A Fully-Integrated 3-Level DC-DC Converter for Nanosecond-Scale DVFS

Wonyoung Kim, *Student Member, IEEE*, David Brooks, *Member, IEEE*, and Gu-Yeon Wei, *Member, IEEE*

***Abstract*—On-chip DC-DC converters have the potential to offer fine-grain power management in modern chip-multiprocessors. This paper presents a fully integrated 3-level DC-DC converter, a hybrid of buck and switched-capacitor converters, implemented in 130 nm CMOS technology. The 3-level converter enables smaller inductors (1 nH) than a buck, while generating a wide range of output voltages compared to a 1/2 mode switched-capacitor converter. The test-chip prototype delivers up to 0.85 A load current while generating output voltages from 0.4 to 1.4 V from a 2.4 V input supply. It achieves 77% peak efficiency at power density of 0.1 W/mm^2 and 63% efficiency at maximum power density of 0.3 W/mm^2. The converter scales output voltage from 0.4 V to 1.4 V (or vice-versa) within 20 ns at a constant 450 mA load current. A shunt regulator reduces peak-to-peak voltage noise from 0.27 V to 0.19 V under pseudo-randomly fluctuating load currents. Using simulations across a wide range of design parameters, the paper compares conversion efficiencies of the 3-level, buck and switched-capacitor converters.**

***Index Terms*—DC-DC conversion, dynamic voltage and frequency scaling, fully integrated converter, switching converter, 3-level.**

I. Introduction

MODERN processors pack increasing amounts of functionality and numbers of components into a single chip, whether in the form of multiple cores or specialized hardware (e.g., memory controllers, hardware accelerators, etc.). Concurrently, power has emerged as a primary concern for both high-performance and battery-driven systems. Current processors employ numerous low-level techniques to reduce power consumption, such as aggressive clock and power gating and dynamic voltage and frequency scaling (DVFS). Current implementations of DVFS track coarse-grained (often OS-level) fluctuations in processor activity to improve energy utilization or boost performance. However, detailed study of processor workloads shows that high levels of activity fluctuation exist at finer timescales that traditional DVFS cannot track [1]. While highly efficient (>90%), the bulky off-chip components and low switching frequencies of conventional off-chip DC-DC converters limit voltage transition times. The high board-level footprint, cost, and package issues related to off-chip converters also hamper the use of multiple voltage domains to improve energy efficiency in multi-core designs. Given these drawbacks of off-chip DC-DC converters, there has been a surge of interest in building on-chip converters [2]–[13] to implement multiple on-chip voltage domains. This paper presents the design and experimental results of a fully integrated 3-level DC-DC converter [13], which enables fast voltage scaling at nanosecond timescales and 77% peak conversion efficiency at 0.3 W/mm^2 maximum power density. (We calculate power density excluding the 10 nF output decoupling capacitors, assuming that the converter uses existing on-chip decoupling capacitors on the processor.)

On-chip DC-DC converter designs range from buck converters to switched-capacitor converters to low-dropout linear regulators. Linear regulators have a maximum efficiency limit given by the ratio of output voltage to input voltage; they suffer from low efficiency at high ratios. In contrast, switching converters can maintain high efficiency across a wide range of output voltages. There are two types of switching DC-DC converters commonly used for low step-down ratios – buck and switched-capacitor (SC) converters. Shown in Fig. 1(a), the buck converter relies on an inductor to generate a step-down voltage on the output capacitor, $\mathrm{C_{OUT}}$. The buck converter creates a square-wave voltage – of varying duty cycles (D) – at the output of the power FETs ($\mathrm{V_X}$). While traditional buck converters rely on single pull-up and pull-down power FETs, series stacks of switches enable use of thin-oxide devices in integrated voltage converters [7]. By adjusting the duty cycle of $\mathrm{V_X}$, buck converters can provide a wide range of $\mathrm{V_{OUT}}$. However, the buck converter requires a large, high-quality inductor, which is difficult to integrate on-chip.

In contrast, the SC converter uses flying capacitors ($\mathrm{C_{FLY}}$), without an inductor, to nominally divide the high input voltage ($\mathrm{V_{IN}}$) by pre-determined integer ratios. For example, the SC converter in Fig. 1(b) divides $\mathrm{V_{IN}}$ by two as it iterates between two phases of capacitor configurations – series-stack and parallel. Although it does not need inductors, this particular configuration of SC can only step $\mathrm{V_{OUT}}$ down to values lower than $\mathrm{V_{IN}}/2$. Additional step-down ratios, such as 1/3 and 2/3, are also possible as demonstrated by Ramadass *et al.* [9] and Le *et al.* [10], in order to extend the range of output voltage conversion. However, the added power switches needed for the additional capacitor configurations can exacerbate conversion loss.

