Improving SoC Power Delivery With Fully Integrated Switched-Capacitor Voltage Regulators

Citation

Tong, Tao. 2015. Improving SoC Power Delivery With Fully Integrated Switched-Capacitor Voltage Regulators. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

http://nrs.harvard.edu/urn-3:HUL.InstRepos:23845472

Terms of Use

Share Your Story

Improving SoC Power Delivery with Fully Integrated Switched-Capacitor Voltage Regulators

A dissertation presented

by

Tao Tong

to

The School of Engineering and Applied Sciences

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Engineering Sciences

Harvard University

Cambridge, Massachusetts

July, 2015

Thesis advisors   Author

Gu-Yeon Wei and David Brooks   **Tao Tong**

Improving SoC Power Delivery with Fully Integrated Switched-Capacitor Voltage Regulators

Abstract

Traditional power delivery solutions f or system-on-chip (SoC) applications rely on off-chip voltage regulators. The off-chip power delivery solution is becoming a bottleneck for SoCs, due to 1) coarse voltage domain management, 2) increased cost as well as complexity of the power delivery network, and 3) high I^2R loss as supply voltages scale down with the fabrication technology. One promising solution is to integrate the voltage regulators in the SoC. While fully integrated voltage regulators (FIVRs) could resolve these problems, their performance is limited by low efficiency and high chip area overhead, especially if the conversion ratio of the converter is high ($\geq$ 4-to-1).

This thesis presents the design and implementation of two fully integrated switched-capacitor (SC) DC-DC voltage regulators. Both regulators are implemented in the SoC along with the microprocessors they deliver power to. I first present a two-stage 4-to-1 SC regulator in a flapping wing micro-robotic bee application. The regulator converts a 3.7V battery voltage down to two lower voltages (~1.8V and ~0.9V) for the rest of the circuits in the SoC. The two-stage topology and the proposed charge recycling technique improve conversion efficiency and provide very fast load regulation to handle the dynamic current fluctuation of the load circuitry. Next, I explore the power delivery architecture at the system level and propose a joint power delivery network that combines SC FIVRs with voltage stacking. Voltage stacking reduces the maximal power that the FIVRs have to provide and "hides" the FIVR conversion loss so that the latter only applies to a portion of the total power consumed by the load. The FIVRs reduce the voltage noise of the stacked voltage domains when the load in the stacked voltage domains consumes a different amount of power. To verify the benefits of this new power delivery system, a fully integrated reconfigurable SC regulator is implemented with 16 Intel microcontroller cores that are stacked in four voltage domains. The SC regulator simultaneously provides power to the four stacked voltage domains (~0.9V) from a single input voltage (~3.6V). The regulator can dynamically change its configuration to optimize its performance according to the current profiles of the stacked load. A hybrid feedback control scheme is implemented to simultaneously regulate the four stacked domains. The proposed power delivery system achieves an average efficiency of 87%

and a peak efficiency of 99%. At the end of this thesis, I present my conclusion and discuss the technologies that could further improve FIVR-based power delivery systems in the future.

Contents

Acknowledgements

Looking back at the graduate school years, I see a young student growing to a determined man who knows what he wants in his life. I am who I am thanks to those surrounding me. I always read the acknowledgements first when I read others' thesis, because I not only learn their research work, but also share the days and nights that they spend in graduate school. Here I write my acknowledgements with gratitude and honor.

I was really fortunate to have Gu and David as my advisors. Gu and David taught me so many things and made me grow better every single day. When I first came to Harvard, I felt extremely challenging when I met with Gu and David, discussing research projects. I didn't know how to present a research idea and there were many questions that I didn't expect or didn't have answers for. I felt depressed. Gu and David encouraged me, and more importantly, kept challenging me. Now, when I meet Gu and David, I am able to lead the meeting and explain the important issues even before they ask. I love being challenged and that makes me who I am today. Gu and David also helped me a lot improve other skills. Before I gave my first conference presentation, Gu and David went through the slides multiple times with me. They gave me incredible suggestions that significantly improved my talk. My discussions with Gu and David went beyond research and projects. We often talked about news in the industry and interesting breakthroughs in other areas. These conversations gave me the perspectives to connect different technologies together and thought about the

applications that could bring the technologies to the next level. I could not have asked more from my advisors.

I also want to thank my committee members. Prof. Zickler was on my qualify exam committee. He was interested in the presentation and asked questions that helped me revisit my research work from a different perspective. He also attended my final defense and gave me suggestions on the thesis. It is my honor to work with him.

I am also very grateful that I am in a wonderful research group. People in our group are interested in a variety of things and I enjoy so much working and talking with them. For most time, we have lunch together. The topics cover technology, entertainment, politics, daily news and so on. I learned so much from Wonyoung when I worked on my first research project. And I enjoyed working with him when I worked at Lion Semi. I always joked with him that we are going to do something together in the future. I hope that will come true. I feel lucky that I worked with Silvia, Saekyu and Mario on various projects. I was often impressed by how sharp they are. Silvia was from the same college that I went to. I felt happy for her that she got a very good faculty position. Saekyu and I co-designed a very interesting silicon chip that I am really proud of. Mario and I were roommate for one year and I was the groomsman when he got married. Xiao, Sophia, Kevin, Simone, Svilen and Mike from the architecture group made my life much easier when I first came to Harvard. They were very patient and helpful. There were also new members joining the lab every year, Siming, Hyunkwang, Simon and Paul joined the circuit lab. Brandon, Bob, Emma, Rafael and Sam joined the architecture lab. These

people gave me a special reason to go to lab every day. I would also like to thank Glenn, our lab manager. He was the "to-go" person whenever I had a question about the servers and simulation tools. I probably couldn't have finished my Ph.D without him. I also remember the random talks we had in the MD kitchen about his life experience.

I am lucky to have many friends outside the research group. I spent so much time together with Yu Lei, Xu Zhang, Ruichao Ma, Lilei Xu, and Kecheng Li, who I knew from HCSSA. I knew Li Yu from college and he was in MIT when I came to Harvard. It was good to have him around. Hongyao joined Harvard one year after me. I love talking with her and I know that she will be a great researcher. Meng Jia worked with me and Yu Lei on interesting projects. Haifei gave me a lot of good advice on career developments. Wei Sun, Ying Xiong, Weiyi Liu, Lu Sun, Yijing Chen, Qiang Zhi, Qiaoying Zhang, Ming Ying, Hui Xu, Ran Li and Weifei shared a lot of memories with me. There are a lot more people that I love being with and I could not list all of them. I am sure the friendship will last in the rest of my life.

I cannot thank my family enough for their endless support over the years. My mom and dad were in China. They were always supportive and were ready to give to advice and love. I regret that I could not have visited them more during my Ph.D.

Chapter 1

Delivering Power to SoCs with Fully Integrated Voltage Regulators: Opportunities and Challenges

1.1 Opportunities for fully integrated power delivery in SoCs

The emergence of mobile computing applications, such as wearable smart devices and micro-robotics, imposes stringent energy and form factor requirements on the electronics. To build an energy-efficient computing system, the industry has moved towards system-on-chip (SoC) solutions that integrate heterogeneous analog and digital circuits, such as digital general-purpose processor cores, hardware specialized functional units, analog-to-digital converters (ADCs), and sensor interfaces [1-66].

As SoCs become more and more complicated, delivering power to them is a challenge. Conventional power delivery solutions use off-chip voltage regulators (VRs). As shown in **Figure 1.1**(a), off-chip VRs connect the SoC and the energy source, such as a Li-ion battery in this mobile application example. The off-chip VRs convert the battery voltage down to lower voltages for different functional units in the SoC. During the past decade, much effort has been devoted to integrating the entire VR solution into the SoC [9-30]. **Figure 1.1**(b) shows a high-level diagram of a power delivery solution based on fully integrated voltage regulators (FIVRs). The FIVRs are integrated in the SoC without using any external board components. A single voltage is delivered to the SoC, and the voltage conversion/regulation is conducted inside the SoC.

Compared to off-chip VRs, FIVRs provide a number of benefits for mobile SoC applications, including:
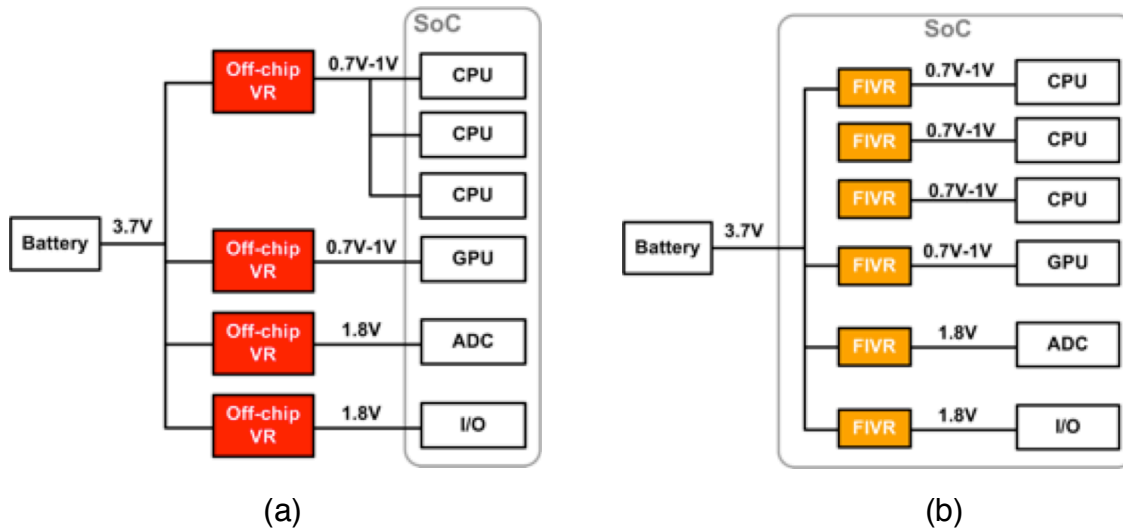
Figure 1.1: Power delivery in a mobile system based on (a) off-chip VRs and (b) FIVRs.

1. **FIVRs enable fine-grain dynamic voltage and frequency scaling (DVFS) for SoCs.** The various units in an SoC, such as processor cores, require dynamic and workload-dependent operating voltages to achieve energy-efficient computing. However, multiple cores in SoCs usually share the same voltage in the off-chip VR solution, due to the difficulty of duplicating off-chip VRs and the complexity of routing a large number of supply rails on a PCB/package. With FIVRs, only one voltage is delivered to the SoC, and the FIVRs in the SoC generate many separate supply rails for the various blocks in the SoC, reducing the overhead associated with sharing voltages. Previous studies have shown significant energy savings from fine-grain DVFS by using FIVRs [77-78].

2. **FIVRs reduce I^2R loss.** SoC supply voltages keep scaling down with the transistor feature size. Even though the peak power consumed by SoCs does not increase much due to thermal issues, the supply currents delivered to an SoC

increase as a result of the reduced supply voltages. The parasitic resistance of the power delivery path, including the PCB trace and the package C4 bump, creates I^2R loss that increases with the supply currents. In the off-chip power delivery solution, low voltages with high currents are delivered to the SoC, creating relatively high I^2R loss. With FIVRs, high voltages with low currents are delivered to the SoC, and the voltage conversion is done in the SoC. This technique significantly reduces the power delivery's I^2R loss [70-74].

3. **FIVRs reduce the occupied area for power delivery as well as the complexity of the PCB.** In mobile applications, the form factor of the device places stringent requirements on the size and weight of the electronics. Integrating FIVRs in the SoC can reduce the number of off-chip discrete components, such as inductors and capacitors. More importantly, since the voltage conversion is performed in the SoC, only a single high supply voltage needs to be delivered to the SoC, which significantly reduces the complexity of the board design.

1.2 Challenges for improving the FIVR-based power delivery

While FIVRs offer a promising solution for power delivery in modern SoCs, there remain a number of challenges for improving the performance of FIVRs and putting the fully integrated power delivery solution into practice.

1. **FIVRs need to be more area efficient.** When FIVRs are integrated in the SoC, they are typically implemented using the same fabrication process as is used for

the SoC. To save power and improve performance, most SoCs are implemented using expensive advanced technologies with 40nm or even smaller feature sizes. The silicon area occupied by an FIVR has to be small compared to the rest of the SoC, due to cost issues. Much progress has been made in improving the power density of FIVRs, defined as the maximal output power of the regulator divided by its area. For example, IBM has implemented switched-capacitor FIVRs with power densities greater than 3W/mm^2 [26, 36]. However, these converters rely on a special on-chip capacitor, a trench capacitor, which has a very high capacitor density (>200nF/mm^2). Trench capacitors are expensive and are not available in most SoC fabrication processes. We need more circuit- and/or system-level innovations to push the frontier further.

2. **FIVRs need to have high conversion efficiency and support high conversion ratios.** Typical off-chip VRs can achieve peak efficiencies of about 90% for a conversion of 3-to-1 or less [79]. The peak conversion efficiencies of FIVRs are limited to about 80% unless special fabrication techniques, such as trench capacitors, are used, as shown in **Figure 1.2**. Moreover, current FIVR research mainly focuses on converters with low conversion ratios ($\leqslant$ 3-to-1). In mobile applications, conversion ratios higher than 3-to-1 are required to deliver power from the ~3.7V battery to the SoC, which may operate under 0.9V. Since the conversion loss increases superlinearly as the conversion ratio increases, it is challenging to design high-performance FIVRs at higher conversion ratios.
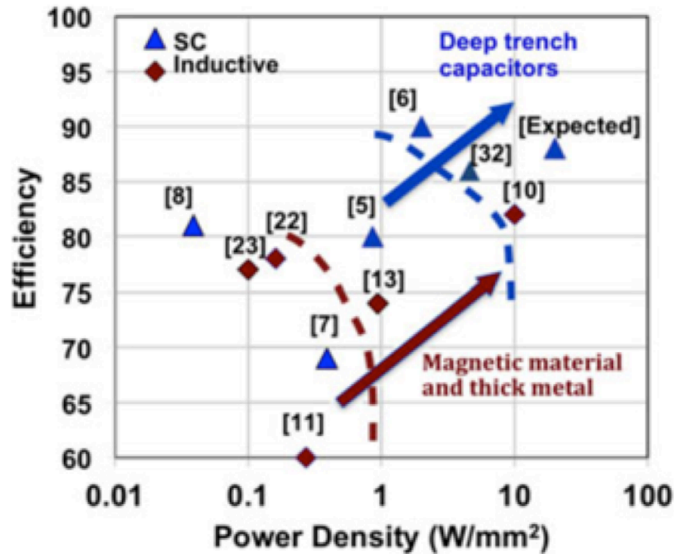
Seth et al, TPEL 2013 [9]

Figure 1.2: Performance of FIVRs in production CMOSs and in processes where extra steps are allowed.

In this thesis, I present techniques to improve the FIVR-based power delivery solution by proposing new FIVR topologies as well as new system-level architectures:

1. I provide an introduction to the basics of FIVRs, specifically, switched-capacitor (SC) fully integrated voltage regulators.

2. I present the design and implementation of a battery-connected 4-to-1 SC FIVR for a micro-robotic flying bee application. The FIVR is integrated in a SoC that works as the "brain" of the robotic bee. The FIVR provides multiple voltage outputs for different blocks in the SoC. I propose a two-stage SC topology and a charge recycling technique to improve the conversion efficiency. A single-bound feedback control loop is also implemented to regulate the output voltages. The FIVR works as expected. However, its performance is limited since high-density

6

capacitors are not available in this fabrication process, which motivates the following proposal to further improve the efficiency and power density of the FIVR.

3. I propose a single-input, multi-output (SIMO) FIVR that can simultaneously deliver power to stacked output domains with a conversion ratio of 4-to-1. Combined with voltage stacked load circuitry, the new power delivery system significantly improves the conversion efficiency and the power density compared to conventional FIVR-based solutions. The entire voltage stacking power delivery system achieves a peak efficiency of 99% and an average efficiency of 87%. The maximal output power of the SIMO FIVR is 4 times that of a conventional FIVR. In my implementation, the FIVR delivers power to 16 Intel Siskyou Peak processor cores that are stacked in 4 voltage domains and implemented on the same chip as the FIVR.

Chapter 2

Switched-Capacitor Fully Integrated Voltage Regulators: Basics and Recent Developments

2.1 Basic FIVR topologies

Switching and linear VRs are two widely used FIVR topologies. **Figure 2.1** shows a high-level block diagram for a typical linear VR. The output (V_{OUT}) is regulated from the input (V_{IN}) using a variable resistor whose value is dynamically changed through a feedback loop that constantly compares the output voltage to a reference voltage (V_{REF}). In practice, the variable resistor is usually implemented with an MOS transistor whose resistance is controlled by varying its gate voltage [49-52]. An output-filtering capacitor (C_{FLY}) is used to reduce the high-frequency output voltage noise. The current source represents the load circuitry, such as a processor. Linear VRs offer several advantages as FIVR solutions: 1) ease of integration in the SoC, 2) relatively small area overhead, and 3) fast load transient response. However, linear VRs have a very low conversion efficiency at a high step-down conversion ratio. As shown in **Figure 2.1**, the same amount of current flows through the VR and the load circuitry, which creates a very large I^2R loss in the resistor. The maximal conversion efficiency of a linear VR is limited by the ratio of V_{OUT} to V_{IN}. For example, if a linear regulator converts a 3.6V input voltage down to an output voltage of 0.9V, its maximal conversion efficiency is as low as 25%.

Switching regulators can be categorized into switched-inductor regulators and switched-capacitor regulators. Figure 2.2 shows examples of a switched-inductor regulator and a switched-capacitor regulator. In the switched-inductor regulator, by turning on one of the two MOS switches, the converter creates a square waveform on
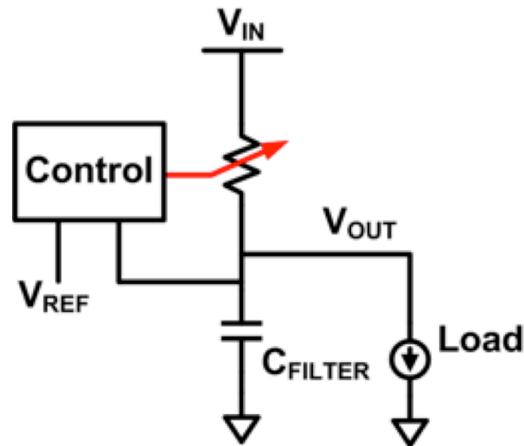
9

Figure 2.1: Linear voltage regulator

V_X. A control loop modifies the on/off time of the two switches to change the duty cycle of the square wave. The inductor (L) and the capacitor (C) work together as a low-pass filter to generate an output voltage (V_{OUT}) whose value is similar to the average voltage of V_X. Switched-inductor converters can regulate a wide range of output voltage levels across a wide range of output currents. However, switched-inductor converters require high-quality inductors to achieve high conversion efficiency—and these are hard to integrate on-chip in commercially available CMOS processes. This challenge makes switched-inductor converters less attractive as FIVR solutions.

SC converters rely on capacitors as passive components. Figure 2.2(b) shows an SC converter that implements a conversion ratio of 2-to-1. By turning on a pair of the switches in a non-overlapping two-phase (Φ1 and Φ2) pattern, the flying capacitor (C_{FLY}) is periodically charged and discharged while the energy is transferred from the input to the output. A control loop (not shown in the figure, for simplicity) compares the output
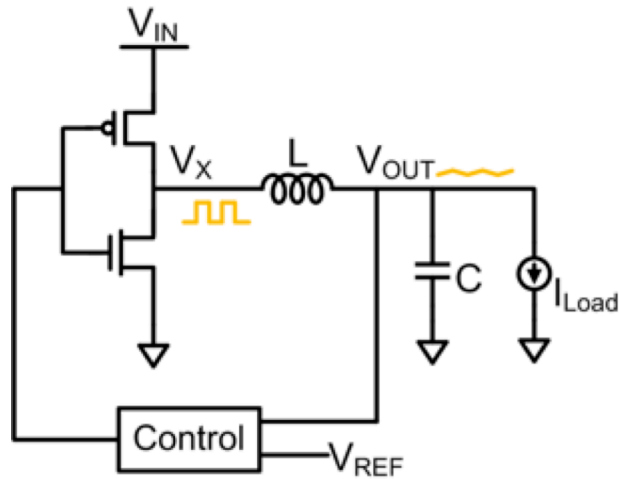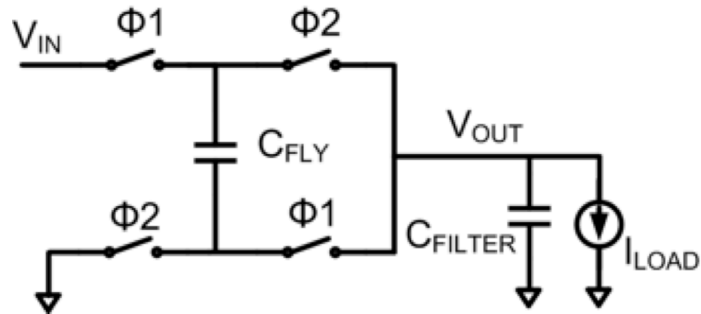
Figure 2.2: (a) Switched-inductor converter and (b) switched-capacitor converter

voltage (V_{OUT}) to a reference and modifies the switching frequency of the switches as needed. By changing the switching frequency, the SC converter changes the amount of power that is delivered from the input to the output. The switches are usually implemented with MOS transistors. An output-filtering capacitor (C_{FILTER}) reduces the output voltage noise. The operations of SC converters are discussed in more detail in Section 2.2. It is much easier to fabricate high-quality on-chip capacitors [9, 10, 27, 28] than inductors. As a result, SC converters are much easier to fully integrate on-chip
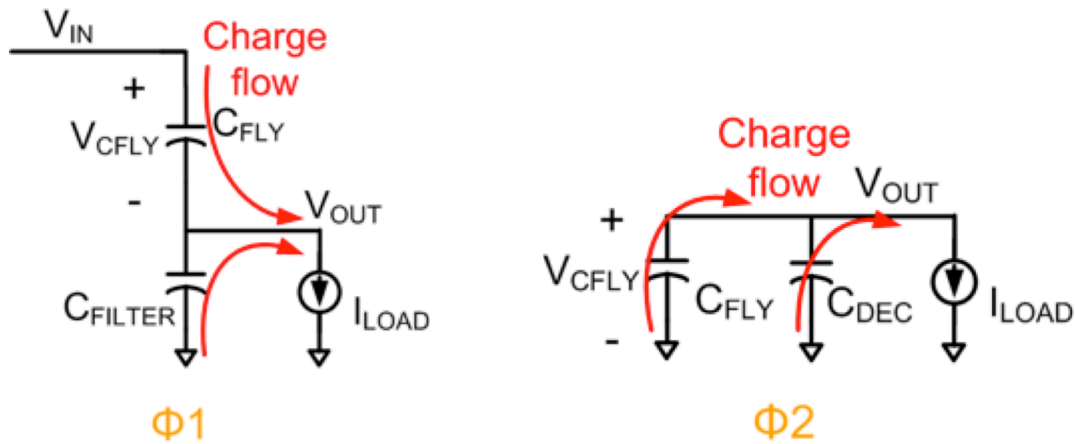
Figure 2.3: Operations of a 2-to-1 SC converter

in SoCs. SC converters also achieve relatively high efficiency (e.g., 90%) at medium

conversion ratios (e.g., 2-to-1). One common concern regarding SC converters is the

challenge of achieving high conversion efficiency across a wide range of conversion

ratios. However, a number of studies [35-38, 41] have proposed reconfigurable

topologies to provide high efficiency across a wide range of output voltages. Overall, SC

converters are promising solutions for FIVRs, compared to linear regulators and

switched-inductor converters.

