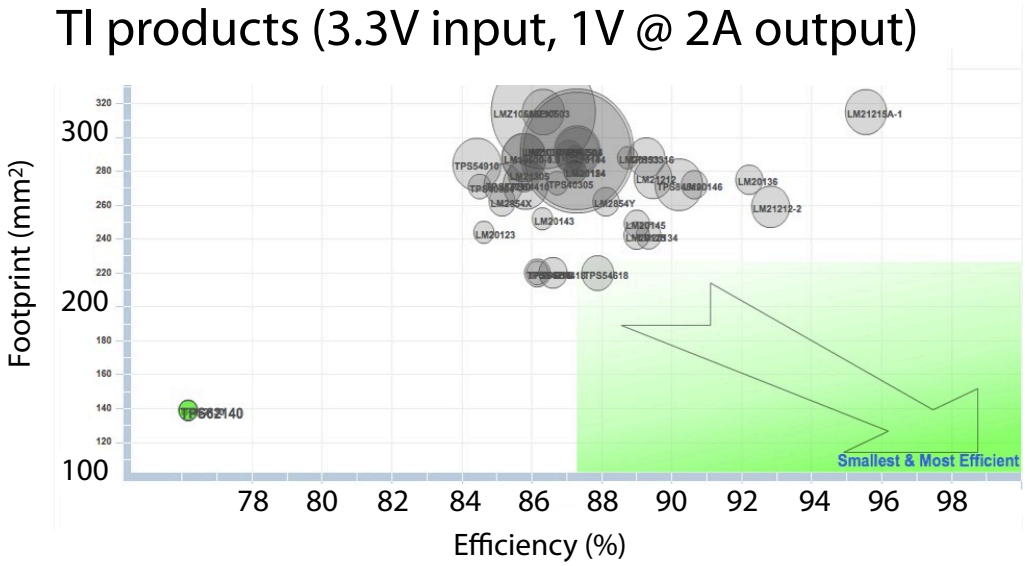

Figure 1.8: Efficiency and footprint data for TI voltage regulator products plotted using TI's WEBENCH, a simulator provided by TI for its voltage regulator products [13]

cially available: With smaller footprints, IVRs can reduce board area, leaving more room for larger batteries in mobile electronics. Figure 1.8 presents efficiency and footprint data for TI's off-chip VR products plotted using TI's WEBENCH, a simulator provided by TI. Existing off-chip VRs from TI occupy 100-300mm^2 with 75-96% efficiency when converting 3.3V to 1V at 2A current. Since regulator footprint is roughly proportional to the current that needs to be delivered to the output, we use current density(A/mm^2) to compare regulator footprints across a wide range of load currents. TI's products in Figure 1.8 presents 0.007-0.02A/mm^2 current densities. In contrast, IBM recently reported an IVR that is 90% efficient with current densities as high as 2A/mm^2, albeit at a lower input voltage of 2V being converted down to 1V [28]. TI's recent 2012 product called MicroSiP, not included in Figure 1.8, is specifically tuned

	TI other products	TI MicroSiP	IBM IVR
efficiency (%)	75-96	80	90
current density (A/mm^2)	0.007-0.02	0.07	2
input voltage (V)	3.3	3.3	2
output voltage (V)	1	1.2	1

Table 1.1: Comparison of existing off-chip voltage regulators offered by TI [13, 12] and IVR published by IBM [28].

towards lower footprint [12]. When converting 3.3V to 1.2V, MicroSiP is 80% efficient with $0.07A/mm^2$, which translates into a more than 20 times larger footprint per delivered current compared to IBM's IVR. Table 1.1 summarizes the specifications of existing off-chip regulators versus the IVR presented by IBM.

3. **Reduce I^2R loss and simplify board-level power distribution:** Cross-country grids deliver electricity at a high voltage and low current to reduce I^2R loss. Similarly, IVRs can reduce I^2R loss on the board by delivery power at a high voltage and converting down to a lower voltage at the point of load (Figure 1.9). This is especially important for high-performance server processors with maximum current exceeding 100A [10]. In these processors, a mere $1m\Omega$ parasitic resistance on the board can add 10W of I^2R loss, assuming 100A delivered at 1V. Using an IVR that instead delivers 50A at 2V, we can reduce

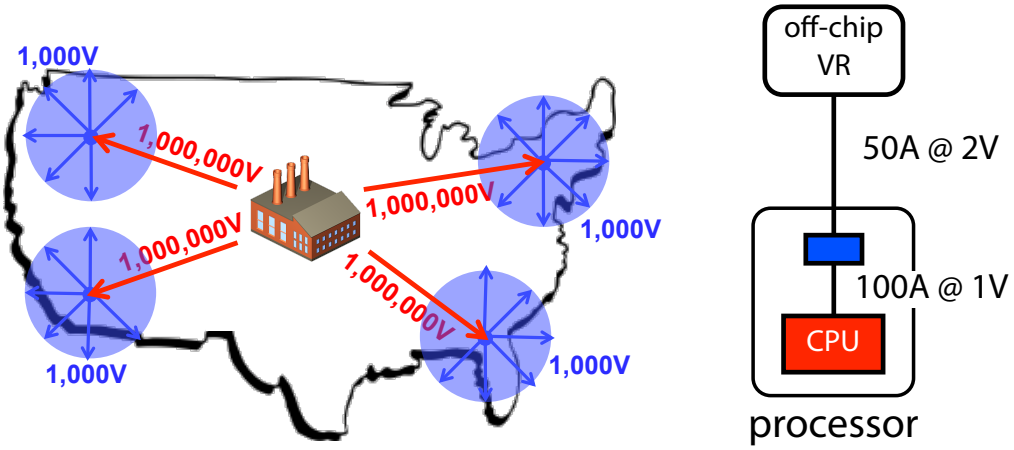

Figure 1.9: IVRs can reduce I^2R loss on the board by delivering power at a high voltage and low current and converting to a lower voltage at the point of load.

this loss to 2.5W. Moreover, IVRs can simplify board-level power distribution and potentially reduce parasitic resistance especially for processors that need large numbers of voltage domains. Revisting Figure 1.5, IVRs let the off-chip VR deliver only a single voltage on the board. Compared to a case where the board is split into multiple voltage planes for delivering multiple voltages, power distribution using IVRs simplifies board-level power distribution and allows the single voltage plane to have less parasitic resistance than split power planes.

To study the benefits of IVRs in delivering power to SoCs, this thesis presents the following points in the next chapters.

1. I provide a brief background on the basics of VR design and prior works on IVRs that this thesis has built upon (Chapter 2).

2. Through a system-level simulation study on the benefits of using IVRs, I show that fast, per-core DVFS can save up to 20% power in a 1.6W 4-core processor (Chapter 3).

3. I present measurement results from two IVR test-chips built using a 3-level topology, which is a hybrid form of an inductor-based buck and a switched-capacitor VR. The 3-level VR reduces inductor size and presents higher efficiencies compared to existing buck VRs (Chapter 4).

4. I discuss what future process technologies can further improve IVR efficiencies and current densities beyond those of IVRs built using standard digital CMOS processes (Chapter 5).

Chapter 2

Basics of Voltage Regulators and Challenges of Integration

Contents

2.1 Basics of step-down voltage regulators 13

2.2 Challenges of Integrating Voltage Regulators 15

2.3 Evolution of IVRs 21

2.1 Basics of step-down voltage regulators

While IVRs facilitate fast, per-core DVFS, they also introduce various overheads compared to existing off-chip VRs. In order to understand these overheads, this section provides an overview of existing step-down off-chip VRs and IVRs.

Switching and linear VRs are two widely-used step-down VR topologies. Linear VRs offer several advantages: ease of on-chip integration, relatively small size, and good response to load current transients [39]. Unfortunately, the maximum achievable power-conversion efficiency of a linear VR is constrained by the ratio of V_{OUT} (the output voltage of the VR) to V_{IN} (the input voltage to the VR). For example, when a linear regulator converts a 1.1V V_{IN} to a 1V V_{OUT}, high power conversion efficiency ($\sim$90%) is possible. However, as V_{OUT} decreases further and deviates away from the input voltage, maximum efficiency degrades linearly. When delivering power to a processor using DVFS, the VR has to deliver a wide range of output voltage levels (e.g., 0.7-1.1V), in which case the efficiency degradation of a linear VR can be prohibitively high at low V_{OUT} levels.

In contrast, a switching VR can regulate a wide range of output voltage levels with higher power-conversion efficiency that is less sensitive to the V_{OUT}/V_{IN} ratio. Hence, switching VRs are better suited for loads employing DVFS [111]. This higher conversion efficiency stems from its reliance on inductors and/or capacitors as low-loss energy-transfer devices between V_{IN} and V_{OUT}, but they can be bulky and consume large area. While there are several types of step-down switching VRs — those using inductors (buck VR), capacitors (switched-capacitor VR) or both (3-level VR) — to transfer energy, we will first study inductor-based buck VRs, which is the most

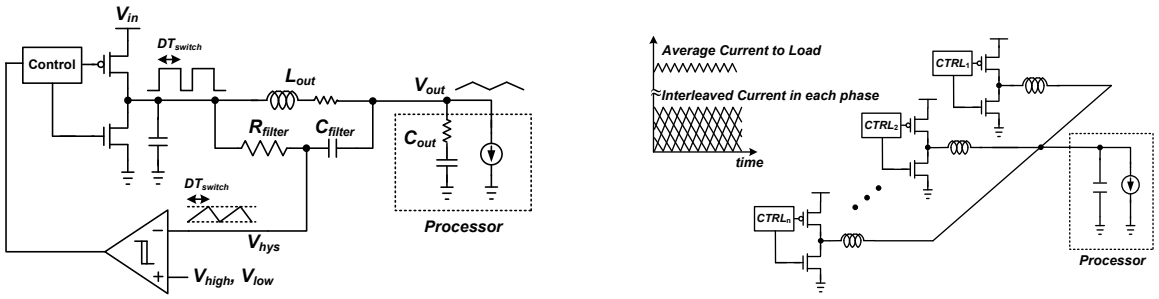

(a) Buck converter with hysteretic control

(b) Multiphase buck converter

Figure 2.1: Buck converter schematics

popular topology for existing off-chip VRs. We will examine the two other topologies, switched-capacitor and 3-level VRs, in more detail later in Chapter 4.

A typical inductor-based buck VR, shown in Figure 2.1(a), consists of three sets of components: switching power transistors, the output filter inductor (L_{OUT}) and capacitor (C_{OUT}), and the feedback control consisting of a hysteretic comparator and associated filter elements (C_{FILTER} and R_{FILTER}) that enhance loop stability. The power transistors can simply be viewed as an inverter that switches on and off at a switching frequency and provides a square wave to the low-pass output filter composed of L_{OUT} and C_{OUT}. The VR output, V_{OUT}, powers the microprocessor load and its voltage is approximately set by the duty cycle of the square wave. This regulated voltage exhibits small ripples since the filter attenuates the high-frequency square wave. The feedback loop is closed by feeding V_{HYS}, which is the output of the filter composed of C_{FILTER} and R_{FILTER}, to the hysteretic comparator. The duty cycle of the square-wave input to the power transistors is set by the hysteretic comparator output. As shown in Figure 2.1(a), the hysteretic comparator has a high threshold (V_{HIGH}) and a low threshold voltage (V_{LOW}). The PMOS power switch turns on when

V_{HYS} drops below V_{LOW}, and the NMOS turns on when the V_{HYS} increases above V_{HIGH}. Since V_{OUT} directly affects V_{HYS}, when V_{OUT} fluctuates in response to load current transients, hysteretic control can react very quickly. While there are several other feedback control schemes one can employ for a buck VR, hysteretic control is one example that offers fast transient response characteristics while keeping design complexity low [69].

The power transistors and inductor shown in Figure 2.1(b) can be interleaved to form a multiphase buck converter. Researchers have proposed multiphase converters for high load current applications [117, 76, 122], since they can reduce the peak current in each inductor. Parallel sets of power transistors and inductors are interleaved and connected to the same load such that current through each inductor is interleaved across even time intervals. Hence these interleaved inductor currents cancel out at the output node and result in an average current that has small ripple. Moreover, this interleaving accommodates the use of small output filter capacitance while meeting small voltage ripple constraints. Since the number of necessary phases increase with load current, VR footprint is roughly proportional to load current.

2.2 Challenges of Integrating Voltage Regulators

While there are various specifications in VR design, conversion efficiency and footprint (or current density) are two of the most important for both off-chip and integrated VRs. Revisiting Figure 1.8, we see that there is a trade-off between footprint and efficiency. As all designs in the figure are buck VRs, the trade-off is present because larger inductors have higher Q, which leads to smaller conversion loss. A

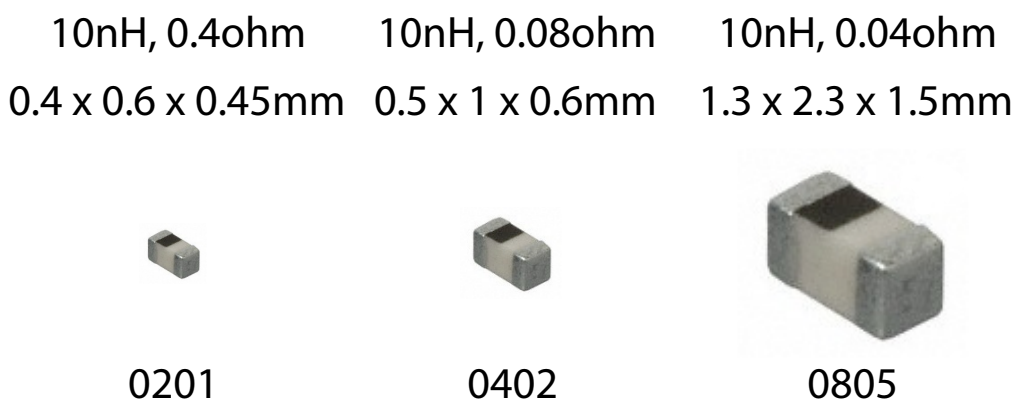

Figure 2.2: Dimensions, DC resistance and inductance of various chip inductors from Coilcraft [4]. Images of inductors is roughly to scale.

big challenge of building IVRs is to make the footprint small enough for integration while achieving high efficiencies. To maximize efficiency (and minimize loss), it is important to understand the sources of VR losses.

Typical buck VRs have the following three main sources of losses.

1. **Capacitive loss of power transistors:** When power transistors switch on and off to generate a square wave on the output (input of inductor), there is a CV^2f loss due to parasitic gate capacitance of the power transistors. This loss is proportional to switching frequency and power switch width.

2. **Resistive loss of power transistors:** As current flows through the power transistors, there is I^2R loss due to on-state parasitic resistance of the power transistors. This loss is inversely proportional to power switch width. As a result, there is a trade-off between capacitive and resistive loss of power transistors [96].

3. **Resistive loss of inductors:** Non-ideal inductors lead to I^2R loss associated

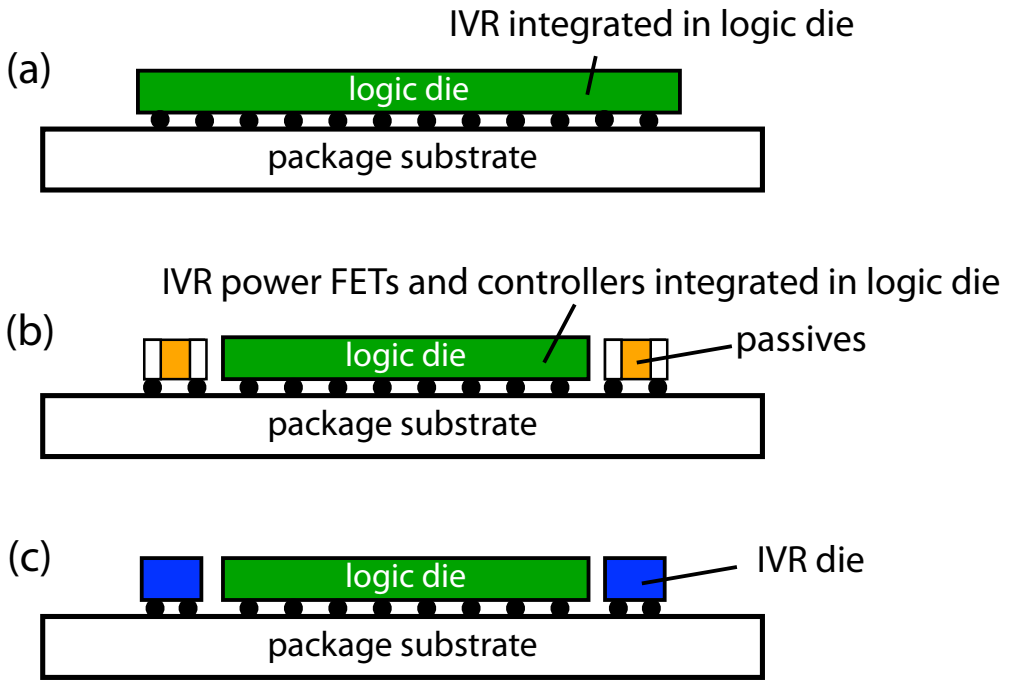

Figure 2.3: IVRs can be integrated on the package-level (a,b) or in the logic IC die (c).

with the inductor coil resistance. Larger inductors have higher Q and lower resistance (Figure 2.2), which leads to a trade-off between conversion efficiency and footprint. Since inductors small enough to be integrated on-die or on-package have lower Q than larger inductors, they tend to have larger parasitic resistance. Assuming a fixed material and process for the inductor, one way to reduce inductor resistance is to reduce the inductance (L), since Q is equal to $2\pi fL/R$ (Note that Q changes with L, so R does not stay proportional to L. The relationship between R and L depends on the structure of the inductor.). To reduce L, we need to increase the switching frequency of the VR to maintain small output voltage ripple. This leads to larger capacitive loss of power transistors, resulting in a trade-off between transistor capacitive loss and inductor

resistive loss.

With unlimited footprint, designers can maximize VR efficiency by using large inductors with high Q, power transistors with high breakdown voltages and switch them at very low frequencies to minimize capacitive losses. However, to implement IVRs with small footprint, designers are forced to use low Q, small L inductors and switch power transistors at high frequencies. Since conventional power transistors with high breakdown voltages are not suitable for high switching frequencies, researchers have proposed using standard digital CMOS transistors as power transistors in IVRs [58, 88]. These transistors cannot sustain high voltages, which is why Figure 1.5 in the previous chapter uses an off-chip VR to convert 12V to an intermediate voltage of 1.8V instead of using IVRs to convert 12V to 1V.

Given these challenges in implementing IVRs, there are several ways to integrate VRs with technologies that exist today (Figure 2.3).

1. **Single-die integration:** Integrating IVRs and logic ICs on a single die offers the highest level of scalability in terms of the number of voltages that can be provided to the logic IC (Figure 2.3(a)). Integrated in the logic die, IVRs could be scattered around the die to provide per-core voltage control even in manycore processors with over 50 cores such as Intel's Xeon Phi [6] or Tilera's Tile 64 [7]. However, the problem is that standard logic process technologies typically do not offer high quality passives and metal layers that are thick enough to deliver large amounts of current. Including these processes in the logic die incurs cost of adding mask layers. Furthermore, IVRs can add large die area, which is especially costly in cutting-edge processes (e.g., 20/22nm,

28/32nm) that are used to fabricate new logic IC products. Amount of IVR die area overhead depends on the power densities of the IVR and processor. For example, Intel's high-end laptop processor (Core i7-3940XM Ivy Bridge) has 55W TDP with 160mm^2 die area, which results in a maximum power density of 0.34W/mm^2 [9]. Assuming an area-efficient IVR with 2W/mm^2 power density, which is one of the highest values reported, the IVR adds 17% additional die area when integrated on the processor die. IVR area overhead can be smaller in processors with lower power consumption. Intel's Core i7-3612QM, which is slightly less powerful than Core i7-3940XM, has 35W TDP with 160mm^2 die area, which results in a maximum power density of 0.22W/mm^2 and 11% IVR die overhead. As a result, single-die integration could be more suitable for low-power SoCs than high-performance servers with high power densities.

2. **Package-level integration with on-package SMT passives:** Designers can integrate power transistors and VR controller blocks in the logic IC die while mounting small SMT chip inductors and capacitors on the package (Figure 2.3(b)). They can take advantage of high-quality passives without adding costly masks to the logic die. Cost of adding on-package passives might be acceptable for high-performance processors since they already have a large number of on-package decoupling capacitors, but the cost might be not as acceptable for mobile SoCs that do not have any discrete deoupling capacitors on-package. Moreoever, SoCs in mobile phones are typically contained in a package-on-package (PoP) in which the package is too thin to mount SMT passives unless the passives are custom-made to be thinner than standard ones. Revisiting

Figure 2.2, it shows that the thickness of an 0201 inductor is 0.45mm, which is too thick to fit in the 0.2mm thick bottom layer of a PoP where the logic package usually sits [3].

3. **Package-level integration with separate IVR dies:** IVRs can be implemented in a separate die using process technologies optimized for high-quality passives, thick metal layers and transistors with low on-state resistance (Figure 2.3(c)). Instead of paying the price of larger die area in expensive, cutting-edge processes, IVRs can be fabricated in a separate die using a process that is not as advanced as those for logic ICs, but more optimized for VR applications. However, known-good-die (KGD) is a problem as is the case in any multi-chip module (MCM) with multiple dies. If there is an error in the relatively cheaper IVR die, the entire package, including the more expensive logic die, is considered faulty since it is very costly to dissemble the MCM and replace the IVR die. To reduce the cost of dealing with faulty IVR dies, it is very important to fully test the IVR die on the wafer-level to guarantee it is a "good die" before integrating in the MCM. However, this is challenging because wafer-level testing is usually more costly and has more restrictions than package-level testing.

There has been various prior works on implmenting IVRs using the aforementioned integration methods. The next section lists prior works on IVRs and how these designs integrated different IVR components.

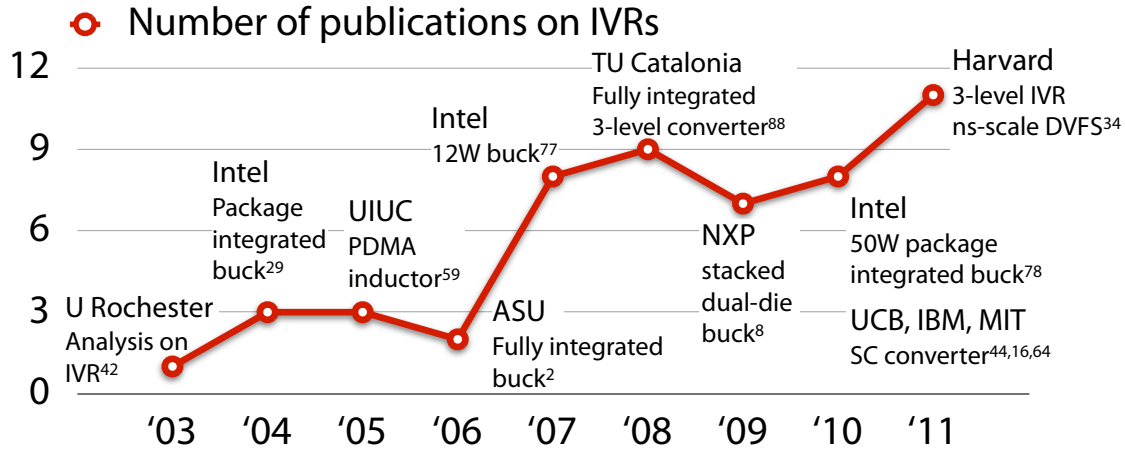

Figure 2.4: There has been increasing interest on IVRs by industry and research communities.

2.3 Evolution of IVRs

IVR publications started to appear in 2003 and have steadily increased, constantly introducing new demonstrations and techniques for IVR design (Figure 2.4).

1. **Before 2003:** PCB mounted buck VRs mainly consisted of power switches, controller chips, inductors and capacitors mounted on the board in separate packages. Power switches built in mature process technologies limited the switching frequency usually to lower than 1MHz, requiring large inductors ($\geq$1μH) and capacitors. Efficiencies reached 95%, but footprints were large and current densities were low in the order of 1-10mA/mm^2.

2. **2003-2007: Package- and chip-level integration:** Following feasibility analyses on integrated buck converters [58, 88], researchers presented buck converters with package- and chip-level integration [41, 91, 42, 71, 70, 74, 19, 112, 100, 73, 107, 89, 62, 21]. Intel used standard digital CMOS transistors cascoded

to sustain higher input voltages switching at over 100MHz. Instead of using low Q on-chip spiral inductors, they mounted small, high-quality SMT chip inductors in the range of 1-20nH on the package. efficiency. Using on-package inductors could be a viable solution for Intel since their processors already have a large number of on-package capacitors and adding several more inductors might not add much cost. However, adding on-package inductors could be a bigger leap for mobile SoCs with no existing on-package capacitance. Other works presented a single-chip solution using on-chip spiral inductors to simplify package design, albeit with lower efficiency due to poor inductor quality.

3. **2008-2011: Fully-integrated switched-capacitor and 3-level converters:** To rely less on low-quality inductors while providing a single-chip solution, other works proposed fully-integrated SC and 3-level converters [82, 102, 81, 60, 28, 83, 52, 121, 48, 36, 26, 59]. We will compare these topologies in more detail in Chapter 4. IBM presented a switched-capacitor converter using deep trench capacitors that are 20 times denser than MOSFET capacitors, saving significant amount of die area and achieiving high current densities. Harvard designed a 3-level converter using a 1nH inductor with capacitors placed underneath to save die area. At the same time, researchers continued to improve both package-integrated [98, 63, 43, 25, 24, 65, 38, 46, 90, 37, 97, 64] and fully-integrated buck converters [108, 101, 113, 22, 20, 99, 72, 54, 110, 68, 55]. Intel presented a buck converter that could deliver 50W [90], paving the way for IVRs to be integrated in high-performance systems. NXP proposed a dual-die solution where a die optimized for high-quality inductors is stacked on top of

another die containing power switches and control circuitry [25].

Now that we have studied the basics of IVR design and what the main challenges are, we take a step back and analyze how much system-level energy savings is possible using IVRs.

Chapter 3

System-Level Energy Savings with Fast, Per-Core DVFS using Integrated Voltage Regulators

Contents

3.1 Prior Works on Fine-Grain DVFS 28

3.2 Potential of Fast and Per-Core DVFS Schemes 29

3.2.1 Simulation Framework 31

3.2.2 Offline DVFS Algorithm 32

3.2.3 Effects of Finer Temporal Resolution 34

3.2.4 Per-Core vs. Chip-Wide DVFS 36

3.3 Characteristics of On-Chip Regulators 40

3.3.1 Model and Simulation of Buck VR 40

3.3.2 Design trade-offs of IVRs 41

3.3.3 Regulator Efficiency 44

3.3.4 Load Transient Response 47

3.3.5 Voltage Scaling Time 49

3.3.6 On-Chip Regulators for Single and Multiple Power Domains 50

3.4 Energy Savings for Per-Core and Chip-Wide DVFS using On-Chip Regulators 53

3.4.1 Comparison of Energy Savings 54

3.4.2 Power Domain Scalability 58

Dynamic voltage and frequency scaling (DVFS) was introduced in the 90's [66], offering great promise to dramatically reduce power consumption in large digital systems by adapting both voltage and frequency of the system with respect to changing workloads [93, 95, 45, 116]. Unfortunately, the full promise of DVFS has been hindered by slow off-chip voltage VRs that lack the ability to adjust to different voltages at small time scales. Modern implementations are limited to temporally coarse-grained adjustments governed by runtime software (i.e. the operating system) [1]. Moreover, the large footprint of off-chip VRs make it difficult to use large numbers of them for per-core or per-IP block voltage control.

This chapter explores the interplay of the promising characteristics and costs of employing IVR designs in modern CMP system architectures. While this study considers CMP designs comprising multiple low-power processor cores within the context of a mobile embedded system, the analysis described can be extended to higher-power processors as well.

Figure 3.1 illustrates three power-supply configurations that this chapter studies.

1. **Slow, Global DVFS:** The first configuration (left) represents a conventional design scenario that only uses an off-chip VR. This VR directly steps the power supply voltage, assumed to be 3.7V provided by a Li-Ion battery, down to a processor voltage ranging from 0.6V to 1V.

2. **Fast, Global DVFS:** The second configuration (middle) implements a two-step voltage conversion scenario. Given an inherent degradation in conversion efficiencies for large step-down ratios, an off-chip regulator performs the initial step-down from 3.7V to 1.8V, which can be shared by other on-board compo-

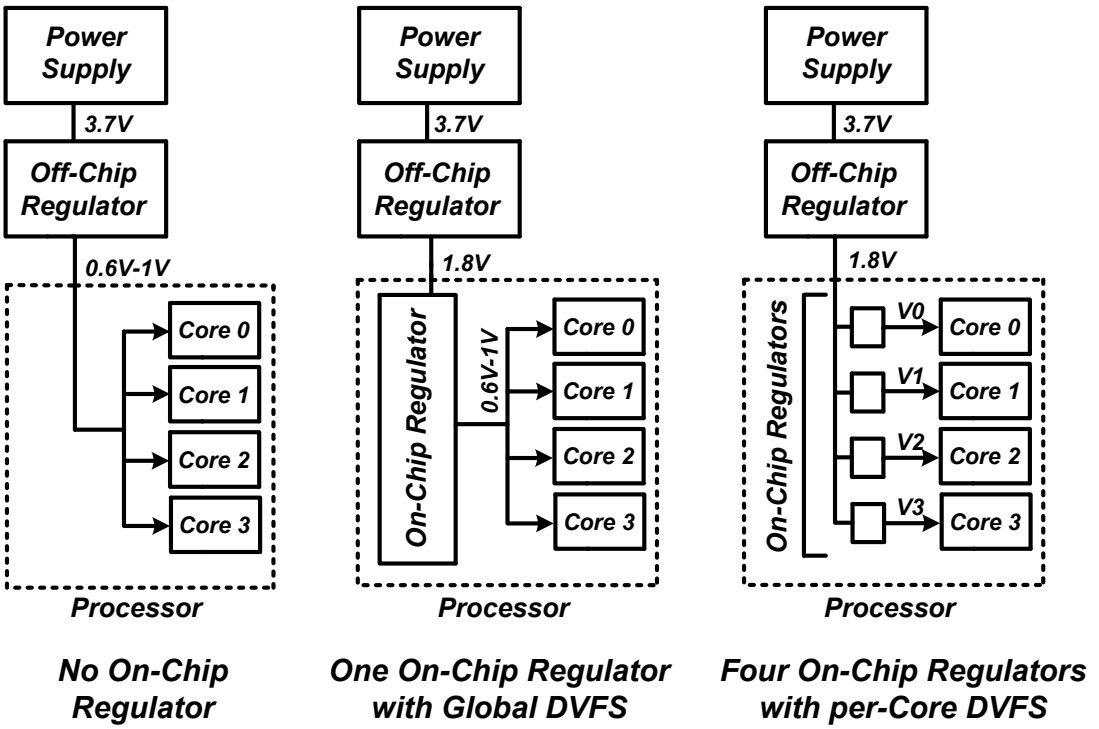

Figure 3.1: Three power-supply configurations for a 4-core CMP.

nents. The 1.8V supply then drives an on-chip voltage regulator that further steps the voltage down to a range of 0.6V to 1V as a single power supply domain distributed across a 4-core CMP.