As shown in Fig. 1(c), a 3-level converter merges characteristics of both inductor-based buck and SC converters to gain the benefits of both [12], [14]. Similar to the buck, the output LC pair of the 3-level converter filters $\mathrm{V_X}$ to generate $\mathrm{V_{OUT}}$ with small ripple. While the $\mathrm{V_X}$ of the buck converter swings between 0 and $\mathrm{V_{IN}}$, the $\mathrm{V_X}$ of the 3-level converter either swings between 0 and $\mathrm{V_{IN}}/2$, or $\mathrm{V_{IN}}/2$ and $\mathrm{V_{IN}}$, to convert $\mathrm{V_{OUT}}$ to voltages under and over $\mathrm{V_{IN}}/2$, respectively. The switching

Manuscript received April 19, 2011; revised June 27, 2011; accepted September 08, 2011. Date of publication November 18, 2011; date of current version December 23, 2011. This paper was approved by Guest Editor Tanay Karnik.

The authors are with Harvard University, Cambridge, MA 02138 USA (e-mail: wonyoung@eecs.harvard.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2011.2169309

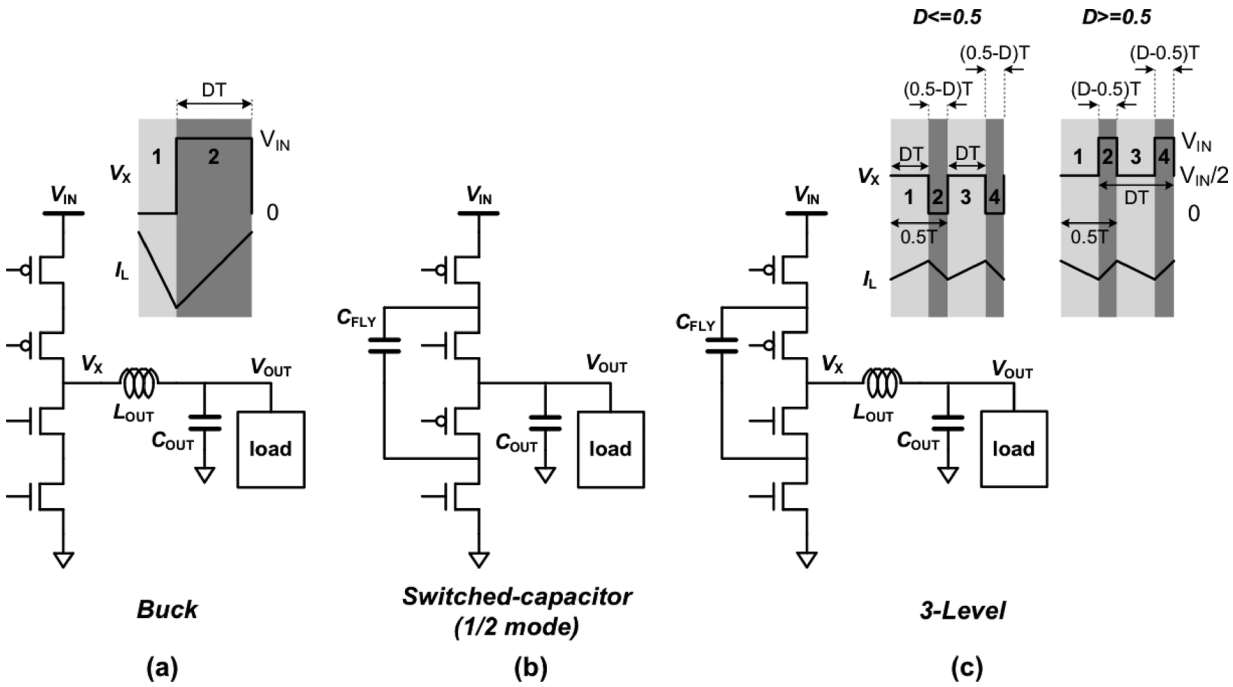

Fig. 1. Power FET and output filters of (a) buck, (b) switched-capacitor, and (c) 3-level converters.

action of the power FETs, combined with the flying capacitor, effectively generates a third voltage, $\mathrm{V_{IN}}/2$ (hence the name 3-level converter), and adjusts D to set $\mathrm{V_{OUT}}$ across a wide range of voltage levels. Notice that $\mathrm{V_X}$ of the 3-level converter swings with half the amplitude and at twice the frequency compared to that of the buck. Both of these attributes enable the 3-level converter to exhibit smaller inductor current ripple and voltage ripple on $\mathrm{V_{OUT}}$ or to use a smaller inductor for the same ripple target.