2.2 Operations and losses in SC converters

2.2.1 Basic operations of SC converters

In this section, I use a 2-to-1 SC converter as an example to discuss the operations

of SC converters and explain how power is delivered from the input to the output. As
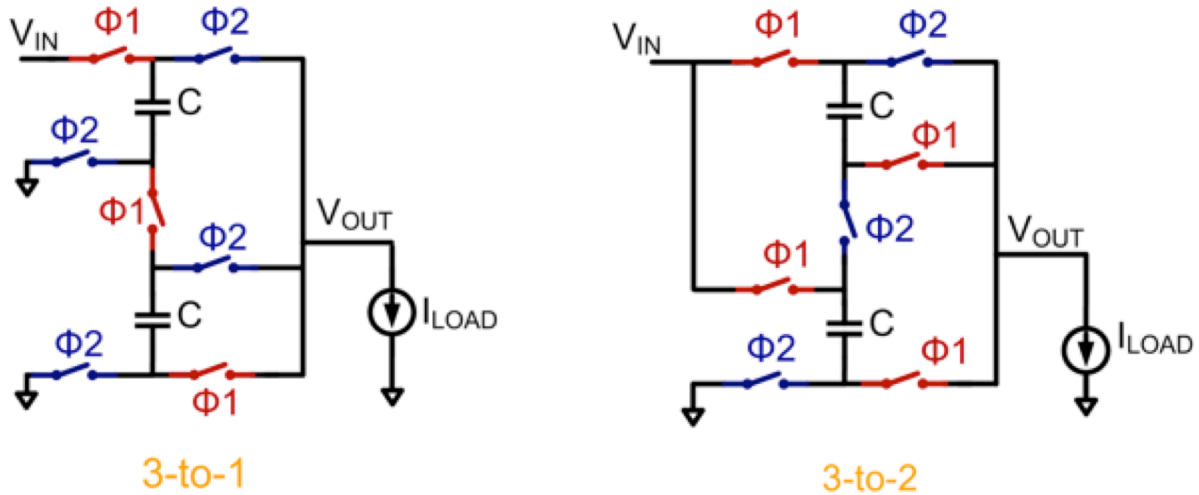
Figure 2.4: Examples of SC converters with different conversion ratios

shown in Figure 2.2, an SC converter operates in two non-overlapping phases. The details of the operation are illustrated in Figure 2.3. During Φ1, C_{FLY} is connected in series with C_{FILTER}. Charges drawn from the input charge the C_{FLY} up and flow to the load. During Φ2, C_{FLY} is connected in parallel with C_{FILTER}. The charges previously stored on C_{FLY} in Φ1 flow to the load. The SC converter periodically switches between Φ1 and Φ2, delivering power from the input to the output.

By connecting more switches and flying capacitors in different patterns, SC converters can be made to achieve different conversion ratios. Figure 2.4 shows two different SC converters implementing different conversion ratios (3-to-1 and 3-to-2). Analyzing how charges flow in an SC converter is a typical method for understanding the behavior of the converter, such as the conversion ratio and the conversion loss. Examples and details of the charge-flow analysis have been well discussed in [42, 43].

2.2.2 Conversion losses in SC converters

Non-idealities in SC converters cause conversion losses [25, 32, 42, 43]. The power switches have non-zero on-resistance and parasitic capacitance. The flying capacitors also have parasitic capacitance.

The conversion loss can be categorized into conductive loss and switching loss. When charges flow through the flying capacitors, the flying capacitors are either charged or discharged, creating ripple on the flying capacitors and resulting in conductive loss. The conductive loss caused by the flying capacitors is sometimes referred to as intrinsic loss or charge redistribution loss. In addition to the capacitor conductive loss, the non-zero on-resistance of the power switches also results in conductive losses when currents flow through them.

The other major loss in SC converters comes from switching the parasitic capacitance of the flying capacitors and the power switches. This switching loss is a function of the parasitic capacitance, the capacitance switching-voltage swing, and the switching frequency of the converter.

More details about the conversion losses and the optimization process for SC converters can be found in [25, 32, 42, 43].

2.3 Recent developments in SC FIVRs

During the past decade, much effort has been devoted to building fully integrated SC converters. In this section, I review some of the interesting work that has been done.

2.3.1 Capacitor fabrication techniques

The flying capacitor is one of the most important components in an SC converter. The performance of an SC converter heavily depends on the quality of the flying capacitor. For FIVRs, a good integrated capacitor should have high capacitor density and low parasitic capacitance/resistance.

In many SC converters, MOS transistors are used to implement the flying capacitors. Although it is technology dependent, the capacitor density of an MOS capacitor in an advanced technology is usually in the range of 5-10nF/mm^2. MOS capacitors also have a relatively high parasitic capacitance that is around 2% to 4% of the intrinsic capacitance.

IBM developed a deep trench capacitor technology [27], as shown in Figure 2.5(a). The vertical structure of this capacitor enables a very high capacitor density and a very small parasitic capacitance. This trench capacitor can achieve a 25 times better capacitor density with a 10 times better parasitic capacitance than MOS capacitors.

Intel designed an SC FIVR based on Metal-Insulator-Metal (MIM) capacitor technology [28]. As shown in Figure 2.5(b), the high density MIM capacitor is implemented using metal layers 8 and 9. The advantage of this MIM capacitor is that it can be placed on top of the load circuitry so that it does not consume extra chip area.
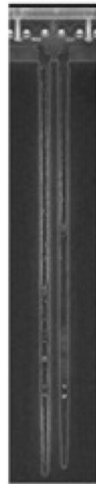
(a) J. Seo, et al., PowerSoC 2012 [27]   (b) R. Jain, et al., JSSC 2014 [28]
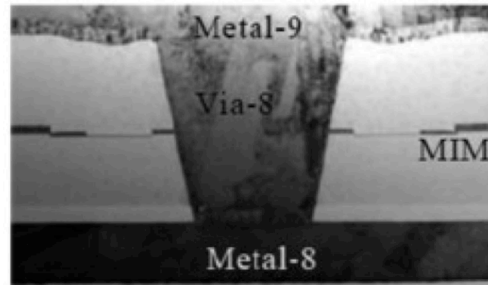
Figure 2.5: (a) Trench capacitor designed by IBM and (b) MIM capacitor designed by Intel

2.3.2 Reconfigurable SC topologies

SC converters with a single conversion ratio cannot efficiently cover a wide range of output voltages. In contrast, reconfigurable SC converters can dynamically change their topology and conversion ratio according to the input and output voltage conditions.

Figure 2.6 shows an example of a reconfigurable SC converter. It operates with two non-overlapping clock phases, similar to other SC converters discussed earlier. By choosing the appropriate switching pattern for all the power switches, the converter can be dynamically reconfigured to achieve different conversion ratios. In each configuration, some power switches are in the "Off" state, which means that these switches are always turned off in that configuration.
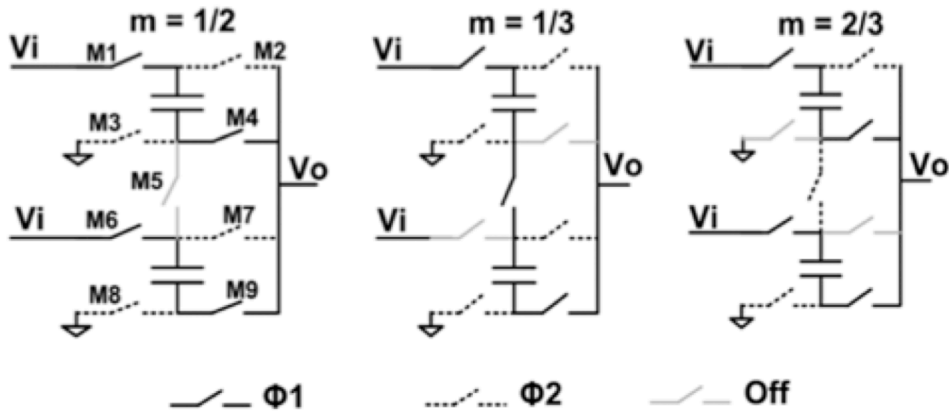
Figure 2.6: A reconfigurable SC converter (*m* is the conversion ratio)

2.3.3 Feedback control loops with nanosecond response time

An SC converter requires a feedback control loop to regulate the output voltages to the appropriate levels when the load current fluctuates. A typical control loop compares the output voltage to a reference voltage and dynamically adjusts the behaviors (e.g., the switching frequency) of the converter.

Figure 2.7 shows a block diagram of a single-bound feedback control loop that is widely used in SC implementations [21, 26, 27, 32, 33]. A digital comparator that is clocked by a ring oscillator compares the output voltage (V_{OUT}) to the reference voltage ($V_{CONTROL}$). If V_{OUT} is larger than $V_{CONTROL}$, the output of the comparator (V_{COMP}) remains low. Whenever V_{OUT} is detected to fall below $V_{CONTROL}$, the comparator generates a pulse on V_{COMP}. A latch creates a switching signal (V_{SWITCH}) from V_{COMP}. V_{SWITCH} is further processed by a non-overlapping phase generator and eventually controls the
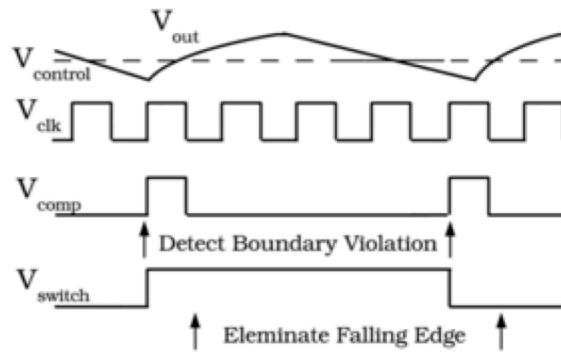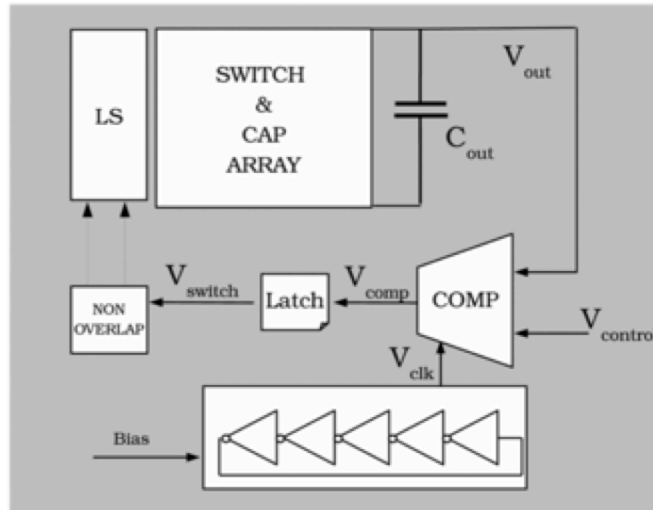
T. Van Breussegem, et al., JSSC 2011 [32]

Figure 2.7: Single-bound feedback control loop in an SC converter

power switches in the SC converter. In the single-bound control scheme, the SC converter does not switch at a known frequency. The converter only switches when the output voltage falls below the reference voltage. In other words, the switching frequency of the converter can change from a very low frequency to the highest frequency very quickly. If operated under a multi-GHz clock, the single-bound feedback control scheme can achieve a response time in the range of nanoseconds [26, 28, 34], reducing the voltage droop caused by the large current step. However, the pulse-skipping nature of

the feedback loop and the non-zero feedback delay also result in large voltage ripples [21].

Figure 2.8 shows an SC converter that implements a pulse-frequency modulation control loop. The control loop consists of a main loop (in black) and an auxiliary loop (in red). The main loop slowly changes the switching frequency of the converter to handle the low-frequency load current fluctuations. The slow main loop results in a large voltage droop when the load current increases from low to high very quickly. The auxiliary loop can bypass the main loop and set the switching frequency of the converter to its highest frequency to reduce voltage droop in the case of a large and fast current-increasing event.

In the main loop, the clocked comparator (Comp1) compares the output voltage (V_O) to the reference (V_{ref}). Based on this comparison, the comparator controls the charge pump integrator to increase or decrease the supply voltage of the VCO in order to gradually change the switching frequency of the SC converter. The comparator (Comp2) in the auxiliary loop uses a different reference voltage (V_{r_low}) that is 30mV lower than V_{ref}. If V_O falls below V_{r_low}, the auxiliary loop will bypass the main loop, using S1, S2, and S3 to speed up the transient response.
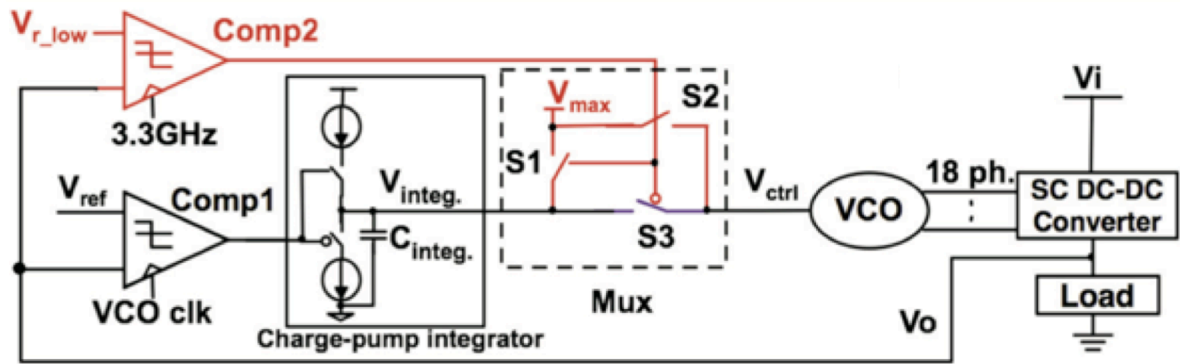
H. Le, et al., ISSCC 2013 [34]

Figure 2.8: Pulse-frequency modulation control in an SC converter

Chapter 3

A Battery-Connected 4-to-1 Fully Integrated SC Converter for Micro-Robotic Bee Applications

3.1 Introduction to Harvard's Robobee and its BrainSoC
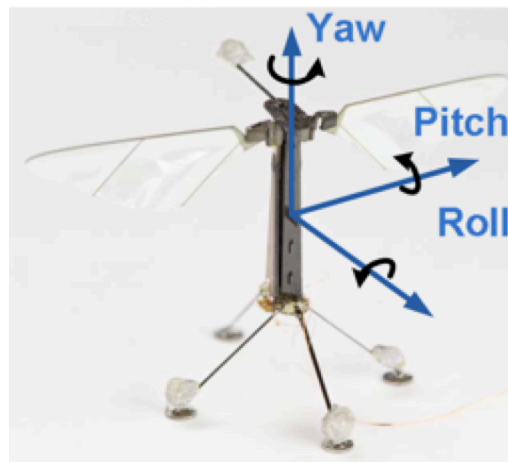
X. Zhang, et al., VLSI 2015 [66]

Figure 3.1: Picture of Harvard's Robobee

Researchers at Harvard University have designed and manufactured an insect-scale flapping wing robot (Harvard's "RoboBee"). Figure 3.1 shows a picture of the robot. Through control of its wings, the robot can be made to perform different movements, including "Yaw", "Pitch", and "Roll", as discussed in [65-67]. Recent research has proved controlled flight—hovering and maneuvering along three axes—that relies on an external motion-capture system, a benchtop high-voltage amplifier to energize piezoelectric (PZT) actuators that flap its wings, and a computer for computation [65].

The ultimate goal of this project is to achieve autonomous flight, which requires replacing all the bulky bench machines with customized components. These customized components have to be small and light enough for the robotic bee to be able to carry them within its extremely tight payload budget.

Figure 3.2: The robotic multiple-chip module consists of the BrainSoC and the power electronic unit (Power IC + tapered-inductor boost converter).

Figure 3.2 presents the electronic components of the robotic bee. An energy-efficient SoC, the BrainSoC, is the central controller of the robot. It processes sensor data and sends wing-flapping control signals to a power electronics unit that generates 200-300V sinusoids to drive a pair of piezoelectric actuators that individually flap each wing [66]. More details about the electronics are available in [66].

The integrated circuits in the BrainSoC operate at two voltages. The digital circuit blocks, such as the ARM Cortex M0 processor core and the memory, work at 0.6V to 0.9V. The mixed-signal blocks, such as the analog-to-digital converter (ADC), operate at around 1.8V.

A Li-ion battery with a voltage around 3.7V is the only source of energy on the bee. A DC-DC voltage regulator is thus necessary to convert the high battery voltage down to

lower voltages and deliver the energy to the BrainSoC. Considering the stringent weight and area requirements, the DC-DC converter must be integrated in the same SoC without using any external components.

3.2 A two-stage 4-to-1 SC converter topology

3.2.1 Different topology options

Among the various options for building an FIVR, linear regulators offer ease of integration and low area overhead. However, they have low efficiencies (~25%) for a high step-down ratio of about 4-to-1 (from ~3.7V to ~0.8V). Many voltage regulators rely on switched-capacitor (SC) converters or switched-inductor converters. The efficiency of a switching converter heavily depends on the quality of the reactive components, inductors, and/or capacitors. It is easier to integrate high-quality on-chip capacitors than on-chip inductors. In the 40nm digital fabrication process that we use to design the converter and the BrainSoC, on-chip MOS capacitors with capacitor densities as high as 10nF/mm^2 are available, making the SC converter a reasonable choice for implementing the fully integrated DC-DC voltage regulator.

Previous work has proposed fully integrated SC converters using series-parallel and ladder topologies [24-44]. However, series-parallel converters suffer from power-switch breakdown issues for voltage step-downs from 3.7V to 0.8V. Cascaded thick-oxide transistors have to be used to implement some of the power switches [25, 34]. The parasitic switching losses of the flying capacitors also increase dramatically as the

conversion ratio increases [25, 34, 29, 31] in series-parallel converters. Ladder converters avoid device breakdown problems, but typically have high equivalent output resistance for a conversion ratio of 4-to-1 [42-43].

This chapter presents a fully integrated two-stage SC regulator. The two-stage topology simplifies the overall converter design and provides the opportunity to optimize the two stages separately. Each stage uses the appropriate flavor of transistors (thin-oxide and thick-oxide transistors) and has different switching frequencies to reduce the conversion loss and improve the load regulation. The design also incorporates a charge recycling technique to mitigate the parasitic switching loss of the flying capacitors. Two separate low-bound feedback control loops regulate each stage's output to the desired levels. Finally, the two-stage topology provides an intermediate voltage (~1.8V) for use by other parts of the micro-robotic bee, such as the image sensors.

3.2.2 A two-stage topology

Figure 3.3 shows a system block diagram for the proposed two-stage SC converter. The design cascades two 2-to-1 SC stages to achieve a conversion ratio of 4-to-1. Each stage is implemented and optimized for different purposes. The first stage connects directly to the battery and converts the 3.7V high voltage down to a 1.8V intermediate voltage (V_{INT}). To handle the 1.8V swing, this stage uses the thick-oxide transistors that are available in the process. The second stage converts the intermediate 1.8V down to ~0.8V for the final output (V_{OUT}), using thin-oxide transistors. Each stage also includes identical, but separate, feedback control loops, which will be discussed below.
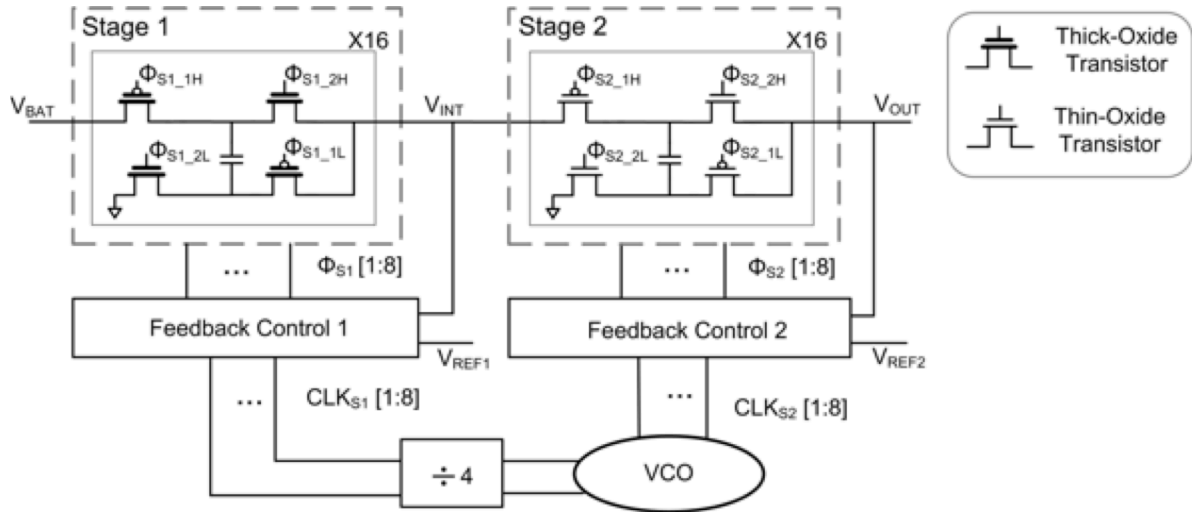
Figure 3.3: Block diagram for the proposed two-stage SC converter

The two SC stages are nearly identical, except for the types of transistors and the sizing. Each SC stage implements a multi-phase topology to reduce voltage ripple. Sixteen modules operate off both edges of eight interleaved clock phases. A multi-phase voltage-controlled oscillator (VCO) generates the clock edges and operates directly off the battery to guarantee proper start-up. To ensure that there is always a balanced number of modules in operation, pairs of modules operate 180° out-of-phase off one shared clock phase. SC converters have two basic phases of operation, which were thoroughly described in Chapter 2. In one phase, energy drawn from the input charges the flying capacitor up and flows to the load. In the other phase, energy stored on the capacitor during the previous phase flows to the load. The power switches operate with stacked voltage domains similar to those in [32]. Taking the first stage as an example, switches driven by Φ_{S1_1H} and Φ_{S1_2H} operate in the high voltage domain (between V_{INT} and V_{BAT}), while switches driven by Φ_{S1_1L} and Φ_{S1_2L} operate in the low

voltage domain (between ground and V_{INT}). Level shifters are implemented to shift the power-switch driving signals to the high voltage domains when necessary.

The two-stage topology provides an opportunity to optimize the two stages separately. The maximal switching frequency of the first stage is designed to be one quarter of that in the second stage. From the perspective of conversion loss, the first stage has a larger voltage swing and uses transistors with higher parasitic capacitance. Switching the first stage at a slow frequency reduces its parasitic switching loss. From the perspective of the voltage conversion and load regulation, the first stage handles the 1.8V voltage step-down but is decoupled from output load transients. The higher switching frequency of the second stage enables a smaller feedback delay to achieve a fast output load transient response. Further details about the optimization are presented in Section 3.2.3.

Cascading two 2-to-1 SC stages offers other advantages. V_{INT} and V_{OUT} can serve as stacked supply voltages for the switch drivers in each stage, so that an additional voltage rail is not required. Moreover, the output of the first stage, V_{INT}, also works as the supply voltage for the ADC in the SoC and off-chip sensors in the robotic bee.

3.2.3 Losses and optimizations of the two-stage converter

When implemented in advanced technologies, the major losses in fully integrated SC converters are usually the conductive loss and the switching loss. The conductive loss is due to the charge redistribution through the flying capacitors and power switches. The switching loss is due to switching the parasitic capacitance of the flying capacitors and

power switches. Other loss mechanisms, such as the digital feedback controller's and VCO's power consumption, are usually much smaller than the conductive and switching losses, especially when the output power of the converter is not very small (i.e., when it is <10mW) [25, 32].

The optimization of the two-stage converter follows the process in [25] and [42], which discuss how to calculate each of the loss mechanisms. Exhaustive searching is an easy way to decide the sizing of all devices and the switching frequency of each stage once the losses are quantified. Instead of going through all the equations, in this section I provide some qualitative intuitions about the optimization of the proposed two-stage structure: The first stage should switch at a relatively low frequency and the second stage should switch at a relatively high frequency to reduce losses and improve load regulation.