3. **Fast, Per-Core DVFS:** The third configuration (right) expands on the second configuration by providing four separate on-chip power domains via individual IVRs. These three configurations constitute the framework through which we compare the costs and benefits of fast, per-core DVFS enabled by IVRs.

The main points of this chapter are as follows:

- We explore the energy savings offered by implementing both temporally fine-grained and per-core DVFS in a 4-core CMP system using an offline DVFS algorithm (Section 3.2).

- We present a buck-type IVR design space analysis that considers key regulator characteristics—DVFS transition times and overheads, load current transient response, and regulator conversion losses (Section 3.3).

- We combine the energy savings with the IVR cost models and come to several conclusions. For a single power domain, on-chip regulator losses offset the gains from fast DVFS for many workloads. In contrast, fast, per-core DVFS can achieve energy savings ($\geq$20%) when compared to conventional, single power domain, off-chip VRs with comparatively slow DVFS (Section 3.4).

3.1 Prior Works on Fine-Grain DVFS

There has been prior work that has focused on exploring the benefits of multiple frequency/power domains in microprocessors compared to a global frequency/voltage. In the area of CMP systems, per-core DVFS has been shown to offer larger energy savings than chip-wide DVFS using four different voltage and frequency levels [45], but this work considered relatively coarse DVFS time intervals and did not consider any of the issues related to power supply regulation. Other works explore multiple clock domain (MCD) architectures, which use globally asynchronous, locally synchronous (GALS) techniques to provide *within-core* energy control. These techniques have demonstrated 17% improvement in energy-delay product compared to using a single domain [93]. An adaptive reaction time scheme for multiple clock domain processors have been proposed [116]. These works focus on the energy savings of the processor using per-core DVFS, and the algorithms associated with it, but do not consider the

practical overheads of integrating multiple on-chip regulators. As this chapter shows, the practical overheads of on-chip regulators must be considered to argue that per-core DVFS actually has large energy savings. At the circuit-level, there have been many works demonstrating on-chip regulators [40, 87, 111, 18], but these works solely analyze the energy conversion efficiency of regulators. These works do not consider any of the system-level overheads (DVFS scaling and voltage transient analysis) or the system-level benefits of on-chip regulators. The contribution of this chapter is the aggregation of ideal energy savings using per-core DVFS with the practical overheads of integrating on-chip regulators within each processor core.

3.2 Potential of Fast and Per-Core DVFS Schemes

Dynamic voltage and frequency scaling can be an effective technique to reduce power consumption in processors. DVFS control algorithms can be implemented at different levels, such as in the processor microarchitecture [67], the operating system scheduler [47], or through compiler algorithms [118, 44]. Most previous work in the domain of DVFS control algorithms focus on coarse temporal granularity, e.g., voltage changes on the order of several microseconds, which is appropriate given slow response times of off-chip VRs. In contrast, on-chip regulators offer much faster voltage transitions as presented in Figure 3.2. This figure, a simulation of the IVR model described in a later section, shows voltage transitions can occur on the order of tens of nanoseconds, several orders of magnitude faster than off-chip VRs. DVFS algorithms implemented at the microarchitecture level provide the finest level of temporal control, hence, are good candidates for the fine-grained approach that we consider.

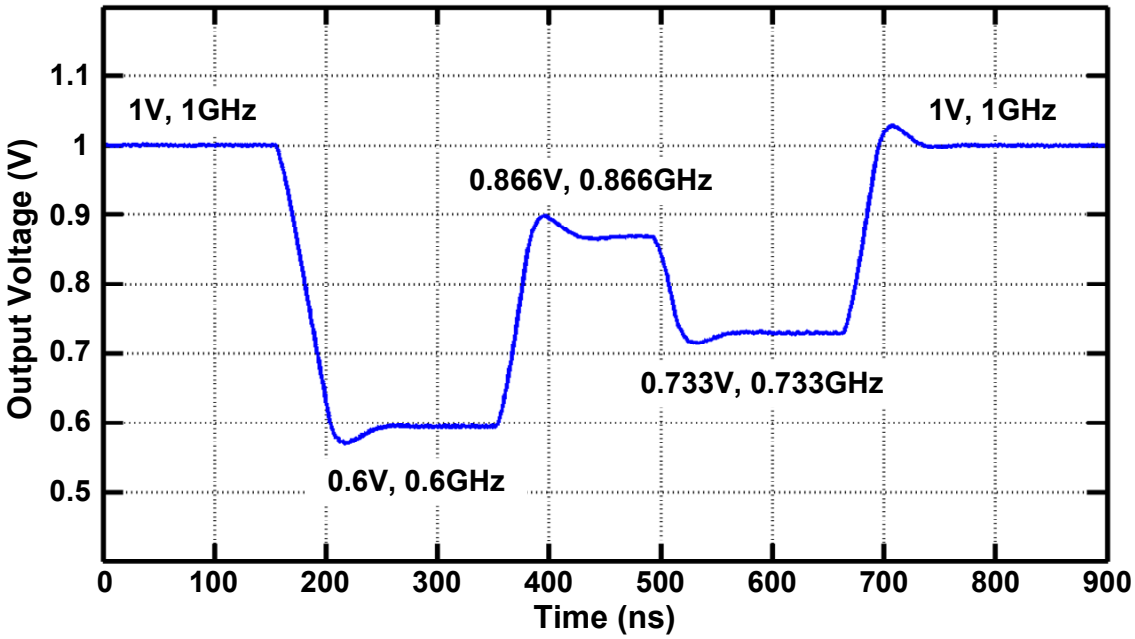

Figure 3.2: DVFS transition times with an IVR

In this section, we explore the benefits of fast DVFS with fine temporal resolution and also highlight the benefits of per-core voltage domains compared to chip-wide DVFS. To explore the benefits and tradeoffs associated with temporally fine-grained and per-core DVFS, we rely on an offline DVFS algorithm that can easily be applied across the wide range of DVFS transition times we consider. Section 3.2.1 provides a brief overview of the simulation framework used in our study, and the methodology of the offline DVFS algorithm is described in Section 3.2.2. We then discuss the effects of finer temporal granularity (Section 3.2.3), and the savings for per-core versus chip-wide DVFS schemes (Section 3.2.4).

Frequency	1GHz @ 65nm	**Vdd**	1 V
Core Area	16mm^2	**Fetch/Issue/Retire**	2/2/2
Branch Penalty	7 cycles Hybrid Branch Predictor	**Branch Predictor**	BTB (1K entries) RAS (32 entries)
Int registers	32	**FP registers**	32
IL1	32KB, 32-way, 32B block Hi/Miss latency 2/1 cycles	**DL1**	32KB, 32-way, 32B block Hi/Miss latency 2/1 cycles MESI-protocol
ITLB entries	64	**DTLB entries**	128
MSHR size	8	**Write Buffer size**	16
L2 size	512 KB	**MSHR size**	16

Table 3.1: Processor configuration and system parameters for SESC.

3.2.1 Simulation Framework

We employ an architectural power-performance simulator that generates realistic current traces. We use SESC [86], a multi-core simulator, integrated with power-models based on Wattch [27], Cacti [94], and Orion [104]. A simple in-order processor model represents configurations similar to embedded processors like Xscale [31]. The per-core current load is 400mA when fully active and 120mA when idle. We model a configuration with a shared-L2 configuration, private-L1 caches in each processor, and a MESI-based coherence protocol. Table 3.1 lists the details of the 4-core processor configuration and system parameters. The simulator was modified to obtain cycle-by-cycle current profiles for each core in the system.

In a CMP-based system, it is important to understand the interactions between the multiple cores. These interactions can be accurately characterized by analyzing a mix of multi-threaded and multi-programmed benchmarks. We use a composite benchmark suite composed of applications from SPEC2K, ALPBench [61], and SPLASH2 [114]. For multi-programmed scenarios, we consider several mixtures of a

Benchmarks	Description	$\frac{\text{Memory Cycles}}{\text{Total Runtime}}$
ocean-con	Large Scale Ocean Simulation	0.47
fft	Fast Fourier Transform	0.4
facerec	CSU Face Recognizer	0.22
cholesky	Cholesky Factorization	0.197
raytrace	Tachyon Ray Tracer	0.058
mcf4	4-high memory-bound (*mcf*)	0.697
mcf3, applu1	3-high memory-bound (*mcf*) and 1-high cpu-bound (*applu*)	0.697 (*mcf*) and 0.051 (*applu*)
mcf2, applu2	2-high memory-bound (*mcf*) and 2-high cpu-bound (*applu*)	
mcf1, applu3	1-high memory-bound (*mcf*) and 3-high cpu-bound (*applu*)	
applu4	4-high cpu-bound (*applu*)	0.051

Table 3.2: Benchmark Suite.

memory-bound benchmark (*mcf*) and a cpu-bound benchmark (*applu*) from SPEC2K. Table 3.2 lists the different benchmarks used in this study along with the ratio of memory cycles to total runtime of the application for each. All benchmarks are run for 400M instructions after fast forwarding through the initialization phase.

3.2.2 Offline DVFS Algorithm

The goal of any DVFS algorithm is to minimize energy consumption of the application within certain performance constraints. This can be done by exploiting the slack due to asynchronous memory events. Scaling down the frequency of the processor slows down cpu-bound operations, but does not affect the time taken by memory-bound operations. We exploit the presence of such memory-bound intervals to reduce the voltage and frequency of the processor. The effectiveness of such a DVFS scheme is directly related to the ratio of memory-bound cycles to cpu-bound

cycles.

As this chapter aims to study the potential system-wide benefits of using on-chip voltage regulators, the offline algorithm is applied to all configurations and it optimizes DVFS settings based on a global view of workload characteristics. We formulate the DVFS control problem as an integer linear programming (ILP) optimization problem, which seeks to reduce the total power consumption of the processor within specific performance constraints (δ). This approach is similar to the one proposed in [118]. We divide the application runtime into N intervals based on different temporal granularities of DVFS. A total of L = 4 voltage/frequency (V/F) levels are considered. For each runtime interval i and frequency j, the power consumption, P_{ij}, is calculated. The delay for each interval and V/F level, D_{ij}, is also calculated. Heuristics for the delay of individual intervals are obtained by calculating the relative memory-boundness of each interval through cache miss behavior. Equations 3.1- 3.3 specify the ILP formulation of our offline algorithm. The overheads associated with switching between different voltage/frequencies settings are not considered in the optimization, but are included later in Section 3.3.

$$min(\sum_{i=1}^{N}\sum_{j=1}^{L} P_{ij}x_{ij}) \tag{3.1}$$

$$(\sum_{i=1}^{N}\sum_{j=1}^{L} D_{ij}x_{ij}) < \delta \tag{3.2}$$

$$\sum_{i=1}^{N}\sum_{j=1}^{L} x_{ij} = N \tag{3.3}$$

We consider an in-order processor with the capability of switching between four voltage settings: 1V, 0.866V, 0.733V, and 0.6V, with proportionally scaled frequencies from 1GHz down to 600MHz. As in Xscale [31], we assume the processor can operate

through voltage transitions by quickly ramping down frequency before the voltage ramps down. Conversely, we ramp up the voltage and only switch the frequency after the voltage has settled to higher levels. Clock synthesis that combines finely-spaced edges out of a delay-locked loop can provide rapid frequency adjustment without PLL re-lock penalties [33].

The offline algorithm finds voltage/frequency settings at each interval to minimize power while maintaining a specified performance constraint. In this study, we consider performance constraints of 1%, 5%, 10%, 15%, and 20%. In order to keep the runtime overheads of the ILP algorithm tractable, we divide the simulation trace into smaller windows of 2M cycles each; finding optimal DVFS assignments within the windows, but not necessarily across the entire trace. The overall power savings presented in this chapter represents the average power savings across all 2M-cycle windows for each application.

3.2.3 Effects of Finer Temporal Resolution

IVRs allow voltage transitions to occur at a rate of tens of nanoseconds as compared to microseconds for off-chip VRs. The fast voltage-scaling capability of IVRs provides the potential for applying DVFS at very fine-grained timescales. A fine-grained DVFS scheme can more closely track different cpu- and memory-bound phases than a coarse-grained scheme and, hence, reduce power consumption without performance degradation. However, the power-saving benefits of a fine-grained technique depend on the distribution of memory misses in the benchmark.

Figure 3.3(a) shows the impact of scaling temporal DVFS resolutions for *mcf*

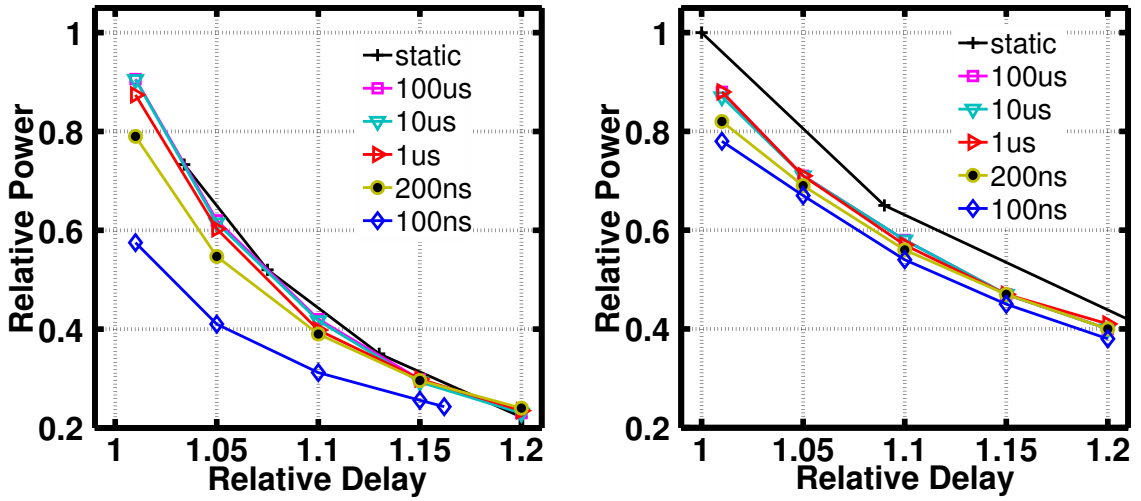

(a) *mcf* (b) *fft*

Figure 3.3: Benefits of fine-grained DVFS scheme for *mcf* and *fft*.

and *fft*. Resolutions in the range of 10-100μs represent the coarse-grained DVFS schemes and 100-200ns represent fine-grained, on-chip DVFS. We also consider a static voltage/frequency scaling scheme (representative of coarse-grained OS-level control) that fixes DVFS settings at one point for the entire benchmark for each performance target. In some cases, the ILP algorithm fails to match the performance constraint and data points may deviate from initial performance targets. As discussed previously, *mcf* is a memory-bound benchmark, with approximately 70% of its runtime spent servicing memory misses. The fine-grained approach can capture these memory-miss intervals and achieve as much as 60% power savings for only 5% performance degradation. In contrast, coarse-resolution windows fail to capture all of these intervals, achieving less power savings for the same performance constraint (between 35-40% savings for the same 5% performance loss). In general,we find that the benefits of fast DVFS depends heavily on the application. For example, fine-grained DVFS is not much

better than the coarse-grained schemes for *fft* (Figure 3.3(b)), but show an 8% power benefit compared to static voltage/frequency scaling.

3.2.4 Per-Core vs. Chip-Wide DVFS

Chip multiprocessor systems running heterogeneous workloads add the dimension of benefiting from per-core DVFS. Isci et al. show multiple power domains offer power savings in CMP systems over a single power domain [45]. However, due to cost and system board area constraints, it may not be practical to implement multiple power domains using off-chip voltage regulators. On the other hand, IVRs can easily be modified to accommodate multiple power-domains with little additional cost (explained in Section 3.3). We refer to chip-wide DVFS as a global setting for voltage/frequency of the entire chip based on the activity of the whole chip, as opposed to each core. In this section we compare per-core and chip-wide DVFS schemes with 100ns transition times for both multi-threaded and multi-programmed workloads.

Figure 3.4 plots the relative power savings for per-core DVFS and chip-wide DVFS schemes across a range of multi-threaded benchmarks and a significant difference can be observed for most of the benchmarks (e.g., *ocean, fft, facerec*). However, benchmarks like *raytrace* yield only slight differences between the two approaches. This can be attributed to the highly cpu-bound behavior of *raytrace*, which offers fewer frequency-scaling opportunities.

Multi-threaded applications can have similar phases (cpu- or memory-bound) of operation across the cores. Figure 3.5(a) shows a snapshot of activity on each core for a four-threaded version of *ocean*. We see similar behavior across all four threads, but

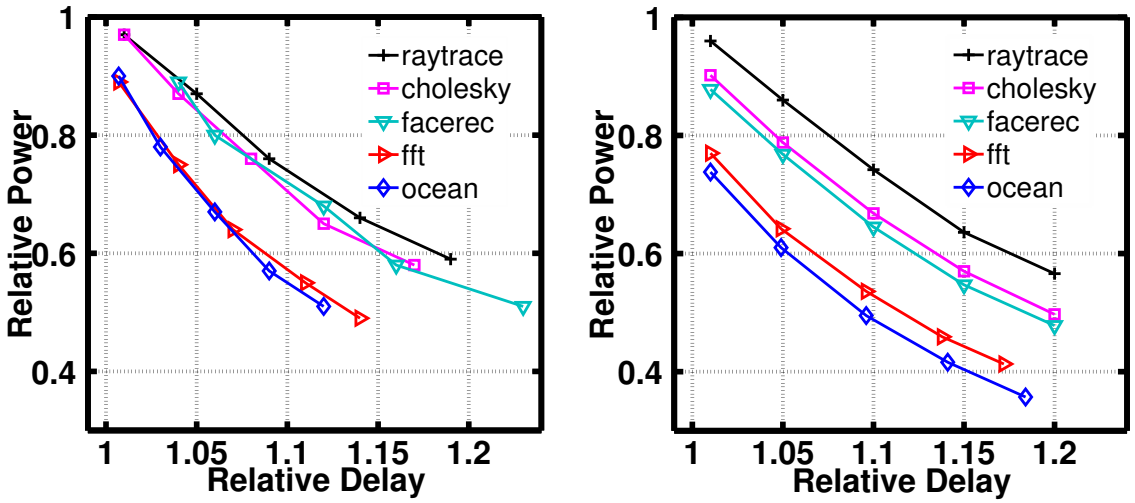

(a) Chip-Wide DVFS (b) Per-Core DVFS

Figure 3.4: Per-Core DVFS for multi-threaded applications.

there is a slight shift in the activity across the cores. While per-core DVFS is able to capture DVFS scaling opportunities in the individual threads, the time windows where the scaling is applied are different. Because of this, a chip-wide DVFS scheme, based on the combined activity of the four threads, finds fewer DVFS scaling opportunities as shown by the global scaling in Figure 3.5(b). In contrast, Figure 3.6(a) presents the activity snapshot for *fft*. We see that the activity profiles of core 0 and core 2 are synchronized in time, as are the activity profiles of core 1 and core 3. This leads to a more effective chip-wide DVFS schedule, demonstrated by the global scaling in Figure 3.6(b). As mentioned in Section 3.2.2, the offline algorithm relies on a global view of each 2M cycle window and, hence, the local voltage/frequency assignments for short intervals shown do not necessarily line up with local activities.

Figure 3.7 plots the relative power vs. delay for multi-programmed scenarios with per-core and chip-wide DVFS. The figure shows different combinations of *mcf* (a

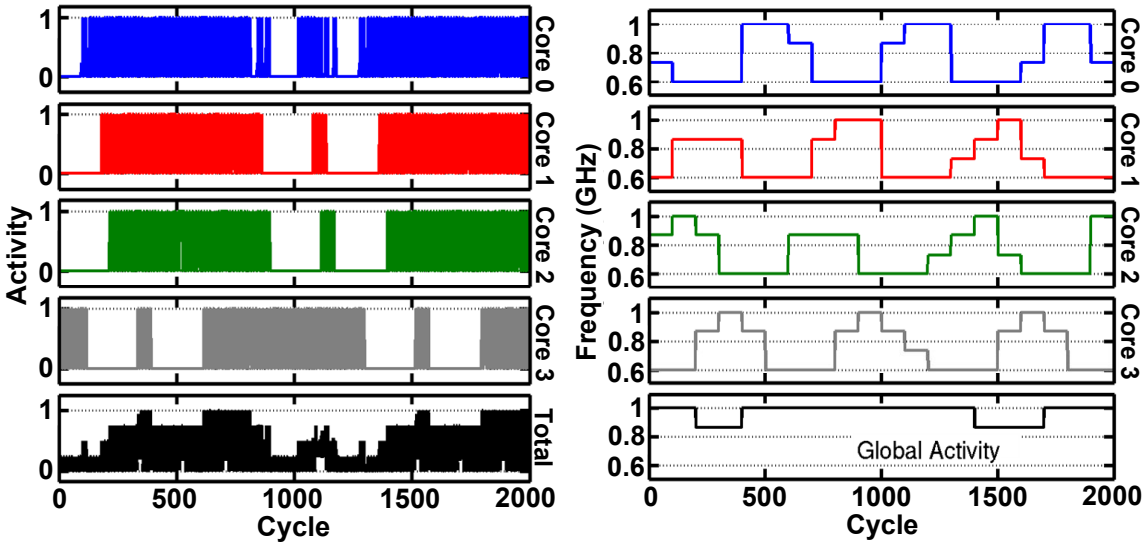

(a) Activity profile of *ocean* (b) Frequency settings

Figure 3.5: Snapshot of *ocean* with per-core and chip-wide DVFS.

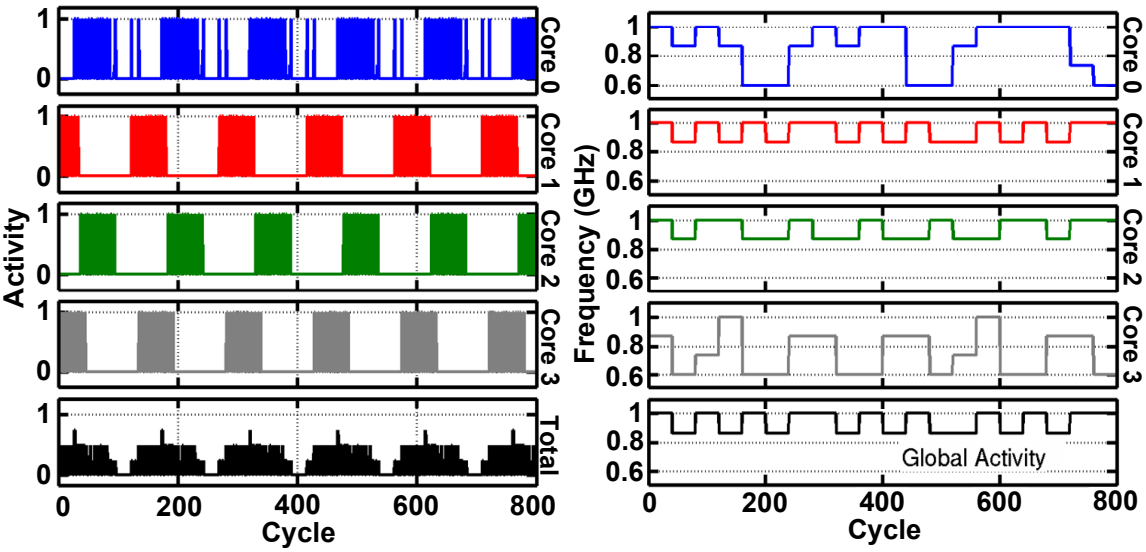

(a) Activity profile of *fft* (b) Frequency settings

Figure 3.6: Snapshot of *fft* with per-core and chip-wide DVFS.

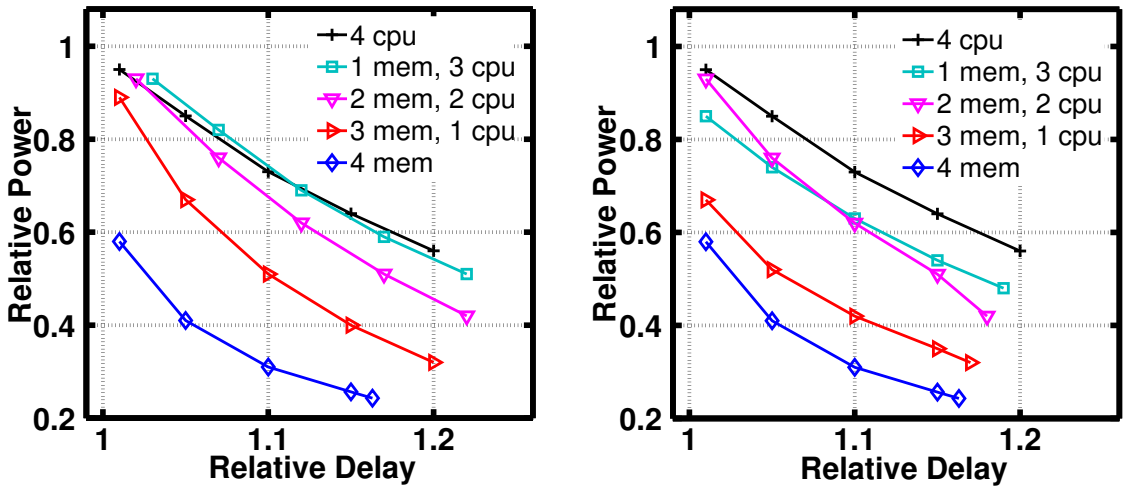

(a) Chip-Wide DVFS (b) Per-Core DVFS

Figure 3.7: Per-core DVFS for multi-programming scenarios.

memory-bound application) and *applu* (a cpu-bound application), ranging from all four cores running *mcf* to all four cores running *applu*. Per-core DVFS achieves similar power savings as chip-wide DVFS for both extremes (all memory-bound or all cpu-bound) as there is little per-core variation to exploit. On the other hand, we observe an additional 18% of power savings for the per-core DVFS scheme over the chip-wide DVFS scheme at a performance degradation of 5% when one copy of *applu* and three copies of *mcf* are run on the 4-core machine.

These results show that depending on the heterogeneity of workload characteristics, per-core DVFS offers substantial additional savings compared to global DVFS schemes by better adapting to the different requirements of each core.

3.3 Characteristics of On-Chip Regulators

3.3.1 Model and Simulation of Buck VR

Before we study the various design trade-offs and overheads of IVRs, this section describes how the off-chip and on-chip VRs are modeled and simulated in this chapter. Figure 3.8 illustrates the overall power delivery network of the example embedded system, from the Li-Ion battery to the processor load, for two VR configurations—with and without an on-chip VR. This is a more detailed version of Figure 3.1, adding in the parasitic elements associated with the power delivery network. This figure shows the parasitic inductors and resistors along the PCB trace and package, and decoupling capacitance added to mitigate voltage fluctuations. This model is derived from the Intel Pentium 4 package model, but scaled to be consistent with our assumptions of power draw in embedded processors [35]. The off-chip VR is modeled as an ideal voltage source, but losses are accounted for by using power-conversion efficiencies extracted from published datasheets [16].

The on-chip VR is modeled in greater detail with parasitics. We assume an on-chip VR using a commercial 65nm CMOS process. Extensive SPICE simulations were run to extract parasitic values that can significantly affect VR efficiency and performance. These parasitics include feedback control path delays, power MOSFET gate capacitance and on-state resistance, and on-chip decoupling capacitor losses.

The inductors required by the on-chip VRs are assumed to be air-core surface-mount inductors [4] attached on-package [40, 87]. The inductors are connected via C4 bumps, which introduce series resistance. The total number of C4 bumps for power

is assumed to be equal for both off-chip and on-chip VRs for fair comparisons. For the on-chip VR, we use 60% of the C4 bumps to connect package-mounted inductors to the die. The remaining bumps are used to connect Vin of the on-chip VR to the PCB. Since the off-chip scheme uses more C4 bumps to connect the processor to the package, it has lower package-to-chip impedance compared to the on-chip scheme. Careful modeling of parasitic losses is required to accurately estimate on-chip VR efficiency, which is found to be consistent with published results [40, 87].

Transient response characteristics also impact the efficacy of using on-chip voltage VRs. Hence, we rely on a detailed Matlab-Simulink model of the on-chip VR to thoroughly investigate the VR's performance given load current transients and voltage transition demands of realistic workloads seen in Section 3.2. The model is built using the SimPowerSystems blockset [15] of Simulink. This Simulink model includes all of the parasitic elements described above since they also impact transient behavior in addition to efficiency.

The next section studies the characteristics of on-chip VRs in more depth with simulation results based on the aforementioned model. The characteristics are presented in comparison to those for an off-chip VR. We also study the tradeoffs associated with different VR characteristics in order to minimize overheads.

3.3.2 Design trade-offs of IVRs

VRs are typically off-chip devices [69, 117, 76, 122] due to the large power transistors and output filter components that are required. However, this VR module can occupy a significant portion of the PCB area, making it costly to utilize multiple

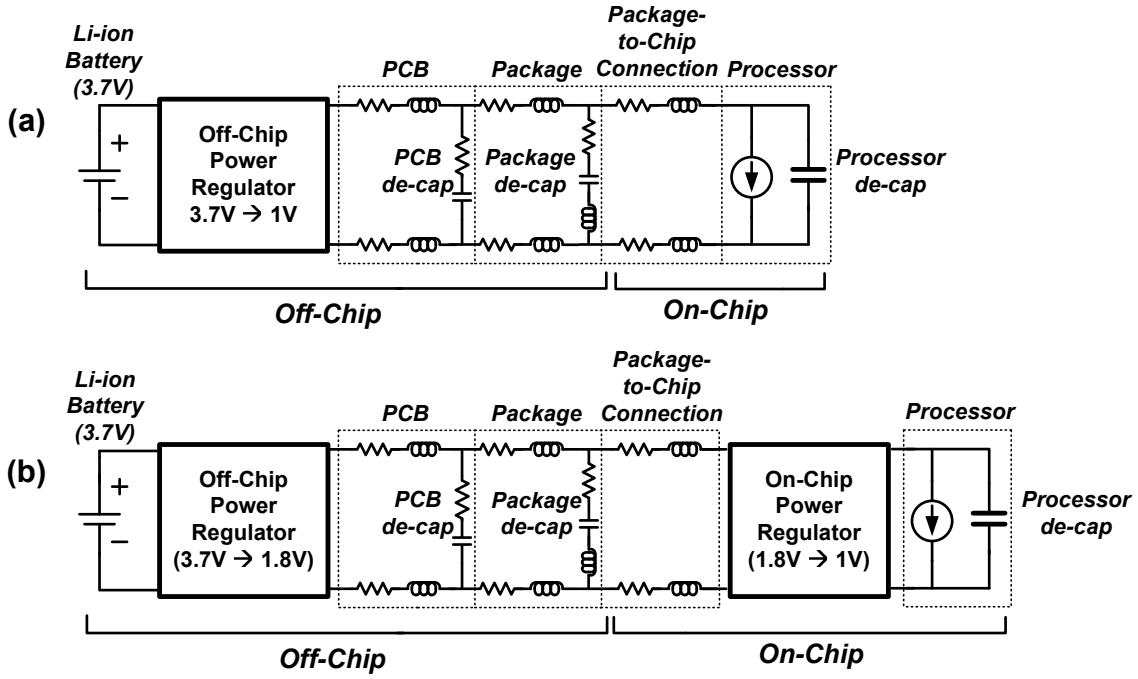

Figure 3.8: Power delivery network using (a) only off-chip and (b) both off-chip and on-chip VRs.