Consider the switching loss first. The first stage uses thick-oxide transistors as power switches. Compared to the thin-oxide power switches in the second stage, thick-oxide transistors have a higher parasitic capacitance. Moreover, the flying capacitor's and power switches' parasitic capacitor in the first stage switches at a swing of about $2V_{OUT}$, while the flying capacitor's parasitic capacitor in the second stage switches between V_{OUT} and the ground, as shown in Figure 3.4. As a result, the first stage would have much greater switching loss if it switched at the same frequency as the second stage. Intuitively, then, it makes sense to switch the first stage at a slower frequency to reduce its switching loss.
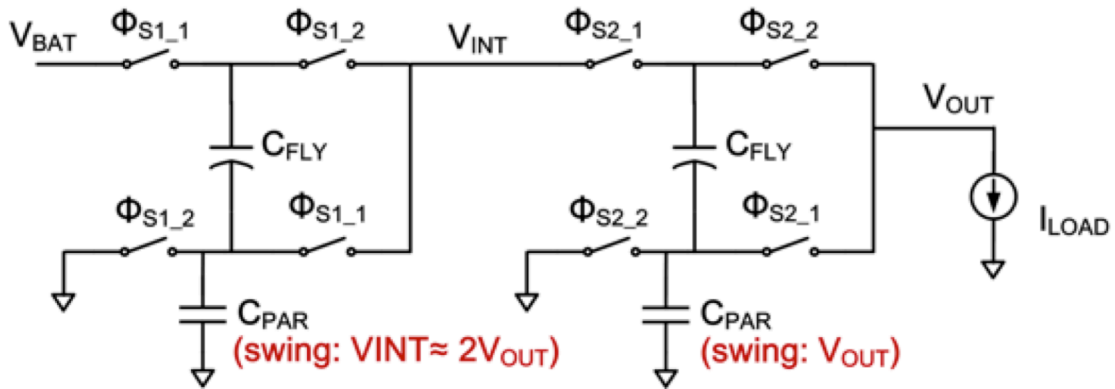
Figure 3.4: Flying capacitors' parasitic capacitors in the converter
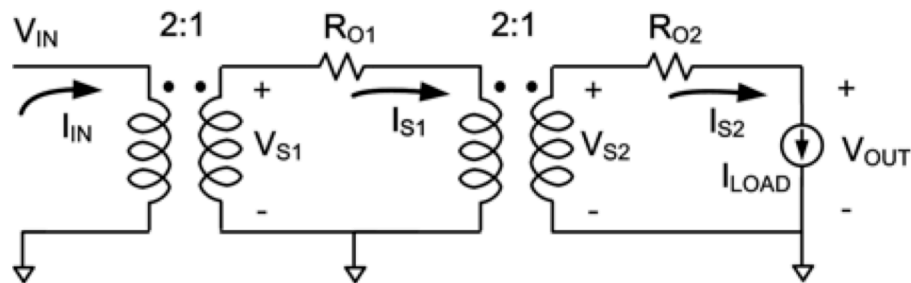
Figure 3.5: Simplified model of the two-stage converter

To better understand the conductive loss, Figure 3.5 presents a simplified model of the two-stage converter that includes ideal transformers and resistors. R_{O1} and R_{O2} are the equivalent output resistances of each stage, which capture the conductive loss. Since each stage has a voltage conversion ratio of 2:1, the current that flows through the first stage is smaller than the current flowing through the second stage. I_{S1} is only half of I_{S2} if we consider conductive loss alone. If the two stages were to contribute similar amounts of conductive loss, R_{O1} should be greater than R_{O2}. Since the equivalent output resistance of an SC converter is inversely related to its switching frequency, the first stage can switch at a lower frequency. Intuitively, switching the first

stage at a lower frequency than the second stage balances the losses from each stage and helps reduce the overall losses.

Switching the second stage at a higher frequency also improves load regulation. In SC converters, the switching frequency is usually related to the response time for handling the load current step [26, 34]. In this application, where the SC converter is fully integrated on the same chip as the load circuitry, there is not much silicon space for output decoupling capacitors. The SC converter needs to respond quickly since the load current of the digital processor changes very quickly. Because the second stage of the SC converter directly interacts with the load circuitry, switching the second stage at a higher frequency provides fast load regulation.

3.3 Implementations of the two-stage converter

This section describes the design techniques that improve the conversion efficiency along with the implementation details for some important converter components, such as the feedback control loop.

3.3.1 Flying Capacitor Parasitic Charge Recycling

Even though high-quality trench capacitors—which have both high capacitor density and low parasitic capacitance—have been used to implement flying capacitors in previous work [27, 36, 41], such high-quality capacitors are not available in many CMOS processes [34, 37-39], including the 40nm technology that we use to implement the SoC. Hence, bulk PMOS or NMOS transistors are often used to implement flying capacitors

because of their relatively higher density. MOS transistors usually have a high parasitic capacitance (~2% in this technology, ~5% in [29]). Switching this parasitic capacitance accounts for one of the dominant losses in SC converters.

In the two-stage converter design, all of the flying capacitors rely on bulk MOS transistors. The flying capacitors in the first stage are implemented with thick-oxide transistors and the flying capacitors in the second stage are implemented with thin-oxide transistors. To reduce the parasitic switching loss of the flying capacitors, each stage implements circuitry that combines two-step charging/discharging with charge recycling. Figure 3.6 illustrates the concept of the two-step charging. If a capacitor C is charged from 0 to V_{DD} by a single voltage source V_{DD}, the energy that is consumed from the voltage source is CV_{DD}^2. Adding another secondary voltage source, $V_{DD}/2$, and charging the capacitor in two steps reduces the required energy by 25%, as shown in equation (3.1). Of the required energy, only $CV_{DD}^2/2$ comes from the voltage source V_{DD}.

$$E_{TOT} = \frac{V_{DD}}{2} \cdot \frac{CV_{DD}}{2} + V_{DD} \cdot \frac{CV_{DD}}{2} = \frac{3CV_{DD}^2}{4} \tag{3.1}$$

However, this two-step charging technique cannot be directly applied to the SC converter. The second stage of the converter will serve as an example. As shown in Figure 3.4, the flying capacitor's parasitic capacitor is charged and discharged between
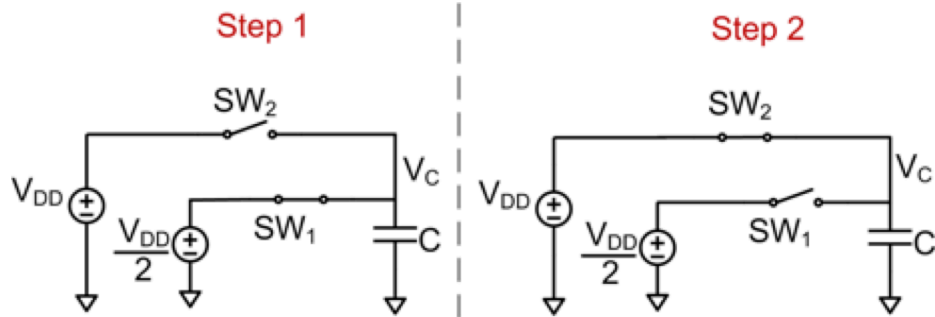
Figure 3.6: Two-step charging technique
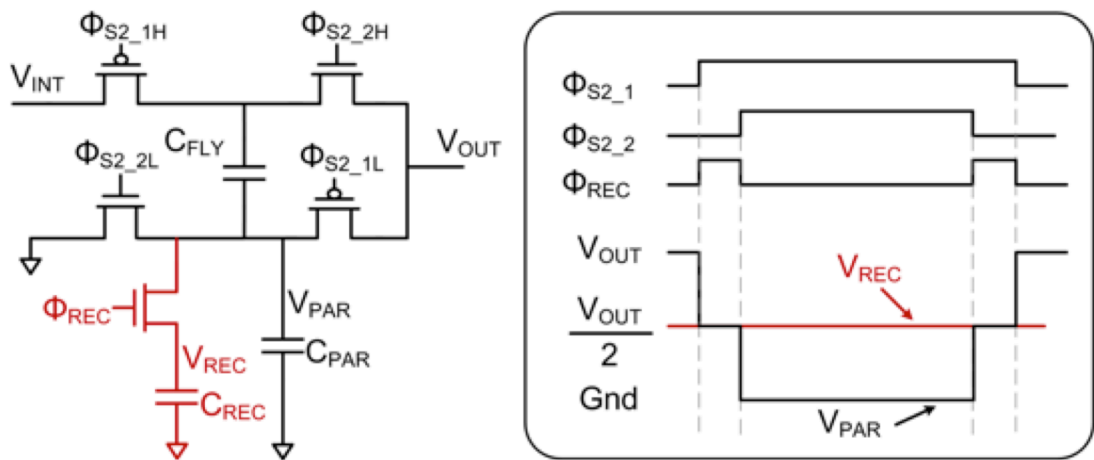
Figure 3.7: Two-step charging/discharging with inherent charge recycling, and timing diagram

gnd and V_{OUT} in every cycle. There is no additional $V_{OUT}/2$ voltage source to take advantage of the two-step charging. To solve this problem, I propose a charge recycling technique and combine it with the two-step charging technique. The new joint technique avoids the secondary voltage source and results in a 50% energy savings, more than two-step charging alone can provide.

Figure 3.7 uses the second stage of the converter as an example to illustrate the proposed technique. C_{PAR} is the parasitic capacitor of C_{FLY}. By adding an additional

recycling capacitor, C_{REC}, the proposed technique avoids using an additional voltage source. The two-step charging/discharging occurs during the converter's dead time to recycle charge, reduce losses, and improve conversion efficiency.

The charge recycling operation is as follows. Assume $C_{REC} \gg C_{PAR}$ and V_{REC} starts out at $V_{OUT}/2$. When discharging C_{PAR}, C_{PAR} first transfers charge to C_{REC} through the additional switch controlled by Φ_{REC}. In this process, C_{PAR} discharges from V_{OUT} to $V_{OUT}/2$. Then, the switch Φ_{REC} turns off and C_{PAR} fully discharges to gnd. The amount of charge transferred from C_{PAR} to C_{REC} is $C_{PAR}V_{OUT}/2$, which is stored on C_{REC} and is recycled in the charging phase. When charging C_{PAR}, C_{PAR} first charges up from gnd to $V_{OUT}/2$ via C_{REC}. In this period, C_{REC} transfers $Q=C_{PAR}V_{OUT}/2$ to C_{PAR}, which is the same amount of charge that C_{REC} gets from C_{PAR} in the discharging process. C_{PAR} then disconnects from C_{REC} and fully charges up to V_{OUT}. From an energy perspective, V_{OUT} only needs to provide $E=C_{PAR}V_{OUT}^2/2$ in this charging process, which is half of the energy otherwise required. It is important to note that V_{REC} eventually settles to $V_{OUT}/2$ regardless of its initial voltage, because this is the only balanced state where the energy stored on C_{REC} when discharging C_{PAR} matches the energy that C_{REC} loses when charging C_{PAR}.

The above recycling process assumes $C_{REC} \gg C_{PAR}$. Thanks to the converter's multi-phase operation, C_{REC} can be shared by all of the phases and C_{REC} only needs to be larger than the parasitic capacitance in one phase. In this work, C_{REC} is only 2% of the total flying capacitance.

The allocated non-overlapping times for the first and the second SC stage are 600ps and 200ps, respectively. The difference in the non-overlapping times is due to the difference in the switching frequencies of the two stages. Within the non-overlapping time, about 300ps and 100ps are used to perform charge recycling for each stage. There is a margin to avoid any short in the circuits. The RC constants related to the charge recycling process are about 120ps and 40ps for the two stages.

3.3.2 Low-Bound Feedback Control

Closed-loop operation regulates V_{OUT} and V_{INT} to desired voltage levels. The VCO generates interleaved clock signals and send them to the interleaved SC modules in both stages. The feedback control logic in each SC module operates off the interleaved clock and eventually controls the switching behavior of each SC module. To ensure there is always a balanced number of modules in operation, pairs of the interleaved modules share separate feedback paths, i.e., there are a total of eight feedback paths in the 2nd stage. All the feedback control paths are based on the same low-bound feedback control scheme, as illustrated in Figure 3.8 [30, 32]. Since the feedback topology is the same in both stages, the following illustration uses the second stage as an example. In each feedback path, two comparators operate off of complimentary clocks generated by the VCO. The comparators compare V_{OUT} with a reference voltage, V_{REF2}, on the rising and falling edges of the clock. If V_{OUT} is smaller than V_{REF2}, V_{LA} switches either from low to high or high to low, depending on its previous state. V_{LA} then propagates through to control the power switches and switch the state of the SC
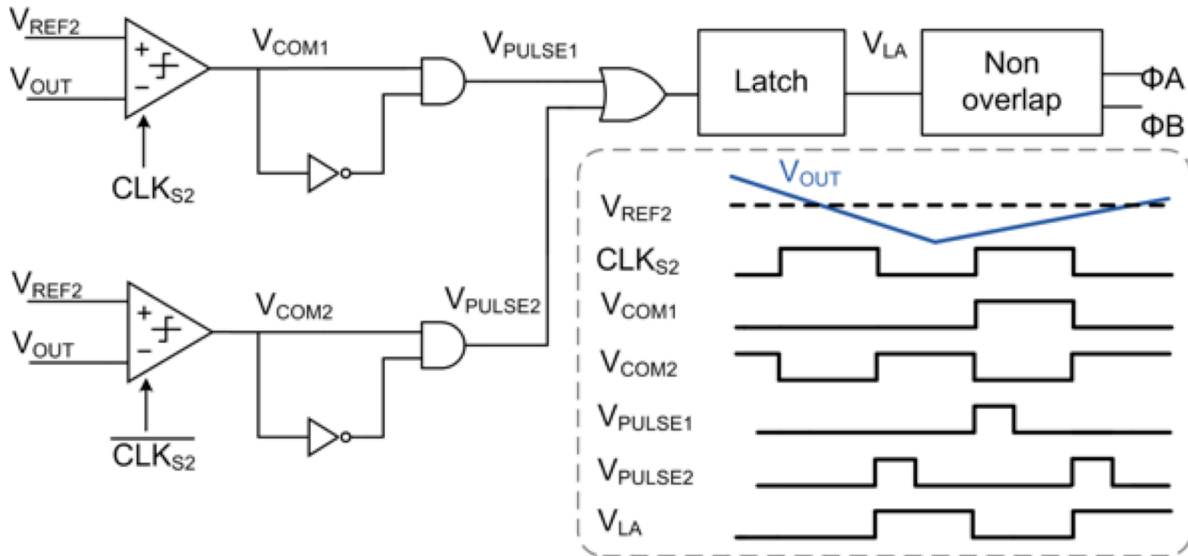
Figure 3.8: Feedback control loop diagram (second stage)

converter. This action increases the output voltage V_{OUT}. If V_{OUT} is larger than V_{REF2}, V_{LA} remains in its previous state. The power switches do not switch and V_{OUT} decreases until the SC converter reacts.

A resistor DAC (R-DAC), shown in Figure 3.9, provides separate reference voltages to the comparators in the first and the second stage of the converter via a switch network that connects each individual comparator to the resistor ladder separately. By doing do, we can use the R-DAC to calibrate comparator offsets. Calibrating comparator offsets improves steady-stage voltage ripple and conversion efficiency.

Figure 3.10 shows the circuit implementation of the digital clocked comparator used in the feedback loop. The comparators in the two stages rely on the same topology but use different transistors (thin-oxide or thick-oxide transistors) to handle different voltage swings. A strong-arm topology similar to the one presented in [69] is used to
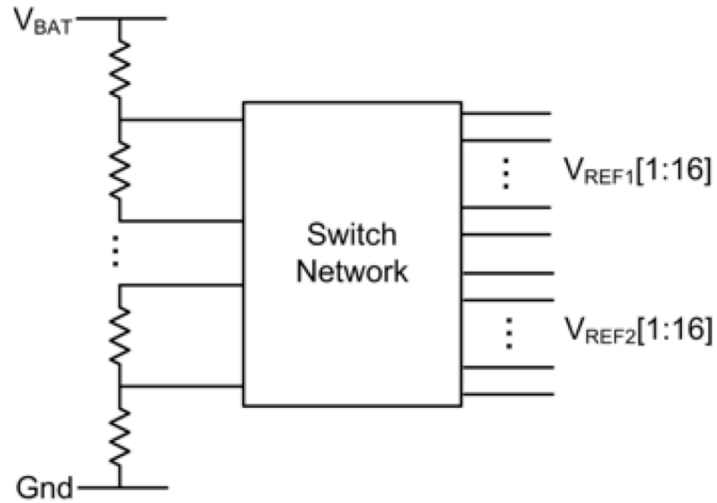
Figure 3.9: Resistor DAC
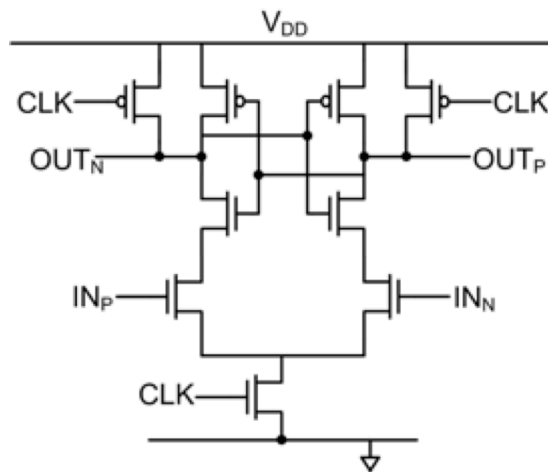
Figure 3.10: Digital clocked comparator

implement the comparators in this converter. When the clock input, CLK, is low, the comparator is in the pre-charge phase. Both the OUT_N and OUT_P nodes are pre-charged to V_{DD}. When CLK transitions from low to high, the comparison occurs. After the comparison, one output node settles at gnd while the other settles at V_{DD}. The outputs of the comparator do not change until the next pre-charge phase.

Figure 3.11: Implementation of the level shifter
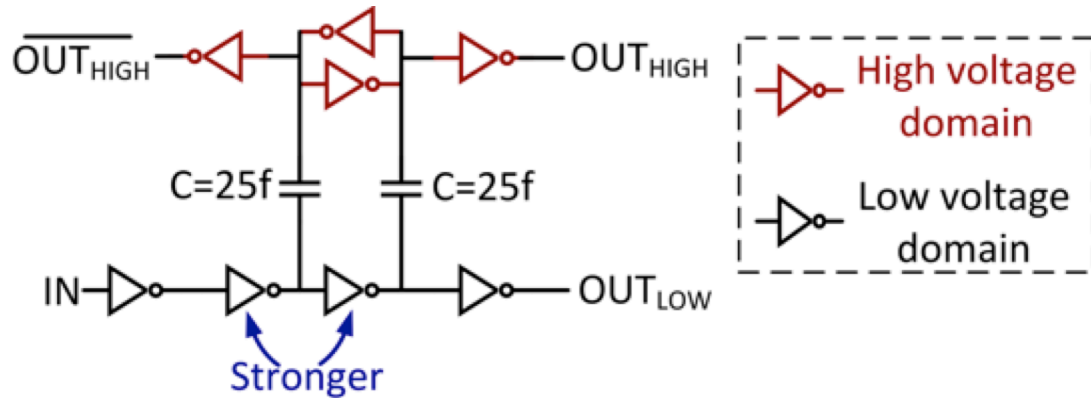
3.3.3 Other SC Converter Components

Figure 3.11 shows the implementation of the level shifter. Level shifters are required in the converter since some of the power switches need driving signals in high voltage domains, such as $V_{OUT} \sim V_{INT}$. Because these level shifters are in the switching signal path, they should have a small delay to reduce the overall feedback loop delay. I chose this capacitor-coupled level shifter since it has a delay as small as 40ps in this 40nm process. The coupling capacitors are implemented using metal-oxide-metal (MOM) capacitors. The inverters are implemented with either thin-oxide or thick-oxide transistors, depending on the voltage across the inverters. Because of the cross-couple inverter pairs in the high voltage domain, the two inverters that are driving them from the low voltage domain need to be sized larger so that they are strong enough to overwrite the top inverter pairs.
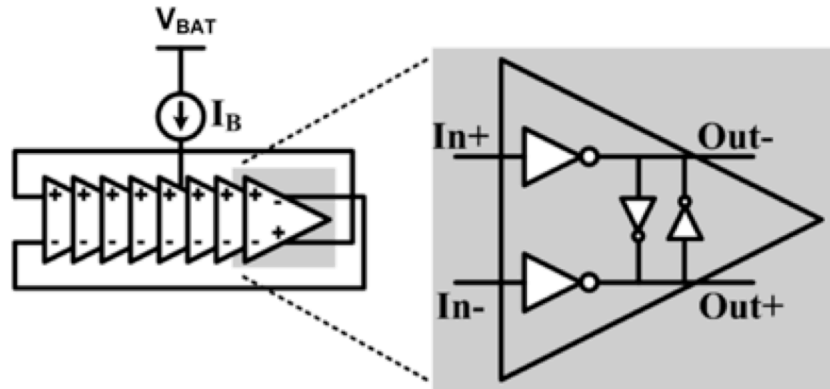
Figure 3.12: Implementation of the VCO

Figure 3.12 shows the implementation of the multi-phase current-starved pseudo-differential VCO that generates the clock edges for the feedback controller. The VCO operates directly off the battery to guarantee proper start-up.

3.4 Measurement results for the two-stage converter

The two-stage SC converter was fabricated in TSMC's 40nm CMOS technology. The chip was tested in two modes: open- and closed-loop operation. In open-loop operation, the output voltage and output power can be tuned by changing the switching frequency，F_{sw}, of the converter via the VCO. And the first stage switching frequency is set to be one quarter of the second stage switching frequency, F_{SW}. In closed-loop operation, the VCO frequency is set to its maximum and the feedback control loop adjusts the effective switching frequency of the converter to regulate the output.
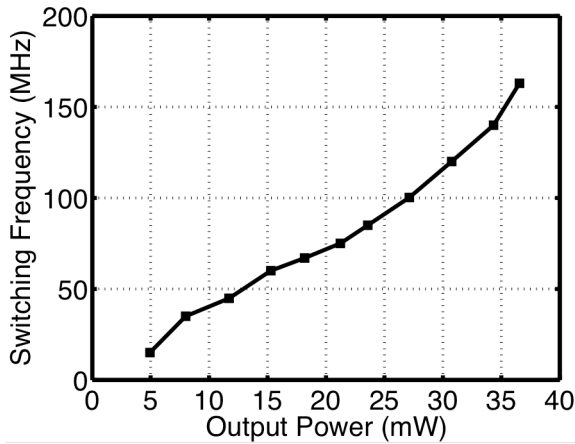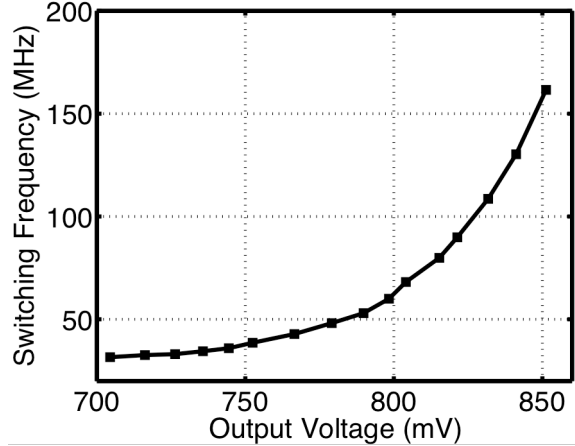
(a)   (b)

Figure 3.13: Open-loop Fsw with V_{BAT}=3.8V for (a) different P_{OUT} @V_{OUT}=~800mV and (b) different V_{OUT} @I_{OUT}=~19mA

In open-loop operation, there is a relationship between the switching frequency and the output voltage/power. Shown in Figure 3.13(a), higher output power requires high switching frequency to deliver energy more frequently. However, when switching frequency increases, there is less time for the switched capacitor circuit to settle in each cycle. Because of this incomplete charge transfer, the energy that is delivered from input to output in each cycle decreases as switching frequency increases. Hence, switching frequency increases super linearly with output power. Switching frequency, and thus switching loss, increases faster than the delivered power. Figure 3.13(b) shows that higher output voltages also require higher switching frequencies. As the output voltage increases, there is less energy that can be delivered from input to output in each cycle [25]. So, switching frequency and switching loss increase faster than V_{OUT} increases.
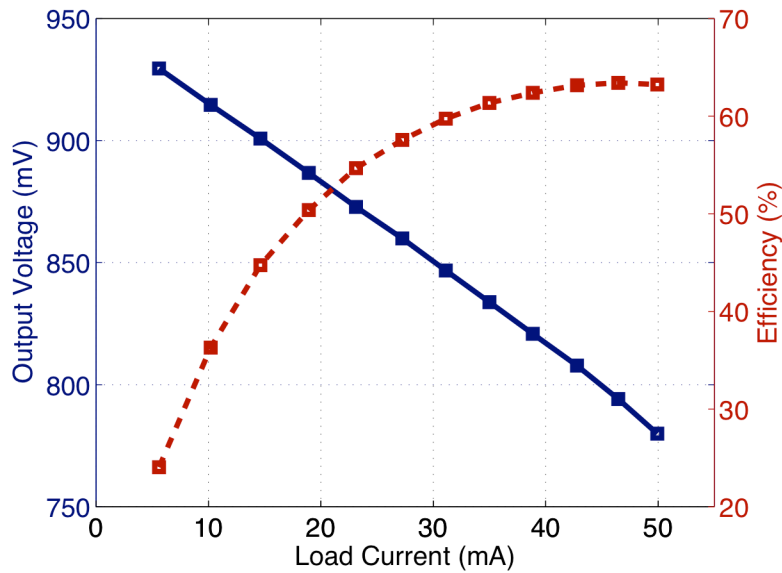
Figure 3.14: Open-loop operation w/ Fsw=160MHz, V_{BAT}=3.8V

Because the SC converter's ability to deliver power depends on its switching frequency, Figure 3.14 shows the maximal amount of power that the two-stage converter can deliver at a peak switching frequency of 160MHz and V_{BAT} of 3.8V. When the load current is very small, the output voltage is close to the ideal output voltage of 950mV. As the load current increases, the output voltage decreases linearly since the voltage drop on the equivalent output resistance of the converter increases. In other words, more current can be delivered to the output of the converter as output voltage decreases because in every switching cycle, the amount of charge delivered to the output increases as the output voltage decreases.

The conversion efficiency is very low at a low output current because of the very high parasitic switching loss. Since the switching frequency remains at 160MHz, the

switching loss does not change much with the load current. However, more power is delivered to the output as the load current increases. So conversion efficiency increases as the load current increases.

The following subsections present other experimentally measured results. First, Section 3.4.1 compares the steady-state voltage ripple for open- and closed-loop modes of operation. Next, Section 3.4.2 presents conversion efficiency results versus V_{OUT} and P_{OUT}. Then the transient responses in the open- and closed-loop modes of operation are discussed in Section 3.4.3. Finally, Section 3.4.4 provides a summary of the test chip characteristics and compares it to prior work.

3.4.1 Voltage ripple

The box plots in Figure 3.15 compare the measured steady-state output voltage ripple across a range of output power conditions for the SC converter in open- and closed-loop operation. In open-lop operation, we manually tuned the VCO frequency to keep V_{OUT} at ~800mV for each power level. In closed-loop operation, the feedback loop keeps the output voltage at ~800mV. Steady-state ripple in open-loop operation is small (~10mV) due to the interleaved design with constant switching frequency. In contrast, closed-loop ripple is generally higher due to the cycle-skipping nature of the feedback topology. In each cycle, the feedback controller must determine whether the converter should switch or not. As a result, the instantaneous switching frequency can vary widely from cycle to cycle. Delay through the feedback loop further exacerbates the ripple, because the control loop must react to the output decreasing below the reference
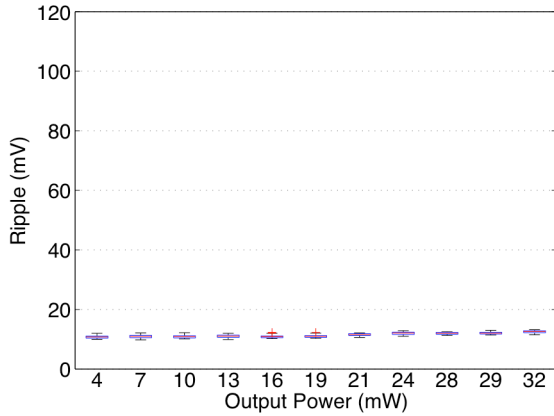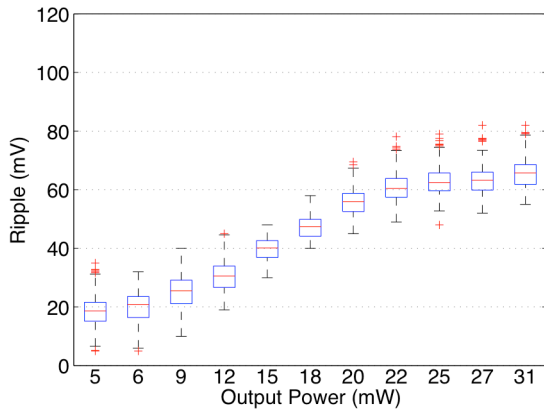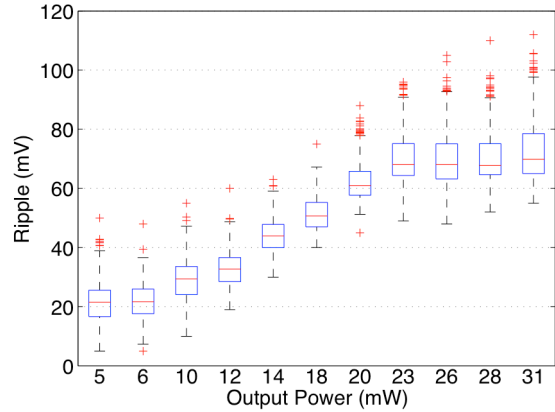
Figure 3.15: Measured output voltage ripple @ V_{BAT}=3.8V, V_{OUT_AVE}= ~800mV in (a) open-loop operation, (b) closed-loop operation with calibrated comparators, and (c) closed-loop operation with uncalibrated comparators

voltage. The longer the feedback delay is, the larger the ripple is. Measurement results show that closed-loop ripple increases with output power since larger load currents discharge the output voltage more quickly. Comparing Figure 3.15(b) and Figure 3.15(c), calibration helps to reduce voltage ripple by minimizing inconsistent switching
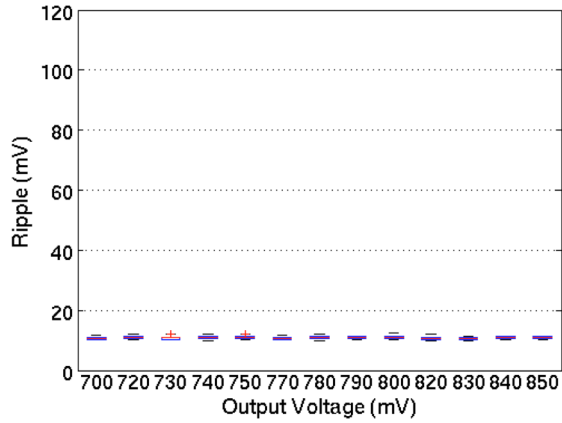
(a)
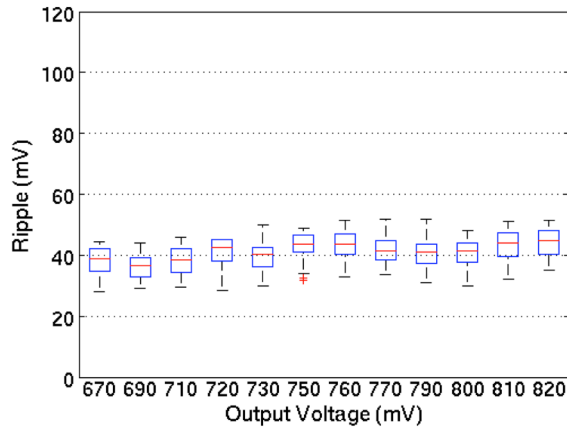
(b)
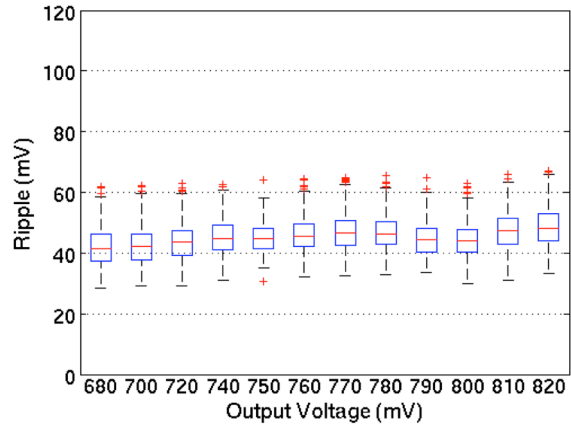
(c)

Figure 3.16: Measured output voltage ripple @ V_{BAT}=3.8V, P_{OUT_AVE}= ~15mW in (a) open-loop operation, (b) closed-loop operation with calibrated comparators, and (c) closed-loop operation with uncalibrated comparators

thresholds across all of the comparators in the multiple feedback paths. In all subsequent plots, the comparators are always calibrated unless noted otherwise.

Figure 3.16 compares the output voltage ripple across output voltages in open- and closed-loop operation. In open-loop operation, I manually tuned the VCO frequency to obtain the desired output voltages. In closed-loop operation, the output voltages are

43

regulated by the feedback loop. The voltage ripple is small in the open-loop operation because of the interleaved design. The ripple is much larger in the closed-loop operation because of the non-ideality of the feedback, as discussed earlier. Calibration helps to reduce the voltage ripple. Overall, the voltage ripple does not have a strong relation to the output voltage because it is mostly caused by the delay and the pulse-skipping nature of the feedback control loop, whose characteristics do not change much with the output voltage of the converter.

3.4.2 Conversion efficiency

In SC converters, the major sources of efficiency loss are linear charge redistribution loss, bottom-plate parasitic loss, other switching losses, and voltage ripple overhead. The minimum output voltage is used to calculate conversion efficiency, because the worst-case speed of the digital load circuits depends on the lowest transient voltage condition.

Figure 3.17 plots efficiency measurements for both open- and closed-loop operation. In Figure 3.17(a), open-loop efficiency reaches a peak of 70% at P_{OUT}=15mW. The efficiency rolls off for higher output power, because switching frequency and switching losses increase faster than the delivered power. Efficiency also rolls off for lower output power, because of static overheads. Comparing Figure 3.17(a) and Figure 3.17(b), closed-loop efficiency is generally lower than open-loop efficiency, because of larger voltage ripple. Figure 3.17 also shows that charge recycling consistently improves conversion efficiency by ~2%. Charge recycling is always on for all subsequent plots.
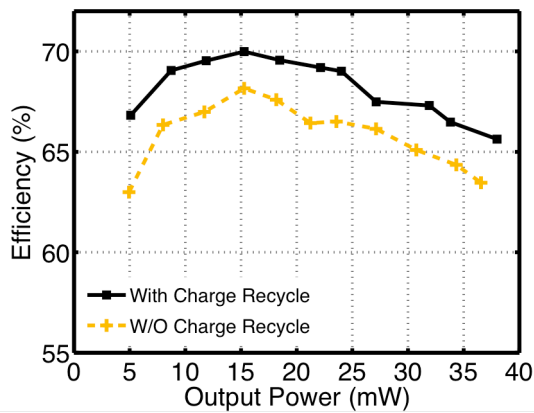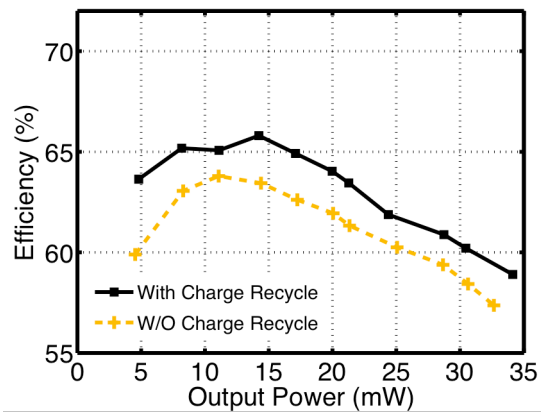
(a) open-loop operation   (b) closed-loop operation

Figure 3.17: Measured efficiency w/ V_{BAT}=3.8V & V_{OUT_MIN}=0.8V

Figure 3.18 plots conversion efficiency across different output voltage levels and exhibits the characteristic efficiency versus voltage curve of SC converters. In open-loop operation, the output voltage is set by tuning F_{SW}. In closed-loop operation, changing the reference voltage regulates the output voltage to different levels. Conversion efficiency rolls off as output voltage decreases due to the increased linear charge redistribution loss and rolls off as output voltage increases due to the higher switching loss.

Comparing the three curves in Figure 3.18, open-loop operation consistently achieves higher conversion efficiency since it has the smallest voltage ripple. Calibration improves efficiency, as expected, since it reduces voltage ripple in closed-loop operation. The efficiency in closed-loop operation peaks at a lower output voltage compared with that in open-loop operation again because of voltage ripple and because the minimum output voltage is used to calculate efficiency.

45
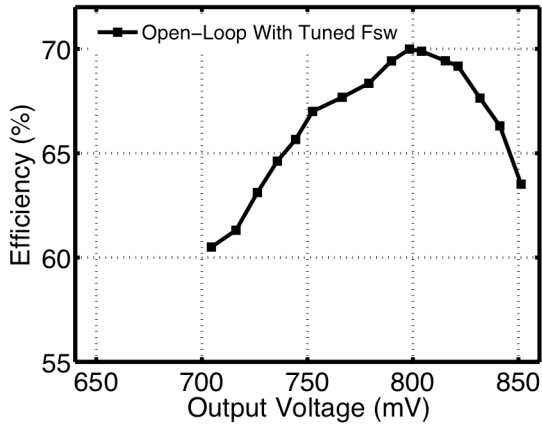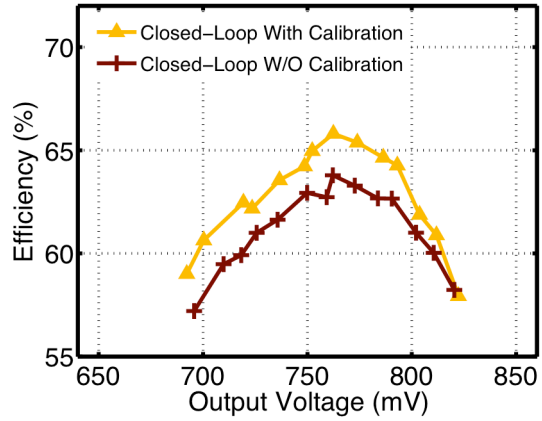
(a) open-loop operation   (b) closed-loop operation

Figure 3.18: Measured efficiency w/ V_{BAT}=3.8V & I_{OUT}=~19mA
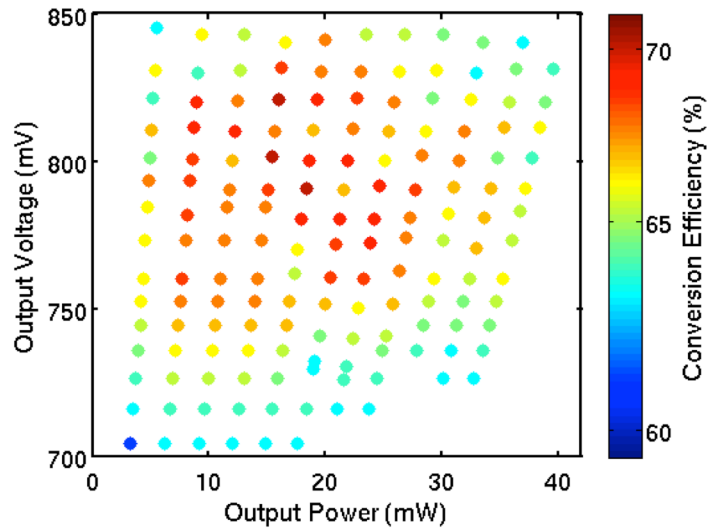
Figure 3.19: Measured open-loop efficiency w/ V_{BAT}=3.8V

The plots in Figure 3.19 and Figure 3.20 summarize the efficiency of the converter across output voltages and output powers in open- and closed-loop operation with V_{BAT}=3.8V. Generally, the efficiency in the open-loop operation is greater than that in the closed-loop operation.

46
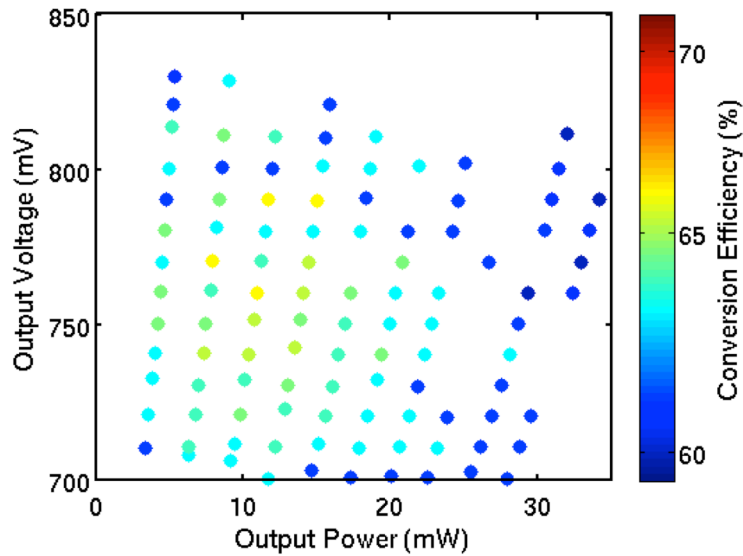
Figure 3.20: Measured closed-loop efficiency w/ V_{BAT}=3.8V

Figure 3.21 and Figure 3.22 characterize the performance of the converter at different battery voltages. Figure 3.21 shows how the output voltage changes with the load current at different battery voltages when the converter is operating in open-loop mode with its switching frequency set to the peak of 160MHz. The output voltage decreases as the load current increases because of the non-zero equivalent output resistance of the converter. The higher V_{BAT} is, the higher the load current the converter can deliver at a certain output voltage. This is because in every switching cycle, more charge can be transferred from input to output during the charge redistribution process at a higher V_{BAT}.

Figure 3.22 summarizes the conversion efficiency versus output voltages for different battery voltages (V_{BAT}). First, conversion efficiency is higher for open-loop operation, consistent with previous results presented above. Second, conversion

47
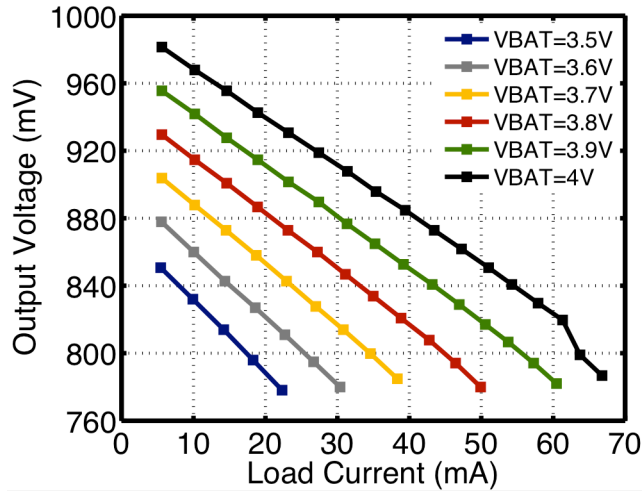
Figure 3.21: Open-loop operation across VBAT w/ Fsw=160MHz

(a) open-loop operation

(b) closed-loop operation
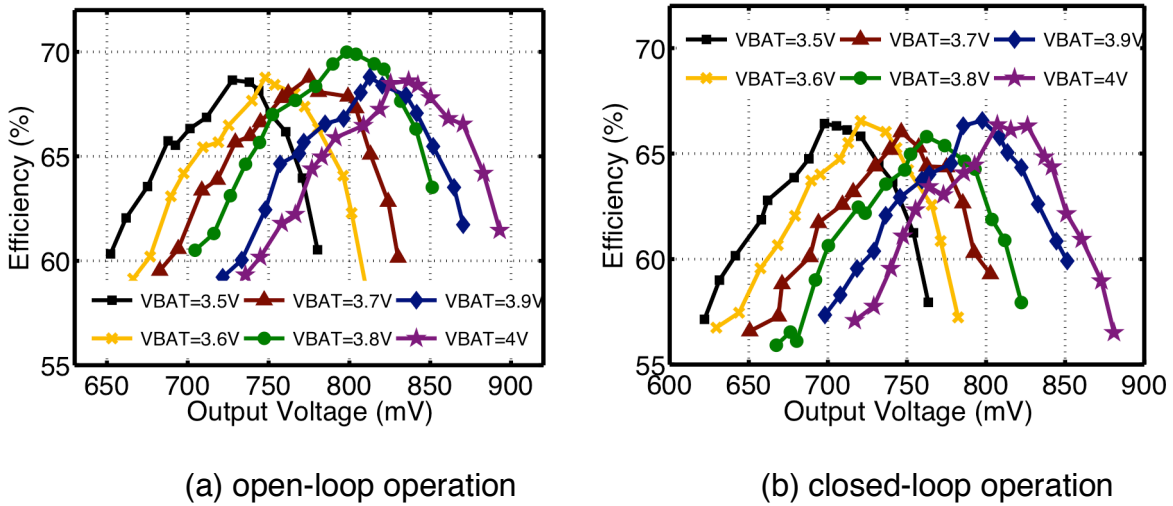
Figure 3.22: Measured efficiency with different V_{BAT} and V_{OUT_MIN}

efficiency peaks at higher output voltages when V_{BAT} is higher since the charge

redistribution loss and switching loss are both related to V_{OUT}/V_{BAT} [25, 43].
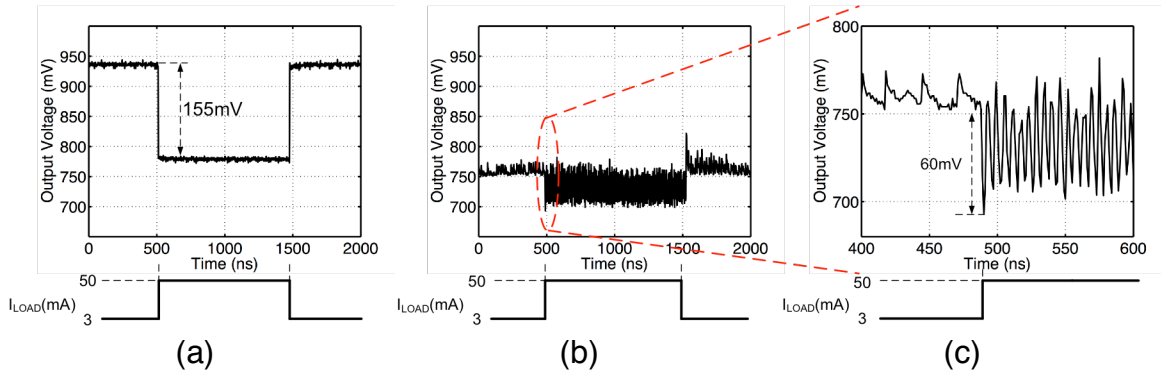
Figure 3.23: Transient response for (a) open-loop with maximum F_{SW}, (b) closed-loop, and (c) zoom-in of (b)

3.4.3 Transient response

Figure 3.23 presents the SC converter's measured response to 47mA output load transients using an on-die load circuit with rise and fall times of ~100ps. As seen in Figure 3.23 (a), when the SC converter runs in open-loop with maximum switching frequency, a 3mA to 50mA load step causes V_{OUT} to drop by 155mV. When running in closed-loop with the nominal output voltage set to 750mV, however, the control loop quickly reacts and the voltage droop caused by the load current step is much smaller. In fact, the ~60mV droop in Figure 3.23 (c) is mostly due to the larger steady-stage voltage ripple previously seen with respect to higher output power. The simulated feedback loop delay is about 1ns.

3.4.4 Test chip summary

The silicon area, shown by the micrograph in Figure 3.24, was not optimized for power density but was governed by the pads and circuitry added for testing. Flying
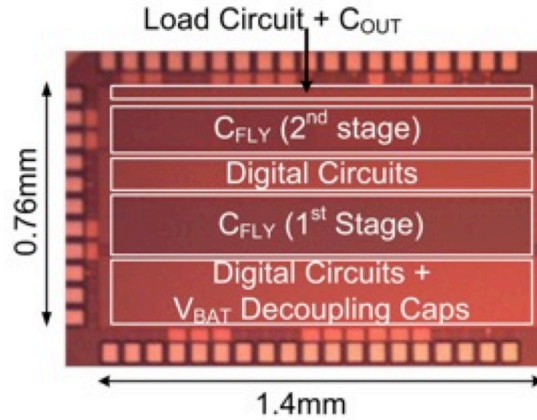
Figure 3.24: Die Photo

	[6]	[8]	[9]	[10]	[12]	This Work
Technology	90nm	65nm	32nm	22nm	250nm	40nm
Input voltage	3V-3.9V	3V-4V	1.8V	1.225V	2.5V	3.5V-4V
Total capacitance	5.2nF	3.88nF	1nF	—	3nF	2.64nF
Efficiency (η)	74%	73%	86%	71%	58%	66%
Conv. ratio @ η	2:1	3:1	2:1	2:1	4:1	4:1
Power density (mW/nF)@η	28.8 (150mW/5.2nF)	31.3 (121mW/3.88nF)	—	—	0.4 (1.2mW/3nF)	13.3 (35mW/2.64nF)
Output ripple @nom. load	—	—	30mV	60mV	—	70mV
Output droop	30mV	76mV	94mV	—	—	60mV

Figure 3.25: Comparison to prior work.

capacitors and output filter capacitors, which occupy half of the overall area, total 2.64nF. Figure 3.25 compares this work to prior art fully integrated SC converters.