VRs for per-core DVFS. Recently, on-chip VRs have been proposed, integrated on the same die as the processor load [40, 87, 111, 18]. By using much higher switching frequencies, the bulky off-chip inductors and capacitors can be reduced in size and moved onto the package and die, respectively. Hence, on-chip VRs offer an interesting solution that can supply multiple power domains in CMPs with per-core DVFS.

In addition to reducing size, on-chip VRs are also capable of fast voltage switching, which again results from higher switching frequencies. The switching frequency of an off-chip VR is typically on the order of hundreds of KHz to single-digit MHz, whereas on-chip VR designs push switching frequency above 100MHz. Unfortunately, the higher frequency switching comes at the cost of degrading the conversion efficiency of

on-chip VRs, lower than that of their off-chip counterparts. Hence, there are tradeoffs between VR size, voltage switching speed, and conversion efficiency.

In order to design an on-chip VR with minimum overheads, we study three important VR characteristics: VR efficiency, load transient response, and voltage switching time. Figure 3.9 summarizes the tradeoffs between these three characteristics. Each dot represents a VR design with different parameters: output filter inductor and capacitor sizes, C_{filter}, R_{filter}, and switching frequency. Voltage variation is the percentage change of the output voltage droops during load transients. VR loss includes both switching power and resistive losses associated with the power transistors in addition to all components of resistive loss throughout the power delivery network. Different colors (or shades) of each dot correspond to how quickly the voltage can transition between 0.6V and 1V. The figure shows that different design parameters can shift VR characteristics. VRs with higher switching frequencies are capable of fast voltage scaling (i.e. short scaling times) and exhibit smaller voltage variations, but incur higher VR loss. Conversely, VRs with lower switching frequencies have lower VR loss, but exhibit larger voltage variations and slower voltage scaling capabilities. By understanding these characteristics, designers can exploit the tradeoffs to minimize overheads depending on the specific needs and attributes of the processor load. For example, if the load can leverage fast DVFS for significant power savings (seen for memory-bound applications), a VR that prioritizes minimization of voltage scaling times may yield the best overall system-level solution. On the other hand, if the load is steady with small current transients, design parameters ought to be chosen to minimize VR loss. To better understand how one can make appropriate design

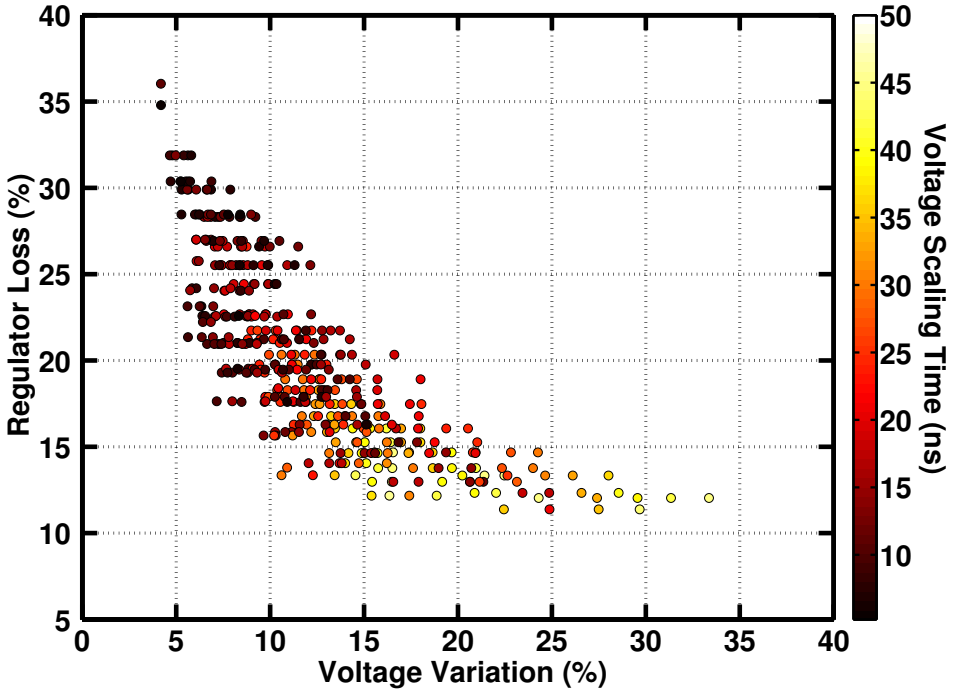

Figure 3.9: Conversion loss, voltage variation, and voltage scaling time of a VR with different parameters.

tradeoffs, the next subsections delve into the VR characteristics in greater detail.

3.3.3 Regulator Efficiency

An ideal VR delivers power from a power source (e.g., battery) to the load without any losses. Unfortunately, the VR itself consumes power while delivering power to a load. Conversion efficiency is an important metric commonly used to evaluate VR performance. It is the ratio of power delivered to the load by the VR to the total power into the VR. VR losses are dominated by switching power and resistive losses, which depend on the size of the switching power transistors, switching frequency, and load conditions (e.g., load current levels). Larger power devices reduce resistive losses at the expense of higher switching power. Higher switching frequencies lead to higher

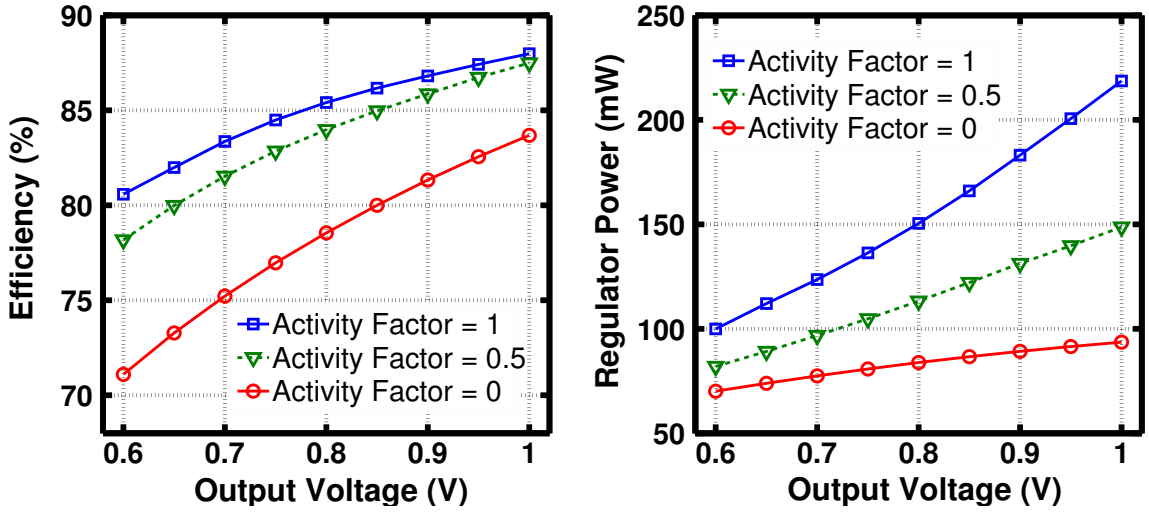

(a) Efficiency (b) Power

Figure 3.10: VR efficiency and power vs. output voltage for different activity factors.

switching power, but can also reduce resistive loss. Hence, it is important to balance these two loss components with respect to different load conditions. Figure 3.10(a) shows that efficiency varies as a function of the output voltage and processor activity, assuming a fixed input voltage. As output voltage scales down, load power scales down with CV^2f and VR power also decreases (Figure 3.10(b)), but not as rapidly. Hence, the efficiency degrades at lower output voltages. Decreasing processor activity also degrades converter efficiency in a similar fashion. Since activity factors differ among benchmarks, VR efficiency changes with benchmarks as well. However, the conversion efficiency metric alone does not appropriately capture the system-level costs and benefits of DVFS. When we later evaluate total system energy consumption and savings, it will be important to combine the on-chip and off-chip VR losses along with DVFS-derived energy savings and overheads. Hence, this chapter presents results in terms of energy (with detailed breakdowns of energy losses) instead of reporting efficiency numbers.

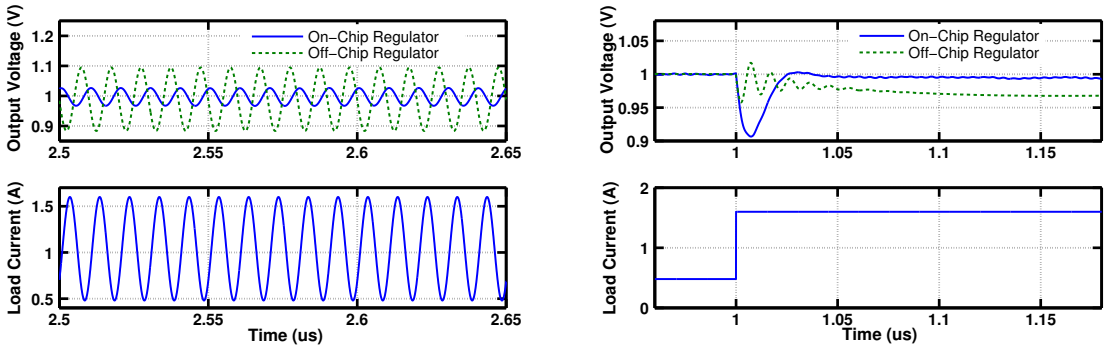

(a) Sine wave load current

(b) Step load current

Figure 3.11: Voltage fluctuation of off-chip and on-chip VRs during step and sine wave load current transient

Although the model treats the off-chip VR as an ideal voltage source, it includes VR power (or loss) based on published efficiency plots found in commercial product datasheets [16]. Based on the peak efficiency values for different output voltages, we calculate the efficiency for our target input and output voltages. Efficiency of the off-chip VR tends to be higher than that of the on-chip VR since they have lower switching frequencies. Recalling Figure 3.8, (a) uses one off-chip VR that converts 3.7V to 1V, and (b) uses an off-chip VR that converts 3.7V to 1.8V and an on-chip VR steps down the 1.8V input to 1V for the processor. Since conversion efficiency varies with output voltage, as shown in Figure 3.10, an off-chip VR can step voltage down from 3.7V to 1.8V with higher efficiency than stepping down to 1V. Besides the losses associated with the VR, we must also consider other losses associated with power delivery. As was observed in Figure 3.8, there are parasitic resistors between the battery and the processor that contribute to loss. Since higher currents flow through this resistive network when delivering power at 1V directly to the processor load from the off-chip VR, I^2R losses are higher. In contrast, using an on-chip VR

that requires a 1.8V input permits lower current flow (∼1/1.8) through the resistive network between the off-chip VR and the chip. This difference in resistive loss is also included when accounting for on-chip and off-chip VR losses.

3.3.4 Load Transient Response

In addition to conversion efficiency, load transient response is another important characteristic that impacts VR performance. Simply put, a VR's load transient response determines how much the voltage fluctuates in response to a change in current. Recalling Figure 3.8, it shows that there are parasitic inductors and resistors along the path between the off-chip VR and the processor. Decoupling capacitors are typically added on the PCB, package, and chip in order to suppress voltage fluctuations. However, these capacitors and inductors can interact to create resonances in the power-delivery network. For a configuration that only relies on the off-chip VR, a mid-frequency frequency resonance occurring in the 100MHz-200MHz range is commonly seen on the chip [35, 79]. Owing to this resonance, load current fluctuations that occur with a frequency near the resonance can lead to large on-chip voltage fluctuations. On the other hand, if the VR is integrated on-chip, most of the parasitic elements fall between the power supply (i.e. battery) and the VR input, as seen in Figure 3.8(b), suppressing this important mid-frequency resonance issue. This can be verified by applying step or sine wave load current patterns and observing how the processor voltage reacts. Figure 3.11 shows that a sinusoidal load current with a frequency at the mid-frequency resonance can cause large on-chip voltage fluctuations due to resonant buildup. In contrast, the on-chip VR does not suffer this resonance

problem and exhibits much smaller voltage fluctuations. Effects of this resonance can also be observed by applying a load current step. The voltage of the off-chip VR rings before settling down, indicative of an under-damped response with resonance. In contrast, the output voltage of the on-chip VR does not ring, but rather reveals a critically-damped system. However, the output voltage of the on-chip VR suffers a different problem. It droops much more in response to the load current step than its off-chip VR counterpart. This is because the on-chip VR relies on the on-chip capacitor for both decoupling and to act as the output filter capacitor. Since this on-chip capacitor is much smaller than the total decoupling and filter capacitance used for off-chip VRs, large load current steps can rapidly drain out the limited charge stored on the capacitor before the VR loop can respond, resulting in a large voltage droop. These plots suggest that the worst-case current trace for the off-chip VR is a sine wave at the resonance frequency, whereas a step change is the worst-case load transient for the on-chip VR.

In order to make a fair comparison between the on-chip and off-chip VRs, two important factors that affect load transient response are kept constant. The total on-chip decoupling capacitance is 40nF (10nF per core) and voltage margin is set to ±10%. The 40nF decoupling capacitance is set such that with the conventional off-chip VR scenario, voltage fluctuations stay within the ±10% voltage margin under worst-case load conditions. This decoupling capacitance value also matches well with the Intel 80200 Processor based on the Xscale Architecture [30]. The 10% voltage margin is also a widely-used value in microprocessors [105, 29]. Unfortunately, the 40nF of on-chip decoupling cannot always guarantee voltage fluctuations stay within

the ±10% margin for on-chip VRs across all load transient conditions.

In order to prevent voltage emergencies, where the on-chip VR's output voltage swings beyond ±10% due to sudden load current steps, we employ a simple architecture-driven mechanism that selectively disables clock gating. Since large load transients can largely be attributed to aggressive clock gating events, disabling some of the gating can reduce the magnitude of load current steps. Figure 3.12 shows voltage traces corresponding to load current transients for two clock gating scenarios. A sudden current increase that occurs after a long stall period causes a voltage emergency and large current steps following the first step also cause subsequent voltage emergencies. By appropriately disabling some of the clock gating (solid line), current transient magnitudes are reduced and the voltage droops can be suppressed to stay within the 10% margin. Since clock gating is used to reduce power consumption, disabling it leads to power overhead that must be accounted for. Hence, this technique is sparingly applied only when there are large current transients due to large fluctuations in processor activity.

3.3.5 Voltage Scaling Time

Voltage scaling time is another important characteristic that affects systems with DVFS. When the VR voltage scales to a new voltage level, it cannot scale immediately, but scales gradually. Figure 3.13 shows voltage, frequency, and current traces for an on-chip VR that drives a single processor core running *fft*. The frequency changes abruptly whereas the voltage scales across tens of nanoseconds. To ensure sufficient timing margins for the processor core, low-to-high frequency transitions are

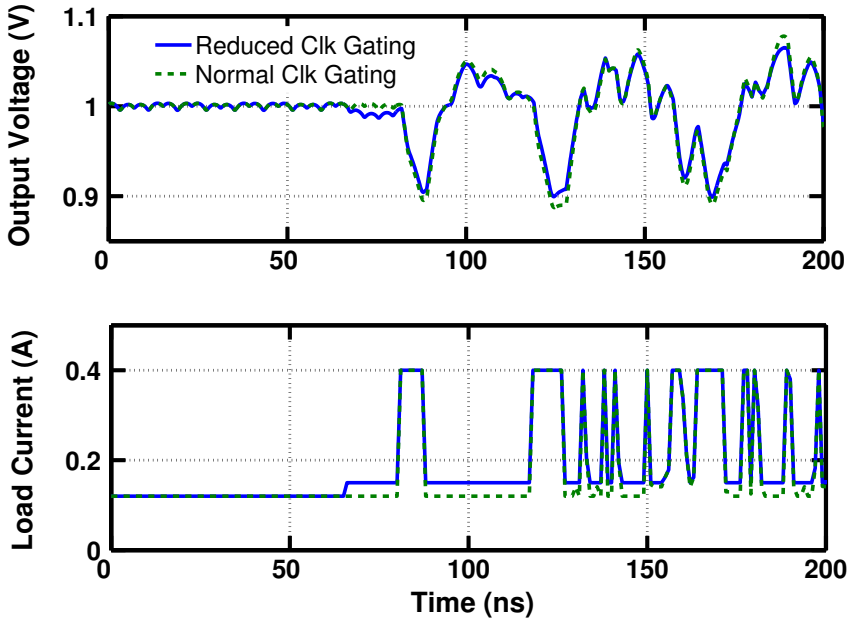

Figure 3.12: Example of reducing voltage fluctuations by selectively disabling clock gating.

allowed after the voltage settles to the higher level. Similarly, high-to-low frequency transitions precede voltage changes. This difference between frequency and voltage transition times leads to energy overhead. We account for this wasted energy as DVFS overhead. Higher switching frequencies and/or smaller output filter component sizes can enable faster voltage scaling to reduce this DVFS overhead, but they introduce penalties of higher VR loss and/or more sensitivity to load current transients.

3.3.6 On-Chip Regulators for Single and Multiple Power Domains

Given their small size compared to off-chip VRs, several on-chip VRs can be integrated on-chip to deliver power to multiple voltage domains. However, there is

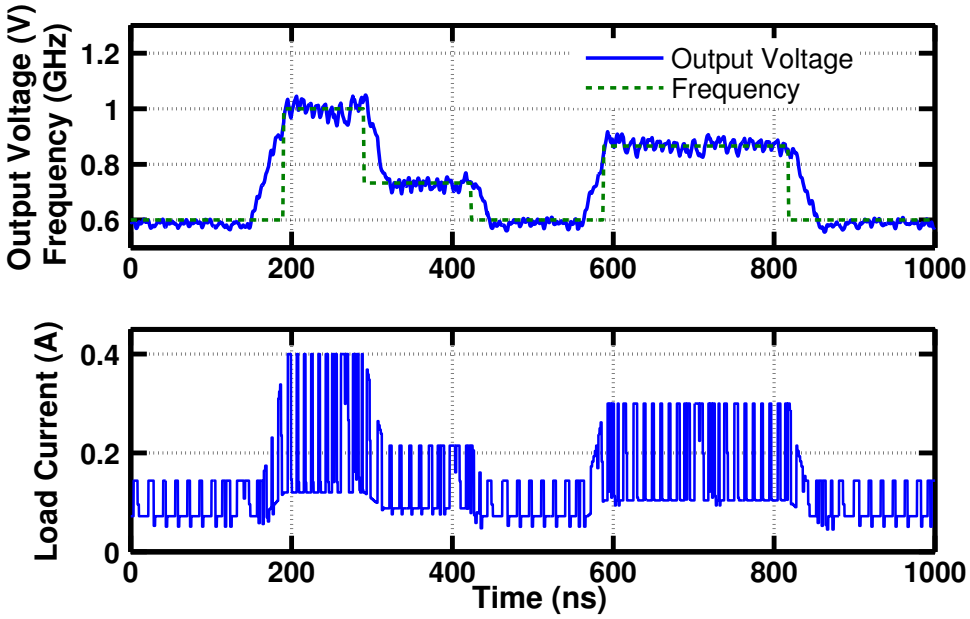

Figure 3.13: Snapshot of output voltage, frequency, and load current traces with DVFS.

a tradeoff between using one voltage domain versus multiple voltages domains. For fair comparison, we assume that the total number of phases for the multiphase on-chip VR we use is constant for single and multiple voltage domain configurations, matching the area overhead. In other words, an 8-phase VR is used to power a single voltage domain, while four 2-phase VRs deliver power to four different voltage domains. Again, we assume that each core has 10nF of on-chip capacitance for each of the 2-phase VRs in the multiple voltage domain scenario and a total capacitance of 40nF for a single 8-phase on-chip VR for the single voltage domain case.

There are several differences related to implementing single versus multiple power domains using on-chip VRs in a 4-core CMP. With four voltage domains, each VR is only sensitive to current transients in its respective core. For a single power domain, the VR sees current transients from all four cores, but also benefits from the larger

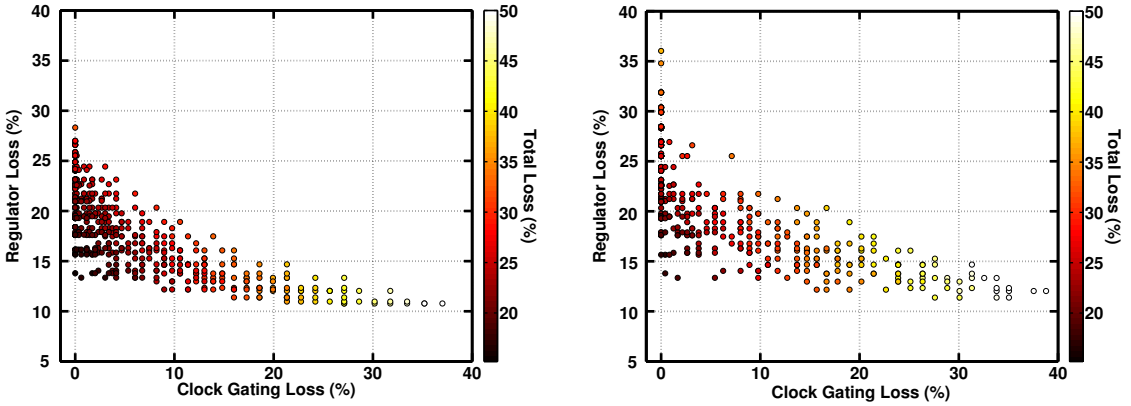

(a) 1 regulator

(b) 4 regulators

Figure 3.14: Total energy overhead with different regulator settings for *facerec*

on-chip capacitance. For a multi-threaded version of *facerec* running on a 4-core CMP, maximum current steps (between idle and full activity) occur over 125K times within 1M cycles for each core. In contrast, with a single voltage domain, the maximum current step (between all four cores idles and all four cores fully active) occurs much less frequently, only 350 times out of 1M cycles. These differences affect the appropriate tradeoffs a designer must make to minimize overheads and maximize energy savings. Given the higher potential for voltage emergencies with multiple power domains, the previously-described technique that disables clock gating may trigger frequently and incur high power penalties. Higher switching frequencies may improve load transient response to reduce overheads in spite of higher switching losses.

Given the tradeoffs between conversion loss, load transient response, and voltage scaling time, we can choose different VR design parameters for both single and multiple voltage domains that minimize energy overhead. Figure 3.14 presents the conversion loss, DVFS overhead, and power overhead of disabling clock gating (labeled Clock Gating Loss) across different VR design parameters for a 4-core CMP

running *facerec* with one and four VRs. These plots are similar to Figure 3.9, but Figure 3.14 combines all of the losses into a total energy overhead represented by different colors for each dot. For both single and four voltage domains, configurations corresponding to dots in the bottom left corner offer the design point with the smallest total energy overheads and losses. Dots extending to the lower right have small conversion loss, but the low switching frequency leads to higher power overhead related to frequently disabling clock gating to limit current swings. Dots in the upper left corner suffer excessive conversion loss. Figure 3.14 also shows that the total loss for the single power domain tends to be smaller than that for four power domains. This can be attributed to the fact that the four power domains have to handle many more worst-case current steps as compared to the single-domain case, in which much of the current hash cancels out. Based on this analysis, the VR design (or dot) that minimizes overhead is chosen for the single and four power domain scenarios. Details of these configuration are list Table 3.3, showing a single power domain scenario has around 2% smaller overhead than implementing four power domains. Similar trends are observed for other benchmarks and so we use the VR design configurations based on the analysis above in subsequent sections of the chapter.

3.4 Energy Savings for Per-Core and Chip-Wide DVFS using On-Chip Regulators

In previous sections, the major benefits (additional DVFS energy-saving opportunities) and overheads (DVFS overheads and VR losses) of on-chip VRs were discussed

	Single Power Domain	Four Power Domains
# of phases for on-chip regulator	8	2 per domain
On-chip regulator switching frequency (MHz)	100	125
Inductance per phase (nH)	13	9.6
Voltage scaling speed (mV/ns)	30	50
Total Energy Overhead (%)	15.49	17.32
Decoupling capacitance (nF)	40	10 per domain
Voltage margin (%)	± 10	

Table 3.3: Characteristics of the on-chip VR (all percentage (%) numbers are relative to the processor energy with DVFS).

in isolation. In this section, we return to Figure 3.1 and evaluate the overall benefits of on-chip VRs compared to traditional, off-chip VRs when considering all of these combined effects. We also extend our analysis to larger numbers of power domains (and on-chip VRs) to understand scalability constraints.

3.4.1 Comparison of Energy Savings

Figure 3.15 provides detailed breakdowns of the DVFS energy savings and the various overheads incurred within a 5% DVFS performance loss constraint. This analysis has been performed for four configurations: an off-chip VR with no DVFS, an off-chip VR with DVFS, an on-chip VR with a single-power domain (global or per-chip DVFS), and an on-chip VR with four power domains (local or per-core DVFS). In this figure, processor energy consumption with no DVFS is set to 100 and the other values are presented relative to this value. The reduced processor

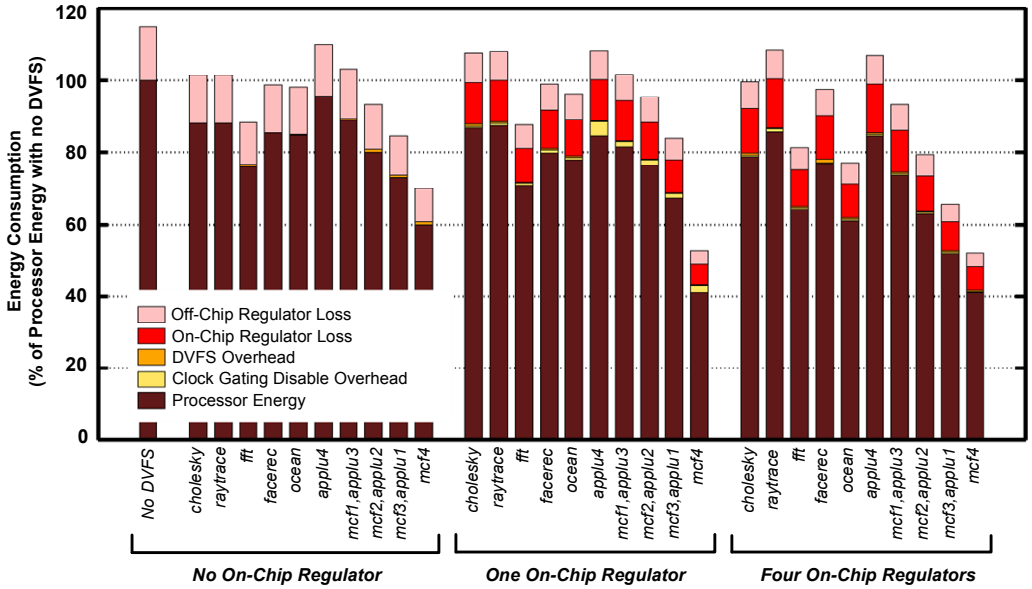

Figure 3.15: Detailed breakdown of energy consumption for the processor and VR for single power domain (global) and multiple domains (per-core) DVFS.

energy results achieved with DVFS represent the best selection of DVFS parameters for each configuration that maximize DVFS-energy savings while minimizing DVFS overheads: the on-chip VR has a 100ns DVFS interval and the off-chip VR has a 100 μs interval. To evaluate the energy savings offered by using on-chip VRs, Figure 3.16 presents a bar graph showing energy savings compared to the off-chip DVFS case for different benchmarks. For each benchmark, the bar on the right corresponds to how much energy savings is possible with fast DVFS, ignoring overheads. The bar on the left presents the relative savings with all of the overheads included. The gap between the left and right bars corresponds to the sum of overheads introduced by using on-chip VRs. Higher bars indicate larger relative energy savings.

These two figures represent several interesting trends in the design space which we discuss in detail below.

Off-chip DVFS vs On-Chip, Single Power Domain: We first compare on-

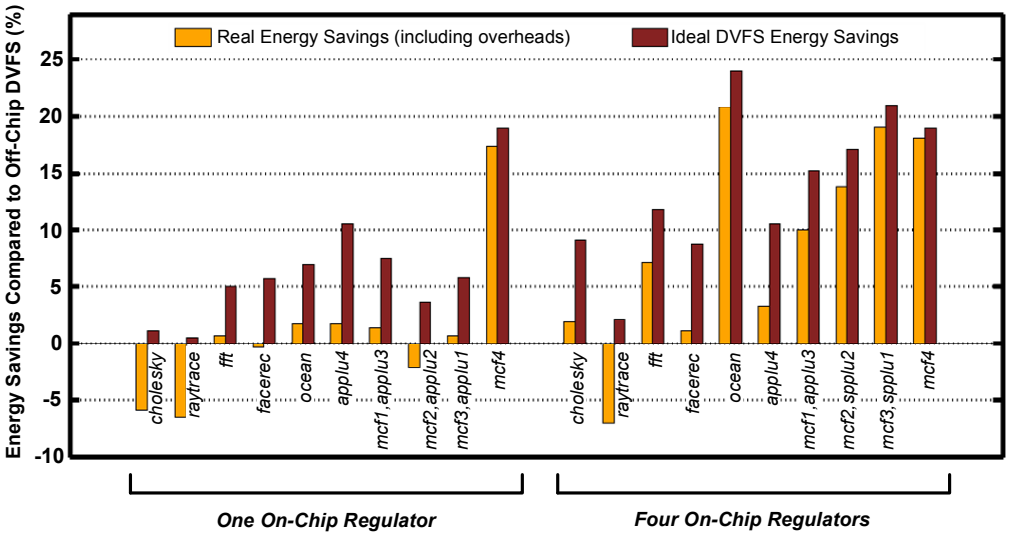

Figure 3.16: Relative energy consumption of on-chip VR configurations compared to a off-chip VR with DVFS.

chip VRs with global DVFS to the off-chip VR. At a high-level, we see that only *mcf4* achieves significant positive energy savings when compared to the off-chip VR with DVFS. The reduction in processor energy, provided by fast DVFS, has the added benefit of reducing conversion losses. Seven of the ten benchmarks are approximately break-even (within ±2%) between the two configurations, which means that the faster DVFS scaling can just offset the additional losses introduced by using an on-chip VR. *Raytrace* and *cholesky* with few opportunities for DVFS, yet still suffering the impact of on-chip VR loss, suffer significant energy overheads. One reason that off-chip DVFS performs well is that the the coarser DVFS intervals lead to less DVFS overhead compared to the on-chip VR which may switch voltage/frequency settings more frequently.