Chapter 4

A Fully Integrated Reconfigurable Switched-Capacitor DC-DC Converter for Voltage Stacking Applications

4.1 Introduction to voltage stacking as a power delivery solution

(a)

(b)
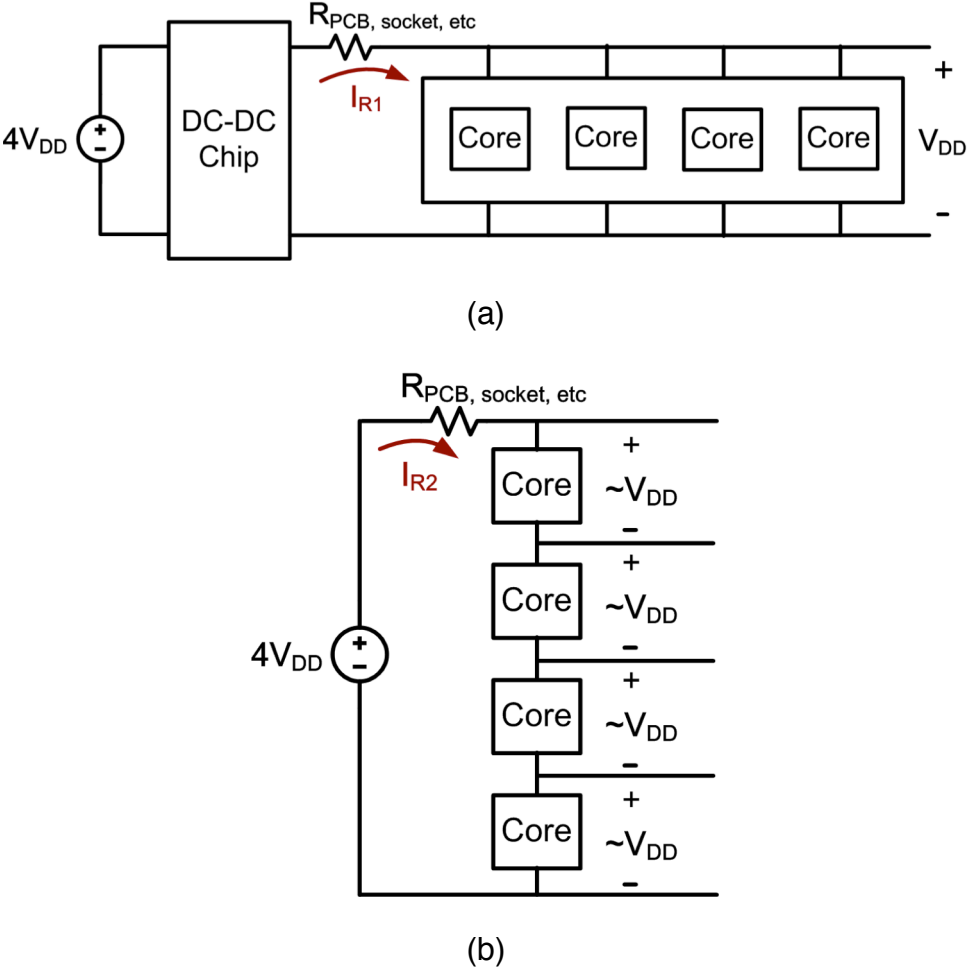
Figure 4.1: (a) Conventional power delivery; (b) Voltage stacked power delivery.

Power delivery has been a challenging issue for multicore SoC applications. The decreasing supply voltages as well as the increasing supply currents of the processors create more losses in the off-chip power delivery networks [70-75]. Figure 4.1(a) shows a diagram of a conventional power delivery network. As the fabrication technology scales down, the processor core supply voltage (V_{DD}) decreases and the current (I_{R1})

delivered to the cores increases. As a result, the off-chip resistance $R_{PCB, SOCKET, etc.}$ creates a huge I^2R loss that increases quadratically with the supply current I_{R1}. Integrating the DC-DC converter with the cores could reduce the I_2R loss. However, fully integrated voltage regulators typically have low efficiencies at high conversion ratios (e.g., 4-to-1) unless ultra-high quality on-chip capacitors or inductors are used [21, 37-38].

Recent work has proposed voltage stacking as an alternative on-chip power delivery solution [70-75]. Rather than delivering current to all of the cores in parallel, voltage stacking vertically connects the cores in serial layers, as shown in Figure 4.1(b). A single high voltage supply ($4V_{DD}$) is delivered to the chip. The supply current (I_{R2}) is reduced compared to the solution in Figure 4.1(a), and it is recycled through the stacked cores in different voltage layers. Consequently, the I^2R loss is dramatically reduced. If the power consumption of all stacked layers is the same, the internal rail voltages should be evenly distributed to around V_{DD}. Unfortunately, a load power mismatch between layers directly translates to inter-layer voltage noise, which calls for a fully integrated voltage regulator to compensate for any load power mismatch between the stacked layers.

4.1.1 Prior work

Prior work has proposed several fully integrated voltage regulators for differential power processing in voltage stacking applications. Push–pull linear regulators are used in [70-71] to provide voltage regulation for stacked outputs. As shown in Figure 4.2, by

53
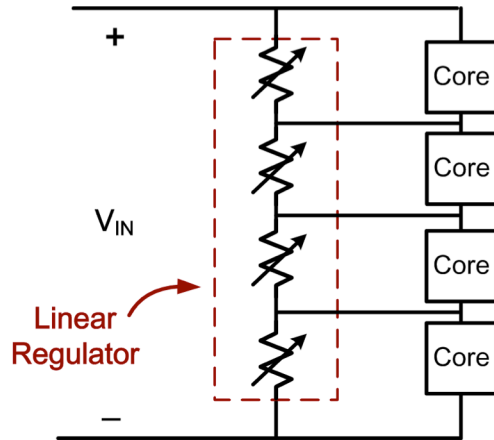
Figure 4.2: Voltage stacked power delivery with linear regulators.

changing the resistance in a linear regulator, the current consumed by each stacked layer is rebalanced when the cores consume different amounts of current. Although the linear regulators have small area overhead and are easy to integrate, the power delivery efficiency of the overall stacking system is limited by the low efficiency of the linear regulator. In the worst case, if cores in only one layer consume current, the efficiency is less than 25% when linear regulators are used.

A 2-to-1 switched-capacitor (SC) converter as shown in Figure 4.3 was implemented in [76] to deliver power to two stacked output layers. To support more than two output layers, multiple 2-to-1 SC converters can be used to regulate the internal rails [74]. For an N-layer voltage stacking system, this multi-stage solution needs a total of N-1 2-to-1 SC converters, resulting in many switches on the power train and making the feedback design complicated. Inductive converters have also been proposed as off-chip solutions for differential power processing [75]. However, it is harder to integrate high-quality on-chip inductors than on-chip capacitors.

ϕ^*_{2-1}

ϕ_{2-1}

ϕ^*_{1-0}

ϕ_{1-0}

Switched-Cap Converter
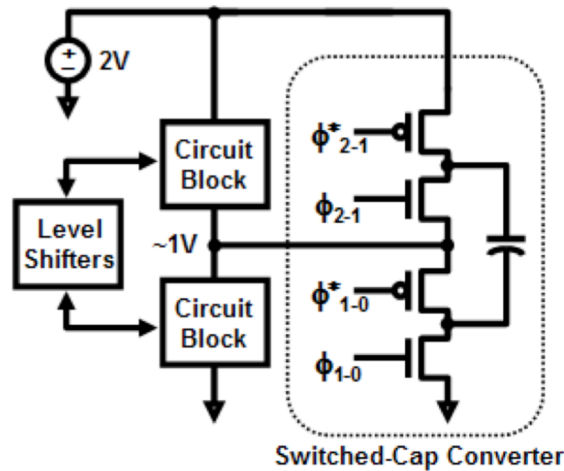
L. Chang, et al., VLSI 2010 [76]

Figure 4.3: 2-way voltage stacking with SC converters.

The remainder of this chapter will present an SC FIVR that I designed and implemented that simultaneously supports four stacked output domains. Section 4.2 describes the basic operations of the SC converter in the voltage stacking application and discusses the optimization of the converter. Section 4.3 presents the important design techniques, such as the flying capacitor reconfiguration and the flying capacitor parasitic charge recycling, that improve the performance of the converter. Section 4.4 discusses the proposed hybrid feedback control scheme. Finally, measurement results from the prototype converter are presented in Section 4.5.

4.2 A symmetric ladder SC converter with stacked output domains

In this thesis, I present a fully integrated 4-to-1 SC converter that absorbs inter-layer load power mismatches and regulates the internal rails of a multicore system that

implements four-layer voltage stacking. The symmetric ladder SC converter (SLSCC) topology [42, 43] is used to implement the converter. By tapping into the internal rails of the symmetric ladder (as shown in Figure 4.4), the SLSCC can neutralize the mismatched load currents. Thanks to the ladder topology, none of the power switches or flying capacitors are exposed to high voltages, and they can be implemented with thin-oxide devices in this technology, which improves the efficiency and power density.

The SLSCC delivers power to 16 Intel microcontroller cores, which are voltaged stacked in 4 layers. I conducted a charge flow analysis of the SLSCC and discovered that the charge flow depends on the layer-to-layer load conditions. Since the conversion loss is related to the charge flow, the optimized SLSCC design also depends on layer-to-layer load conditions. To optimize its performance, the proposed SLSCC dynamically allocates valuable flying capacitor resources according to different load conditions, which improves conversion efficiency and allows greater power mismatches between the layers. Conversion losses only apply to inter-layer mismatched power, and recycled current flows efficiently through the entire stack. The average power delivery efficiency of the entire voltage stacking system is as high as 87%.

I also propose a new hybrid feedback control scheme that regulates the four stacked layers simultaneously and reduces voltage ripple for high levels of power mismatch, a condition that exacerbates voltage ripple in conventional SC converters.
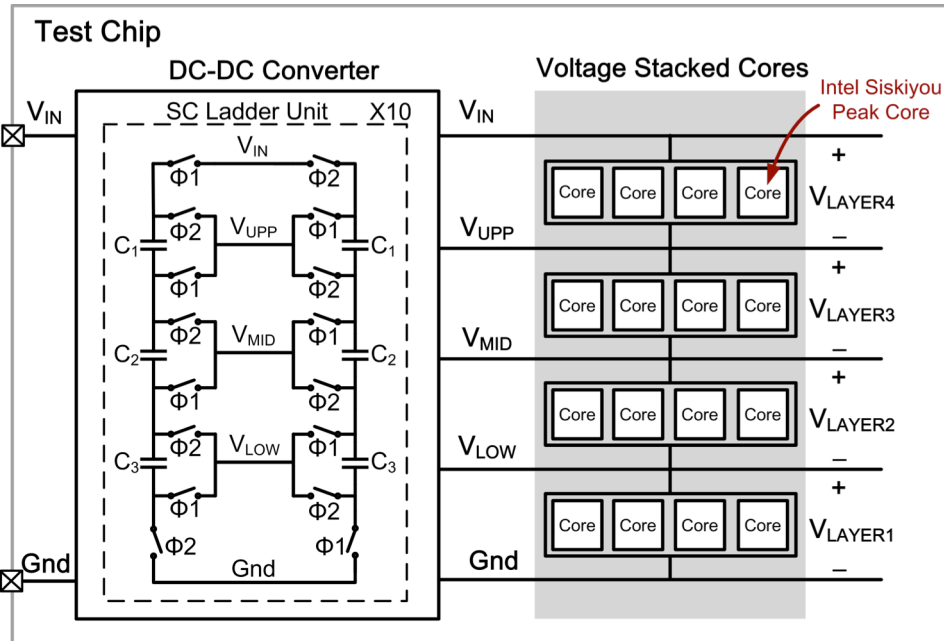
Figure 4.4: Voltage stacking system diagram showing the SLSCC and the 16 four-layer stacked cores.

The SLSCC operates off a 3.6V input voltage and supports four stacked layers. Thus, the nominal voltage of each output layer is 900mV, which is the nominal operating voltage of the transistors in the 40nm technology chosen for this test chip. The maximal supported power mismatch is between 20mW and 30mW for each output layer, limited by the available chip area. Better capacitor technology, such as high-density trench capacitors, can reduce the area of the SLSCC and improve the supported output power. Even though I chose the symmetric ladder SC topology and four-layer stacking, many of the conclusions and findings of this research can be applied to other SC designs with different topologies and different numbers of stacked layers. For example, the analysis that I conducted to study the charge flow and the conversion loss can be applied to

other SC converters that target voltage stacking applications. The techniques that I propose to enable multi-layer regulation and improve conversion efficiency can also be applied in other designs.

4.2.1 Basic operations of the SLSCC

Figure 4.4 presents an overview of the voltage stacked system that I implemented in a 40nm digital process. A total of 16 Intel Siskiyou Peak microcontroller cores are configured in a 4x4 stacked array. A fully integrated SLSCC is implemented on the same chip with the cores to support the power mismatch and regulate the internal rails of the stacked system. Depending on which layer consumes more current, the SLSCC pushes current to the stacked cores or pulls current from them through the V_{UPP}, V_{MID}, and V_{LOW} rails. Connected to a 3.6V input voltage, V_{IN}, the SLSCC simultaneously regulates the four stacked output layers, each nominally at 900mV. The SLSCC consists of 10 SC ladder units, each controlled by one of the 10 interleaved switching signals. Clock interleaving reduces the voltage ripple.

The SC ladder unit operates with respect to two non-overlapping clocks Φ_1 and Φ_2, as shown in Figure 4.5. In phase Φ_1, the left capacitor ladder is connected to the input voltage while the right capacitor ladder is connected to the ground. In phase Φ_2, two capacitor ladders are connected in a symmetric fashion. The left ladder is connected to the ground while the right ladder is connected to V_{IN}. The operation of this SLSCC is similar to that of other SC-based converters [24-39]: the flying capacitors are charged in one phase and discharged in the other while the current flows from the input to the

58
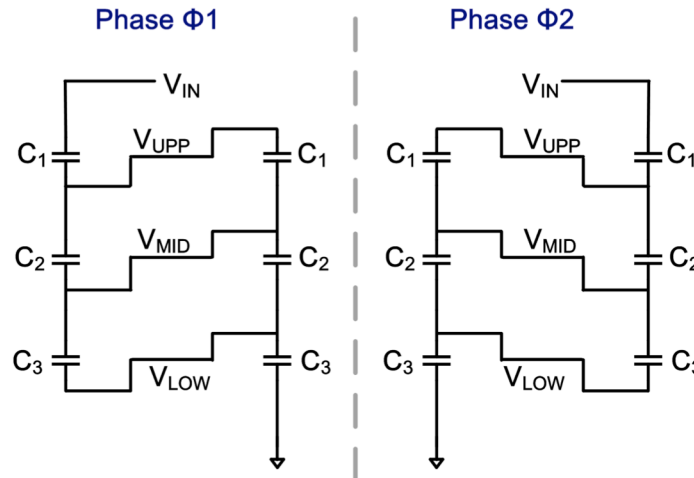
Figure 4.5: Two-phase operation of the SLSCC.

output through the capacitors and the power switches. As an example, assume that there is load current only in the bottom layer (from V_{LOW} to gnd). In Φ_1, energy drawn from the input charges the flying capacitors in the left capacitor ladder and flows to the load. At the same time, the capacitors in the right capacitor ladder are discharged, transferring energy to the load. In Φ_2, the energy stored in the flying capacitors in the left ladder during the previous phase flows to the load. At the same time, energy drawn from the input charges the flying capacitors in the right ladder and flows to the load.

The overall voltage stacking system depicted in Figure 4.4 was shown to work in [73]. This thesis focuses on the implementation and the measurement results of the SLSCC itself. I will not discuss the interaction between the SLSCC and the stacked cores. In [73], I explored the system-level performance, such as the energy–delay product and the computing throughput of the stacked cores. I also compared different clocking

strategies for the cores, including fixed frequency clocking and adaptive frequency clocking, in [73].

4.2.2 SLSCC losses and optimizations

Losses in the SC converter can be categorized as switching losses or conductive losses. Switching losses come from the switching of the parasitic capacitance in the circuit, mainly the parasitic capacitance of the flying capacitors and of the power switches. Conductive losses are associated with charge redistribution through the flying capacitors and power switches.

The way that SLSCC delivers power to the load in this voltage stacking application is very different from a conventional system, because the load circuits are spread out in multiple voltage layers. I found that the charge flow in the SLSCC depends heavily on the layer-to-layer load conditions and that the SLSCC has different performance characteristics, such as efficiency and maximal supported power mismatches, for delivering power to different output layers. To explain these interesting characteristics of this voltage stacking system, Figure 4.6 presents examples of charge flow analyses for different load conditions. For a fair comparison, the total load currents in all scenarios are the same. For simplicity, Figure 4.6 only shows the charge flow during Φ_1. Since the operation of the SLSCC is symmetric, the situation in Φ_2 is similar to that in Φ_1. Charge-flow analysis is a good tool for analyzing conductive loss, which comes from both the flying capacitors and the power switches. The loss depends on the switching frequency of the converter; it is dominated by the flying capacitors when the switching frequency is
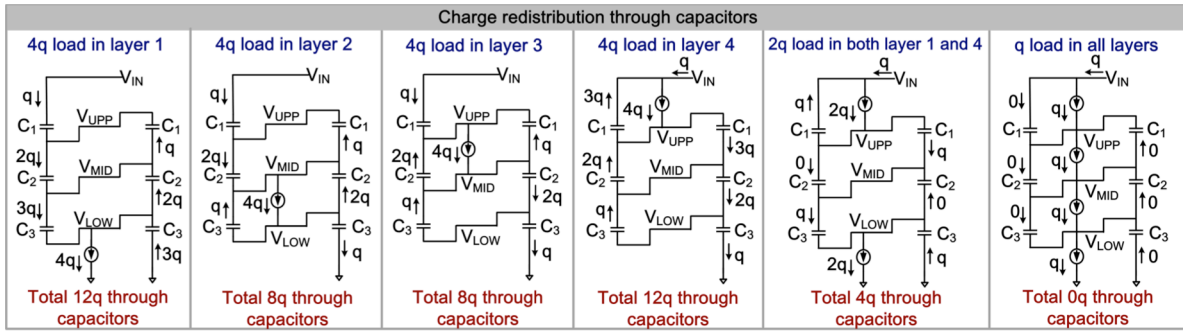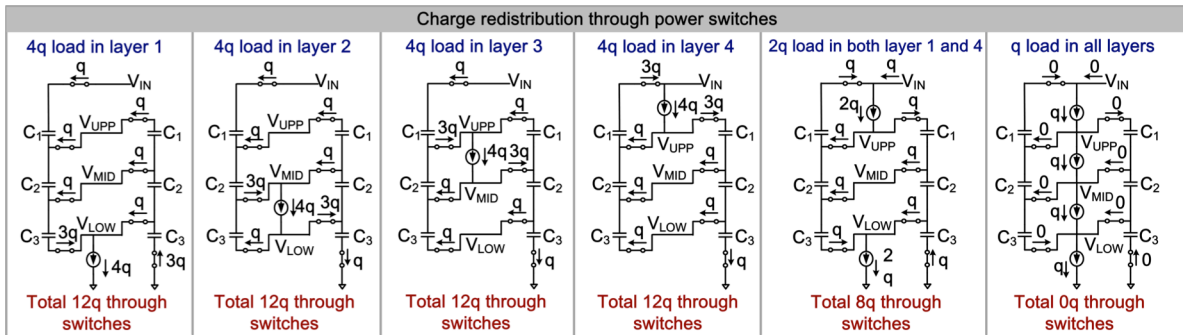
Figure 4.6: Internal charge flow diagrams in the SC ladder for different load conditions. (a) Charge flow through flying capacitors; (b) Charge flow though power switches.

low and by the power switches when the switching frequency is high. Figure 4.6 shows the charge flowing through both the flying capacitors and the power switches.

Figure 4.6(a) shows how charges flow through the flying capacitors in the SLSCC. The load currents are represented by the current sources. Among the six different scenarios, the first four scenarios show situations where the load current is extremely imbalanced and only the output layer consumes current. These are the worst cases in the voltage stacking application because all of the power needs to be delivered by the SLSCC. Comparing the first four scenarios, we see that charges flow through the

capacitors depending on which layer is consuming current when only one layer is doing so. Typically, the greater the charges flowing through the flying capacitors, the higher the conductive loss [25, 43]. Thus, the SLSCC will have better performance when it delivers power to the middle layers than to the top and bottom layers. The load currents are spread out to multiple layers in the last two scenarios. Charges flowing through the flying capacitors are reduced compared to the previous scenarios. In the last case, where load currents are perfectly balanced, all charge flows through the load circuitry, with no charge flowing through the SLSCC.

Figure 4.6(b) shows the charge flow though the power switches. The total amount of charge that flows through the power switches is the same for the first four scenarios. Spreading out the load currents helps reduce the charges flowing through the power switches.

Based on the charge flow analysis in Figure 4.6, we can conclude that (1) the charge flows through the capacitors and power switches may be different when the SLSCC delivers power to different layers, and (2) spreading the load current from a single layer to multiple layers reduces the total amount of charge flowing through the capacitors and power switches, which reduces the losses.

The charge transfer flow also provides a guideline for optimizing the efficiency of the SLSCC and its maximal supported power mismatch. A typical optimization process for an SC converter involves optimizing the size of each flying capacitor and power switch; this is thoroughly discussed in [25, 43]. To minimize the loss, the capacitors are sized
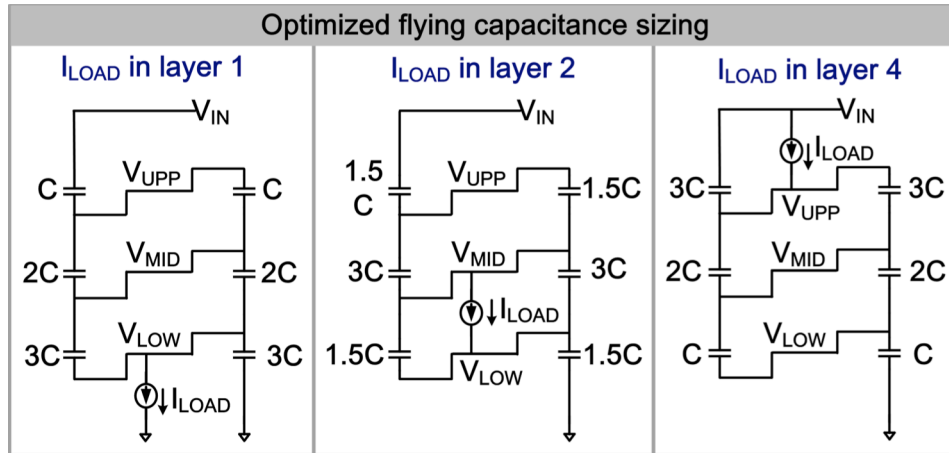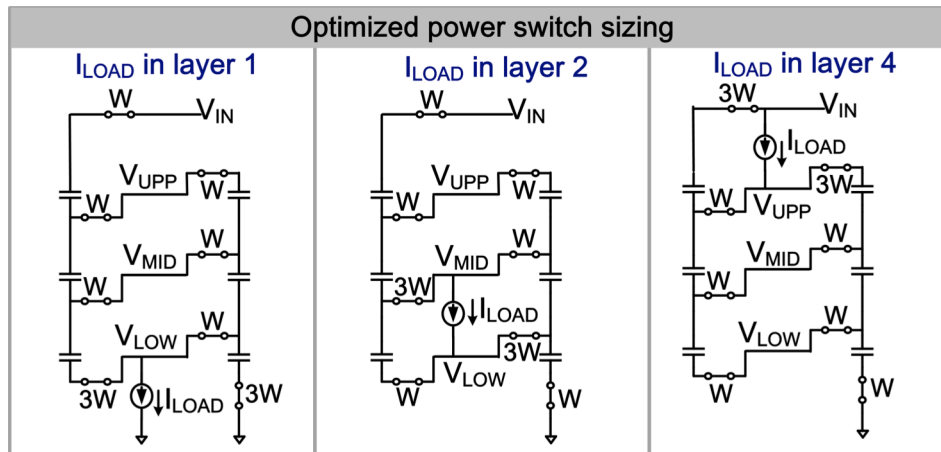
Figure 4.7: Optimized SLSCC for different load conditions with (a) optimized flying capacitors; (b) optimized power switches.

proportional to the charge flowing through them. The optimal power switch width is also proportional to the charge flowing through it. The total flying capacitance is limited by the available chip area, while the switch width is determined by trade-offs between conductive loss and switching loss.

In this voltage stacking application, the optimized SLSCC design depends on the load current profiles, because charge flow depends on the layer-to-layer load conditions.

Figure 4.7 presents the optimized flying capacitor size and the optimized power switch size at three different load conditions. Because of the symmetric operation of the SLSCC, Figure 4.7 only shows the converter in phase Φ_1. The optimized capacitor size is proportional to the charge that flows through each capacitor. Similarly, the optimized switch width is also proportional to the charge flowing through each switch. These optimizations allow for better utilization of flying capacitors and power switches, minimizing loss.

In this particular design, the SLSCC operates in the so-called "slow-switching limit" (SSL) mode [43], where the flying capacitors rather than the power switches usually dominate the conductive loss.

The switching loss in this SLSCC is similar to that of a typical SC converter in conventional applications where only the bottom layer consumes load current. The switching loss has been thoroughly explored in [25, 43]. I do not discuss switching loss in this thesis, although I considered it when designing and measuring the SLSCC.

4.3 Implementation of the SLSCC: Open-loop operation

4.3.1 Implementation of the SC ladder

Figure 4.8 shows the implementation of the SC ladder unit. The power switches in the main SC ladder are implemented with thin-oxide NMOS or PMOS transistors. The flying capacitors rely on thin-oxide PMOS transistors, which have a 20% smaller capacitor density but create only 1/6 of the leakage current of NMOS transistors.
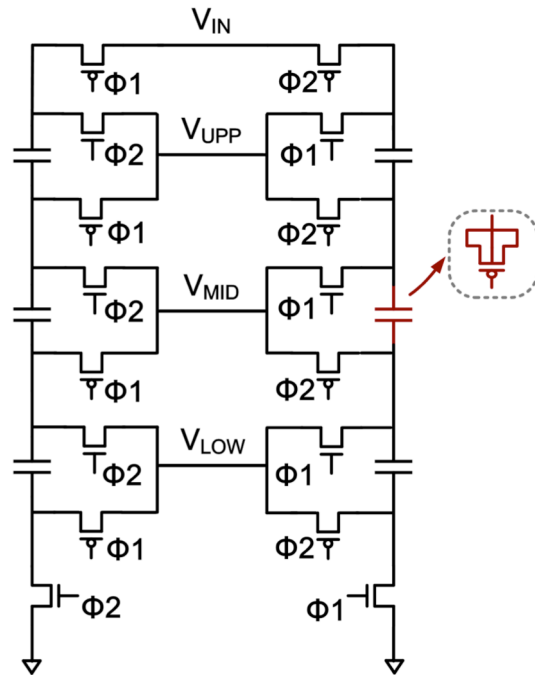
Figure 4.8: Transistor implementation of the SC ladder.
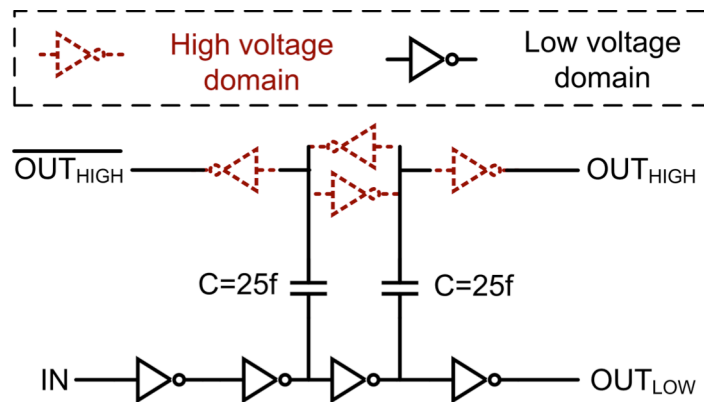
Figure 4.9: Implementation of the level shifter.

The thin-oxide power switches in the SLSCC work in different voltage domains.

Level shifters are required to shift the switching signals from low voltage domains to

higher domains. Figure 4.9 presents the implementation of the level shifters. Since the level shifters are on the switching signal path, their delays would add up to the delay of the feedback loop, degrading the SLSCC's closed-loop performance. The capacitor-coupled level shifter topology was chosen in this design; it has a delay of less than 200 ps. The nominal voltage across the inverters is about 0.9V. Therefore, thin-oxide CMOS transistors were used to implement the inverters. The voltage across the coupling capacitor depends on the output voltage domain. For the highest voltage domain (V_{UPP}~V_{IN}), the voltage across the capacitors is about 2.7V, much higher than the breakdown voltage of the transistors that are available in this process. The capacitors in the level shifters are therefore implemented using metal-oxide-metal (MOM) capacitors. The sizing of the coupling capacitor involves a trade-off between the power consumption, reliability, and occupied area of the level shifter. Bigger coupling capacitors mean stronger coupling, but at the same time they have higher parasitic capacitance and also occupy more chip area. The level shifters occupy a total of 2.5% of the die area in this SLSCC.

4.3.2 Flying capacitor bottom-plate charge recycling

The flying capacitor parasitic switching loss is one of the major losses in fully integrated SC converters. A charge recycling technique similar to those in [21, 36] is implemented in this design, which can reduce the bottom-plate loss by about 50%.

Figure 4.10 shows the switching behaviors of the flying capacitors' parasitic capacitors in this SLSCC. Different parasitic capacitors switch between different rails
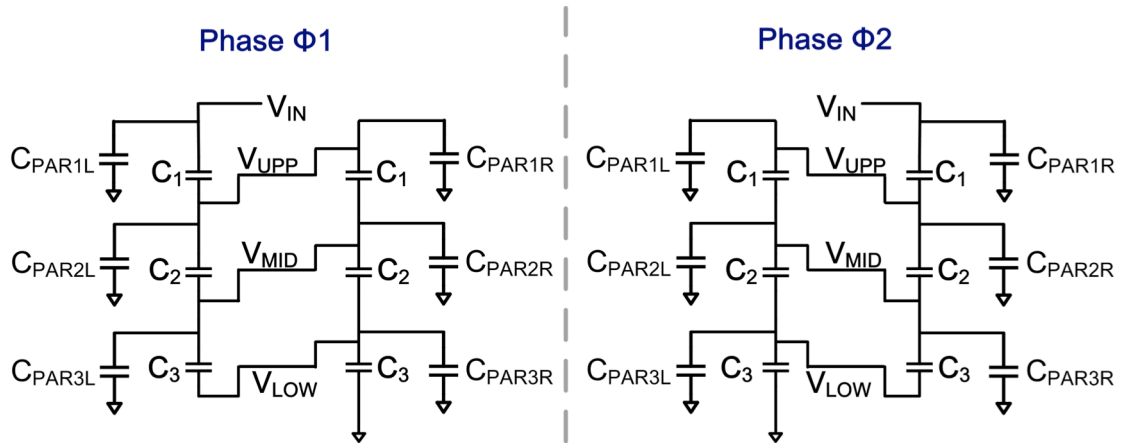
Figure 4.10: Flying capacitor parasitic switching loss in the SLSCC.

Figure 4.11: Implementation of the charge recycling technique.

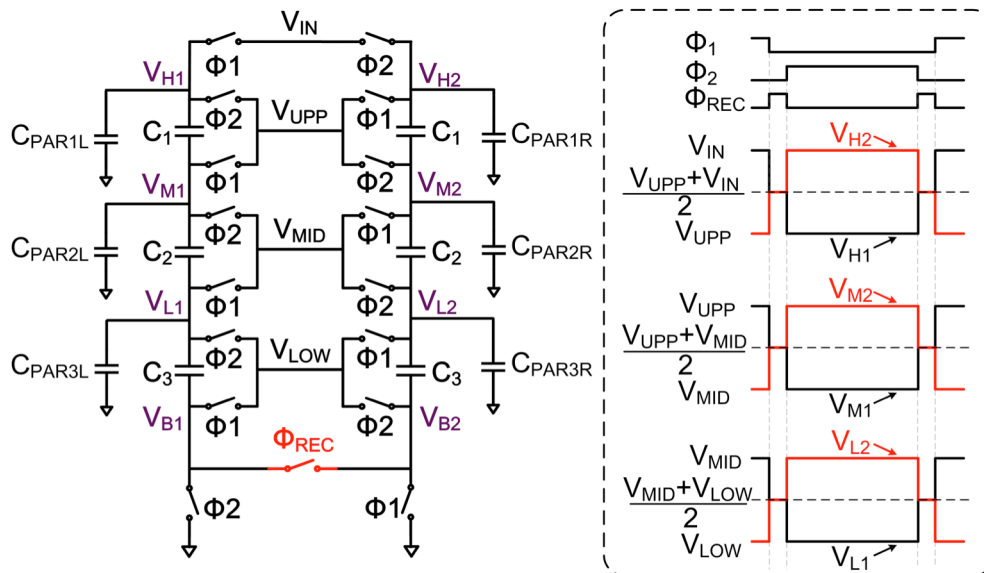with switching swings of about $V_{IN}/4$. Pairs of parasitic capacitors switch between the same rails. For example, in phase Φ_1, C_{PAR3L} connects to V_{MID} while C_{PAR3R} connects to V_{LOW}. In phase Φ_2, C_{PAR3L} is discharged to V_{LOW} while C_{PAR3R} is charged up to V_{MID}.

Figure 4.11 presents a charge recycling technique similar to those proposed in [21, 36]. Whenever the SLSCC switches between Φ_1 and Φ_2, it goes through an additional recycling phase, Φ_{REC}. One additional recycling switch is added at the bottom of the symmetric ladder, controlled by Φ_{REC}. During Φ_{REC}, all the power switches are turned OFF. Only the recycling switch controlled by Φ_{REC} turns ON. Charges redistribute between the parasitic and flying capacitors. As a result, some amount of energy that would otherwise be wasted in charging/discharging the parasitic capacitors is recycled. Take the bottom pair of parasitic capacitors, C_{PAR3L} and C_{PAR3R}, as an example to explain the recycling process. When Φ_1 is ON, C_{PAR3L} is charged up to V_{MID} while C_{PAR3R} is discharged to V_{LOW}. After Φ_1 is turned OFF, Φ_{REC} is turned ON before Φ_2 turns ON. During Φ_{REC}, the charge that was previously stored on node C_{PAR3L} partially transfers to C_{PAR3R}. This charge would otherwise be wasted without this charge recycling scheme. Since the parasitic capacitors are much smaller than the flying capacitors, the voltage across C_{PAR3L} and C_{PAR3R} is around $(V_{MID}+V_{LOW})/2$. After the recycling phase Φ_{REC}, Φ_2 turns ON. C_{PAR3L} is fully discharged to V_{LOW}. C_{PAR3R} is charged from $(V_{MID}+V_{LOW})/2$ to V_{MID}, rather than from V_{LOW} to V_{MID}. A similar process occurs for the other parasitic capacitors. This recycling process reduces the bottom-plate loss by about 50%, as previous work has demonstrated [21, 36].

4.3.3 Flying capacitance reconfiguration scheme

Since the optimized flying capacitance allocation depends on layer-to-layer load conditions in this voltage stacking system, it is preferable for the SLSCC to be able to
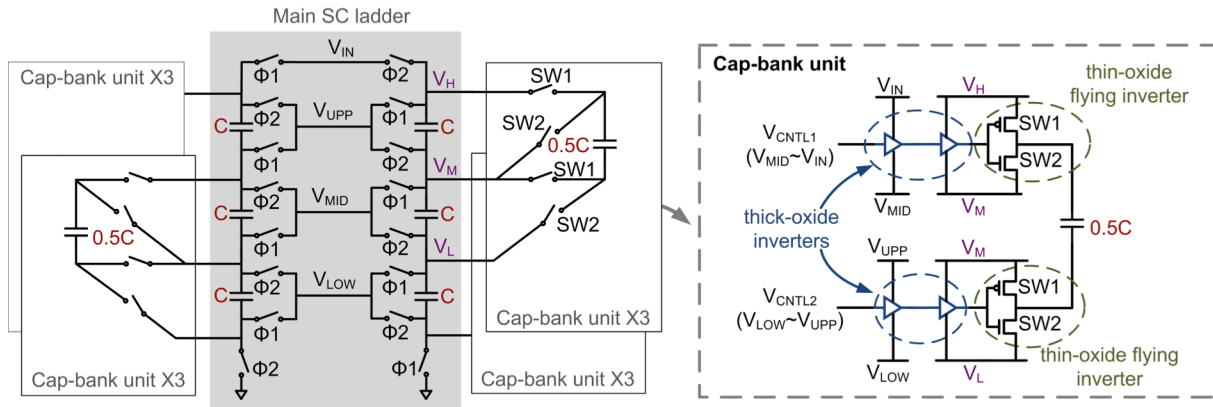
Figure 4.12: Implementation of the reconfigurable SC ladder.

dynamically modify the sizes of its flying capacitors according to load current conditions. In this section, I present the implementation of a capacitance reconfiguration scheme for the SLSCC.

Figure 4.12 shows the implementation of the reconfigurable SC ladder. The overall SC ladder consists of one non-configurable main SC ladder (shown with a grey background in Figure 4.12) and four sets of reconfigurable cap-bank units that are connected to the main SC ladder. Each cap-bank set contains three identical cap-bank units that can individually configure their connections. Each of the reconfigurable capacitors in the cap-bank units can be connected in parallel with different capacitors in the main ladder by closing either SW1 or SW2. In this way, the capacitance is reallocated dynamically according to the load conditions. The switches (SW1 and SW2) in the cap-bank units only switch ON/OFF when the capacitors need to be reconfigured.

Implementing the switch network in the cap-bank units was challenging. These additional switches in the cap-bank units add conductive loss and switching loss to the

converter. To minimize the loss, I propose using a pair of thin-oxide flying inverters to implement the reconfigurable switches SW1 and SW2, as shown in the right half of Figure 4.12. In each paired SW1 and SW2, the gates are connected together, driven by another small flying inverter. Using this design rather than connecting the gates to a fixed voltage to turn SW1 and SW2 ON/OFF, the gate voltage switches together with V_H, V_M, and V_L when the main SC ladder is switching. Thus, SW1 and SW2 can be implemented using thin-oxide transistors to reduce the associated conductive and switching loss. Since load conditions fluctuate at a much lower rate than the main switching frequency of the converter, the switching losses associated with SW1 and SW2 are small and justified by the efficiency improvements that configurability offers. A total of 12C is used in each of the 10 interleaved SC ladder units, where 1C equals 37.5pF, for a total capacitance of 4.5nF.

4.4 Ripple-reduced hybrid feedback control

Figure 4.13 shows the system block diagram for the proposed SLSCC. A hybrid feedback control circuitry operates off the clock from a voltage-controlled ring oscillator (VCRO). The feedback circuitry monitors the voltages across each output layer and generates a 10-phase interleaved switching signal SW_{HYBRID}. The power switch control signal generator then creates the switching signals, Φ_1, Φ_2, and Φ_{REC}, for the 10 interleaved reconfigurable SC ladder units. The level shifters shift these switching signals to the correct voltage domains and eventually drive the switches in the reconfigurable SC ladder.
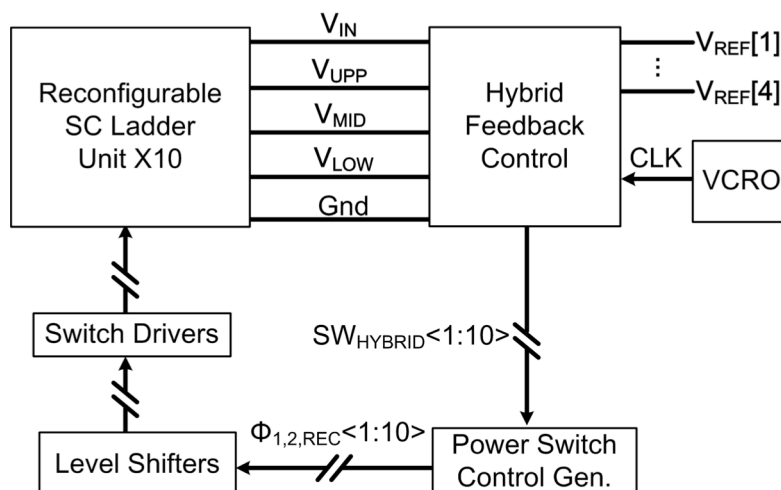
Figure 4.13: Block diagram for the SLSCC.

The feedback control loop in this converter is composed of a primary single-bound control loop that simultaneously regulates each of the voltage layers and another secondary proactive loop, which helps reduce voltage ripple for heavily mismatched load conditions.

4.4.1 Primary single-bound control loop

In this design, the primary feedback loop tries to keep all output layer voltages above the reference voltages, as opposed to regulating the layer voltages to the reference voltages. In a voltage stacking application, all the output layer voltages add up to the input voltage, which is 3.6V in this design. If the load power is the same for all layers, the output voltage is 900mV across all layers. If one of the layers consumes more current than the rest of the layers, its voltage decreases to below 900mV while some of the other layer voltages increase to above 900mV, to maintain Kirchhoff's current law
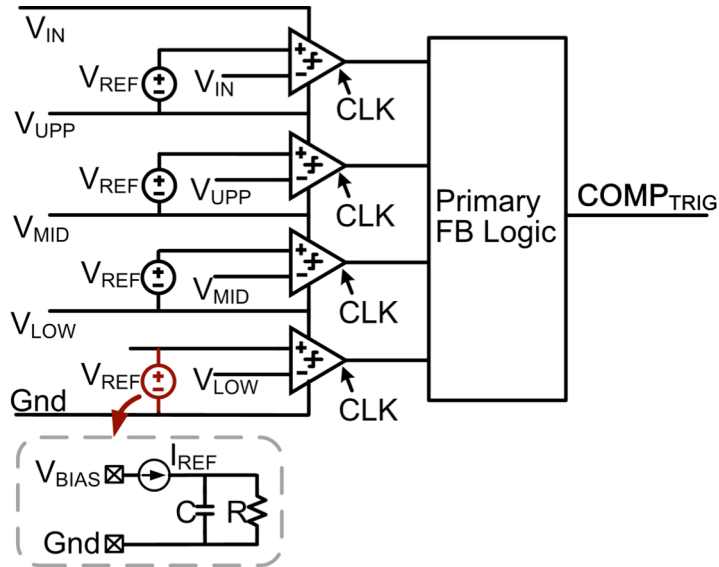
Figure 4.14: Implementation of the primary single-bound control scheme.

(KCL). If the mismatch current is too large, its voltage will fall below the reference voltage. In those scenarios, the feedback loop in the SLSCC detects that one of the layer voltages is lower than the references and tells the SC units to switch and restore the voltage levels. The switching behavior of the SLSCC tries to keep all layer voltages above the reference voltages.

Figure 4.14 illustrates the implementation and operation of the primary feedback control loop. Four 2.5GHz digital-clocked comparators compare the voltages across the layers with the corresponding reference voltages generated on-chip. If the voltage of any layer falls below the reference, the associated comparator generates a pulse. The primary feedback control logic combines the outputs of all the comparators and generates a high frequency switching signal, $COMP_{TRIG}$. If all output-layer voltages are above their reference voltages, $COMP_{TRIG}$ stays low. If the voltage in any output layer is

detected to fall below the references, a pulse is created on COMP$_{TRIG}$. COMP$_{TRIG}$ is further processed by the secondary proactive loop (discussed latter) and is turned into 10 interleaved slow switching signals by a barrel shifter. The interleaved switching signals eventually drive the switches in the SC ladder, as shown in Figure 4.13. Each interleaved SC ladder unit switches at a maximum frequency of 250MHz.

One of the challenges in designing this primary single-bound control loop was to create the reference voltages for the comparators in each layer. In this design, V$_{REF}$ is created using a current flowing through a resistor, as shown in Figure 4.14, so that

$$V_{REF} = I_{REF} \cdot R \tag{4.1}$$

I_{REF}, which is created using an on-chip current mirror, is tuned by changing the V_{BIAS} that is controlled off-chip. At dc, IR creates a stable reference voltage. But the slow slew rate limits its performance when there is high-frequency noise in the rails (Gnd, V$_{LOW}$, V$_{MID}$, V$_{UPP}$). Another coupling capacitor C is added in parallel with R to couple the high-frequency noise of the internal rails to the inputs of the comparators.

4.4.2 Secondary ripple-reduced proactive loop

One of the major advantages of the single-bound control is its fast response for handling large current steps. It can change the effective switching frequency of an SC converter from a very low frequency to its maximum frequency within a few nanoseconds [21, 26, 28]. However, due to the non-zero feedback latency and the pulse-skipping nature of the control loop, single-bound control loops typically result in

73
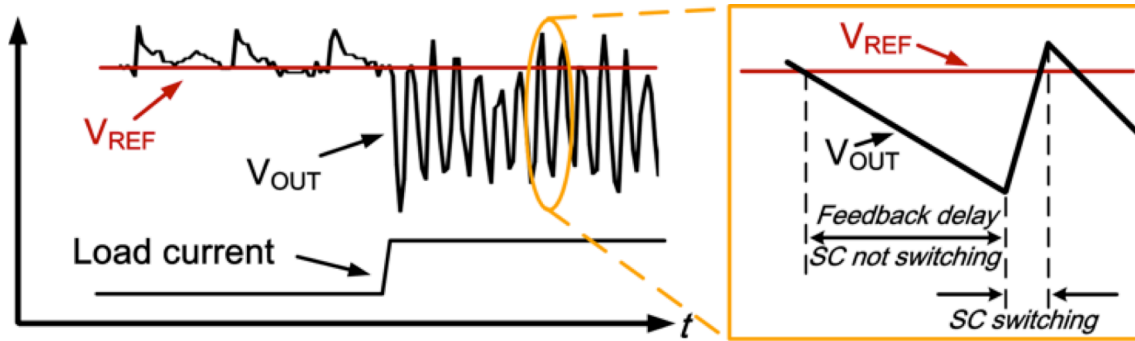
Figure 4.15: Typical voltage noises in an SC converter using single-bound control.

much larger static voltage ripples compared to voltage-controlled-oscillator (VCO) based pulse frequency modulation (PFM) loops [28, 34].

Techniques such as resistance modulation [39] have been proposed to reduce the voltage ripple at light loads. Interleaved designs can also reduce voltage ripple at light loads. However, the ripple at heavy loads can also be very large [21]. Figure 4.15 illustrates a typical output voltage waveform for an SC converter that relies solely on single-bound control. At heavy load conditions, the load current quickly discharges V_{OUT} before the loop can detect this and react, creating a large ripple. The larger the delay is, the larger the ripple will be. In this design, the simulated feedback delay is about 1.5ns.