Off-chip DVFS vs. On-Chip, Four Power Domains: The next comparison that we perform investigates the benefits of per-core DVFS scaling (on top of the fast voltage transition times) compared to the off-chip configuration which only provides

a single voltage domain. This comparison provides very encouraging results for the on-chip VR design: all of the benchmarks except *raytrace* achieve energy savings, and several by significant amounts with *ocean* achieving 21% savings. The multiple power domain configuration allows even more savings through DVFS than the single domain, but needs more VR power to deal with the additional load current hash that each core introduces. When we compare the two cases that both use on-chip VRs, Figure 3.15 shows that on-chip VR loss is consistently higher by a small amount in the four domain case, but this is clearly overshadowed by the additional DVFS energy savings. There is another interesting effect that can be observed. Since VR losses scale with load power, the gap between adjacent bars that correspond to total overheads reduces for several benchmarks, in Figure 3.16, since more energy savings is possible with fast, per-core DVFS. Thus, applications that significantly benefit from DVFS to reduce processor energy can also benefit from the synergistic reduction of VR overheads.

From this analysis, we can form several conclusions regarding the impact of on-chip VRs on system design.

- Systems architects who plan to utilize on-chip voltage regulation must carefully account for energy-efficiency costs when calculating projected benefits. This requires a detailed understanding of many of the costs and overheads that on-chip VRs incur.

- DVFS scaling algorithms must adapt to take advantage of the fast, fine-grained nature of on-chip VRs. Future DVFS scaling algorithms will likely require significant microarchitectural control, rather than traditional OS-based control,

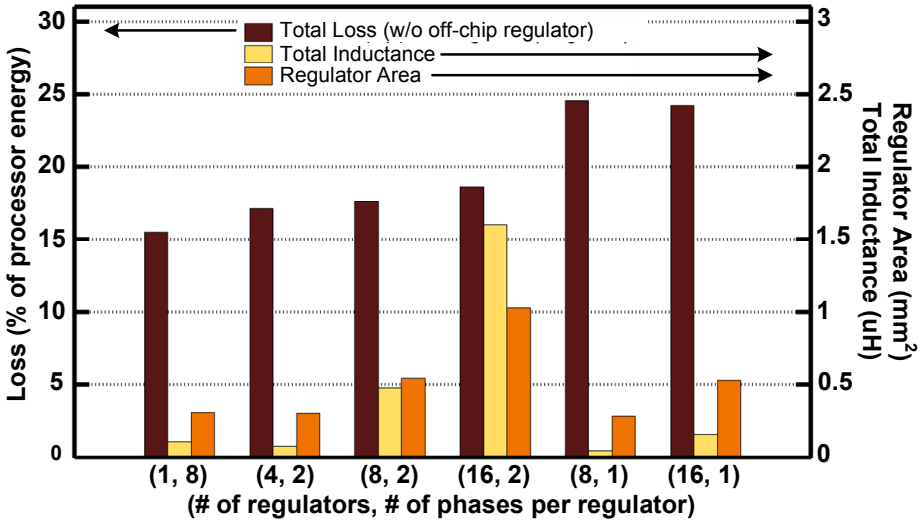

Figure 3.17: Loss, inductor size, and area of on-chip VRs for different numbers of power domains.

and must carefully take into the DVFS scaling overheads.

- On-chip VRs provide significant benefits to designers of CMP systems and we expect that future systems will be developed to capture this potential. The power scalability of on-chip VRs is a key future research question to extend this analysis to high-performance CMP systems with four to eight cores.

3.4.2 Power Domain Scalability

The previous analysis shows that multiple power domains using DVFS with finer granularity allow large energy savings. However, there is a limit to the number of on-chip power domains that can be implemented due to various overheads. This subsection compares different overheads related to implementing 1, 4, 8, and 16 power domains, equal to the total number of VRs since one VR is used per power domain. Figure 3.17 shows simulation results for *facerec* with the energy loss, area overhead,

and total inductance of on-chip VRs assuming these power domain scenarios in a 4-core CMP. With a total maximum power of 1.6W, 1, 4, 8, and 16 power domains consume 1.6W, 0.4W, 0.2W, and 0.1W per domain, respectively. The total loss corresponds to the sum of on-chip VR loss, DVFS overhead, and power overhead from the architectural mechanism that disables clock gating to limit current swings, as a percentage of the processor energy. The chart also shows the total sum of inductance, indicating the number of inductors mounted onto the package scales up rapidly. The two main components that occupy significant on-die area are the power transistors and feedback circuits. Power transistor sizes are obtained using Simulink/Matlab simulations, and the values from a recently built on-chip VR [40] are used for the feedback circuits including the hysteretic comparator, C_{filter}, and R_{filter}. This does not include the area consumed by on-chip decoupling capacitors. The total decoupling capacitance is again fixed to 40nF, which means more power domains get smaller, equally divided units of decoupling capacitance per domain. For each scenario, the VR design is optimized to minimize energy overheads using design parameter sweeps similar to those shown in Figure 3.14.

The results in Figure 3.17 again suggest basic tradeoffs between the number of power domains and associated overheads. The first four sets of bars show that loss only increases slightly with the number of power domains. There is roughly a 3% difference between the loss for 1 domain and 16 domains. However, more power domains occupy significantly larger area, both on the package and on the die. The main reason for this is the increasing number of VR phases. Since power transistor size scales with load current, power transistor area remains relatively constant. However,

the area occupied by the feedback circuit grows proportionally with the number of phases used in the VRs. The area corresponding to 1 and 4 domains are the same, because the total number of phases used in the VRs are fixed to 8 for fair comparison as shown previously in Table 3.3. For 8 and 16 domains with 2-phase VR designs, the area-increases are two- and four-fold over the 4 domain case, respectively. In addition to increases in on-die area, the total inductance increases rapidly because the number of inductors increase with more phases. Moreover, the inductance per phase increases in order to minimize energy loss associated with lower load currents. This increase in total inductance leads to higher costs and packaging complexity to mount all of the inductors. One can offset these increasing costs for 8 and 16 domains by implementing single-phase VRs at the expense of incurring more loss.

Systems that seek to use a large number of power domains with a multitude of on-chip VRs to implement DVFS with finer spatial granularity must carefully consider all of the related losses, overheads, and costs. The ideal benefits of very fine-grained DVFS may be lost or difficult to justify.

Now that we studied the system-level energy savings of SoCs using IVRs, the next chapter compares three different types of IVR topologies and presents implementation and chip measurement results of a 3-level VR.

Chapter 4

Fully-Integrated 3-Level Voltage Regulators

Contents

4.1 Buck, Switched-Capacitor and 3-Level IVRs 62

4.2 3-Level Voltage Converter 65

4.2.1 Design Parameters for 3-Level Converters 65

4.2.2 Comparison to Buck and SC Converters 69

4.3 3-Level Implementation: Open-Loop 74

4.3.1 Power FETs 76

4.3.2 Driver circuits 78

4.3.3 Passive elements 80

4.3.4 Feedback loop and shunt regulator 82

4.4 Measurement: Open-Loop 83

4.5 3-Level VR: Closed-Loop 94

4.1 Buck, Switched-Capacitor and 3-Level IVRs

IVR designs range from buck VRs to switched-capacitor VRs to low-dropout linear regulators. Linear regulators have a maximum efficiency limit given by the ratio of output voltage to input voltage; they suffer from low efficiency at high ratios. In contrast, switching VRs can maintain high efficiency across a wide range of output voltages. There are two types of switching VRs commonly used for low step-down ratios - buck and switched-capacitor (SC) VRs. Shown in Figure 4.1(a), the buck VR relies on an inductor to generate a step-down voltage on the output capacitor, C_{OUT}. The buck VR creates a square-wave voltage – of varying duty cycles (D) – at the output of the power FETs (V_X). While traditional buck VRs rely on single pull-up and pull-down power FETs, series stacks of switches enable use of thin-oxide devices in integrated voltage VRs [89]. By adjusting the duty cycle of V_X, buck VRs can provide a wide range of V_{OUT}. However, the buck VR requires a large, high-quality inductor, which is difficult to integrate on-chip.

In contrast, the SC VR uses flying capacitors (C_{FLY}), without an inductor, to nominally divide the high input voltage (V_{IN}) by pre-determined integer ratios. For example, the SC VR in Figure 4.1(b) divides V_{IN} by two as it iterates between two phases of capacitor configurations – series-stack and parallel. Although it does not need inductors, this particular configuration of SC can only step V_{OUT} down to values lower than $V_{IN}/2$. Additional step-down ratios, such as 1/3 and 2/3, are also possible as demonstrated by Ramadass, et al. [82] and Le, et al. [59], in order to extend the range of output voltage conversion. However, the added power switches needed for the additional capacitor configurations can exacerbate conversion loss.

As shown in Figure 4.1(c), a 3-level VR merges characteristics of both inductor-based buck and SC VRs to gain the benefits of both [119, 102]. Similar to the buck, the output LC pair of the 3-level VR filters V_X to generate V_{OUT} with small ripple. While the V_X of the buck VR swings between 0 and V_{IN}, the V_X of the 3-level VR either swings between 0 and $V_{IN}/2$, or $V_{IN}/2$ and V_{IN}, to convert V_{OUT} to voltages under and over $V_{IN}/2$, respectively. The switching action of the power FETs, combined with the flying capacitor, effectively generates a third voltage, $V_{IN}/2$ (hence the name 3-level VR), and adjusts D to set V_{OUT} across a wide range of voltage levels. Notice that V_X of the 3-level VR swings with half the amplitude and at twice the frequency compared to that of the buck. Both of these attributes enable the 3-level VR to exhibit smaller inductor current ripple and voltage ripple on V_{OUT} or to use a smaller inductor for the same ripple target.

Although the three VRs look similar in schematic, the loss mechansims are different, which leads to interesting design decisions. Buck and 3-level VRs have an inductor that forces current to always flow through the power switches and the inductor. As a result, a large part of the conversion loss comes from I^2R losses on the switches and inductors. In contrast, conversion loss on SC VRs come from charge redistribution loss between C_{FLY} and C_{OUT} and this loss is not dependent on on-state switch resistance to a certain extent [50, 92]. A simple example to explain this loss mechanism is when two capacitors (capacitance C) with 0V and 2V are connected through a switch. After the charge transfer is complete, both capacitors have 1V and have lost energy equal to C compared to the initial state. As long as the charge transer completes, the on-state resistance of the switch does not affect how much

energy is wasted in the redistribution process. Similarly, conversion loss of SC VR is not affected by the power switches' on-state resistance, as long as the switching period is long enough for the SC to complete the charge redistribution process. This indicates that SC VR could be more suitable for older process nodes with higher switch resistance than modern process nodes.

Prior 3-level VR designs include an off-chip VR for envelop tracking [119] and an integrated 3-level VR with 27nH bondwire inductors [102]. We build upon these works and present a fully-integrated 3-level VR with 1nH on-chip spiral inductors. 1nH inductors, placed on top of flying capacitors to minimize area overhead, enable voltage transition across 1V within 20ns, which is 100 times faster than previously published data [90]. The VR can be externally programmed to adjust design parameters (switching frequency, number of phases and power FET size) to study trade-offs associated with different design parameters. We also add fast shunt regulation to the VR to reduce voltage noise.

The next section studies how design parameters affect conversion loss in 3-level VRs and compares the conversion efficiencies of 3-level to those of buck and SC VRs. Then Section 4.3.4 presents a detailed, circuit-level description of the 3-level VR design that was implemented in a test-chip prototype using a 130nm CMOS process technology. Experimental results from the test-chip, in Section 4.4, demonstrate fast voltage scaling and high conversion efficiency across a wide range of output voltages. In Section 4.5, I present the design and measurement results of a second version 3-level regulator test-chip that fixes a couple of drawbacks — limited duty-cycle resolution and inefficient shunt regulation — of the first version test-chip.

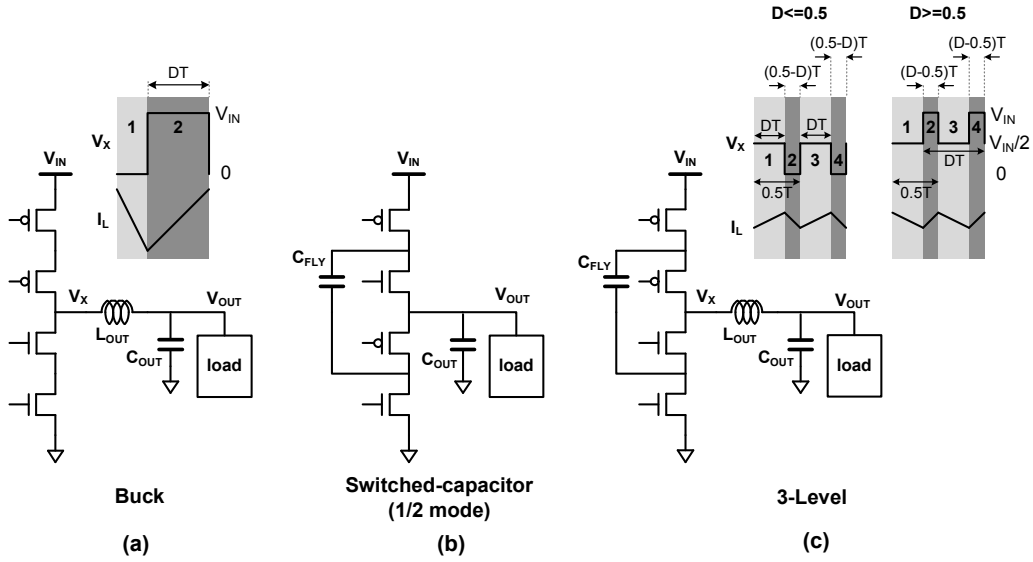

Figure 4.1: Power FET and output filters of (a) buck, (b) switched-capacitor, and (c) 3-level VRs

4.2 3-Level Voltage Converter

There are multiple sources of conversion loss in the 3-level VR. Understanding how VR design parameters affect different sources of losses is important for achieving maximum efficiency. We first study the different design parameters of the 3-level VR and then compare its efficiency to those of buck and SC VRs.

4.2.1 Design Parameters for 3-Level Converters

Three design parameters of the 3-level VR significantly affect conversion loss - switching frequency, number of phases and power FET size. For maximum efficiency, the choice of design parameters should take output voltages and load currents into account.

Figure 4.2 presents simulated conversion efficiencies of a 3-level VR running in

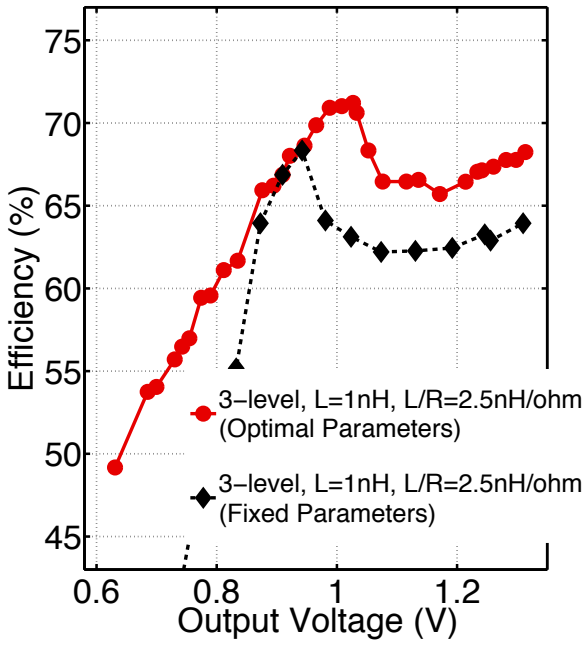

input voltage	2.4V
switching frequency	10-500MHz
number of phases	1,2,4 (buck, 3-level) 8 (SC)
Total power FET width	8-256mm
output voltage (V_{OUT})	0.6 - 1.35
load current (A)	0.55* V_{OUT}^2

Figure 4.2: Simulated conversion efficiencies of 3-level VRs with fixed and optimal design parameters. Table shows the range of design parameters used in simulations.

continuous conduction mode (CCM) acquired using a fast circuit simulator HSIM, set to the highest simulation accuracy level. As specified in the table in Figure 4.2, the VR operates with DC load current ranging from 0.2A to 1A for output voltages ranging from 0.6 to 1.35V. Load current scales quadratically with output voltage to mimic a processor operating with DVFS. Simulations sweep design parameters to find the maximum efficiency for each output voltage value. The VR uses 1nH inductors with 400mΩ series resistance. Up to 4 copies of power FETs and inductors can be interleaved to form multi-phase VRs [42] to distribute current flow and reduce output voltage ripple. Figure 4.12 presents an example of a 4-phase VR that can dynamically change the number of operating phases according to load levels. Figure 4.2 shows that optimizing design parameters significantly improves conversion efficiency compared

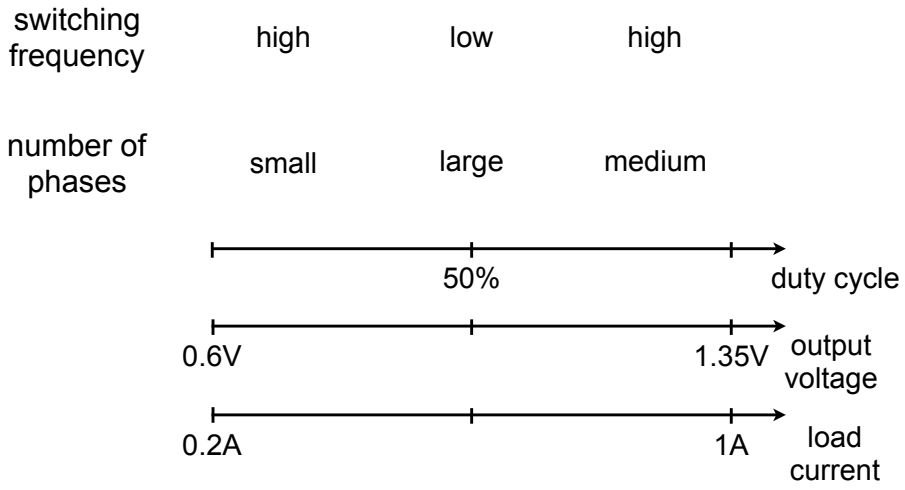

Figure 4.3: Design parameters that maximize efficiencies across duty cycle, output voltage and load current ranges.

to a VR using fixed parameters (100MHz frequency, 2 phases, 48mm total power FET width).

Figure 4.3 shows how to determine switching frequency and number of phases to maximize efficiency. When duty cycle is in the vicinity of 50%, a VR needs to operate at low switching frequency with maximum number of phases. As duty cycle deviates from 50%, the VR needs to increase switching frequency and reduce the number of phases. The selection of design parameters aim to balance different sources of losses. High switching frequencies increase switching loss (CV^2f), but reduce resistive loss ($I^2_{RMS}R$) caused by inductor current ripple ($\Delta I_{L,PP}$). Assuming a VR operating under CCM with a triangular wave for the inductor current (IL), Equation 4.1 shows that both DC value and peak-to-peak ripple of the inductor current contribute to $I^2_{RMS}R$ loss.

$$\Delta I^2_{L,RMS} = I^2_{L,DC} + (\Delta I^2_{L,PP})/12 \tag{4.1}$$

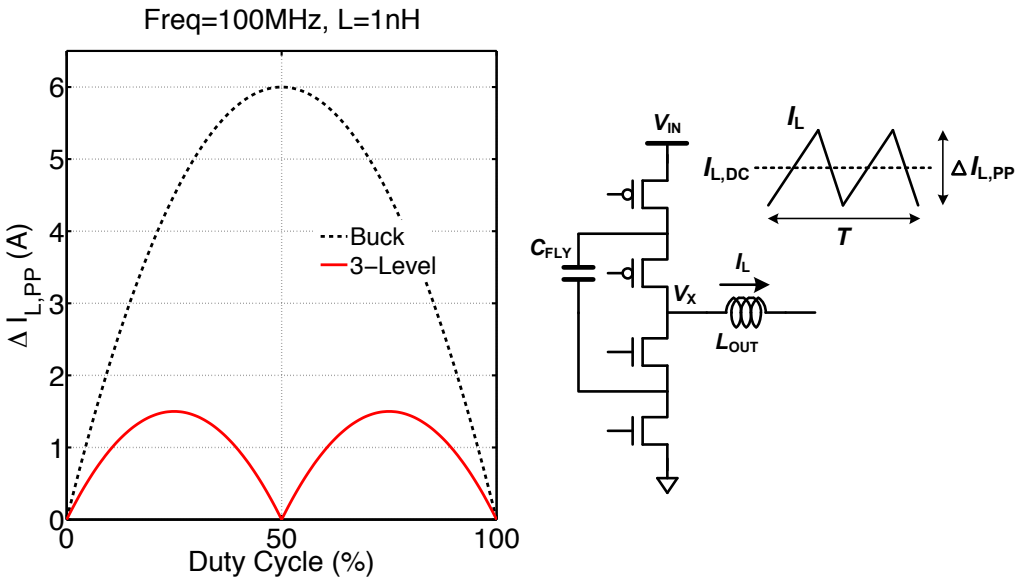

Figure 4.4: Simulated peak-to-peak inductor current ripple ($\Delta I_{L,PP}$) of 3-level and buck VRs in continuous conduction mode (CCM).

Shown in Figure 4.4, $\Delta I_{L,PP}$ of the 3-level VR reaches minimum at 50% duty cycle, increases as duty cycle deviates from 50% and decreases again when duty cycle goes below 25% or over 75%. Taking advantage of small $\Delta I_{L,PP}$ at duty cycles near 50%, the 3-level VR minimizes switching loss by running at low frequencies. As $\Delta I_{L,PP}$ grows at duty cycles away from 50%, the VR runs at higher frequencies to suppress $I_{RMS}^2 R$ loss, albeit with larger switching loss. Increasing switching frequency at light loads contradicts the conventional wisdom of using pulse frequency modulation (PFM) in buck VRs to reduce frequency at light loads. As duty cycle deviates from 50%, $\Delta I_{L,PP}$ of the 3-level increases while that of the buck VR decreases. This allows the buck to reduce frequency at light loads, while forcing the 3-level VR to increase frequency.

To study how the number of phases affects conversion loss, Equation 4.2 expands Equation 4.1 to a multi-phase 3-level VR, which consists of multiple interleaved copies

of a single phase VR.

$$\begin{aligned} \Delta I_{L,RMS}^2 &= (I_{L,DC}^2 + \Delta I_{L,PP}^2/12) \times N_{PH} \\ &= I_{LOAD}^2/N_{PH} + \Delta I_{L,PP}^2/12 \times N_{PH} \end{aligned} \tag{4.2}$$

$$\begin{aligned} N_{PH} &: \textit{number of phases} \\ I_{L,DC} &: \textit{DC inductor current per phase} \\ \Delta I_{L,PP} &: \textit{inductor current ripple per phase} \end{aligned}$$

Equation 4.2 shows that using larger number of phases reduces loss due to DC current, while increasing loss caused by $\Delta I_{L,PP}$. At light loads, the VR uses a single phase because $\Delta I_{L,PP}$ is a larger source of loss compared to DC current. Near 50% duty cycle, the VR uses all 4 phases since $\Delta I_{L,PP}$ is small. At high load currents, the VR uses 2 out of 4 phases to balance the losses due to $\Delta I_{L,PP}$ and DC current, contradicting conventional wisdom that increases the number of phases at full loads to minimize loss due to DC current. Again, the difference is due to increasing $\Delta I_{L,PP}$ as duty cycle deviates from 50% at full loads. Moreover, reducing the number of phases allows a portion of C_{FLY} to stay idle, resulting in smaller loss due to bottom-plate parasitic capacitance.

4.2.2 Comparison to Buck and SC Converters

Remaining simulation plots (Figures 4.5, 4.6, 4.7, 4.8) present efficiencies with optimized design parameters using ranges specified in Figure 4.2. Figure 4.5 presents a similar efficiency versus output voltage plot of the buck VR for different inductance

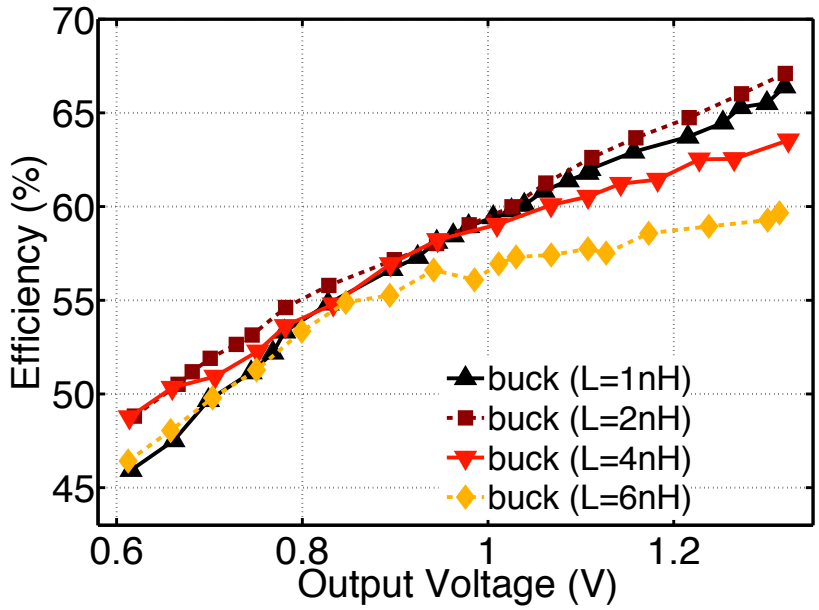

Figure 4.5: Simulated conversion efficiencies of buck VRs across inductance values (L/R = 2.5nH/Ω).

values, assuming CCM operation across all load conditions. Simulations use a buck VR design similar to one proposed in [89] with series stacks of power FETs using thin-oxide devices. For the same inductor quality (L/R = 2.5nH/Ω), larger inductance reduces $\Delta I_{L,PP}$ while increasing inductor series resistance (R_L). At low load currents, 2nH and 4nH inductors achieve higher efficiencies than 1nH and 6nH, which suffer from large $\Delta I_{L,PP}$ and R_L, respectively. At high load currents, R_L significantly affects conversion loss, allowing 1nH and 2nH to achieve higher efficiencies than 4nH and 6nH. We choose 2nH for further comparisons to 3-level VRs.

Since a 3-level VR adds flying capacitors on-die, it occupies larger die area than a buck using the same inductor. Assuming the buck VR can use additional die area to implement larger, higher quality inductors, Figure 4.6 compares conversion efficiencies of 3-level and buck VRs, providing similar or higher quality inductors to buck VRs.

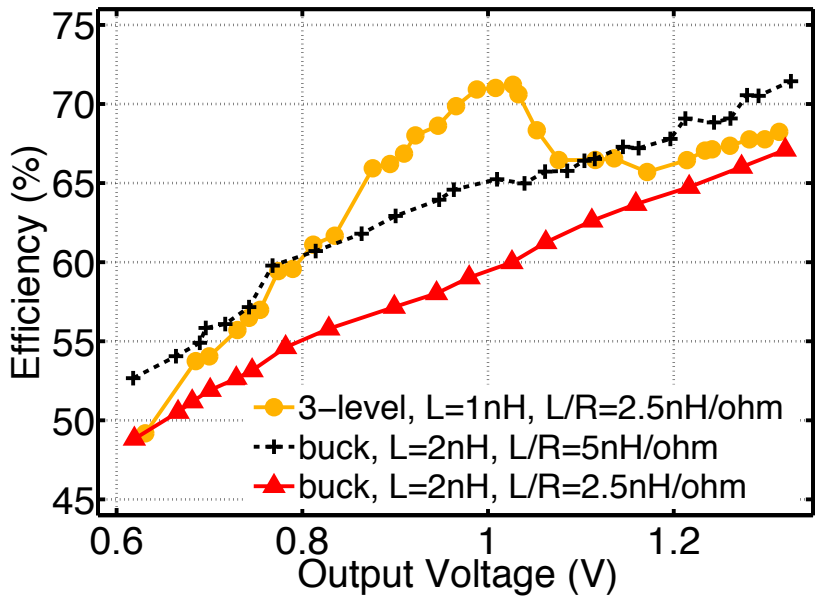

Figure 4.6: Simulated conversion efficiencies of 3-level and buck VRs across inductor qualities.

The 3-level VR uses 16nF of C_{FLY}, and both buck and 3-level VRs use 10nF of C_{OUT}, operating with up to 4 phases. To make a fair comparison between VRs with different V_{OUT} ripple characteristics, as proposed in [59], we calculate conversion efficiency using the minimum value of V_{OUT} ripple, instead of the average V_{OUT} value. For the same inductor quality (L/R = 2.5nH/Ω), the 3-level VR exhibits higher efficiency than the buck VR. Both VRs suffer from degrading efficiencies at low voltages, but the slope of 3-level is steeper than that of the buck. This is because $\Delta I_{L,PP}$ of the 3-level increases as duty cycle deviates from 50%, while that of the buck decreases. Using a higher quality inductor (L/R = 5nH/Ω) allows the buck to achieve higher efficiencies than the 3-level VR at low and high loads.

Figure 4.7 compares the conversion efficiency of the 3-level VR to a reconfigurable SC VR that can switch between three modes – 1/3, 1/2 and 2/3. Simulations use a

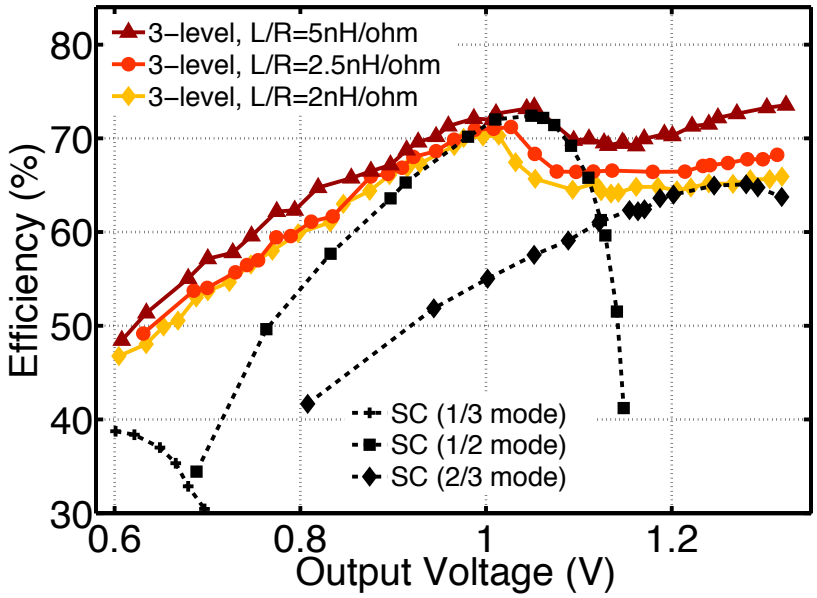

Figure 4.7: Simulated conversion efficiencies of 3-level and switched-capacitor VRs across inductor qualities.

SC VR design similar to one in [59] with series stacks of thin-oxide devices to support high input voltage. While the 3-level VR has 16nF of C_{FLY} and 10nF of C_{OUT}, SC can use C_{FLY} as an output decoupling capacitor, obviating additional C_{OUT}. Assuming the same die area for the two VRs, the SC VR can use 26nF of C_{FLY} without any C_{OUT}. For the 3-level VR, we assume that C_{FLY} is MOS capacitors placed underneath the inductor to avoid additional area overhead (as explained later in Figure 4.13). Since 16nF of MOS capacitance occupies 1.6mm^2 in UMC 130nm technology, four 0.4x0.4mm inductors occupying 0.64mm^2 can fit on top of C_{FLY}. In contrast to the 3-level and buck VRs, the SC VR does not need a thick metal layer for high quality inductors. For fair comparison, we present conversion efficiencies across inductor qualities that represent different metal characteristics. Again, efficiency is calculated using minimum V_{OUT} instead of average V_{OUT}.

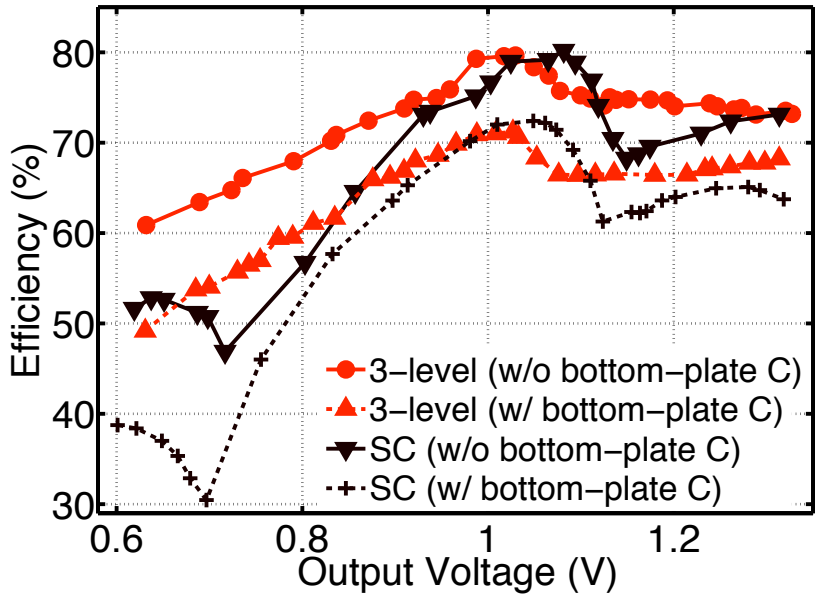

Figure 4.8: Simulated conversion efficiencies of 3-level and SC VRs with and without bottom-plate parasitic capacitance.

Assuming an inductor built with two metal layers in parallel using the digital logic process in UMC 130nm (L/R=2nH/Ω), the SC VR in 1/2 mode achieves higher efficiency than the 3-level VR at the center where duty cycles are in the vicinity of 50%, while the 3-level VR exhibits higher efficiencies at light loads than the SC VR in 1/3 mode. The trend is similar assuming an inductor built with two metal layers (one 2μm thick layer) using the RF process in UMC 130nm (L/R=2.5nH/Ω). The 3-level has the potential for even higher efficiencies when assuming an ultra-thick metal available in modern process technologies that enables an even higher quality inductor (L/R=5nH/Ω), albeit with higher cost. Although the inductor adds series resistance, the 3-level VR has the following benefits when operating at 50% duty cycle. First, the inductor allows the 3-level to have a lower per-phase peak current than the SC VR, reducing resistive loss [78]. Second, the inductor reduces loss caused by charge

redistribution in the 3-level VR. As mentioned in Section 4.1, whenever capacitors switch between series stack and parallel configurations in the SC VR, the resulting charge redistribution between C_{FLY} and C_{OUT} increases conversion loss [50, 92]. In contrast, the inductor in the 3-level VR sits between C_{FLY} and C_{OUT} to store a portion of the charge otherwise lost to charge redistribution.

The next section provides an in-depth explanation of 3-level VR operation and circuit details found in a multi-sector, multi-phase regulator test-chip prototype, which we evaluate in 4.4.

4.3 3-Level Implementation: Open-Loop

Figure 4.9 presents an overall block diagram comprising a set of thin-oxide transistors used as power FETs for power conversion, drive circuitry for the power FETs, a flying capacitor, an on-die LC filter, and control circuitry for voltage regulation. A relatively slow digital feedback loop sets the signals out of the digital pulse-width modulator (DPWM) that feed drivers to switch the 3-level converter[1] with appropriate duty cycles (D). In parallel, a fast shunt regulator [23] on the output reacts to sudden load current transients to maintain a steady voltage. The overall design target is to minimize conversion loss, on-die area overhead, voltage fluctuations, and dynamic voltage scaling time. This section further studies the components in Figure 4.9 and looks at circuit implementations in detail.

The 3-level converter uses four power FETs, a flying capacitor, and an output

[1] I call this particular design a 3-level "converter", instead of a 3-level "voltage regulator" because this design has to operate in open-loop due to a design mistake. I explain this mistake in more detail in Section 4.3.4.

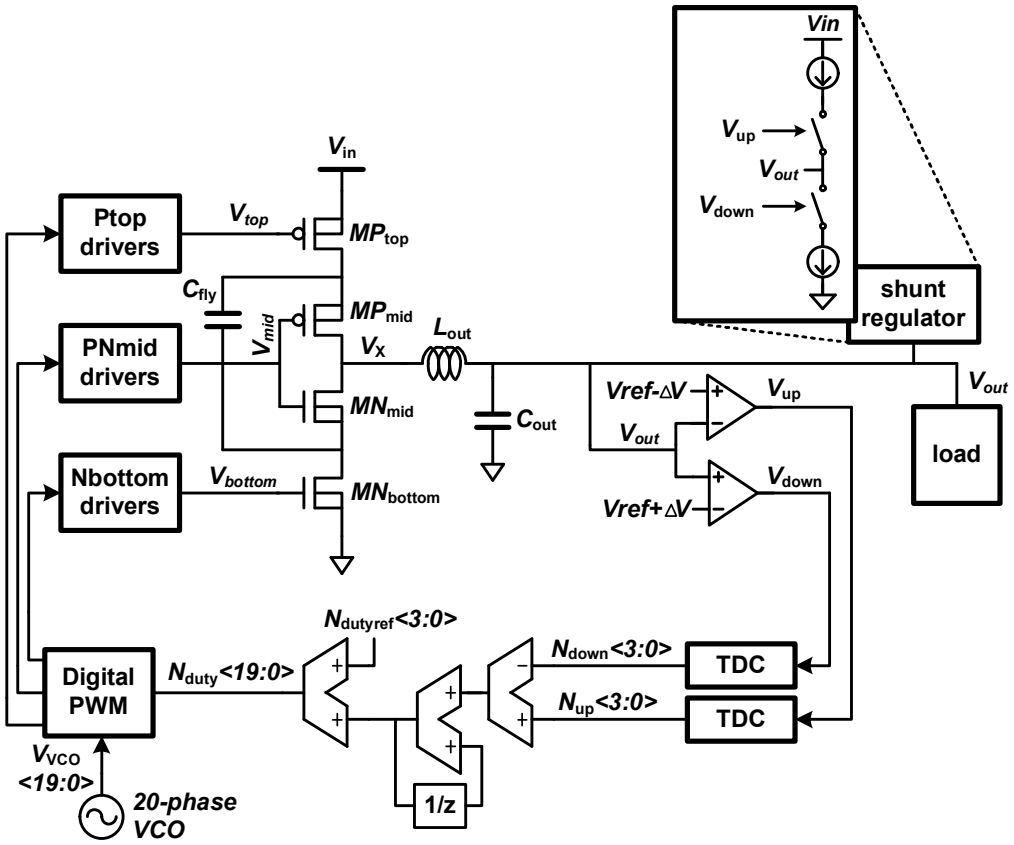

Figure 4.9: Block diagram of 3-level converter with slow digital feedback control and fast shunt regulation. Finer duty cycle control is necessary to avoid limit-cycling.

LC filter to generate a wide range of output voltages. Figure 4.10 illustrates the converter's operation via signal waveforms associated with the power FETs (MP_{TOP}, MN_{BOTTOM}, MP_{MID}, and MN_{MID}) and the output inductor for two scenarios: $0.5 \leq D$ and $D \leq 0.5$. As previously described, node V_X can swing between three voltage levels by iterating through four steps per switching period (T) that control the power FETs and C_{FLY}.

For $D \geq 0.5$, step 1 turns on MN_{BOTTOM} and MP_{MID}, placing C_{FLY} between V_X and 0. In step 3, MP_{TOP} and MN_{MID} turn on, placing C_{FLY} between V_{IN} and

V_X. As in a SC converter, where two capacitors alternate between series-stack and parallel configurations, steps 1 and 3 generate VCfly and V_{IN} - VCfly, respectively, on V_X. Assuming the ideal case where VCfly is equal to $V_{IN}/2$, V_X stays at $V_{IN}/2$ in steps 1 and 3. In steps 2 and 4, V_X connects to V_{IN} through MP_{TOP} and MP_{MID}. By adjusting the time spent in each step, the converter can generate any voltage between V_{IN} and V_{IN} /2 at V_{OUT}.

Conversely, for $D \leq 0.5$, steps 2 and 4 connect V_X to ground through MN_{MID} and MN_{BOTTOM}. Steps 1 and 3 operate in the same manner as described above for $D \geq 0.5$, generating V_{IN} /2 at V_X. Again, by adjusting D, the converter can generate any voltage between $V_{IN}/2$ and 0 at V_{OUT}. For the special case when $D = 0.5$, steps 1 and 3 in the above descriptions effectively disappear and the 3-level converter operates much like a conventional SC converter.

To understand what input signals these power FETs need to iterate across the different steps, we investigate the operation and design requirements for the four power FETs.

4.3.1 Power FETs

The power FETs use thin-oxide devices in a stacked structure to support input voltages (V_{IN}) up to twice the maximum gate-source voltage allowed by the process technology. Compared to thick-oxide devices for I/O, the thin-oxide counterparts exhibit lower conversion loss due to lower parasitic resistance and capacitance. They also require lower voltages to operate, which reduces switching loss. To minimize ON-state resistance, each of the middle transistors MP_{MID} and MN_{MID} connects its

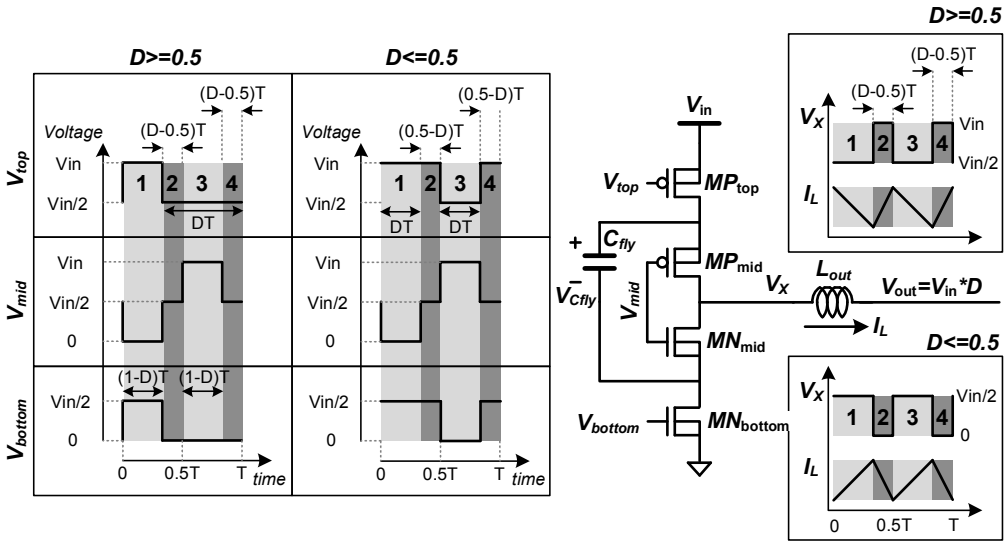

Figure 4.10: Schematic of the proposed 3-level power converter. Signal timing diagrams illustrate different operating modes.

body node to its source instead of to V_{IN} or ground (either of which is possible with triple-well devices).

Again referring to Figure 4.10, the stacked structure using thin-oxide devices requires voltage stress across each device to be limited to $V_{IN}/2$. Input signals to the power FETs need to be carefully set in order to meet this requirement in each step. For this purpose, the input signal to MP_{TOP} (V_{TOP}) swings between V_{IN} and $V_{IN}/2$, while V_{BOTTOM} swings between $V_{IN}/2$ and 0. To limit voltage stress on the middle FETs (MP_{MID} and MN_{MID}), their input (V_{MID}) swings across three voltage levels, V_{IN}, $V_{IN}/2$ and 0. In step 1 for $D \geq 0.5$, V_{MID} is set to 0 to simultaneously turn MP_{MID} on and turn MN_{MID} off. In step 2, both MP_{MID} and MN_{MID} remain in their respective on and off states from step 1. However, as V_X goes up to V_{IN}, V_{MID} must increase to $V_{IN}/2$ to meet voltage stress requirements on MP_{MID} and MN_{MID}. In step 3, V_{MID} is set to V_{IN} to turn MN_{MID} on and turn MP_{MID} off. Step 4 sets V_{MID} to

$V_{IN}/2$ again to alleviate voltage stress as seen in step 2. When the converter operates with $D \leq 0.5$, similar voltage stress constraints must be observed.

This circuitry requires an additional voltage, $V_{IN}/2$, to generate inputs for power FETs that switch between two sets of supply rails (V_{IN} and $V_{IN}/2$ or $V_{IN}/2$ and ground). To generate $V_{IN}/2$, we use an external power source with 20uF on-board and 660pF on-chip decoupling capacitance. Since the pFETs switching between the top supply rails (V_{IN} and $V_{IN}/2$) are larger than nFETs between the bottom rails ($V_{IN}/2$ and ground), current usually flows into the power source that provides $V_{IN}/2$. An integrated linear regulator [39] can replace the external source by bleeding in current caused by the imbalance between top and bottom rails without adding significant power loss.

4.3.2 Driver circuits

Creating appropriate signals to limit voltage stress on the power FETs requires careful design of the circuitry that generates V_{TOP}, V_{MID}, and V_{BOTTOM}. Figure 4.11 presents schematics of the drivers for the four power FETs and associated signal waveforms for the case when $D \geq 0.5$. A digital pulse-width modulator (DPWM) block generates signals based on a digital, thermometer coded representation of the desired converter duty cycle, NDUTY[19:0], using a 20-phase VCO. The DPWM consists of digitally controlled switches that choose two VCO phases that determine the duty cycle of the output signal. While inverters can generate V_{BOTTOM} from the DPWM output signal ($V_{DPWMbottom}$), V_{TOP} requires a level-shifter [80] to shift the DPWM output ($V_{DPWMtop}$), which swings between $V_{IN}/2$ and 0, up to swing between

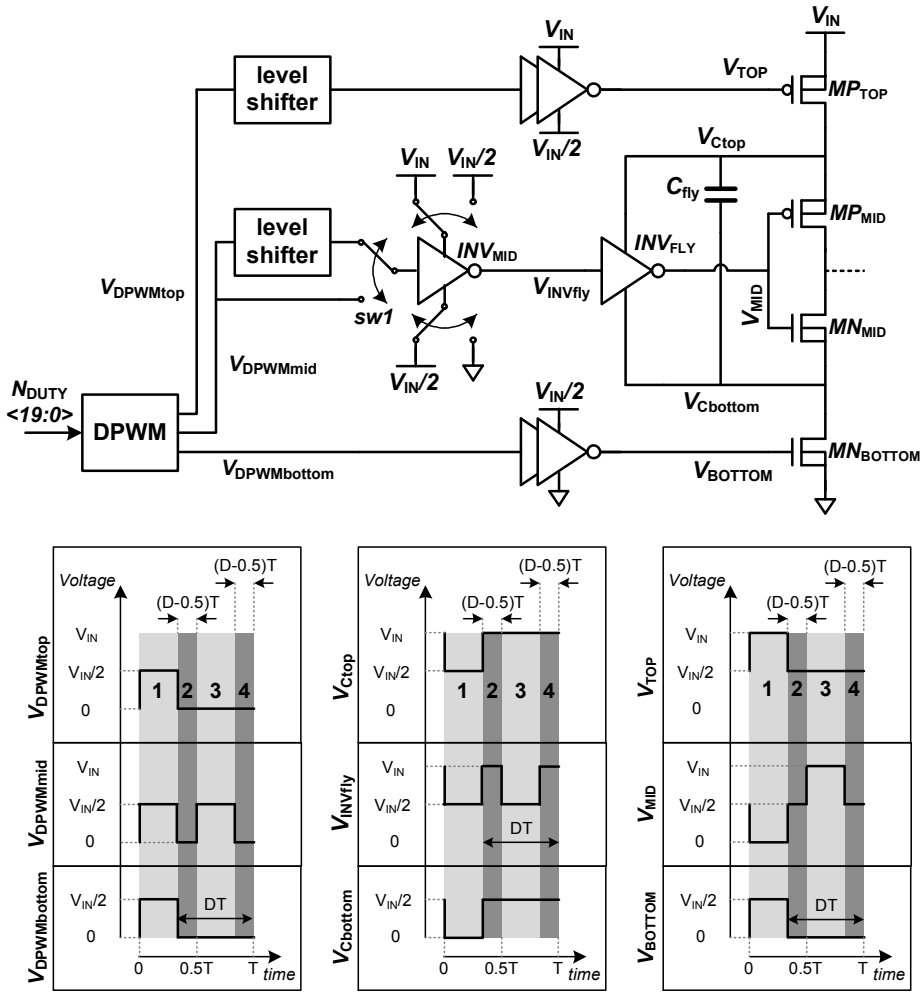

Figure 4.11: Schematic and waveforms that drive power FETs when duty cycle is over 50%.

V_{IN} and $V_{IN}/2$.

The middle FETs, MP_{MID} and MN_{MID}, need a special driver to generate V_{MID} that swings across three different voltages, V_{IN}, $V_{IN}/2$ and 0. The buffer, INV_{FLY}, needs to dynamically switch between two configurations—sitting between V_{IN} and $V_{IN}/2$ and sitting between $V_{IN}/2$ and 0. Since C_{FLY} alternates between the same two configurations, one way to implement INV_{FLY} is to place it between the top (V_{Ctop})

and bottom plate ($V_{Cbottom}$) of C_{FLY}, creating a *flying inverter* [60]. In step 1, INV_{FLY} follows C_{FLY} to sit between $V_{IN}/2$ and 0. Input to INV_{FLY} (V_{INVfly}) is set to $V_{IN}/2$ to generate 0 at V_{MID}. In steps 2 to 4, INV_{FLY} sits between V_{IN} and $V_{IN}/2$ with V_{INVfly} swinging between V_{IN} and $V_{IN}/2$ to generate V_{MID}. While this is the case for $D > 0.5$, V_{INVfly} needs to swing between $V_{IN}/2$ and 0 for $D < 0.5$ (V_{INVfly} is fixed at $V_{IN}/2$ for $D = 0.5$). To accommodate both cases, $D \geq 0.5$ and $D \leq 0.5$, the buffer, INV_{MID}, that generates V_{INVfly} sits between V_{IN} and $V_{IN}/2$ for $D \geq 0.5$, while it sits between $V_{IN}/2$ and 0 for $D \leq 0.5$. INV_{MID} switches between these two configurations using power switches. The switch (sw1) that connects the input to INV_{MID} is an analog 2:1 mux built with thick-oxide devices to accommodate input signals ranging from 0 to V_{IN}.

4.3.3 Passive elements

For high efficiency, it is crucial to design high quality passive elements while not incurring excessive on-die area overhead. Table 4.1 shows specifications for the spiral inductor implemented using top two metal layers in parallel to reduce series resistance. To save on-die area, the flying capacitor resides under the inductor. Since the flying capacitors can potentially inject noise into the inductor, a patterned ground shield protects the inductor from noise coupling [120]. The flying capacitor is implemented with a MOS gate capacitor, because of its higher density compared to metal wire capacitors. However, both sides of the flying capacitor swing by $V_{IN}/2$, which impacts the design of the MOS capacitor.

While a triple-well nFET offers slightly higher capacitor density, a pFET incurs

Inductance	1nH
Series Resistance	400mΩ (@200MHz)
Area	400x400μm
# of turns	1.25
Trace Width	80um
Metal Layers	M7 and M8 (top 2 layers)
Capacitor Density	10fF/μm^2
Bottom-plate Capacitance	0.3fF/μm^2

Table 4.1: Specifications of on-chip spiral inductors modeled using ASITIC [11] and MOS capacitors.

less area overhead associated with the surrounding wells. Hence, we opted to implement the MOS cap using a pFET with drain, source, and body all tied together. A major overhead of this choice comes from the junction capacitance between the P-substrate and N-well, which adds large bottom-plate parasitic capacitance that exacerbates switching loss. Figure 4.8 presents simulated conversion efficiencies of SC and 3-level converters including and excluding bottom-plate parasitic capacitance. Both converters benefit from a 10% efficiency gain across a wide range of loads when bottom-plate parasitic capacitance is eliminated. This motivates using a process technology with high density capacitors with less bottom-plate parasitic capacitance.

4.3.4 Feedback loop and shunt regulator

Building on the previous blocks that generate an output voltage with respect to different duty cycles, we now turn our attention to the relatively slow digital feedback loop and shunt regulator loop that regulate V_{OUT} to a desired level, especially under load fluctuations. Revisiting Figure 4.9, both loops share a pair of fast voltage comparators with hysteresis to sense whether the output voltage is above or below a desired reference level, V_{REF}. In the digital loop, a pair of simple time-to-digital converters (TDC) generates 4-bit thermometer codes, $N_{UP}[3:0]$ and $N_{DOWN}[3:0]$, whose difference corresponds to the V_{OUT}-V_{REF} error within each switching cycle. Accumulating the difference between $N_{UP}[3:0]$ and $N_{DOWN}[3:0]$, and adding it to a reference duty cycle, $N_{DUTYREF}[3:0]$, results in a digital code, $N_{DUTY}[19:0]$, that feeds the DPWM described above. $N_{DUTYREF}[3:0]$ can be programmed externally and changes only when the converter needs to dynamically scale the output voltage. Simultaneously changing V_{REF} and $N_{DUTYREF}[3:0]$ together enables nanosecond-scale voltage scaling, as opposed to only adjusting V_{REF} and slowly accumulating error through the digital loop. $N_{DUTY}[19:0]$ can generate a range of duty cycles between 25% and 75%, in 5% steps, which leads to 120mV output voltage resolution for a 2.4V V_{IN}. The coarse resolution hinders the feedback from providing tight regulation, often resulting in steady-state limit-cycling [77]. Because of this design mistake, all of the measurement results in Section 4.4 are made in open-loop operation and hence call the test-chip a 3-level "converter", instead of a "voltage regulator(VR)". Finer-grain duty cycle control, possible using a VCO with a larger number of phases, is necessary to achieve tighter regulation. I present another 3-level VR test-chip in Section 4.5

that incorporates these changes and operates in closed-loop.

Since the digital loop cannot easily track sudden load current transients, there is a supplemental shunt regulator that suppresses output voltage fluctuations by detecting when V_{OUT} crosses low or high thresholds and injecting or extracting current [23]. Based on the V_{UP} and V_{DOWN} signals from the two comparators, the shunt regulator can either turn on pFETs sitting between V_{IN} and V_{OUT} to inject current to V_{OUT}, or turn on nFETs between V_{OUT} and 0V to extract current from V_{OUT}. Since V_{OUT} varies widely, the shunt regulator uses thick-oxide devices for pFETs sitting between V_{IN} and V_{OUT}. In contrast, maximum voltage stress is 1.4V for nFETs sitting between V_{OUT} and 0, allowing for thin-oxide devices.

4.4 Measurement: Open-Loop

To demonstrate the benefits of the 3-level converter, we designed a test-chip prototype in a 130nm Mixed-Mode/RF CMOS process from UMC with a 2μm thick top metal layer. Figure 4.12 shows the high-level architecture of the test-chip prototype that consists of a pair of 2-phase, 3-level converters arranged as two identical sectors. The two phases share a single output capacitor to reduce ripple on V_{OUT}. Low-impedance, on-chip switches can connect the two sectors together to create a single 4-phase converter with each phase offset by 90 degrees. Otherwise, the test chip implements two independent 2-phase converters. An ability to disable power FETs further enables multiple 3-level converter configurations consisting of one to four phases. A programmable load in each sector facilitates experimental measurements by sinking up to 0.5A in 25mA steps as steady or pseudorandom patterns of

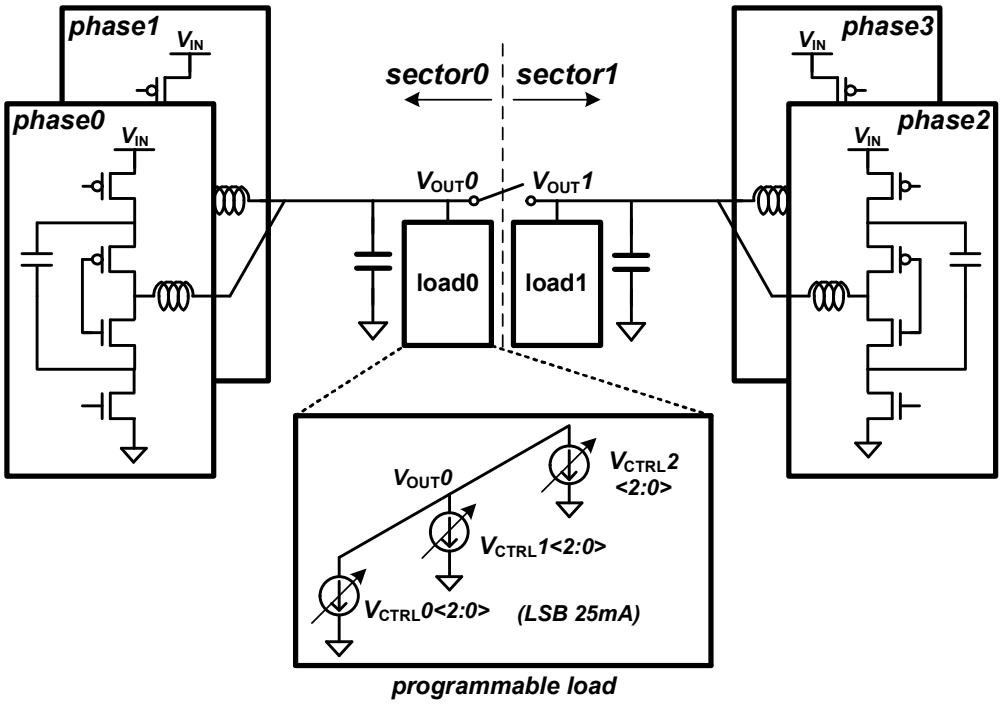

Figure 4.12: High-level architecture of the 3-level converter test-chip prototype.

current.

Measurement results demonstrate that the 3-level converter can generate a wide range of output voltages using 1nH integrated inductors. The converter presents nanosecond-scale voltage transition times and peak conversion efficiency of 77%. Figure 4.13 presents a die micrograph and a list of specifications for the test chip.

Data captured from a real-time oscilloscope (plotted in Figure 4.14) demonstrates the converter can generate output voltages across a wide range – from 0.4 to 1.4V when the input voltage is 2.4V – and rapidly scale V_{OUT} by 1V within 20ns. Such high-speed voltage transitions at nanosecond time scales enable complex digital systems to leverage temporally fine-grained DVFS and improve system-wide energy efficiency

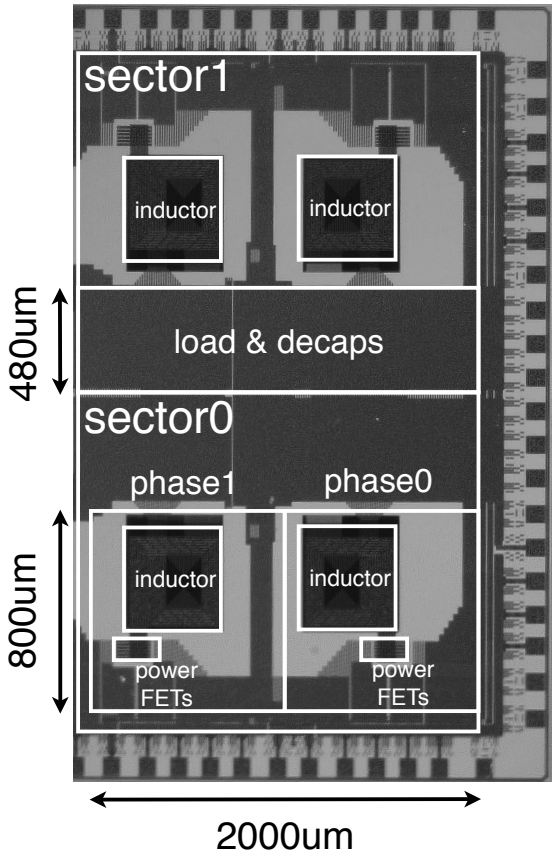

Technology	130nm CMOS
Load Power	0.2-1W
Input Voltage	2.4V
Output Voltage	0.4-1.4V
Inductor per phase	1nH
Total Flying Capacitance	18nF
Total Output Capacitance	10nF
Total Input Capacitance	1nF
Switching Frequency	50-250MHz
Peak Efficiency	77%

Figure 4.13: Die micrograph of the converter with dimensions of main blocks. Flying capacitors are placed underneath the inductors to reduce area overhead. The table shows converter specifications.