In this voltage stacking application, the SLSCC only processes the mismatched power between the output layers. I propose a secondary proactive loop to reduce voltage ripple for heavily mismatched load conditions. To support the ripple reduction feature, all of the SC ladder units are dynamically divided into two groups. One group is controlled by the primary single-bound control loop. The other ladder units, which I call *proactive units*, are controlled by the secondary loop. By detecting load conditions, the
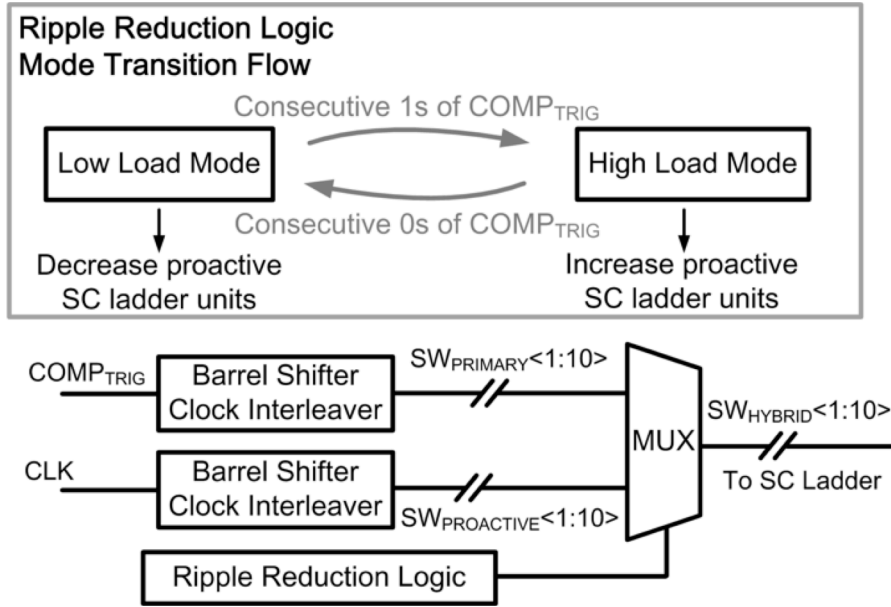
Figure 4.16: Implementation of the proactive feedback control.

secondary proactive loop tells the proactive units to always switch at the maximum rate, which reduces voltage ripple. As shown in Figure 4.16, ripple reduction logic monitors consecutive 1s and 0s in $COMP_{TRIG}$ to dynamically allocate SC ladder units between single-bound and proactive control. If several consecutive 1s are detected, more SC ladder units become proactive units. If several consecutive 0s are detected, the number of proactive units is reduced. In a very heavy load condition, most of the SC ladder units are proactive units switching at peak frequency, periodically delivering power to the output and reducing the voltage ripple.

The ripple reduction scheme implemented in this voltage stacking application can be applied to other SC converters with single-bound control, especially in conventional applications where the load circuitry operates in a single output layer.
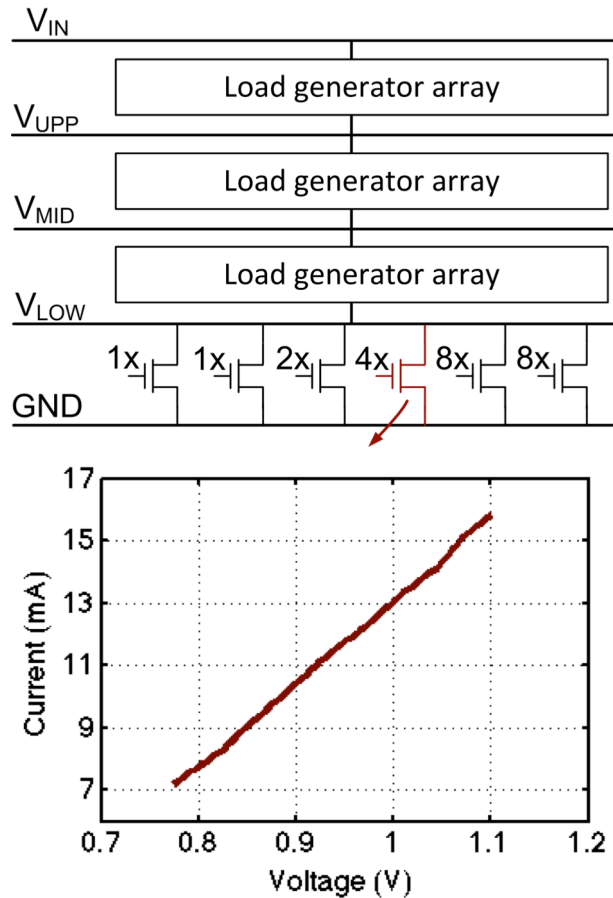
Figure 4.17: Implementation and characterization of the on-chip load generator.

4.5 Measurement Results

This section presents measurement results for the test chip prototype. The test setup has off-chip capacitors to bypass V_{IN}, but no external capacitors connect between the internal rails. The Intel Siskyou Peal processor cores are turned off during the measurements. Discussions of the processor cores and their measurement results can be found in [73].

I used on-chip load current generators to create different load conditions in the stacked layers in order to measure the conversion efficiency and transient response of the SLSCC. Figure 4.17 shows the implementation of the load generators and their measurement results. Each layer has an identical but individually configurable load generator array. Each array consists of six binary weighted NMOS transistors, implemented using triple-well technology to avoid breakdown issues. These transistors create load currents in the stacked layers when they are turned on. Each array is controlled by a programmable LFSR that is also implemented on-chip. The current of the load generator is a function of the transistors that are turned on as well as the voltage across them. The bottom part of Figure 4.17 shows the measured load current of a 4X transistor across its supply voltage.

4.5.1 Conversion efficiency

As discussed in previous sections, the SLSCC only processes differential power consumed by the load. In the worst case, only the load circuitry in one of the four stacked output layers consumes currents. In such cases, all the power that is consumed is delivered by the SLSCC. Figure 4.18 plots the measured efficiencies of the converter when only one layer consumes current. In Figure 4.18(a), the proposed SLSCC (with dynamic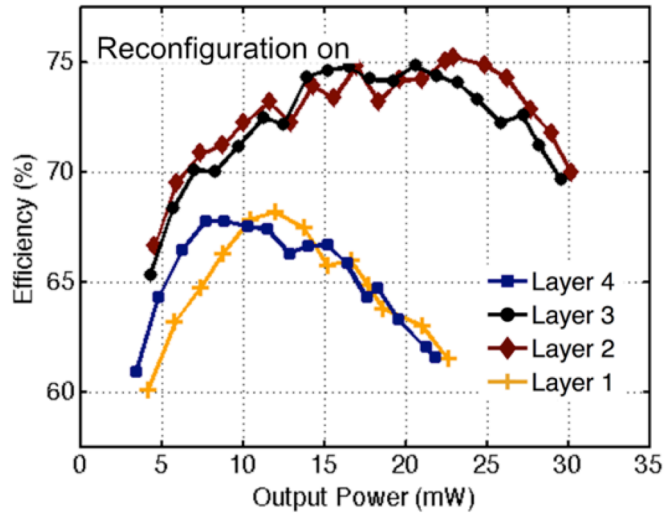 flying capacitor allocation) achieves higher efficiency and supports higher mismatched power (output power) for the two middle layers. Analysis of the internal SLSCC charge flow in Figure 4.6 shows that the losses are smaller (i.e., lower conductive loss) when delivering current to the middle layers. In the measurements

77
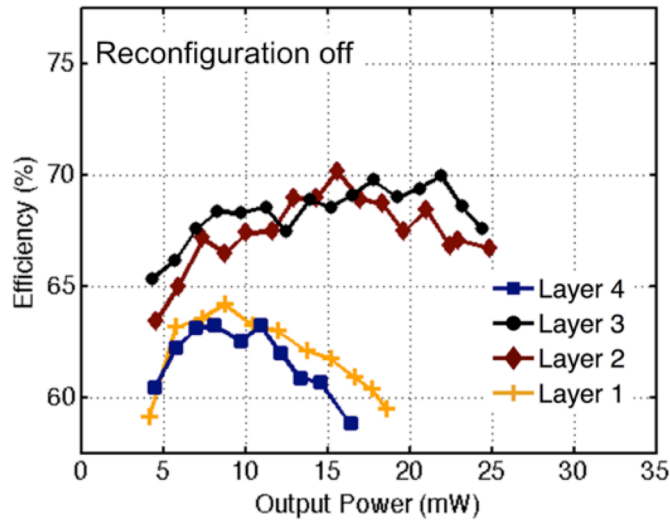
Figure 4.18: Measured efficiencies of the SLSCC when only one layer consumes current, with reconfiguration (a) on and (b) off. (VIN = 3.6V; all reference voltages are set to 800mV.)

presented in Figure 4.18(b), the flying capacitor reconfiguration is turned off. The flying

capacitance resource is equally distributed. Comparison of Figure 4.18(a) with Figure

4.18(b) confirms the benefits of reconfiguring capacitor allocations. Reconfiguration
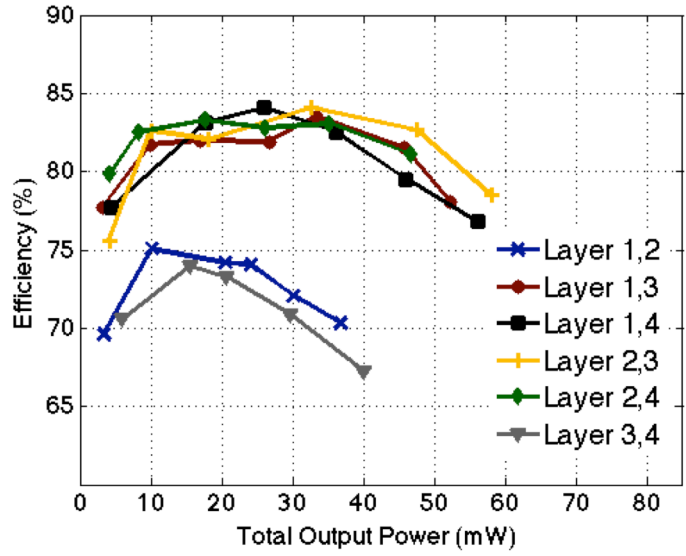
(a)
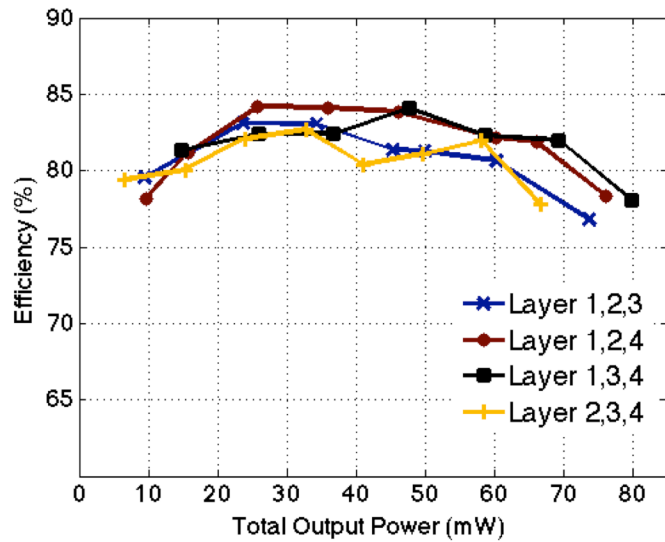
(b)

Figure 4.19: Measured efficiencies of the SLSCC when multiple layers consume currents. (VIN = 3.6V; all reference voltages are set to 800mV.)

improves conversion efficiency as well as maximal supported mismatched power. For all measurements that follow, reconfiguration is always turned on unless stated otherwise.

As shown in Figure 4.18, the efficiency of this SLSCC is not very high when delivering power to only one output layer. However, the efficiency improves significantly when more than one layer consumes current, which is the more common case for this voltage stacking application. Figure 4.19 presents the measured efficiencies of the SLSCC when more than one layer consumes currents. In Figure 4.19(a), two layers consume currents. The load generators are set in such a way that the same size load-creating transistors are turned on in the two layers. Generally, the efficiencies are higher compared with those in Figure 4.18. The maximal supported total output power also increases over that in Figure 4.18. These results confirm the charge-flow analysis in Section II. Figure 4.19(a) also shows that the efficiency will be different depending on which two layers consume current. This is because the charge flow depends on the layer-to-layer load conditions. For the same total power, the efficiency is higher if a smaller amount of charge flows through the flying capacitors and power switches. Figure 4.19(b) shows the efficiencies when three output layers consume currents. Both efficiency and maximal output power show much improvement over the scenarios where only one layer consumes current.

Figure 4.20 presents the power delivery efficiency of the overall voltage stacking system for a diverse collection of load current conditions. Each output layer consumes a random amount of current. Efficiency is computed as the total power consumed by all loads versus the total power supplied by V_{IN} at 3.6V. The average efficiency is as high as 87%, confirming the benefits of the voltage stacking system. SLSCC losses only

80
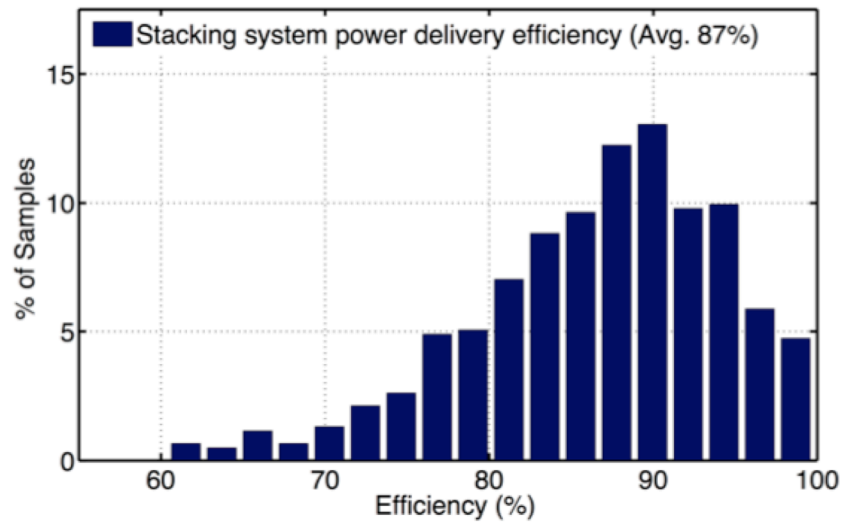
Figure 4.20: A histogram of the measured power delivery efficiency of the overall voltage stacking system. (VIN = 3.6V; all reference voltages are set to 800mV.)

apply to inter-layer power mismatches. When the load power across all layers matches well, voltage stacking evenly distributes internal voltage levels and SLSCC losses are small, leading to high system-level efficiency. All of the converter's reference voltages are 800mV, and the SLSCC might not need to switch at all unless one or more of the layer voltages fall below 800mV.

4.5.2 Voltage ripple and transient response

Figure 4.21 shows the static transient waveforms and histograms of V_{LAYER1} when only the bottom layer (Layer 1) consumes currents, with the proposed hybrid control turning on/off. In both plots, V_{LAYER1} stays around 800mV, regulated by the feedback loop. When the hybrid control is turned off, there is a static peak-to-peak voltage ripple
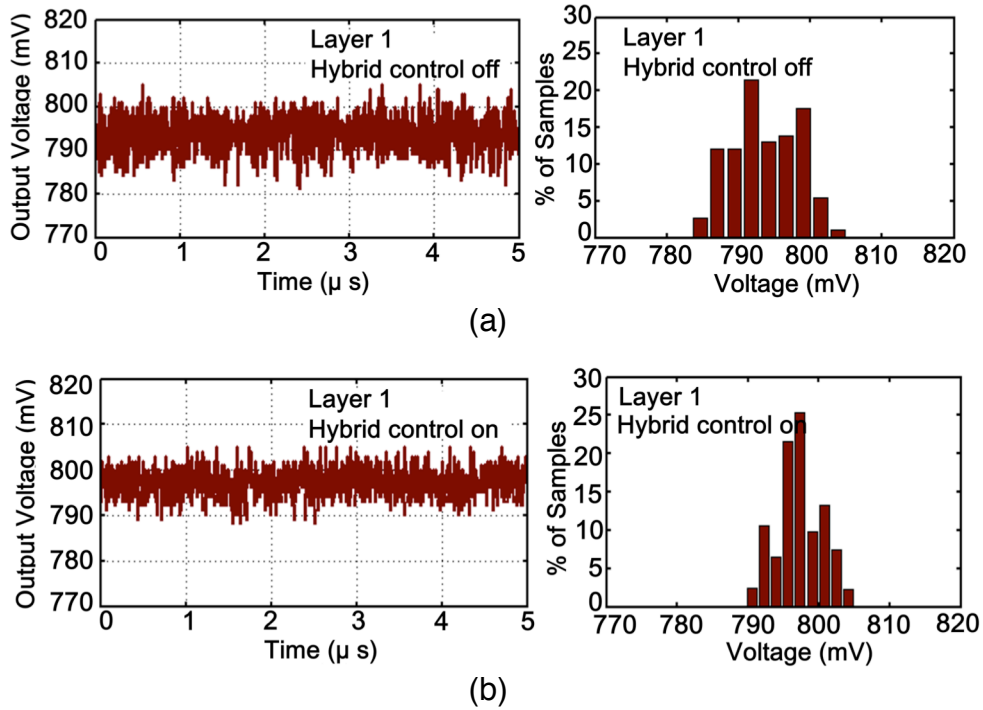
Figure 4.21: Measured transient waveforms and histograms of V$_{LAYER1}$ with load current only in Layer 1: (a) hybrid control turned on; (b) hybrid control turned off. (P$_{OUT}$≈16mW; all reference voltages are set to 800mV.)

of 25mV, shown in Figure 4.21(a). The proposed hybrid feedback control scheme reduces the voltage ripple by 30%, to 18mV, as shown in Figure 4.21(b).

To further explore the functionality of the hybrid control, the static voltage noise was measured when the SLSCC was delivering power to all four output layers, and the results were compared with the hybrid control scheme turning on/off. The measurement results are presented in Figure 4.22. The voltage noise distributions are presented in box plots.
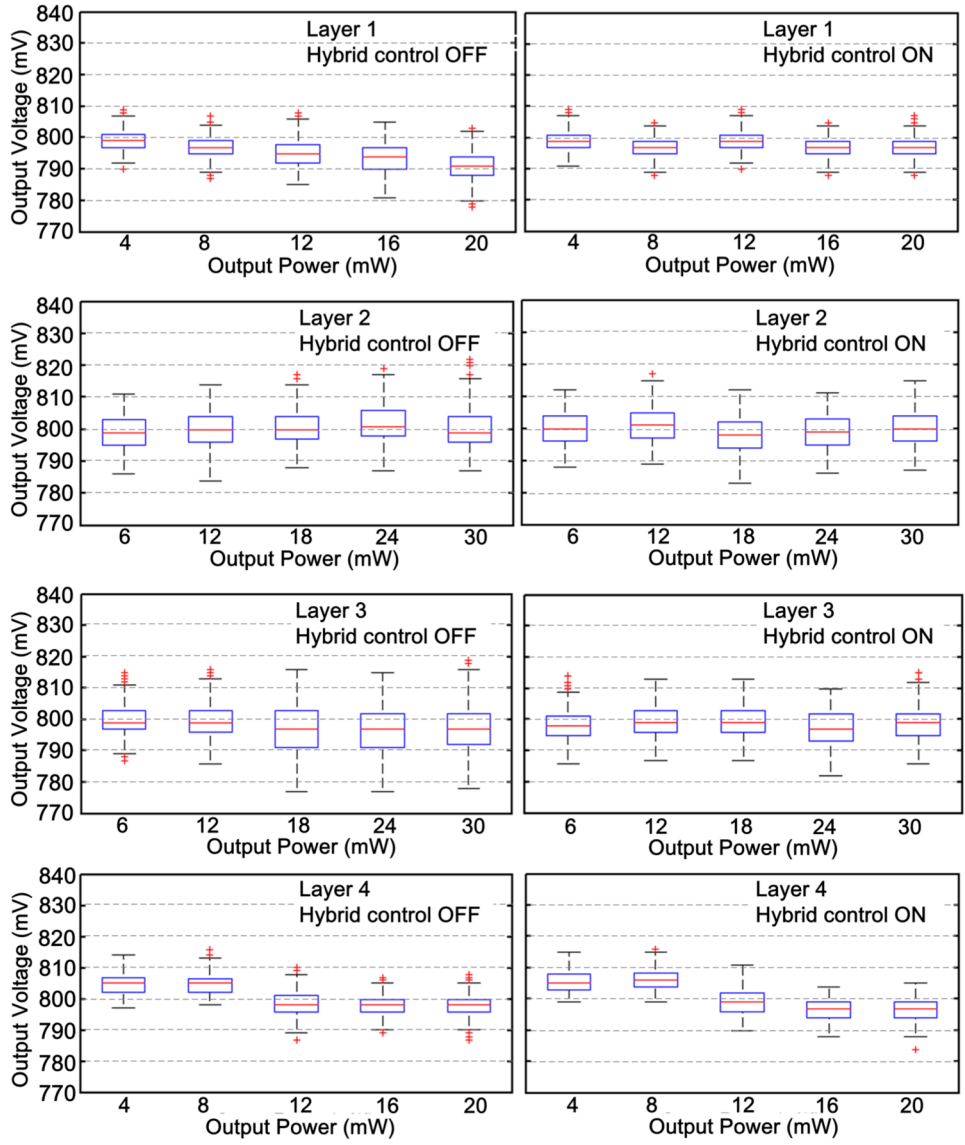
Figure 4.22: Measured voltage noise distribution when only one layer consumes current, with hybrid control turned (a) on, (b) off. (All reference voltages are set to 800mV.)

Generally, when the hybrid control is turned off, the static voltage ripple increases as the output power increases. Because of the 10-phase interleaved design, the voltage ripple is small when the output power is low. As the output power increases, the load current discharges the output faster, increasing the voltage noise, as discussed in

Section 4.4. Comparing the results with the hybrid control turned on and off, I found that the hybrid control scheme does not change the noise characteristics very much at light load conditions, but reduces the voltage ripple at heavy load conditions. At light load conditions, the SLSCC switches at a very low frequency. Most of the interleaved units are controlled by the primary single-bound control loop. As the load increases, the layer voltages fall below the reference voltage more frequently and the SLSCC needs to switch at high frequencies to deliver power to the load. As a result, more interleaved units become proactive units controlled by the secondary ripple-reduction loop. The hybrid control helps reduce the voltage ripple by 20% to 40% at a heavy load. The hybrid control is always turned on for the following measurements, unless noted otherwise.

Figure 4.23 presents the transient response of the SLSCC, verifying the functionality of the feedback control loop. Figure 4.23 shows the output voltage of all four layers. As I discussed earlier, the load current depends on the current-creating NMOS transistors that are turned on and the voltages across them. In this plot, the layer current is labeled using the nominal current at 900mV. Overall, the SLSCC ensures a minimum voltage of about 800mV for all layers, set by the reference voltage. From 0 to 5μs, all load generators are turned off. All layers settle at about 900mV, a quarter of the 3.6V input voltage, as expected. The slight differences among the layer voltages are a result of the mismatch in leakage currents from the processor cores. The load current in the bottom layer increases to 25mA at t=5μs. Voltage stacking redistributes the layer voltages to be
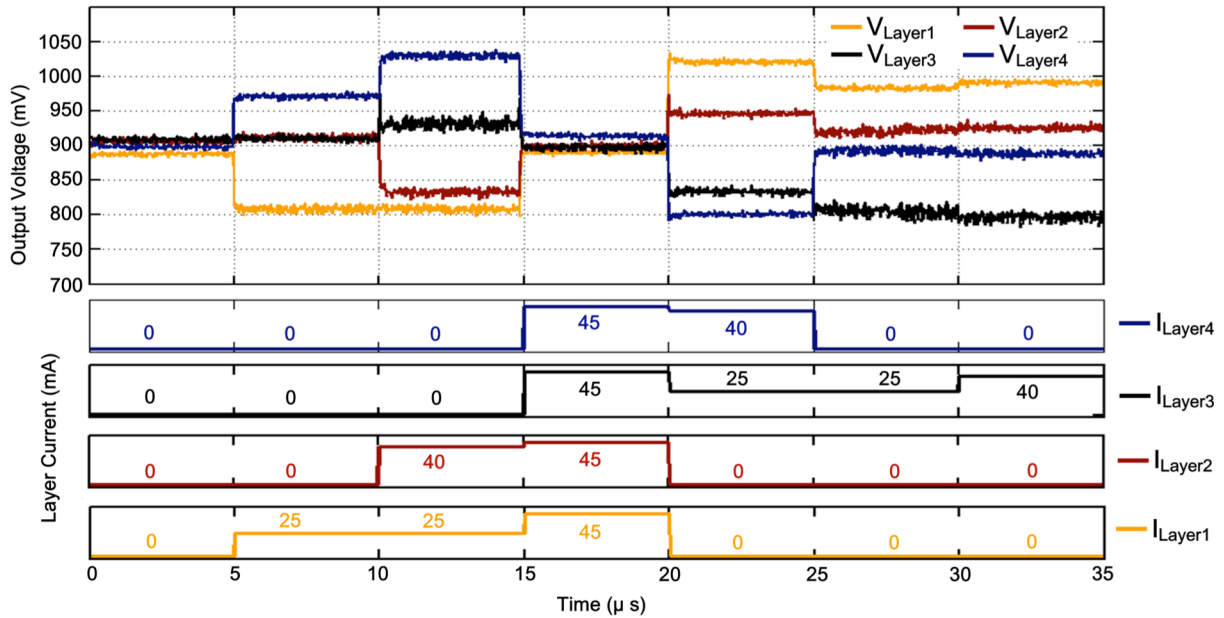
Figure 4.23: Measured transient responses with dynamic load currents in multiple layers. (V_{IN}=3.6V; all reference voltages are set to 800mV.)

lower for the bottom layer. The SLSCC maintains a minimal voltage of V_{LAYER1} around 800mV. I_{LAYER2} increases to 40mA at t=10µs. As a result, V_{LAYER2} decreases. All the layer currents increase to 45mA at t=15µs. The voltage across all layers redistributes to a balanced state around 900mV. At t=20µs, the load currents in the bottom two layers go down to 0. The SLSCC regulates to maintain a minimal voltage of about 800mV across the top two layers. After t=25µs, only the third layer consumes current. I_{LAYER3} increases from 25mA to 40mA at t=30µs. The SLSCC regulates the layer voltages so that V_{LAYER3} is around 800mV.