[51].

Figure 4.15 summarizes the conversion efficiency measurements made on the test chip in CCM mode. The converter operates in open-loop with fixed duty cycles ranging from 40% to 65% in 5% steps to facilitate measurements across a wide range of conditions. Two converter sectors can also operate with duty cycles that differ by 5% to implement finer steps. Since duty cycle is fixed during open-loop measurements, IR drop due to parasitic resistance causes a spread in output voltages with respect to load currents for the same duty cycle. IR drop is larger than expected due to

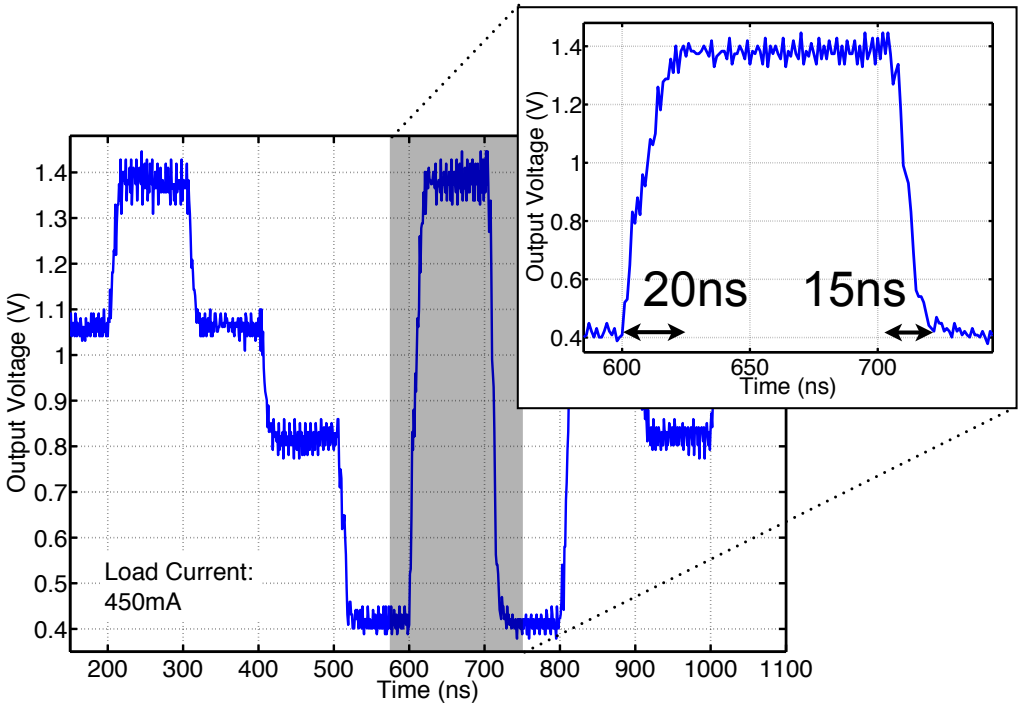

Figure 4.14: Measured snapshot of fast dynamic voltage scaling of the converter operating in open-loop. Voltage scales from 1.4V to 0.4V and vice versa within 20ns.

parasitic resistance on the external power supply, bond-wires, and metal trace. Figure 4.15(a) aggregates all of the measured efficiencies collected across a range of static load current conditions (0.3 to 0.8A), duty cycles (40 to 65%), switching frequencies (50 to 160MHz), and number of phases (1 to 4). Efficiency peaks at 77% for low load current conditions (0.1W/mm^2) at 50% duty cycle. Figure 4.15(b) compares measured data for 50% duty cycle operation using 2 and 4 phases. IR losses increase as load current increases, increasing further for the 2-phase configuration. Higher switching frequency can also degrade efficiency at low load currents due to higher switching losses. Figure 4.15(c) plots the upper range of efficiency measurements for the 4-phase configuration by picking the best efficiency data across different duty cycle settings. Trend line overlays again illustrate the spread in output voltages due to

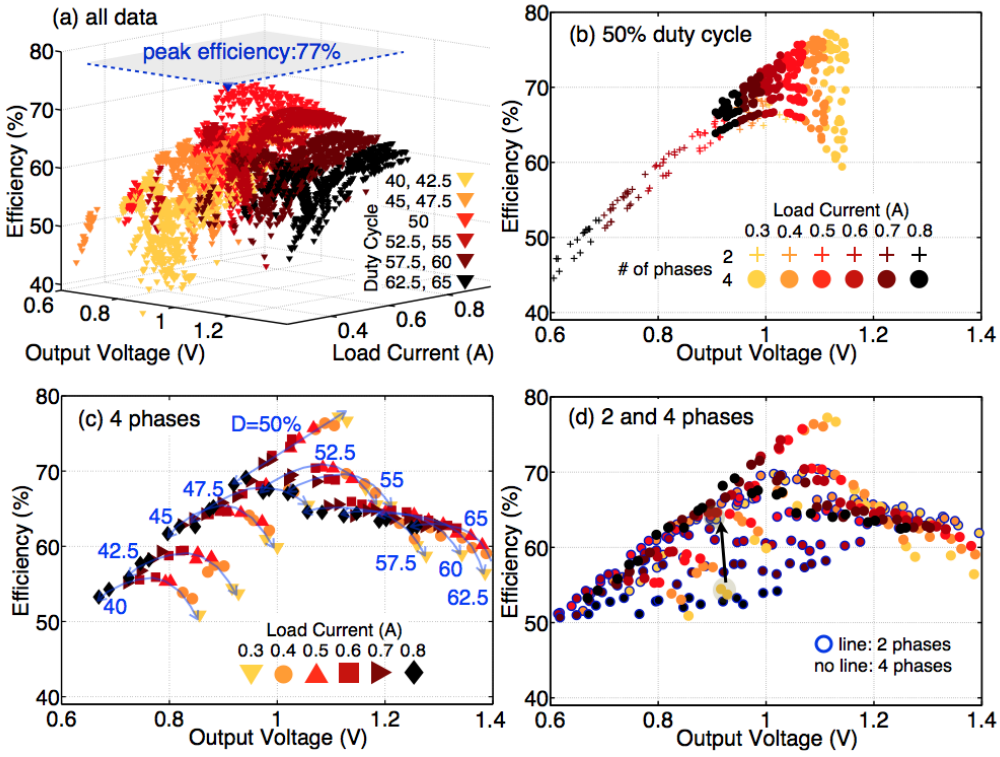

Figure 4.15: Measured efficiency of converter operating in open-loop.

IR drop. Efficiency peaks for 50% duty cycle owing to small inductor current ripple as explained in Section 4.2.2. As duty cycle deviates from 50%, inductor current ripple grows and the corresponding increase in resistive losses degrades conversion efficiency. Figure 4.15(d) adds results for the 2-phase configuration (symbols with outlines) to show that fewer phases can improve efficiency at duty cycles away from 50%.

Using data from Figure 4.15, Table 4.2 presents the breakdown of conversion loss for three different design points. At low loads (point 1), the 3-level converter runs at 152MHz with a single phase to reduce loss due to inductor current ripple. At mid-loads (point 2) where duty cycle is 50%, the number of phases increases to 4 and switching frequency decreases, matching the analysis in Section 4.2.2. However, at

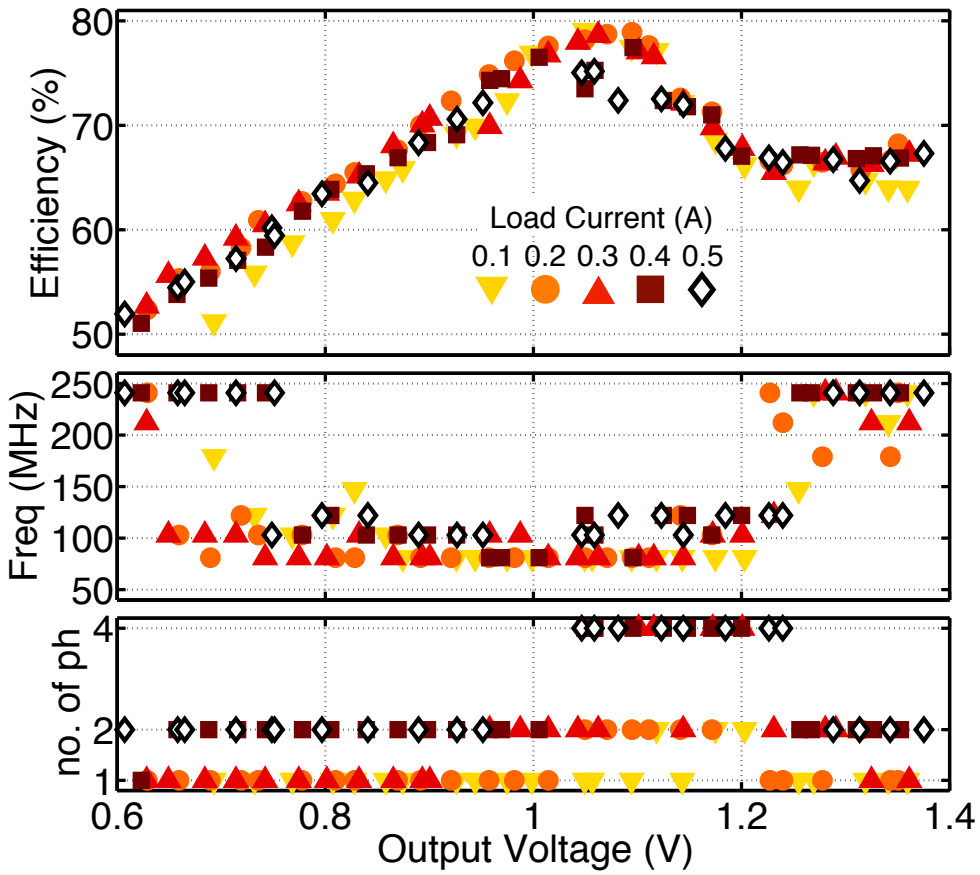

Figure 4.16: Measured conversion efficiency with optimal switching frequencies and number of phases.

high loads, the number of phases does not decrease to 2, but stays at 4. Contrary to the analysis in Section 4.2.2, using 4 phases exhibits higher efficiency than using 2 phases at high loads because a 2-phase converter suffers larger parasitic resistance in the power delivery wires due to floor-plan issues in the test-chip. The die micrograph in Figure 4.13 shows that pads are placed close to each phase of the converter, allowing all phases to have low-impedance connections to power/ground pads. When the converter operates with 2 phases, it has low-impedance connection to about half of the pads that are close to the 2 phases that are turned on. The rest of the pads that

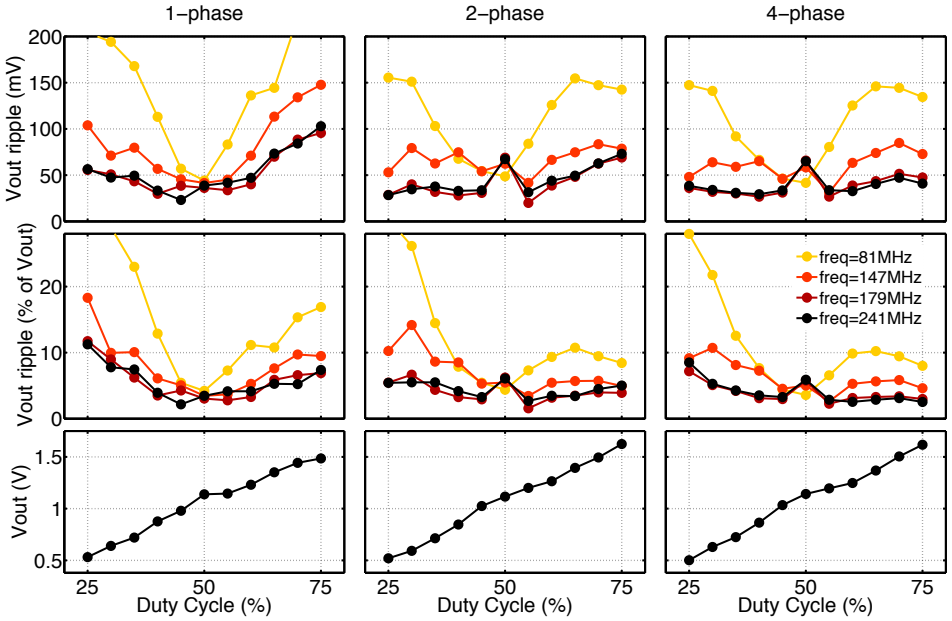

Figure 4.17: Open-loop measurement of peak-to-peak output voltage ripple of the 3-level converter with DC load current. Ripple changes across duty cycles, switching frequencies and number of phases.

are farther away from the 2 phases provide a higher impedance connection with larger parasitic resistance. Compared to a 4-phase converter with short distance to most of the pads, a 2-phase converter suffers from loss due to larger parasitic resistance on the power delivery path.

To further study the effect of frequency and number of phases on efficiency, we measured a second chip across a wider range of switching frequencies. Figure 4.16 presents maximum efficiencies for each load current from 0.1A to 0.5A plotted across output voltages. As shown in the second subplot, frequency reaches a minimum at the center and increases as duty cycle deviates from 50%, following a U-shaped curve. The optimal number of phases, presented in the bottom subplot, also matches the aforementioned trend of 1 phase at low load, 4 phases at the center and 2 phases at high loads. Since the maximum load current is 0.5A, lower than 0.8A in Figure 4.15,

	V_{OUT}	I_{LOAD}	Freq	No. of ph	conduction	switch-ing	bottom-plate	Efficiency
1	0.71V	0.26A	152MHz	1	28%	11%	9%	52%
2	1.03V	0.57A	82MHz	4	12%	7%	8%	73%
3	1.25V	0.78A	152MHz	4	20%	10%	8%	62%

Table 4.2: Breakdown of conversion loss of the 3-level converter for three design points.

larger parasitic resistance on the power delivery path has less impact on conversion efficiency, favoring 2 phases over 4 phases at high loads.

Figure 4.17 presents peak-to-peak output voltage ripple across duty cycles for 1, 2 and 4-phase configurations and different switching frequencies. In this measurement, the converter operates with DC load current ranging from 0.1A to 0.7A that scales linearly with output voltage. As seen in the top row, voltage ripple reaches a minimum at 50% duty cycle for all cases, and increases symmetrically as duty cycle deviates from 50%, matching the trends of $\Delta I_{L,PP}$ in Figure 4.2. Although the absolute magnitude of ripple is roughly symmetric, ripple grows larger as a percentage of V_{OUT} at low output voltages (second row). Interleaving larger numbers of phases helps reduce voltage ripple, especially at extreme duty cycles far from 50%. By increasing the frequency as duty cycle deviates from 50%, and operating with 2 or 4 phases, the converter can maintain 5% peak-to-peak (+/-2.5%) ripple at duty cycles ranging from 30% to 75%, which covers a wide 0.6-1.5V output voltage range.

Compared to steady-state voltage ripple, rapidly changing load current further increases voltage fluctuation. Figure 4.18 presents histogram plots created by sampling the output voltage of the converter. We measure voltage noise due to pseudoran-

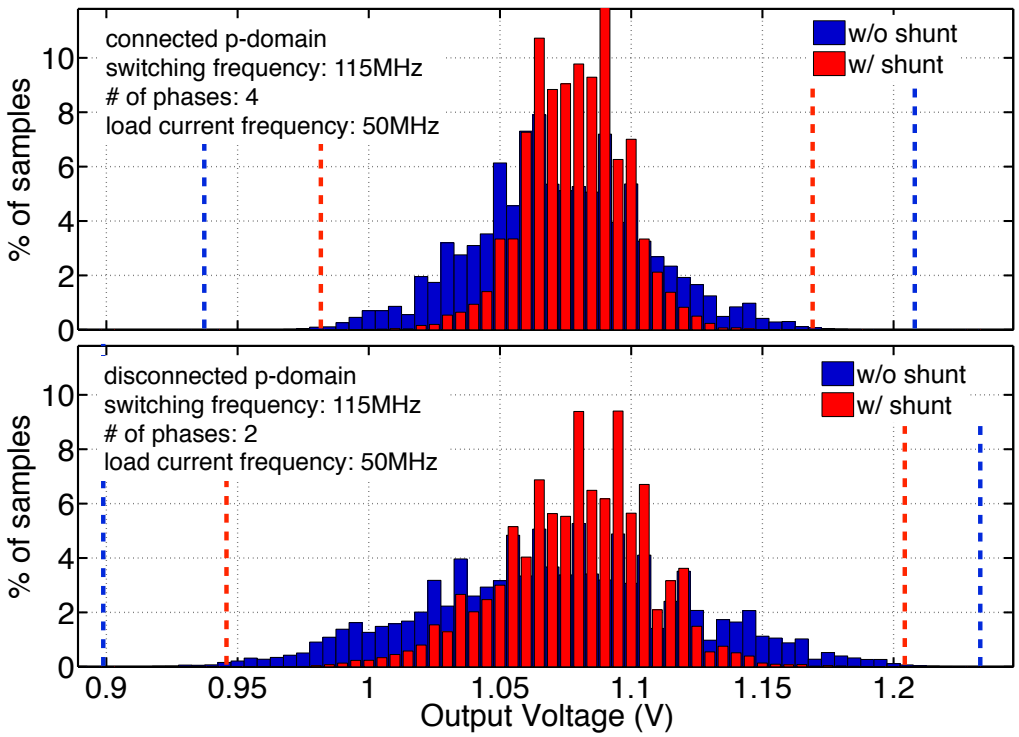

Figure 4.18: Histogram of voltage noise measured in open-loop with and without shunt regulator for connected and disconnected power domains of two sectors.

dom current patterns generated by the programmable loads, with and without the supplemental shunt regulator turned on. The simulated ramp time of load current is 1.5mA/ps. With connected sectors (top plot), the shunt regulator is able to reduce peak-to-peak voltage noise from 0.27V to 0.19V. These results verify that the shunt regulator can appreciably squeeze the noise distribution together and reduce peak-to-peak voltage excursions, shown in dotted lines. Moreover, connecting the power domains reduces voltage noise as a result of larger output capacitance and some canceling of the pseudorandom load currents.

While the shunt regulator – reacting to threshold crossings – reduces voltage fluctuations, it has two drawbacks. First, internal circuit delays limit how quickly this feedback loop can sense and react. Second, simply relying on thresholds provides

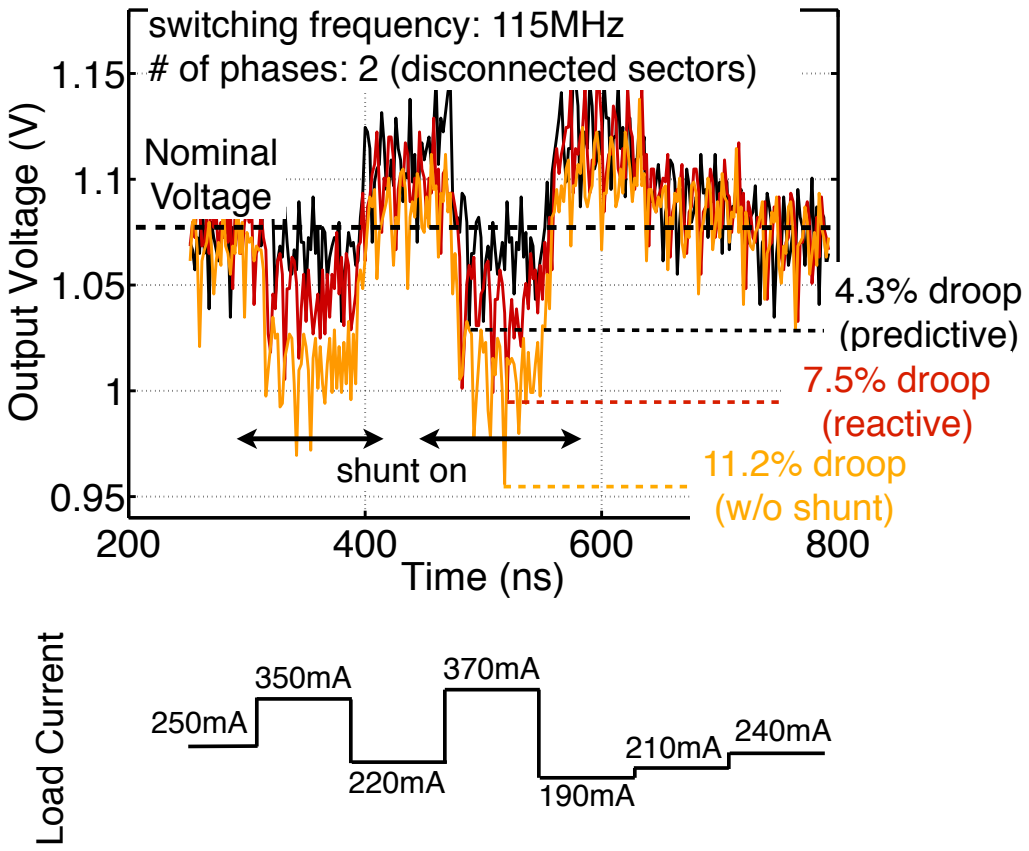

Figure 4.19: Comparison of open-loop measurement of on-die voltage noise without shunt regulator, with reactive shunt, and with predictive shunt.

limited information as to the magnitude of voltage noise and the appropriate response needed to suppress it. One solution is to use a prediction-based shunt regulator that leverages microarchitecture-level information to reliably predicts upcoming voltage droops [85]. The processor can track the history of microarchitecture events using a memory structure to predict events that lead to a surge in load current.

To demonstrate the potential of predictive shunt regulation, we use pulse signals generated externally to turn on the shunt regulator, mimicking signals provided by a processor that predict upcoming voltage droops. Figure 4.19 presents snapshots of measured voltage droops due to two consecutive 80ns wide current pulses of 100mA

and 150mA. Predictive current shunting reduces the maximum voltage droop by over 40% compared to simply reacting to threshold crossings.

Lastly, Table 4.3 compares recently published IVRs using chip-integrated or package-integrated passive elements. Since the published test-chips use different process technologies, input/output voltage ranges and inductor technologies, it is difficult to make a fair comparison across all of them. The test-chip that is most similar to this work is a buck converter built in 130nm using on-chip spiral inductors with 2-2.6V input and 1.1-1.5V output voltage ranges [109]. Compared to this buck converter, our 3-level converter uses a 4x smaller inductor and exhibits 15 percentage points higher efficiency at comparable power densities.

Measurement and analysis from a 130nm test-chip prototype demonstrate the benefits of a fully-integrated 3-level converter. Merging the characteristics of the buck and SC converters, the 3-level converter offers a wide output voltage range using a small 1nH inductor that is suitable for on-chip integration. For a 2.4V to 0.6-1.4V conversion, the converter achieves 79% peak efficiency and voltage scaling across 1V within 20ns, which is 100 times faster than previously presented converters using on-package inductors [90]. Process technologies with smaller bottom-plate parasitic capacitance and thick metal layers offer the potential to further increase the conversion efficiency of future 3-level converter designs.

Building upon this test-chip, the next section dicusses a second-version 3-level VR test-chip that aims to improve upon the first design.

	[2]	[3]	[4]	[8]	[9]	[10]	[11]	[12]	This work
Year	2004	2008	2008	2008	2010	2010	2010	2008	2011
Process Tech (nm)	90 bulk	130 bulk	130 bulk	130 bulk	45 bulk	32 SOI	45 SOI	250 bulk	130 bulk
Topology	buck	buck	stacked inter-leaved	buck	SC	SC	SC	3level	3level
Inductor Capacitor	Air-core on-pkg	Fe-core on-pkg	on-chip spiral	on-chip spiral	MOS cap	MOS cap	Trench cap	bond-wire L	on-chip spiral
Vin	1.2	3.3	1.2	2-2.6	1.8	2	2	3.6	2.4
Vout	0.9	0-1.6	0.9	1.1-1.5	0.8-1	0.5-1.1	0.95	1	0.4-1.4
Freq (MHz)	233	60	170	225	30	-700	100	37.3	50-200
No. of phases	4	16	1	4	No info	32	No info	2	4
L per ph (nH)	6.8	No info	2	3.9	N/A	N/A	N/A	26.7	1
Cfly (nF)	N/A	N/A	N/A	N/A	0.534	No info	0.2	5.07	18
Cout (nF)	2.5	No info	5.2	12.2	0.7	0	No info	25.9	10
Max power (W)	0.27	120	0.32	0.8	0.008	0.3	0.0026	0.1	1
Area (mm2)	1.26	37.6	1.5	3.8	0.16	0.378	0.0012	5.1	5
Power density (W/mm2)	0.21	3.19	0.21	0.213	0.05	0.55	2.19	0.02	0.2*
Efficiency (at power density above, %)	82.5	No info	76	48	No info	81	90	No info	63
Efficiency (peak, %)	83.2	88	77.9	58	69	84	90	69.7	77

* power density includes output decoupling capacitance

Table 4.3: Comparison with prior IVR designs.

4.5 3-Level VR: Closed-Loop

I implemented a second version test-chip to solve the following two problems of the first design.

1. **Prevent limit-cycling in closed-loop operation:** As mentioned in the previous section, closed loop operation was not possible in the first design because coarse duty cycle resolution (5%) resulted in limit-cycling.

2. **Inefficient shunt regulator:** The shunt regulator, operating like an inefficient

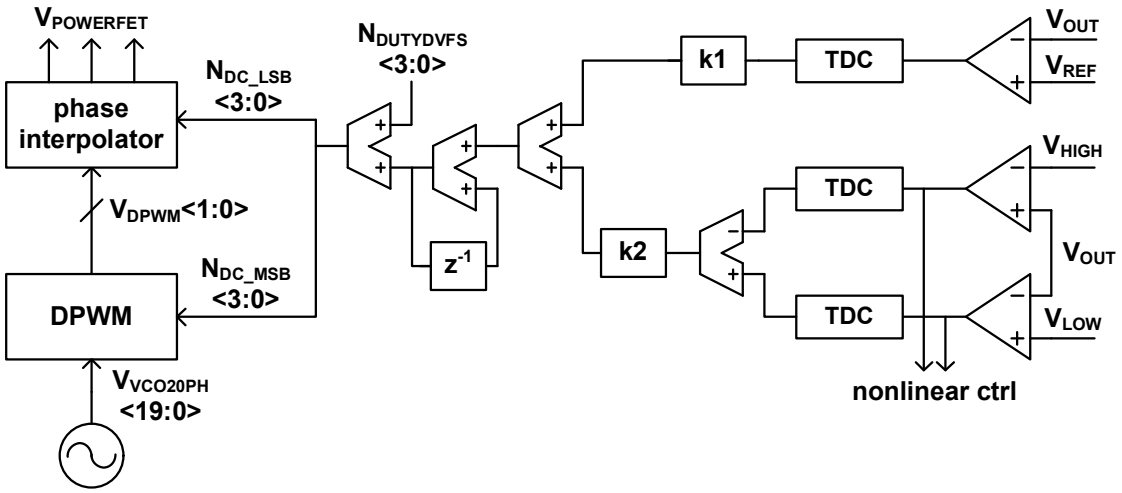

Figure 4.20: High-level diagram of the feedback control in the second version 3-level VR test-chip.

linear regulator, added large conversion loss whenever V_{OUT} crossed low and high thresholds because the charge was dumped directly from V_{IN}.

Figure 4.20 shows a high-level diagram of the feedback control in the second version 3-level VR. The rest of the VR design is similar to the one in Figure 4.9 except that the second version does not have a shunt regulator. To prevent limit-cycling, this feedback adds a digital phase interpolator in addition to the DPWM [106]. Since 5% resolution provided by the DPWM is too coarse, the digital phase interpolator blends two signals with 5% duty cycle difference coming from the DPWM and selects among 16 different edge positions based on N_{DC_LSB}. This results in a duty-cycle resolution of 5/16, or 0.3%, which translates to 5.4mV V_{OUT} resolution when V_{IN} equals 1.8V. This is much finer than the 120mV resolution of the first test-chip.

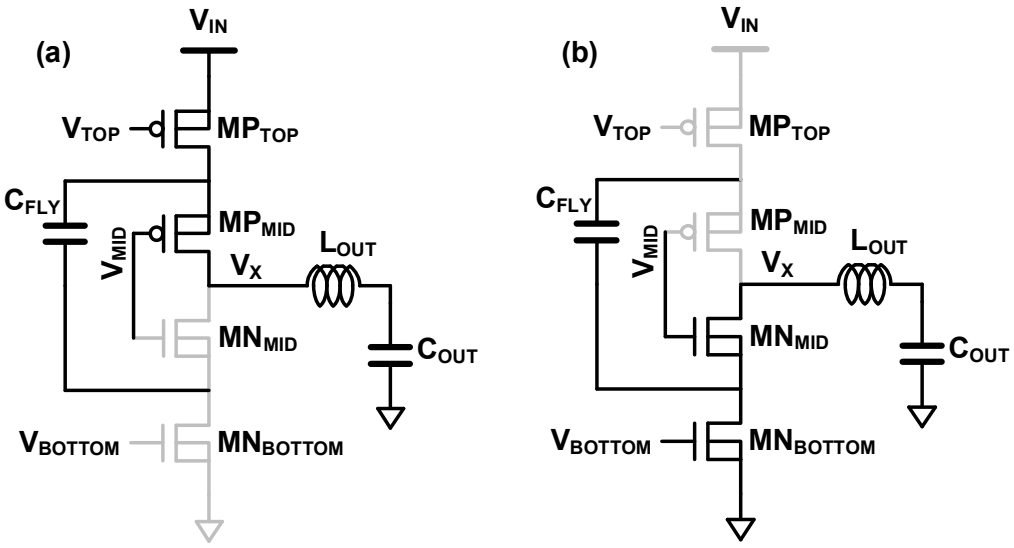

Figure 4.21: Illustration of how the nonlinear control works. Both PFETs turn on whenever V_{OUT} drops below V_{LOW} (a), while both NFETs turn on whenever V_{OUT} spikes above V_{HIGH} (b).

To replace the inefficient shunt regulator, the feedback circuitry includes a nonlinear control that quickly provides charge to V_{OUT} through the inductor, which is more efficient than providing charge directly to V_{OUT}. Revisiting Figure 4.20, a comparator compares V_{OUT} against V_{REF} and the output is sampled using a TDC, which is similar to the feedback in the first test-chip. With a small gain k1, this is a slow feedback path that ensures V_{OUT} always settles to V_{REF}, but does not help in reacting to rapid load transients. There are two additional comparators that sense when V_{OUT} crosses thresholds V_{LOW} and V_{HIGH}, which are voltages that can be externally programmed to be 40-60mV lower and higher than V_{REF}, respectively. Gain k2 is set high so that the duty cycle can quickly change when V_{OUT} deviates from V_{REF}. However, this feedback cannot react quickly enough to respond to sudden load current transients. To react more quickly, the output of the two comparators trigger

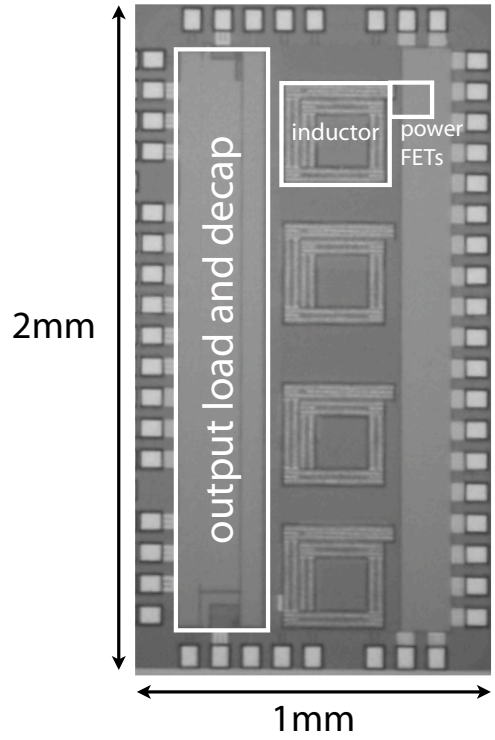

Technology	40nm CMOS
Max Load Power	1W
Input Voltage	1.8V
Output Voltage	0.6 - 1.2V
Inductor per phase	1nH
Total Flying Capacitance	4.4nF
Total Output Decap	5nF on-die 10nF on-board
Switching Frequency	200-400MHz

Figure 4.22: Die photo of the second version 3-level VR test-chip. Similar to the first version, flying capacitors are placed under the inductors to save die area.

a nonlinear control that bypasses the digital blocks in the feedback and directly controls the power switches (Figure 4.21). When V_{OUT} drops below V_{LOW}, power transistors MP_{TOP} and MP_{MID} are forced on to connect V_X to V_{IN}, which provides more charge to V_{OUT} through the inductor. Similarly, when V_{OUT} spikes above V_{LOW}, power transistors MN_{MID} and MN_{BOT} are forced on to connect V_X to 0V, discharging V_{OUT} through the inductor. This control is not as fast as the shunt regulator since the inductor limits how rapidly current can be delivered to V_{OUT}, while the shunt regulator can instantaneously deliver charge to V_{OUT}. However, this control is more efficient since charge is delivered through an inductor, while the shunt regulator acts like a linear regulator that delivers charge to V_{OUT} from a higher V_{IN}.

Incorporating the new feedback design, we implemented a test-chip in 40nm CMOS process that has similar specifications as the first test-chip (Figure 4.22). The differences are the following.

1. Process node advanced from 130nm to 40nm.

2. Input voltage is 1.8V, down from 2.4V, because we stack two 0.9V-rated transistors instead of 1.2V-rated ones.

3. Inductor per phase is 1nH, which is same as the first chip. However, inductor resistance increased from 400mohm to 750mohm, because we used a process with thinner metals.

4. Reduced total flying capacitance from 16nF to 4.4nF to save die area and increase power density. Due to smaller flying capacitance, we increased switching frequency, which was possible due to a more advanced process technology.

5. Added 10nF on the board for output decoupling capacitance because there was just enough die area to put 5nF of on-die decoupling capacitance, which was not enough to keep voltage ripple small.

Measured efficiency of the 3-level IVR was lower than expected. The efficiency degradation is due to parasitic resistance caused by three different wires.

1. **Large resistance on wires connecting C_{FLY} due to ground patterning:** To place C_{FLY} under inductors without degrading inductor Q, wires connecting the flying capacitors are drawn in a ground pattern that is perpendicular to the direction of the wires for the spiral inductor. Because of this limitation, the

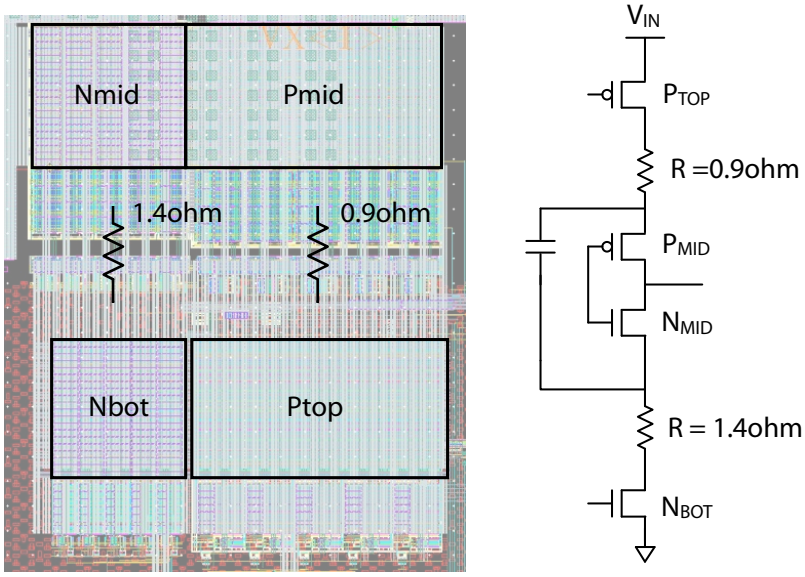

Figure 4.23: Snapshot of layout showing that two connections betweeen power FETs — N_{BOT}-N_{MID} and P_{BOT}-P_{MID} — have high resistance, which significantly degrades conversion efficiency. This is due to a mistake of using too narrow and long paths to connect different power FETs.

impedance of wires connecting C_{FLY} is higher compared to the case where C_{FLY} is connecting by a grid of wires. pattern.

2. **Large resistance on wires connecting nFET and pFET power switches:** I made a layout mistake of using too narrow and long wires to connect different power transistors. Figure 4.23 shows a layout snapshot of four power transistors. GND-N_{BOT}, V_{IN}-P_{TOP} and P_{MID}-N_{MID} connections have little resistance, but the estimated resistance values for N_{BOT}-N_{MID} and P_{BOT}-P_{MID} are 1.4ohm and 0.9ohm, respectively. Considering that this is a single phase in a 4-phase VR delivering up to 1A, maximum 0.25A can flow through 1.4ohm and 0.9ohm.

Adding the losses from all four phases, the parasitic resistance can waste up to 0.3W.

A fundamental problem when designing power ICs in modern process nodes is that transistors are getting smaller, leaving less room for routing wires, while current per transistor width is increasing. As a result, current density increases with process scaling. Furthermore, sheet resistance of metals grow larger due to thinner metals, increasing the risk of significant I^2R loss due to metal resistance. One way to tackle this issue is to leave space between the power transistors to provide more area for routing wires. For example, in Figure 4.23, we could increase the space between N_{BOT} and P_{TOP}, and P_{MID} and N_{MID}, so that wires connecting N_{BOT}-N_{MID} and P_{TOP}-P_{MID} can be wider.

3. **Bondwire resistance:** Bondwire resistance of GND and V_{IN} were 75mohm and 83mohm, respectively. Again, the resistance values seem minor at first glance, but these parasitics can significantly degrade efficiencies when flowing currents as high as 1A.

Figure 4.24 presents the measured and simulated conversion efficiencies for several cases based on which parasitic resistance is included in the simulation. The following lists what each legend in the figure represents.

- **L/R=10nH/ohm (no parasitic):** Simulated efficiency assuming a 0201 inductor with L=1nH, R_L=100mΩ. This case does not include any wire parasitic resistance mentioned above.

- **L/R=1.3nH/ohm (no parasitic):** Simulated efficiency assuming an on-chip

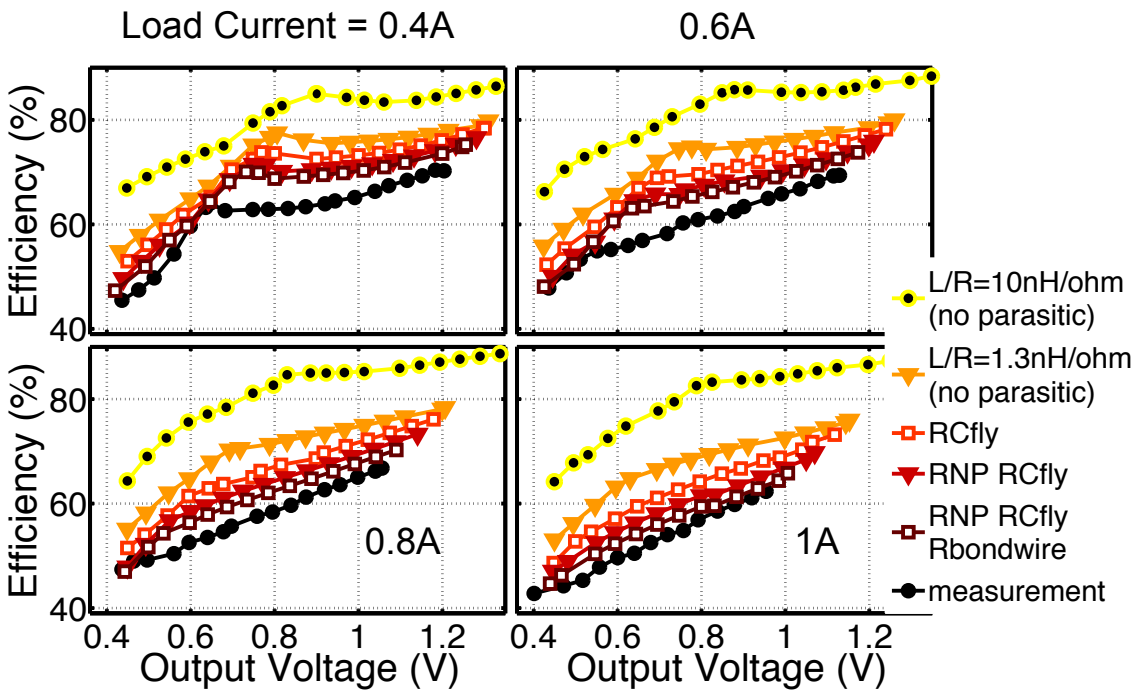

Figure 4.24: Measured efficiencies are lower than expected due to parasitic resistance caused by wires connecting C_{FLY} (RCfly), wires connecting nFET and pFET power switches (RNP) and bondwires (Rbondwire). Simulated efficiencies including three parasitics match well with measured efficiencies. Higher efficiencies are possible with better inductors with higher Q.

spiral inductor with L=1nH, R_L=750mΩ, which is modeled after the inductor used in this IVR test-chip. All of the following cases also assume this inductor. This case does not include any wire parasitic resistance mentioned above.

- **RCfly:** Simulated efficiency including parasitic resistance of wires connecting C_{FLY} (no. 1 in the list of wire parasitics above).

- **RNP RCfly:** Simulated efficiency including parasitic resistance of wires connecting C_{FLY} and wires connecting power switches (no. 1 and 2 in the list of parasitics above, respectively).

- **RNP RCfly Rbondwire:** Simulated efficiency including parasitic resistance of wires connecting C_{FLY}, wires connecting power switches and bondwires (no. 1, 2 and 3 in the list of parasitics above, respectively).

- **measurement:** Measured efficiency with on-chip spiral inductor modeled as L=1nH, R_L=750mΩ. While there is a large gap between measurement and simulation without parasitics, including parasitics help match simulated results and measured efficiencies more closely.

Figure 4.24 shows that measured efficiency matches well with simulated efficiency including all three sources of parasitic resistance, which indicates that the wire parasitic resistance is the reason behind the lower-than-expected efficiencies. Without those parasitics, simulation results show that efficiencies increase by 5-10%p across output voltages and load currents and peak efficiency reaches 80%. Efficiencies can increase further when we use high-Q inductors built using ultra-thick metal [49] or discrete 0201 inductors instead of low-Q on-chip spiral inductors built using thin metal layers. The top line in the figure shows that better inductors increase efficiencies by 5-10%p and can achieve 90% peak efficiency. This shows that 3-level IVRs have the potential to achieve much higher efficiencies with better inductors and smaller parasitics.

Figure 4.25 shows measured V_{OUT} across duty cycles when the converter operates in open-loop with 0A load current at two different switching frequencies. As explained in Figure 4.20, a 20-phase VCO provides 5% duty cycle LSB and a 4-bit digital phase interpolator further divides that 5% into 16, providing 0.31% LSB. Each "tick" in the x-axis of Figure 4.25 is spaced by 5%, which is the LSB of the 20-phase VCO. The

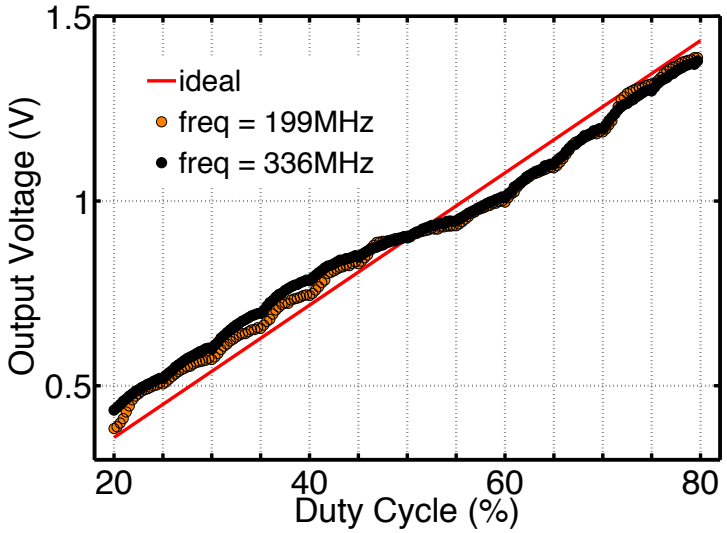

Figure 4.25: Measured output voltage across duty cycles at 0A load current in open-loop.

digital phase interpolator generates finer duty cycles within that 5% range. There is a repeating pattern every 5% due to the non-linearity of the digital phase interpolator.

Revisiting Figure 4.20, $N_{DUTYDVFS}[3:0]$ speeds up voltage transition by instantaneously changing the duty cycle during voltage transitions. Without $N_{DUTYDVFS}$, the slow feedback loop has to change the duty cycle, leading to slow voltage transition. Instead, $N_{DUTYDVFS}$ can change the duty cycle to a value close to the new target duty cycle, and the small adjustment can be handled by the feedback. Figure 4.26 confirms this by showing a comparison of measured voltage traces when the regulator is operating in open-loop, closed-loop and closed-loop with $N_{DUTYDVFS}$. Voltage transition is fast in open-loop since the converter changes the duty-cycle instantaneously, whereas it is slow (4mV/ns) in closed-loop since the feedback is slow to adjust to a new duty-cycle. With $N_{DUTYDVFS}$, closed-loop DVFS is almost as fast as open-loop, presenting a 30mV/ns slew rate. During the first part of voltage transition, the voltage trace of

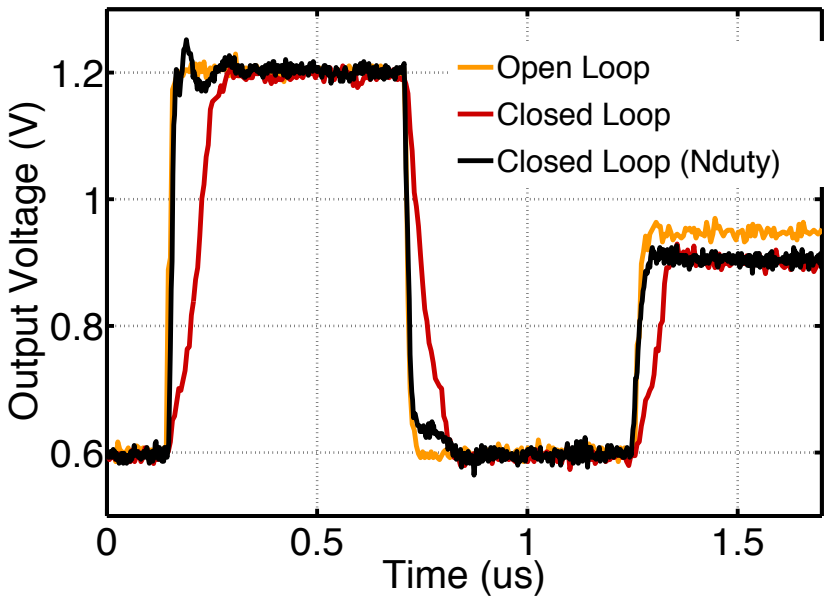

Figure 4.26: Measured voltage traces show that using $N_{DUTYDVFS}$ in closed-loop operation allows the voltage to scale faster. Load current ranges between 0.33A (at 0.6V output voltage) and 0.38A (1.2V output voltage).

closed-loop ($N_{DUTYDVFS}$) follows that of the open-loop since duty-cycles change instantaneously in both cases. However, after the instantaneous change in duty-cycle, the slow feedback starts to handle voltage regulation in closed-loop ($N_{DUTYDVFS}$), which is why the voltage trace deviates from that of open-loop towards the end of the voltage transition. The voltage levels of open-loop and closed-loop at the right side of the figure are slightly different due to limited duty-cycle resolution in open-loop operation. This is also the case in Figure 4.27 that shows measured trasces of voltage scaling across multiple levels in open- and closed-loop (without $N_{DUTYDVFS}$), again with 0.33-0.38A load currents.

Figure 4.28 compares voltage fluctuations in open- and closed-loop operations using load current steps of different magnitudes. The load current increases at 0.5us and decreases at 1.5us and step magnitudes are labeled in each subplot. All load

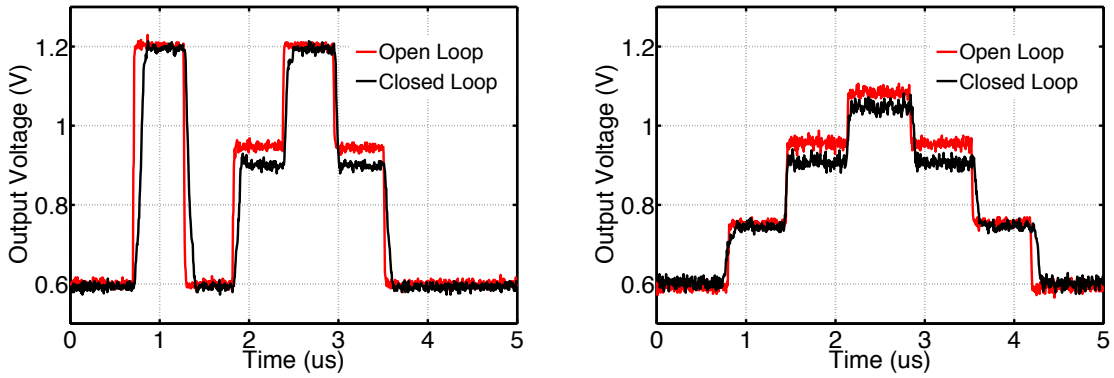

Figure 4.27: Measurement shows that voltage scaling is slower in closed-loop than in open-loop. Both operate with 0.33-0.38A load current.

current transitions occur within 50ps. Duty-cycle is fixed in open-loop, so the voltage levels change significantly when load current changes. In contrast, closed-loop operation maintains the voltage at 1V by adjusting duty cycles except when the voltage droops/spikes due to load current steps. To show how nonlinear control reduces voltage noise, Figure 4.29 compares voltage fluctuations in closed-loop operation with and without nonlinear control (explained in Figure 4.21) using the same load current steps. Measured voltage traces show that nonlinear control reduces voltage droop/spike by up to 90mV.

The frequency of load current steps affect the magnitude of voltage droops/spikes. Figures 4.30 and 4.31 present voltage fluctuations across a range of load current step frequencies, which are labeled in each subplot, for regulators operating in closed-loop at 200MHz and 100MHz switching frequencies, respectively. Voltage fluctuation decreases as load step frequencies increase beyond 25MHz. After a voltage spike/droop is caused by a load current step, the next load step occurs before the voltage settles to the nominal 1V value, leading to smaller voltage fluctuation.

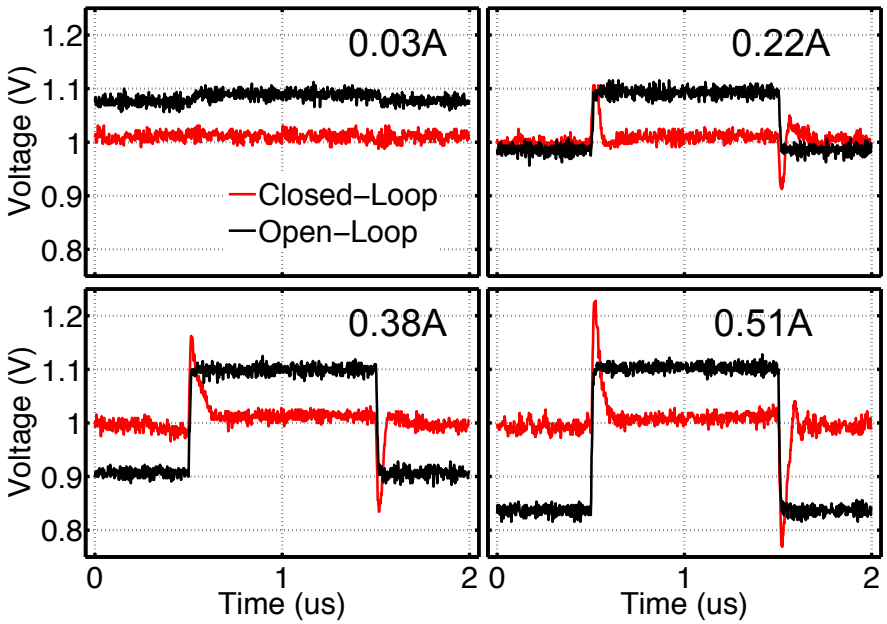

Figure 4.28: Measured voltage fluctuation with load current steps of various magnitudes (labeled in each subplot) when converter operates in open- and closed-loop.

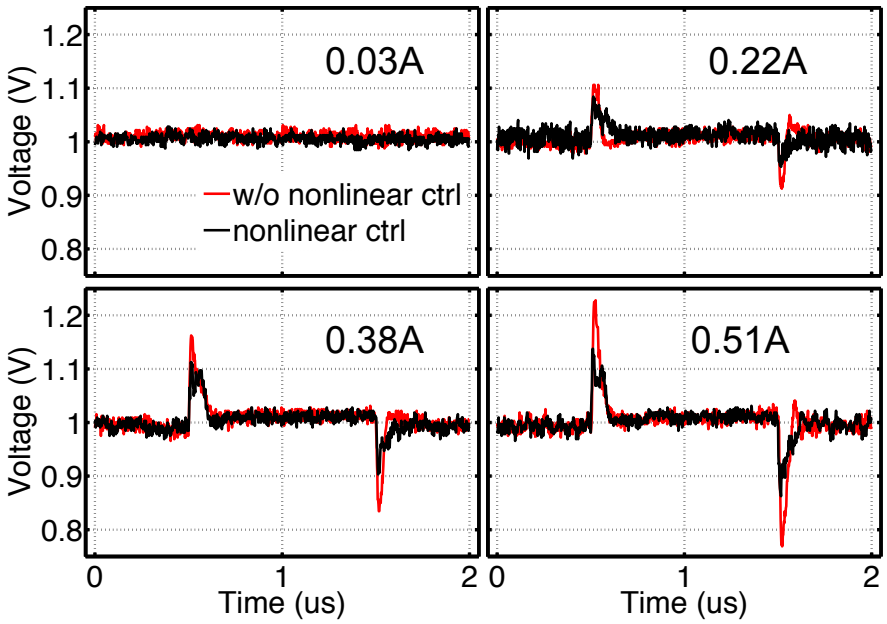

Figure 4.29: Measured voltage traces show that nonlinear control reduces voltage droops/spikes during load current steps in closed-loop operation.

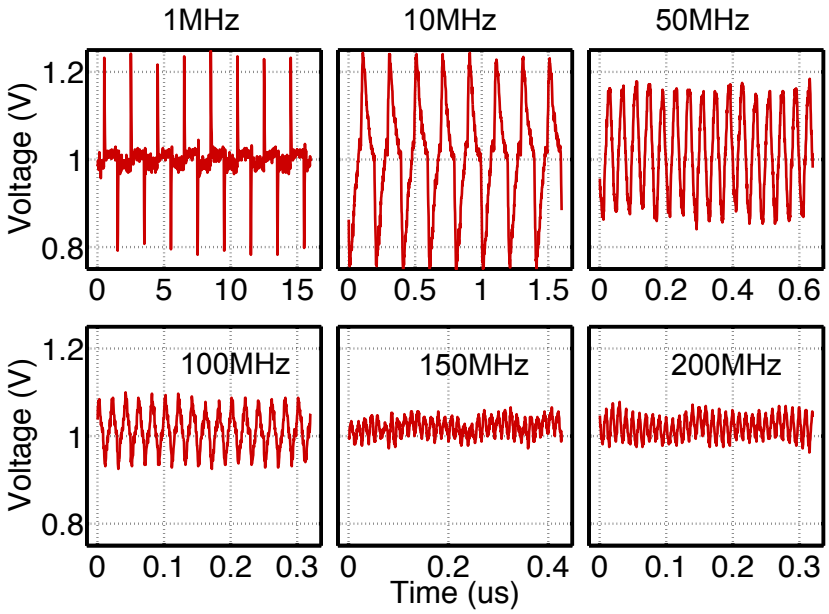

Figure 4.30: Measurement results show that magnitude of voltage fluctuation changes as frequencies of load current steps change. Regulator operates at 200MHz switching frequency. Magnitude of load current step is 0.51A and the current transition occurs within 50ps based on simulation.

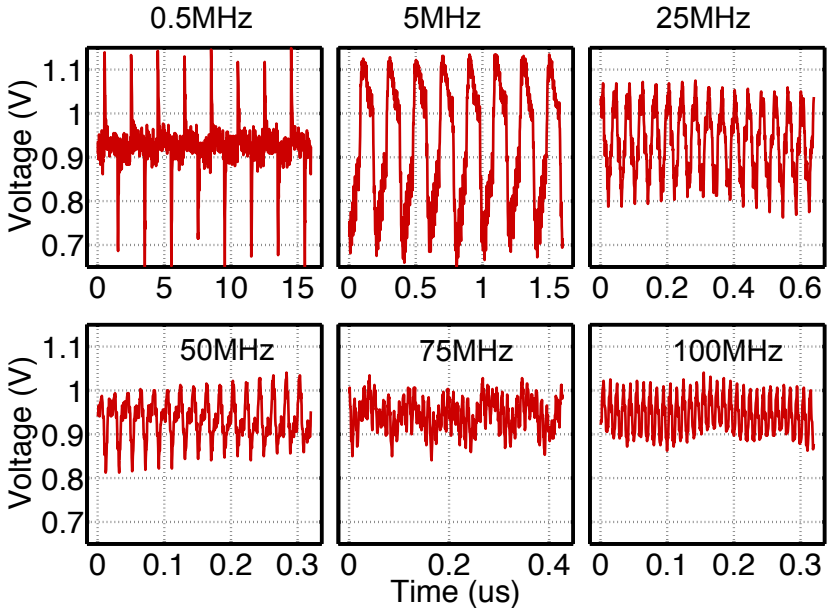

Figure 4.31: Measurements with same settings as Figure 4.30 except that regulator operates at 100MHz instead of 200MHz.

Chapter 5

Technologies on the Horizon

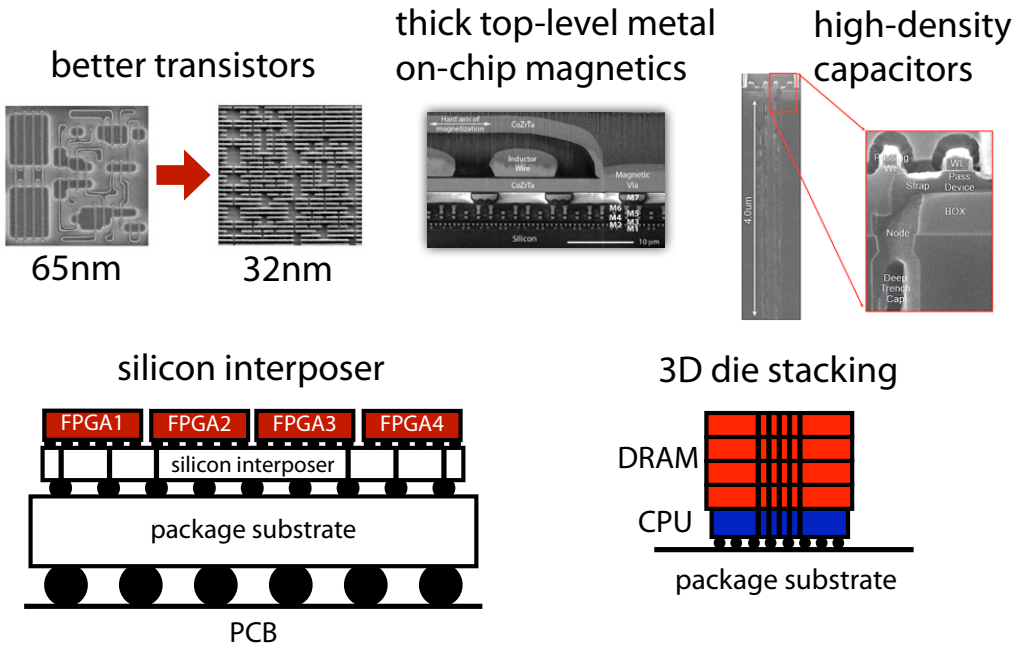

Figure 5.1: Technologies that will impact IVR designs include better transistors [56], thick metal layers and integrated magnetics [34], dense capacitors [103], 2.5D silicon interposers [17] and 3D die stacking [57].

Although IVR designs have come a long way since the early 2000s, most of the IVR publications present lower efficiencies compared to off-chip regulators. However, there are various technologies migrating into commercial products that can be used to design better IVRs with higher conversion efficiency and smaller die area (Figure 5.1).

- **Better CMOS technology:** IVR inefficiencies are caused by parasitic resistance and capacitance of power switches. With cutting-edge process technologies with small parasitic RC, IVRs can operate at switching frequencies high enough ($\geq$50MHz) to enable small inductors ($\leq$50nH) and deliver high current while minimizing resistive loss.

- **Thick top-level metal and integrated magnetics:** Inductor series resistance is a significant source of IVR loss. Thick top-level metal similar to the 8μm thick metal used at Intel [75] or 20μm thick metal at TSMC [49] can reduce series resistance of on-chip spiral inductors. Researchers have also presented on-chip inductors using integrated thin-film magnetic materials to further boost the quality of on-chip inductors.