	[4] VLSI 10	[8] ISSCC 14	[9] ISSCC 13	[10] ISSCC 13	[15] JSSC 14	This Work
Technology	45nm	0.25um	180nm	65nm	22nm	40nm
Capacitor technology	Trench	MIM	On-Chip	MOS	MIM	MOS
Total capacitance	—	3nF	2.24nF	3.88nF	—	4.5nF
V_{IN}	2V	2.5V	3.4V-4.3V	3V-4V	1.23V	3.6V
V_{OUT}	0.95V-1.05V	0.1V-2.18V	0.9V-1.5V	1V	0.45V-1V	0.8V-1V
Voltage stacking	2-Way	None	None	None	None	4-Way
Quoted efficiency (η)	90%	60%	72%	73%	70%	65% *
Conv. ratio @ η	2:1	4:1	4:1	3:1	2:1	4:1
P_{OUT} @ η	—	1mW	0.27mW	122mW	6.4mW	17.5mW *
Power density $(mW/mm^2)@η$	2185	0.215	0.16	190	—	21.1 *

* For consistency, the SLSCC only delivers power to the bottom layer (Layer 1).

Figure 4.24: Performance summary.
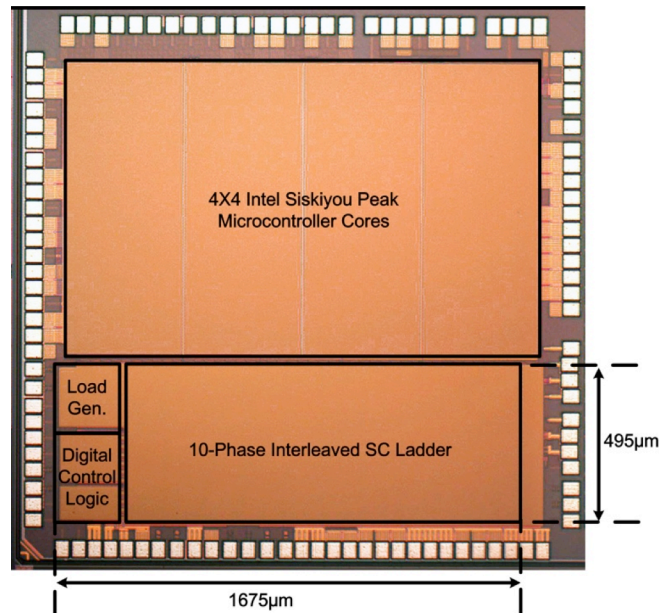
Figure 4.25: Die micrograph of the SLSCC and the stacked processor.

For consistency, Figure 4.24 compares this work to prior work assuming power delivery to a single layer, but note that both conversion efficiency and power density improve when power is delivered to multiple layers, as required by voltage stacking.

A chip micrograph of the 0.829mm^2 SLSCC is shown in Figure 4.25.

Chapter 5

Conclusions and Technologies on the Horizon

In this thesis, I have presented the design and implementation of two fully integrated SC DC-DC voltage regulators. The experimental results verified the effectiveness of many proposed design techniques, such as flying capacitor charge recycling and hybrid feedback control. By combining the switching regulator with voltage stacking, the new power delivery method significantly reduces the overall cost of the SoC power delivery circuitry, including the conversion loss and area overhead.

As mobile SoCs become more and more complicated to fulfill the increasing performance requirements of consumer electronics, delivering power to these SoCs is and will remain to be a very challenging issue. As discussed in this thesis, fully integrated SC converters provide a promising solution to reduce the power delivery system's complexity, footprint, and loss. I expected that if we combine the circuit and system techniques proposed in this thesis with advanced fabrication technologies, such as trench capacitors, we can build even better fully integrated power delivery solutions.

Looking ahead, there are various fabrication, circuit, and system techniques that can improve the performance of FIVR-based power delivery solutions:

1. **Better fabrication technologies.** Most mobile SoCs are fabricated with very advanced (i.e., 22nm or better) digital processes. These processes are usually very expensive, and the cost is a huge concern in consumer electronics. Economically, the area overhead for the FIVRs needs to be very small so that they can be integrated with commercial SoCs. Even though the power density of FIVRs has been much improved in the past decade, there is still a long way to go.

Better and cheaper on-chip magnetics and capacitors are needed to further improve the power density of FIVRs. Better passive components will also reduce the conversion loss, which is very important, especially in SoCs, due to thermal concerns.

2. **Hybrid converter topologies.** Most of the published FIVRs rely on either switched-capacitor or switched-inductor topologies. In the future, I expect more effort to combine the two topologies, exploiting the advantages of both. Ideally, a hybrid topology will lead to an FIVR with a high power density and high efficiency over a wide input/output voltage/current range. There have been some recent endeavors in this direction. For example, Wonyoung Kim has built a fully integrated 3-level converter [53] that uses switched-capacitor circuits to create three fixed voltage levels and switched-inductor circuits to provide finer voltages between the fixed voltages.

3. **Regulator-processor co-design.** Even though FIVRs have attracted much attention from both academia and industry, there are not many examples that integrate FIVRs with commercial SoCs. With the improvements in FIVR implementations, more study is needed to explore the performance trade-offs of integrating FIVRs in commercial SoCs with real workloads. This will not only provide insights about the applications run in the SoCs, but will also guide FIVR designs. For example, SoC designers are worried that FIVR loss could increase the heat in SoCs and consequently lead to more difficult thermal management.

Such problems can only be explored and resolved when FIVRs and the rest of the SoC are studied and co-designed as a complete system. In the case of voltage stacking, more research is needed to explore methodologies to schedule the workload on the stacked cores so that the power consumed by the load in all stacked layers is balanced.

With all of the potential benefits that FIVRs bring, it will take time for this technology to penetrate into various applications. Wearable electronics and micro-robots may be suitable applications that will embrace FIVRs early, since battery life and device footprint are very important product specifications that can be improved through FIVRs.

Bibliography

1. N. Kurd, M. Chowdhury, et al., "Haswell: A Family of IA 22nm Processors," *IEEE ISSCC Dig. Tech. Papers*, pp. 112-113, Feb. 2014

2. H. Mair, G. Gammie, et al., "A Highly Integrated Smartphone SoC Featuring a 2.5GHz Octa-Core CPU with Advanced High-Performance and Low-Power Techniques," *IEEE ISSCC Dig. Tech. Papers*, pp. 424-425, Feb. 2015

3. R. Islam, A. Sabbavarapu, et al., "Next Generation Intel® ATOM™ Processor Based Ultra Low Power SoC for Handheld Applications," *IEEE ASSCC*, pp. 1-4, Nov. 2010

4. Qualcomm Snapdragon mobile SoC processor, online documents, website: https://www.qualcomm.com/products/snapdragon, accessed on May, 17th, 2015

5. Online documents, website: http://www.macrumors.com/2013/09/27/a-closer-look-at-apples-a7-chip-from-the-iphone-5s/, accessed on May, 18th, 2015

6. Online documents, website: http://www.xperiablog.net/2014/04/07/qualcomms-new-64-bit-snapdragon-810-and-808-chipsets-to-land-in-2015/, accessed on May, 18th, 2015

7. H. Lakdawala, M. Schaecher, et al., "A 32 nm SoC With Dual Core ATOM Processor and RF WiFi Transceiver," *IEEE J. Solid-State Circuits,* vol. 48, no. 1, pp. 91–103, 2013

8. D. Jeon, Y. Chen, et al., "An Implantable 64nW ECG-Monitoring Mixed-Signal SoC for Arrhythmia Diagnosis," *IEEE ISSCC Dig. Tech. Papers*, pp. 416-417, Feb. 2014

9. S. Sanders et al., "The road to fully integrated DCDC conversion via the switched-capacitor approach," *IEEE Trans. Power Electronics*, vol. 28, pp. 4146–4155, 2013.

10. G. Pique, et al., "Survey and Benchmark of Fully Integrated Switching Power Converters: Switched-Capacitor Versus Inductive Approach," *IEEE Transactions on Power Electronics*, Vol. 28, No. 9, Sep. 2013

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

12. K. Rangan, G. Wei, and D. Brooks. "Thread motion: fine-grained power management for multi-core systems," *ACM SIGARCH Computer Architecture News*. Vol. 37. No.

3. ACM, 2009.

13. A. Sinkar, H. Wang, and N. S. Kim. "Workload-aware voltage regulator optimization for power efficient multi-core processors." *IEEE Design, Automation & Test in Europe Conference & Exhibition (DATE), 2012*, 2012.

14. S. Eyerman, and L. Eeckhout. "Fine-grained DVFS using on-chip regulators," *ACM Transactions on Architecture and Code Optimization (TACO)*, 2011

15. P. Hammarlund, et al. "Haswell: The fourth-generation Intel core processor," *IEEE Micro*, pp. 6-20, 2014

16. B. Edward, et al. "FIVR—Fully integrated voltage regulators on 4th generation Intel® Core™ SoCs," *IEEE Applied Power Electronics Conference and Exposition (APEC)*, 2014.

17. X. Wang, et al. "Characterizing power delivery systems with on/off-chip voltage regulators for many-core processors," *Proceedings of the conference on Design, Automation & Test in Europe*, 2014.

18. J. Gjanci, "On-Chip Voltage Regulation for Power Management in System-on-Chip," Diss. University of Illinois at Chicago, 2008.

19. X. Zhang, et al. "Supply-noise resilient adaptive clocking for battery-powered aerial microrobotic System-on-Chip in 40nm CMOS," *IEEE Custom Integrated Circuits Conference (CICC)*, 2013.

20. P. Zhou, et al. "Exploration of on-chip switched-capacitor DC-DC converter for multicore processors using a distributed power delivery network," *IEEE Custom Integrated Circuits Conference (CICC)*, 2011

21. T. Tong, et al. "A fully integrated battery-connected switched-capacitor 4: 1 voltage regulator with 70% peak efficiency using bottom-plate charge recycling," *IEEE Custom Integrated Circuits Conference (CICC)*, 2013.

22. R. Jevtic, et al. "Per-Core DVFS With Switched-Capacitor Converters for Energy Efficiency in Manycore Processors," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 23, Issue. 4

23. E. Alon and M. Horowitz, "Integrated regulation for energy-efficient digital circuits," *IEEE J. Solid-State Circuits*, Vol. 43, No. 8, 2008.

24. G. Pique, "A 41-Phase Switched-Capacitor Power Converter with 3.8mV Output Ripple and 81% Efficiency in Baseline 90nm CMOS," *IEEE ISSCC Dig. Tech. Papers*, pp. 98-99, Feb. 2012

25. H.-P Le, S. Sanders, and E. Alon, "Design techniques for fully integrated switched-capacitor dc-dc converters," *IEEE J. Solid-State Circuits,* vol. 46, no. 9, pp. 2120–2131, 2011

26. T. M. Andersen, F. Krismer, J. Walter et al., "A Sub-ns Response On-Chip Switched-Capacitor DC-DC Voltage Regulator Delivering 3.7W/mm2 at 90% Efficiency Using Deep-Trench Capacitors in 32nm SOI CMOS," *IEEE ISSCC Dig. Tech. Papers*, pp. 90-91, Feb. 2014

27. J.Seo, et al., "Deep Trench Capacitors for Switched Capacitor Voltage Converters," presented at 2012 PowerSoC conference, online documents, website: http://powersoc2012.org/session-4/4.5%20_j.seo.pdf, accessed on May, 20th, 2015

28. R. Jain, B.M. Geuskens, S.T. Kim, et al., "A 0.45–1 V Fully-Integrated Distributed Switched Capacitor DC-DC Converter With High Density MIM Capacitor in 22 nm Tri-Gate CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 4, pp. 917–927, 2014.

29. Y. K. Ramadass and A. P. Chandrakasan, "Voltage scalable switched capacitor DC-DC converter for ultra-low-power on-chip applications," in *Proc. Power Electronics Specialists Conf.*, 2007, pp. 2353–2359.

30. M. Seeman and R. Jain, "Single-Bound Hysteretic Regulation of Switched Capacitor Converters," U.S. Patent 2011/0 074 371, Mar. 31, 2011.

31. Y. K. Ramadass, A. Fayer, and A. P. Chandrakasan, "A Fully-Integrated Switched-Capacitor Step-Down DC-DC Converter With Digital Capacitance Modulation in 45 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2557–2565, 2010.

32. T. Van Breussegem and M. Steyaert, "Monolithic capacitive DC-DC converter with single boundary–multiphase control and voltage domain stacking in 90 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1715-1727, July 2011.

33. R. Jain and S. Sanders, "A 200 mA switched capacitor voltage regulator on 32 nm CMOS and regulation schemes to enable DVFS," *IEEE European Conf. Power Electronics,* pp. 1–10, 2011.

34. H.-P. Le, J. Crossley, S.R. Sanders, and E. Alon, "A Sub-ns response fully integrated battery-connected switched-capacitor voltage regulator delivering 0.19W/mm2 at 73% efficiency," *IEEE ISSCC Dig. Tech. Papers*, pp. 372-373, Feb. 2013

35. D. El-Damak et al., "A 93% efficiency reconfigurable switched-capacitor DC-DC converter using on-chip ferroelectric capacitors," *IEEE ISSCC Dig. Tech. Papers*, pp. 374–375. 2013.

36. T. M. Anderson, F. Krismer, and J. W. Kolar, "A 4.6 W/mm2 power density 86% efficiency on-chip switched capacitor DC-DC converter in 32 nm SOI CMOS," in *Proc. IEEE Appl. Power Electron. Conf.*, Mar. 2013.

37. L.G. Salem P.P. Mercier, "An 85%-Efficiency Fully Integrated 15-Ratio Recursive Switched-Capacitor DC-DC Converter with 0.1-to-2.2V Output Voltage Range," *IEEE ISSCC Dig. Tech. Papers*, pp. 88-89, Feb. 2014

38. S. Bang et al., "A Fully Integrated Successive-Approximation Switched-Capacitor DC-DC Converter with 31mV Output Voltage Resolution," *ISSCC* Dig. Tech. Papers, pp. 370-371, Feb. 2013.

39. S. S. Kudva and R. Harjani, "Fully-integrated capacitive DC-DC converter with all digital ripple mitigation technique," *IEEE J. Solid-State Circuits*, vol. 48, no. 9, pp. 1910–1920, Sep. 2013.

40. S. Kim, et al. "Enabling wide autonomous DVFS in a 22nm graphics execution core using a digitally controlled hybrid LDO/switched-capacitor VR with fast droop mitigation," *ISSCC* Dig. Tech. Papers, pp. 154-155, Feb. 2015.

41. T. Andersen, et al. "A deep trench capacitor based 2: 1 and 3: 2 reconfigurable on-chip switched capacitor DC-DC converter in 32 nm SOI CMOS," *IEEE Applied Power Electronics Conference and Exposition (APEC)*, 2014.

42. M. Seeman and S. Sanders, "Analysis and optimization of switched-capacitor DC–DC converters," *IEEE Trans. Power Electron.*, vol. 23, no. 2, pp. 841–851, Mar. 2008.

43. M. Seeman, "Analytical and Practical Analysis of Switched-Capacitor DC-DC Converters," Technical Report No. UCB/EECS-2006-111, September 1, 2006.

44. G. V. Pique and E. Alarcon, "CMOS Integrated Switching Power Converters: A Structured Design Approach," 1st ed. Berlin, Germany: Springer, 2011.

45. C. Mathúna, et al. "Review of integrated magnetics for power supply on chip (PwrSoC)," *IEEE Transactions on Power Electronics, Vol. 27, No.11* 2012.

46. K. Luria, J. Shor, M. Zelikson, and A. Lyakhov, "Dual-use low-drop-out regulator/power gate with linear and on-off conduction modes for microprocessor on-

die supply voltages in 14nm," *IEEE ISSCC Dig. Tech. Papers*, pp. 156-157, Feb. 2015

47. Z. Toprak-Deniz, et al. "Dual-loop system of distributed microregulators with high DC accuracy, load response time below 500ps, and 85mV dropout voltage," *IEEE Symposium on VLSI Circuits (VLSIC)*, pp. 274-275, Jun. 2011

48. E. Fluhr., et al. "The 12-Core POWER8™ Processor With 7.6 Tb/s IO Bandwidth, Integrated Voltage Regulation, and Resonant Clocking," (2015). *IEEE J. Solid-State Circuits,* vol. 50, no. 1, pp. 10–23, 2015

49. S. S. Chong, and P. K. Chan. "A Sub-1V Transient-Enhanced Output-Capacitor less LDO Regulator With Push–Pull Composite Power Transistor," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 22, No. 11, pp. 2297-2306, 2014

50. Y. Lam and W. Ki, "A 0.9V 0.35um Adaptively Biased CMOS LDO Regulator with Fast Transient Response," *IEEE ISSCC* Dig. Tech. Papers, pp. 442-443, Feb. 2008

51. M. Wieckowski, G. K. Chen, M. Seok, D. Blaauw, and D. Sylvester, "A hybrid DC-DC converter for sub-microwatt sub-1V implantable applications," *IEEE Symposium on VLSI Circuits (VLSIC)*, pp. 166-167, Jun. 2009.

52. Y. Lu, W.-H. Ki, and C. P. Yue, "A 0.65ns-response-time 3.01ps FOM fully-integrated low-dropout regulator with full-spectrum power-supply-rejection for wideband communication systems," *IEEE ISSCC Dig. Tech. Papers*, pp. 306-307, Feb. 2014

53. W. Kim, D. Brooks, G.Y. Wei, "A Fully-Integrated 3-Level DC-DC Converter for Nanosecond-Scale DVFS," *IEEE J. Solid-State Circuits*, vol. 47, no. 1, pp. 206–219, 2011.

54. W. Kim, D. Brooks, G.Y. Wei, "A Fully-Integrated 3-Level DC/DC Converter for Nanosecond-Scale DVS with Fast Shunt Regulation," *IEEE ISSCC Dig. Tech. Papers*, Feb. 2011

55. N. Sturcken, M. Petracca, S. Warren, et al., "A Switched-Inductor Integrated Voltage Regulator With Nonlinear Feedback and Network-on-Chip Load in 45 nm SOI," *IEEE J. Solid-State Circuits*, vol. 47, no. 8, pp. 1935–1945, 2012.

56. K Kesarwani, R Sangwan, J.T. Stauth, "A 2-Phase Resonant Switched-Capacitor Converter Delivering 4.3W at 0.6W/mm2 with 85% Efficiency," *IEEE ISSCC Dig. Tech. Papers*, pp. 86-87, Feb. 2014

57. M. Belloni, E. Bonizzoni, and F. Maloberti, "High Efficiency DC-DC Buck Converter with 60/120-MHz Switching Frequency and 1-A Output Current," *IEEE European Solid-State Circuits Conference*, 2009.

58. N. Sturcken, E. O'Sullivan, et al, "A 2.5D integrated voltage regulator using coupled-magnetic-core inductors on silicon interposer de- livering 10.8 A/mm2," *IEEE ISSCC Dig. Tech. Papers*, Feb. 19–23, pp. 400–402, 2012.

59. H. J. Bergveld, K. Nowak, et al, "A 65-nm-CMOS 100-MHz 87% efficient DC-DC down converter based on dual-die System-in- Package integration," *IEEE Energy Conversion Congress and Exposition*, 2009.

60. P. Hazucha, G. Schrom, et al., "A 233MHz, 80-87% Efficient, Integrated, 4-Phase DC-DC Converter in 90nm CMOS," *IEEE Symp. VLSI Circuits*, 2004.

61. Harvard Robobee, Online, website: http://wyss.harvard.edu/viewpage/457

62. M. Karpelson, J.P. Whitney, G.-Y. Wei, and R.J. Wood, "Energetics of flapping-wing robotic insects: towards autonomous hovering flight," *IEEE/RSJ Int. Conference on Intelligent Robots and Systems*, October 2010.

63. P. Chirarattananon, K. Ma, and R.J. Wood, "Adaptive Control of a Millimeter-Scale Flapping-Wing Robot," *Bioinspiration & Biomimetics*, vol. 9, no. 2, 2014.

64. M. Karpelson, R.J. Wood, and G.-Y. Wei, "Low Power Control IC for Efficient High-Voltage Piezoelectric Driving in a Flying Robotic Insect," *IEEE Symp. VLSI Circuits*, 2011.

65. K. Ma, P. Chirarattanon, S. Fuller, and R.J. Wood, "Controlled Flight of a Biologically Inspired, Insect-Scale Robot, " *Science*, vol. 340, pp. 603-607, 2013.

66. X. Zhang, M. Lok, et al., "A Multi-Chip System Optimized for Insect-Scale Flapping-Wing Robots," *IEEE Symp. VLSI Circuits, Jun.* 2015

67. M. Lok, D. Brooks, et al., "Design and analysis of an integrated driver for piezoelectric actuators," *IEEE Energy Conversion Congress and Exposition (ECCE),* 2013.

68. X. Zhang, T. Tong, D. Brooks and G.-Y. Wei, "Evaluating Adaptive Clocking for Supply-Noise Resilience in Battery-Powered Aerial Microrobotic System-on-Chip", in *IEEE Trans. Circuits Syst*. I, Reg. Papers, vol. 61, no. 8, pp. 2309–2317, Aug. 2014.

69. D. Draxelmayr et al., "A 6b 600MHz 10mW ADC array in digital 90nm CMOS," *IEEE ISSCC, Dig. Tech. Papers*, pp. 264- 264, Feb. 2004

70. S. Rajapandian et al., "High-voltage power delivery through charge recycling," *IEEE J. Solid-State Circuits*, vol. 41, no. 6, pp. 1400–1410, Jun. 2006.

71. S. Rajapandian et al., "Implicit DC–DC Downconversion Through Charge-Recycling," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 846–852, Apr. 2005.

72. S.K. Lee et al., "Evaluation of Voltage Stacking For Near-Threshold Multicore Computing," *ACM/IEEE ISLPED*, pp. 373–378, July 2012.

73. S.K. Lee et al., "A 16-core voltage-stacked system with an integrated switched-capacitor DC-DC converter," *IEEE Symp. VLSI Circuits*, June 2015.

74. K. Mazumdar and M. Stan. "Breaking the power delivery wall using voltage stacking," *In Proceedings of the Great Lakes Symposium on VLSI*, pp. 51–54, May 2012.

75. K. Kesarwani, C. Schaef, C. R. Sullivan, and J. T. Stauth, "A MultiLevel Ladder Converter Supporting Vertically Stacked Digital Voltage Domains," *in Proc. Twenty-Eigth Annual IEEE Applied Power Electronics Conf. and Exposition (APEC)*, pp. 429–434, 2013.

76. L. Chang, et al. "A Fully-Integrated Switched-Capacitor 2:1 Voltage Converter with Regulation Capability and 90% Efficiency at 2.3A/mm2," IEEE Symp. VLSI Circuits, pp. 55-56, June 2010.

77. W. Kim, et al. "System level analysis of fast, per-core DVFS using on-chip switching regulators." *IEEE International Symposium on High Performance Computer Architecture* (HPCA), 2008.

78. G. Yan, et al. "AgileRegulator: A hybrid voltage regulator scheme redeeming dark silicon for power efficiency in a multicore architecture," *IEEE International Symposium on High Performance Computer Architecture* (HPCA), 2012.

79. Texas Instruments product guide, website:http://www.ti.com/lsds/ti/power-management/dc-dc-switching-regulator-products.page, accessed on June, 10th, 2015