- **Deep trench caps or other super caps:** High-density capacitors developed for embedded DRAMs offer 20 times higher density than MOSFET capacitors, reducing die area of IVRs. In addition, capacitors with small bottom-plate parasitic can reduce conversion loss of SC and 3-level converters [103].

- **2.5D interposers and 3D stacking:** 2.5D interposers and 3D die stacking are slowly being introduced to commercial products. Xilinx uses 2.5D interposers to connect 4 separate FPGA dies that require fine-pitch interconnects [17]. IBM presented a prototype of an eDRAM die stacked on top of a logic die connected with TSVs at ISSCC 2012 [115]. Using 2.5D and 3D stacking, IVRs can be separate dies connected to the load with silicon interposers or TSVs. The IVR die can use a process technology optimized for the IVR, equipped with the technologies listed above. However, due to the longer distance between the IVR and load dies, voltage regulation could be potentially worse than a single-chip solution.

Based on these technologies, there are a couple of additional applications where IVRs can provide various benefits.

- **Mitigate power/performance degradation due to process variation:** In addition to saving power with fast, per-core DVFS, IVRs can tackle process variation with fine-grain voltage domains. As process variation grow worse and as the number of cores increase, some cores can have slow transistors while others have faster transistors. With a shared voltage, the slowest core determines the voltage for the entire chip. In contrast, IVRs can adjust the voltage of each core separately depending on the process corner of each core.

- **Medical, robotics and defense related applications:** There are various applications in medical, robotics and defense where it is crucial to reduce form factor and number of discrete components. One example is Harvard's Robobee project that aims to create a microrobotic bee that can fly by itself [2]. A 3.7V Li-Ion battery has to power a processor that can calculate where the robotic bee should fly to. Due to stringent requirements on weight and footprint, the voltage regulator has to be integrated in a single die.

IVRs offer the potential for significant energy savings by providing additional knobs for advanced power management in logic ICs. However, that advantage can stay only when IVRs can maintain high efficiency and, at the same time, be small enough to be duplicated many times to support a large number of voltage domains. Going forward, it is going to be increasingly more important to take advantage of advanced process technologies that are co-optimized with novel VR topologies.

Bibliography

[1] Mobile Pentium® III processors® Intel SpeedStep® Technology.

[2] [Online] http://robobees.seas.harvard.edu.

[3] [Online] http://techon.nikkeibp.co.jp/article/HONSHI/ 20100727/184585/?SS=imgview&FD=3561930.

[4] [Online] http://www.coilcraft.com.

[5] [Online] http://www.datacenterknowledge.com/archives/2009/11/27/facebook-follows-google-to-data-center-savings/.

[6] [Online] http://www.intel.com/content/www/us/en/high-performance-computing/high-performance-xeon-phi-coprocessor-brief.html.

[7] [Online] http://www.tilera.com/products/processors/TILE64.

[8] [Online] iFixit - http://www.ifixit.com.

[9] [Online] Intel Ivy Bridge Core i7-3940XM - http://www.notebookcheck.net/Intel-Core-i7-3940XM-Notebook-Processor.80057.0.html.

[10] [Online] Intel Sandy Bridge TDP - http://ark.intel.com/products/64622.

[11] [Online] ASITIC - Analysis and simulation of inductors and transformers for integrated circuits. http://www.eecs.berkeley.edu/ niknejad/asitic.html.

[12] [Online] TI MicroSiP - http://www.ti.com/general/docs/lit/getliterature.tsp?literatureNumber=slvsai0f&fileType=pdf.

[13] [Online] TI WEBENCH http://www.ti.com/ww/en/analog/webench/ index.shtml.

[14] [Online] Xilinx Virtex-7 FPGA.

[15] SimPowerSystems, The MathWorks, Inc.

[16] Low Voltage, 4A DC/DC uModule with Tracking, 2007.

[17] Xilinx WP380: Xilinx Stacked Silicon Interconnect Technology. 2010.

[18] S. Abedinpour, B. Bakkaloglu, and S. Kiaei. A Multi-Stage Interleaved Synchronous Buck Converter with Integrated Output Filter in a 0.18um SiGe Process. In *IEEE International Solid-State Circuits Conference*, February 2006.

[19] Siamak Abedinpour, Bertan Bakkaloglu, and Sayfe Kiaei. A Multi-Stage Interleaved Synchronous Buck Converter with Integrated Output Filter in a 0.18μm SiGe Process. In *International Solid-State Circuits Conference*, 2006.

[20] Mehdi Alimadadi, Samad Sheikhaei, Guy Lemieux, Shahriar Mirabbasi, William G. Dunford, and Patrick R. Palmer. A Fully Integrated 660 MHz Low-Swing Energy-Recycling DC-DC Converter. In *IEEE Transactions on Power Electronics*, June 2009.

[21] Mehdi Alimadadi, Samad Sheikhaei, Guy Lemieux, Shahriar Mirabbasi, and Patrick Palmer. A 3GHz Switching DC-DC Converter Using Clock- Tree Charge-Recycling in 90nm CMOS with Integrated Output Filter. In *International Solid-State Circuits Conference*, 2007.

[22] Mehdi Alimadadi, Samad Sheikhaei, Guy Lemieux, Patrick Palmer, Shahriar Mirabbasi, and William Dunford. A 660MHz ZVS DC-DC Converter Using Gate-Driver Charge-Recycling in 0.18μm CMOS with an Integrated Output Filter. In *Power Electronics Specialists Conference*, 2008.

[23] Elad Alon and Mark Horowitz. Integrated regulation for energy-efficient digital circuits. In *IEEE Journal of Solid-State Circuits*, August 2008.

[24] Massimiliano Belloni, Edoardo Bonizzoni, and Franco Maloberti. High Efficiency DC-DC Buck Converter with 60/120-MHz Switching Frequency and 1-A Output Current. In *European Solid-State Circuits Conference*, 2009.

[25] Henk Jan Bergveld, Katarzyna Nowak, Ravi Karadi, Sebastien Iochem, Jorge Ferreira, Sophie Ledain, Eric Pieraerts, and Mickael Pommier. A 65-nm-CMOS 100-MHz 87% efficient DC-DC down converter based on dual-die System-in-Package integration. In *Energy Conversion Congress and Exposition*, 2009.

[26] Tom M. Van Breussegem and Michiel S. J. Steyaert. Monolithic Capacitive DC-DC Converter With Single Boundary Multiphase Control and Voltage Domain Stacking in 90 nm CMOS. In *IEEE Journal of Solid-State Circuits*, July 2011.

[27] David Brooks, Vivek Tiwari, and Margaret Martonosi. Wattch: a Framework for Architectural-level Power Analysis and Optimizations. In *27th Annual International Symposium on Computer Architecture*, 2000.

[28] Leland Chang, Robert K. Montoye, Brian L. Ji, Alan J. Weger, Kevin G. Stawiasz, and Robert H. Dennard. A Fully-Integrated Switched-Capacitor 2:1 Voltage Converterwith Regulation Capability and 90% Efficiency at 2.3A/mm2. In *Symposium on VLSI Circuits*, 2010.

[29] Ching-Te Chuang, P.F. Luo, and C.J. Anderson. SOI for digital CMOS VLSI: Design Considerations and Advances. In *Proceedings of the IEEE*, 1998.

[30] L. Clark, M. Morrow, and W. Brown. Reverse-Body Bias and Supply Collapse for Low Effective Standby Power. In *IEEE Transactions on VLSI Systems*, September 2004.

[31] L. T. Clark and et al. An Embedded 32-b Microprocessor Core for Low-Power and High-Performance Applications. *IEEE J. Solid-State Circuits*, 36(11):1599–1608, November 2001.

[32] Lawrence T. Clark, Eric J. Hoffman, Jay Miller, Manish Biyani, Yuyun Liao, Stephen Strazdus, Michael Morrow, Kimberley E. Velarde, and Mark A. Yarch. An Embedded 32-b Microprocessor Core for Low-Power and High-Performance Applications. In *IEEE Journal of Solid-State Circuits*, November 2001.

[33] T. Fischer, J. Desai, B. Doyle, S. Naffziger, and B. Patell. A 90-nm variable frequency clock system for a power-managed Itanium architecture processor. *IEEE Journal of Solid State Circuits*, 41:218–228, January 2006.

[34] Donald S. Gardner, Gerhard Schrom, Fabrice Paillet, Brice Jamieson, Tanay Karnik, and Shekhar Borkar. Review of On-Chip Inductor Structures With Magnetic Films. In *IEEE Transcations on Magnetics*, October 2009.

[35] Meeta S. Gupta, Jarod L. Oatley, Russ Joseph, Gu-Yeon Wei, and David Brooks. Understanding Voltage Variations in Chip Multiprocessors using a Distributed Power-Delivery Network. In *Proceedings of DATE'07*, 2007.

[36] Tom Van Breussegem Hans Meyvaert and Michiel Steyaerte. A Monolithic 0.77W/mm2 Power Dense Capacitive DC-DC Step-Down Converter in 90nm Bulk CMOS. In *European Solid-State Circuits Conference*, 2011.

[37] Takayuki Hashimoto, Tetsuya Kawashima, Tomoaki Uno, Noboru Akiyama, Nobuyoshi Matsuura, and Hirofumi Akagi. A System-in-Package (SiP) With Mounted InputCapacitors for Reduced Parasitic Inductancesin a Voltage Regulator. In *IEEE Transactions on Power Electronics*, March 2010.

[38] Takayuki Hashimoto, Masaki Shiraishi, Noboru Akiyama, Tetsuya Kawashima, Tomoaki Uno, and Nobuyoshi Matsuura. System in Package (SiP) With Reduced Parasitic Inductance for Future Voltage Regulator. In *IEEE Transactions on Power Electronics*, June 2009.

[39] P. Hazucha, T. Karnik, B.A. Bloechel, C. Parsons, D. Finan, and S. Borkar. Area-Efficient Linear Regulator With Ultra-Fast Load Regulation. *IEEE Journal of Solid State Circuits*, 40(4), April 2005.

[40] P. Hazucha, G. Schrom, H. Jaehong, B.A. Bloechel, P. Hack, G.E. Dermer, S. Narendra, D. Gardner, T. Karnik, V. De, and S. Borkar. A 233-MHz 80%-87% Efficiency Four-Phase DC-DC Converter Utilizing Air-Core Inductors on Package. In *IEEE Journal of Solid-State Circuits*, 2005.

[41] Peter Hazucha, Gerhard Schrom, Jae-Hong Hahn, Bradley Bloechel, Paul Hack, Greg Dermer, Siva Narendra, Donald Gardner, Tanay Karnik, Vivek De, and Shekhar Borkar. A 233MHz, 80-87% Efficient, Integrated, 4-Phase DC-DC Converter in 90nm CMOS. In *Symposium on VLSI Circuits*, 2004.

[42] Peter Hazucha, Gerhard Schrom, Jaehong Hahn, Bradley A. Bloechel, Associate, Paul Hack, Gregory E. Dermer, Siva Narendra, Donald Gardner, Tanay Karnik, Senior, Vivek De, and Shekhar Borkar. A 233-MHz 80%-87% Efficient Four-Phase DC-DC Converter Utilizing Air-Core Inductors on Package. In *IEEE Journal of Solid-State Circuits*, April 2005.

[43] R. Karadi H.J. Bergveld and K. Nowak. An Inductive Down Converter System-in- Package for Integrated Power Management in Battery-powered Applications. In *Applied Power Electronics Conference*, 2008.

[44] C. Hsu and U. Kremer. The Design, Implementation, and Evaluation of a Compiler Algorithm for CPU Energy Reduction. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'03)*, June 2003.

[45] Canturk Isci, Alper Buyuktosunoglu, Chen-Yong Cher, Pradip Bose, and Margaret Martonosi. An Analysis of Efficient Multi-Core Global Power Management Policies: Maximizing Performance for a Given Power Budget. In *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*, 2006.

[46] Koichi Ishida, Koichi Takemura, Kazuhiro. Baba, Makoto Takamiya, and Takayasu Sakurai. 3D Stacked Buck Converter with 15um Thick Spiral Inductor on Silicon Interposer for Fine-Grain Power-Supply Voltage Control in SiPs. In *International 3D System Integration Conference*, 2010.

[47] Tohru Ishihara and Hiroto Yasuura. Voltage Scheduling Problem for Dynamically Variable Voltage Processors. In *International Symposium on Low Power Electronics and Design*, 1998.

[48] Rinkle Jain and Seth Sanders. A 200mA Switched Capacitor Voltage Regulator on 32nm CMOS and regulation schemes to enable DVFS. In *European Conference on Power Electronics and Applications*, 2011.

[49] Alex Kalnitsky, Y.W. Tseng, T.H. Chien, C.Y. Chang, and Felix Tsui. 1 milli Ohm/square Bondable Post-Passivation Interconnect for Power Management Technologies. In *International Workshop on Power Supply on Chip*, November 2012.

[50] W. H. Ki, F. Su, and C. Y. Tsui. Charge redistribution loss consideration in optimal charge pump design. In *International Symposium on Circuits and Systems*, 2005.

[51] W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks. System level analysis of fast, per-core DVFS using on-chip switching regulators. In *14th International Symposium on High-Performance Computer Architecture (HPCA-14)*, 2008.

[52] Wonyoung Kim, David Brooks, and Gu-Yeon Wei. A Fully-Integrated 3-Level DC/DC Converter for Nanosecond-Scale DVS with Fast Shunt Regulation. In *International Solid-State Circuits Conference*, 2011.

[53] Wonyoung Kim, David Brooks, and Gu-Yeon Wei. A Fully-Integrated 3-Level DC/DC Converter for Nanosecond-Scale DVFS. In *IEEE Journal of Solid-State Circuits*, January 2012.

[54] Sudhir S. Kudva and Ramesh Harjani. Fully Integrated On-Chip DC-DC Converter with a 450x Output Range. In *Custom Integrated Circuits Conference*, 2010.

[55] Sudhir S. Kudva and Ramesh Harjani. Fully Integrated On-Chip DC-DC Converter with a 450x Output Range. In *IEEE Journal of Solid-State Circuits*, August 2011.

[56] K. Kuhn. Variation in 45nm and Implications for 32nm and Beyond? International CMOS Variability Conference (keynote presentation). In *International Electron Devices Meeting*, 2009.

[57] K. Kujala. Presentation. In *Semicon Taiwan*, 2010.

[58] V. Kursun, S. G. Narendra, V. K. De, and E. G. Friedman. Analysis of Buck Converters for On-Chip Integration With a Dual Supply Voltage Microprocessor. In *IEEE Transactions on VLSI Systems*, June 2003.

[59] Hanh-Phuc Le, Seth R. Sanders, and Elad Alon. Design Techniques for Fully Integrated Switched-Capacitor DC-DC Converters. In *IEEE Journal of Solid-State Circuits*, September 2011.

[60] Hanh-Phuc Le, Michael Seeman, Seth R. Sanders, Visvesh Sathe, Samuel Naffziger, and Elad Alon. A 32nm Fully Integrated Reconfigurable Switched- Capacitor DC-DC Converter Delivering 0.55W/mm2 at 81% Efficiency. In *International Solid-State Circuits Conference*, 2010.

[61] Man-Lap Li, Ruchira Sasanka, Sarita V. Adve, Yen-Kuang Chen, and Eric Debes. The ALPBench Benchmark Suite for Complex Multimedia Applications. In *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC-2005)*, 2005.

[62] Pengfei Li, Rizwan Bashirullah, Peter Hazucha, and Tanay Karnik. A Delay Locked Loop Synchronization Scheme for High Frequency Multiphase Hysteretic DC-DC Converters. In *Symposium on VLSI Circuits*, 2007.

[63] Pengfei Li, Deepak Bhatia, Lin Xue, and Rizwan Bashirullah. A 90-240MHz Hysteretic Controlled DC-DC Buck Converter with Digital PLL Frequency Locking. In *Custom Integrated Circuits Conference*, 2008.

[64] Pengfei Li, Deepak Bhatia, Lin Xue, and Rizwan Bashirullah. A 90-240 MHz Hysteretic Controlled DC-DC Buck Converter With Digital Phase Locked Loop Synchronization. In *IEEE Journal of Solid-State Circuits*, September 2011.

[65] Pengfei Li, Lin Xue, Peter Hazucha, Tanay Karnik, and Rizwan Bashirullah. A Delay-Locked Loop Synchronization Scheme for High-Frequency Multiphase Hysteretic DC-DC Converters. In *IEEE Journal of Solid-State Circuits*, November 2009.

[66] P. Macken, M. Degrauwe, M. Van Paemel, and H. Oguey. A voltage reduction technique for digital systems. In *IEEE International Solid-State Circuits Conference*, pages 238–239, February 1990.

[67] D. Marcalescu. On the Use of Microarchitecture-Driven Dynamic Voltage Scaling. In *Workshop on Complexity-Effective Design*, 2000.

[68] Hans Meyvaert, Tom Van Breussegem, and Michiel Steyaert. A 1.65W Fully Integrated 90nm Bulk CMOS Intrinsic Charge Recycling Capacitive DC-DC Converter: Design & Techniques for High Power Density. In *Energy Conversion Congress and Exposition*, 2011.

[69] Rais Miftakhutdinov. An Analytical Comparison of Alternative Control Techniques for Powering Next-Generation Microprocessors.

[70] Surya Musunuri and Patrick L. Chapman. Design of Low Power Monolithic DC-DC Buck Converter With Integrated Inductor. In *Power Electronics Specialists Conference*, 2005.

[71] Surya Musunuri, Patrick L. Chapman, Jun Zou, and Chang Liu. Design Issues for Monolithic DC-DC Converters. In *IEEE Transactions on Power Electronics*, May 2005.

[72] Jinhua Ni, Zhiliang Hong, and Bill Yang Liu. Improved On-Chip Components for Integrated DC-DC Converters in 0.13μm CMOS. In *European Solid-State Circuits Conference*, 2009.

[73] Kohei Onizuka, Kenichi Inagaki, Hiroshi Kawaguchi, Makoto Takamiya, and Takayasu Sakurai. Stacked-Chip Implementation of On-Chip Buck Converter for Distributed Power Supply System in SiPs. In *IEEE Journal of Solid-State Circuits*, November 2007.

[74] Kohei Onizuka, Hiroshi Kawaguchi, Makoto Takamiya, and Takayasu Sakurai. Stacked-chip Implementation of On-Chip Buck Converter for Power-Aware Distributed Power Supply Systems. In *Asian Solid-State Circuits Conference*, 2006.

[75] P. Packan, S. Akbar, M. Armstrong, D. Bergstrom, M. Brazier, H. Deshpande, K. Dev, G. Ding, T. Ghani, O. Golonzka, W. Han, J. He, R. Heussner, R. James, J. Jopling, C. Kenyon, S-H. Lee, M. Liu, S. Lodha, B. Mattis, A. Murthy, L. Neiberg, J. Neirynck, S. Pae, C. Parker, L. Pipes, J. Sebastian, J. Seiple, B. Sell, A. Sharma, S. Sivakumar, B. Song, A. St. Amour, K. Tone, T. Troeger, C. Weber, K. Zhang, Y. Luo, and S. Natarajan. High Performance 32nm Logic Technology Featuring 2nd Generation High-k + Metal Gate Transistors. In *International Electron Devices Meeting*, 2009.

[76] Y. Panov and M.M. Jovanovic. Design Considerations for 12-V/1.5-V, 50-A Voltage Regulator Modules. *IEEE Transactions on Power Electronics*, 16(6), November 2001.

[77] A. V. Peterchev and S. R. Sanders. Quantization Resolution and Limit Cycling in Digitally Controlled PWM Converters. In *IEEE Transactions on Power Electronics*, January 2003.

[78] R. C. N. Pilawa-Podgurski, D. M. Giuliano, and D. J. Perreault. Merged Two-stage Power Converter Architecture with Soft Charging Switched-Capacitor Energy Transfer. In *Power Electronics Specialists Conference*, 2008.

[79] Michael Powell and T. N. Vijaykumar. Exploiting Resonant Behavior to Reduce Inductive Noise. In *Int'l Symp. on Computer Architecture*, Jun 2004.

[80] S. Rajapandian, K. L. Shepard, P. Hazucha, and T. Karnik. High-Voltage Power Delivery Through Charge Recycling. In *IEEE Journal of Solid-State Circuits*, June 2006.

[81] Yogesh Ramadass, Ayman Fayed, Baher Haroun, and Anantha Chandrakasan. A 0.16mm2 Completely On-Chip Switched-Capacitor DC-DC Converter Using Digital Capacitance Modulation for LDO Replacement in 45nm CMOS. In *International Solid-State Circuits Conference*, 2010.

[82] Yogesh K. Ramadass and Anantha P. Chandrakasan. Voltage Scalable Switched Capacitor DC-DC Converter for Ultra-Low-Power On-Chip Applications. In *Power Electronics Specialists Conference*, 2007.

[83] Yogesh K. Ramadass, Ayman A. Fayed, Senior, and Anantha P. Chandrakasan. A Fully-Integrated Switched-Capacitor Step-Down DC-DC Converter With Digital Capacitance Modulation in 45 nm CMOS. In *IEEE Journal of Solid-State Circuits*, December 2010.

[84] Krishna Rangan, Gu-Yeon Wei, and David Brooks. Thread Motion: Fine-Grained Power Management for Multi-Core Systems. In *International Symposium on Computer Architecture*, June 2009.

[85] V. J. Reddi, M. S. Gupta, G. Holloway, Gu-Yeon Wei, M. D. Smith, and D. Brooks. Voltage emergency prediction: Using signatures to reduce operating margins. In *International Symposium on High-Performance Computer Architecture*, February 2009.

[86] Jose Renau, Basilio Fraguela, James Tuck, Wei Liu, Milos Prvulovic, Luis Ceze, Smruti Sarangi, Paul Sack, Karin Strauss, and Pablo Montesinos. SESC simulator, January 2005. http://sesc.sourceforge.net.

[87] G. Schrom, P. Hazucha, J. Hahn, D.S. Gardner, B.A. Bloechel, G. Dermer, S. Narendra, T. Karnik, and V. De. A 480-MHz, Multi-Phase Interleaved Buck DC-DC Converter with Hysteretic Control. In *IEEE Power Electronics Specialist Conference*, 2004.

[88] G. Schrom, P. Hazucha, J.-H. Hahn, V. Kursun, D. Gardner, S. Narendra, T. Karnik, and V. De. Feasibility of Monolithic and 3D-Stacked DC-DC Converters for Microprocessors in 90nm Technology Generation. In *International Symposium on Low Power Electronics and Design*, 2004.

[89] G. Schrom, P. Hazucha, F. Paillet, D. J. Rennie, S. T. Moon, D. S. Gardner, T. Kamik, P. Sun, T. T. Nguyen, M. J. Hill, K. Radhakrishnan, and T. Memioglu. A 100MHz Eight-Phase Buck Converter Delivering 12A in 25mm2 Using Air-Core Inductors. In *Applied Power Electronics Conference*, 2007.

[90] G. Schrom, F. Paillet, and J. Hahn. A 60MHz 50W Fine-Grain Package-Integrated VR Powering a CPU from 3.3V. In *Applied Power Electronics Conference*, 2010.

[91] Gerhard Schrom, Peter Hazucha, Jaehong Hahn, Donald S. Gardner, Greg Dermer Bradley A. Bloechel, Siva G. Narendra, Tanay Karnik, and Vivek De. A 480-MHz, Multi-Phase Interleaved Buck DC-DC Converter with Hysteretic Control. In *Power Electronics Specialists Conference*, 2004.

[92] Michael Douglas Seeman. *A Design Methodology for Switched-Capacitor DC-DC Converters*. PhD thesis, EECS Department, University of California, Berkeley, May 2009.

[93] Greg Semeraro, Grigorios Magklis, Rajeev Balasubramonian, David H. Albonesi, Sandhya Dwarkadas, and Michael L. Scott. Energy-efficient processor design using multiple clock domains with dynamic voltage and frequency scaling. In *International Symposium on High-Performance Computer Architecture*, 2002.

[94] Premkishore Shivakumar and Norman P. Jouppi. Cacti 3.0: An integrated cache timing, power, and area model. Technical report, Western Research Labs, Compaq, 2001.

[95] Tajana Simunic, Luca Benini, Andrea Acquaviva, Peter Glynn, and Giovanni De Micheli. Dynamic Voltage Scaling and Power Management for Portable Systems. In *Design Automation Conference*, 2001.

[96] A. J. Stratakos, S. R. Sanders, and R. W. Brodersen. A Low-Voltage CMOS DC-DC Converter for a Portable Battery-Operated System. In *Proc. IEEE Power Electronics Specialists Conference*, pages 619–626, June 1994.

[97] N. Sturcken, M. Petracca, S. Warren, L. P. Carloni, A. V. Peterchev, and K. L. Shepard. An Integrated Four-Phase Buck Converter Delivering 1A/mm2 with 700ps Controller Delay and Network-on-Chip Load in 45-nm SOI. In *Custom Integrated Circuits Conference*, 2011.

[98] S. Sugahara, K. Yamada, M. Edo, T. Sato, and K. Yamasawa. Low Power Consumption and High Power Density Integrated DC-DC Converter for Portable Equipments. In *Asian Solid-State Circuits Conference*, 2008.

[99] Jian Sun, David Giuliano, Siddharth Devarajan, Jian-Qiang Lu, T. Paul Chow, and Ronald J. Gutmann. Fully Monolithic Cellular Buck Converter Design for 3-D Power Delivery. In *IEEE Transactions on VLSI Systems*, March 2009.

[100] Jian Sun, Jian-Qiang Lu, David Giuliano, T. Paul Chow, and Ronald J. Gutmann. 3D Power Delivery for Microprocessors and High-Performance ASICs. In *Applied Power Electronics Conference*, 2007.

[101] Prabal Upadhyaya, Nan Shi, Sean Bradburn, and Herbert L. Hess. A High Power Density 1.75 mm2 Fully Integrated Closed-loop Buck Converter with Varactor Control Scheme. In *Applied Power Electronics Conference*, 2008.

[102] Gerard Villar and Eduard Alarcn. Monolithic Integration of a 3-Level DCM-Operated Low-Floating-Capacitor Buck Converter for DC-DC Step-Down Conversion in Standard CMOS. In *Power Electronics Specialists Conference*, 2008.

[103] G. Wang, D. Anand, N. Butt, A. Cestero, M. Chudzik, J. Ervin, S. Fang, G. Freeman, H. Ho, B. Khan, B. Kim, W. Kong, R. Krishnan, S. Krishnan, O. Kwon, J. Liu, K. McStay, E. Nelson, K. Nummy, P. Parries, J. Sim, R. Takalkar, A. Tessier, R. M. Todi, R. Malik, S. Stiffler, and S. S. Iyer. Scaling deep trench based eDRAM on SOI to 32nm and Beyond. In *International Electron Devices Meeting*, 2009.

[104] Hangsheng Wang, Xinping Zhu, Li-Shiuan Peh, and Sharad Malik. Orion: A power-performance simulator for interconnection networks. In *Proceedings of MICRO 35*, 2002.

[105] J.D. Warnock, J.M. Keaty, J. Petrovick, J.G. Clabes, C.J. Kircher, B.L. Krauter, P.J. Restle, B.A. Zoric, and C.J. Anderson. The circuit and physical design of the POWER4 microprocessor. *IBM Journal of Research and Development*, 46(1), 2002.

[106] Gu-Yeon Wei. *Energy-Efficient I/O Interface Design With Adaptive Power-Supply Regulation*. PhD thesis, EECS Department, Stanford University, June 2001.

[107] Mike Wens, Koen Cornelissens, and Michiel Steyaert. A Fully-Integrated 0.18μm CMOS DC-DC Step-Up Converter, Using a Bondwire Spiral Inductor. In *European Solid-State Circuits Conference*, 2007.

[108] Mike Wens and Michiel Steyaert. A Fully-Integrated 130nm CMOS DC-DC Step-Down Converter, Regulated by a Constant On/Off-Time Control System. In *European Solid-State Circuits Conference*, 2008.

[109] Mike Wens and Michiel S. J. Steyaert. An 800mW Fully-Integrated 130nm CMOS DC-DC Step-Down Multi-Phase Converter, with On-Chip Spiral Inductors and Capacitors. In *Energy Conversion Congress and Exposition*, 2009.

[110] Mike Wens and Michiel S. J. Steyaert. A Fully Integrated CMOS 800-mW Four-Phase Semiconstant ON/OFF-Time Step-Down Converter. In *IEEE Transactions on Power Electronics*, February 2011.

[111] J. Wibben and R. Harjani. A High Efficiency DC-DC Converter Using 2nH On-Chip Inductors. In *IEEE Symposium on VLSI Circuits*, 2007.

[112] Josh Wibben and Ramesh Harjani. A High Efficiency DC-DC Converter Using 2nH On-Chip Inductors. In *Symposium on VLSI Circuits*, 2007.

[113] Josh Wibben and Ramesh Harjani. A High-Efficiency DC-DC Converter Using 2nH Integrated Inductors. In *IEEE Journal of Solid-State Circuits*, April 2008.

[114] Steven Cameron Woo, Moriyoshi Ohara, Evan Torrie, Jaswinder Pal Singh, and Anoop Gupta. The splash-2 programs: Characterization and methodological considerations. In *Proceedings of the 22nd International Symposium on Computer Architecture*, 1995.

[115] M. Wordeman, J. Silberman, G. Maier, and M. Scheuermann. A 3D System Prototype of an eDRAM Cache Stacked Over Processor-Like Logic Using Through-Silicon Vias. In *International Solid-State Circuits Conference*, 2012.

[116] Qiang Wu, Philo Juang, Margaret Martonosi, and Douglas W. Clark. Voltage and Frequency Control With Adaptive Reaction Time in Multiple-Clock-Domain Processors. In *11th International Symposium on High-Performance Computer Architecture*, 2005.

[117] W. Wu, N.C. Lee, and G. Schuellein. Multi-Phase buck Converter Design with Two-Phase Coupled Inductors. In *IEEE Applied Power Electronics Conference and Exposition*, 2006.

[118] Fen Xie, Margaret Martonosi, and Sharad Malik. Compile-time Dynamic Voltage Scaling Settings: Opportunities and Limits. In *PLDI '03: Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation*, 2003.

[119] V. Yousefzadeh, E. Alarcon, and D. Maksimovic. Three-level Buck Converter for Envelope Tracking Applications. In *IEEE Transactions on Power Electronics*, March 2006.

[120] F. Zhang and P. R. Kinget. Design of Components and Circuits Underneath Integrated Inductors. In *IEEE Journal of Solid-State Circuits*, October 2006.

[121] 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*, 2011.

[122] X. Zhou, P.L. Wong, P. Xu, F.C. Lee, and A.Q. Huang. Investigation of Candidate VRM Topology for Future Microprocessors. In *IEEE Applied Power Electronics Conference and Exposition*, 1998.