Prior 3-level converter designs include an off-chip converter for envelop tracking [14] and an integrated 3-level converter with 27 nH bondwire inductors [12]. We build upon these works and present a fully-integrated 3-level converter with 1 nH on-chip spiral inductors. One-nanoHertz inductors, placed on top of flying capacitors to minimize area overhead, enable voltage transition across 1 V within 20 ns, which is $100\times$ faster than previously published data [3]. The converter can be externally programmed to adjust design parameters (switching frequency, number of phases and power FET size) to study trade-offs associated with different design parameters. We also add fast shunt regulation to the converter to reduce voltage noise.

The next section studies how design parameters affect conversion loss in 3-level converters and compares the conversion efficiencies of 3-level to those of buck and SC converters. Then Section III presents a detailed, circuit-level description of the 3-level converter design that was implemented in a test-chip prototype using a 130 nm CMOS process technology. Experimental results from the test-chip, in Section IV, demonstrate fast voltage scaling and high conversion efficiency across a wide range of output voltages. Section V summarizes our findings.

II. 3-Level Converter Topology

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

A. Design Parameters for 3-Level Converters

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

Fig. 2 presents simulated conversion efficiencies of a 3-level converter running in continuous conduction mode (CCM) acquired using a fast circuit simulator HSIM, set to the highest simulation accuracy level. As specified in the table in Fig. 2, the converter operates with DC load current ranging from 0.2 A to 1 A for output voltages ranging from 0.6 to 1.35 V. Load current scales quadratically with output voltage to mimic a processor operating with DVFS. Simulations sweep design parameters to find the maximum efficiency for each output voltage value. The converter uses 1 nH inductors with 400 mΩ series. Up to four copies of power FETs and inductors can be interleaved to form multi-phase converters [2] to distribute current flow and reduce output voltage ripple. Fig. 11 presents an example of a 4-phase converter that can dynamically change the number of operating phases according to load levels. Fig. 2 shows that optimizing design parameters significantly improves conversion efficiency compared to a converter using fixed parameters (100 MHz frequency, two phases, 48 mm total power FET width).

Fig. 3 shows how to determine switching frequency and number of phases to maximize efficiency. When duty cycle is in the vicinity of 50%, a converter needs to operate at low switching frequency with maximum number of phases. As duty cycle deviates from 50%, the converter needs to increase

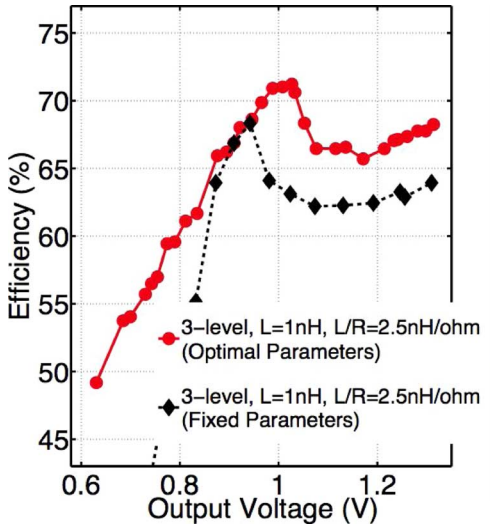

input voltage	2.4V
switching frequency	10-500MHz
number of phases	1,2,4 (buck, 3-level) 8 (SC)
Total power FET width	8-256mm
output voltage (V_{OUT})	0.6 - 1.35
load current (A)	0.55* V_{OUT}^2

Fig. 2. Simulated conversion efficiencies of 3-level converters with fixed and optimal design parameters. Table shows the range of design parameters used in simulations.

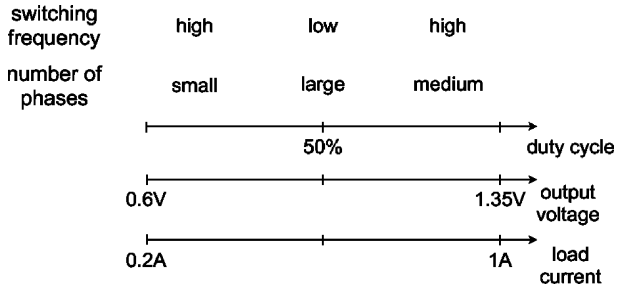

Fig. 3. Design parameters that maximize efficiencies across duty cycle, output voltage, and load current ranges.

switching frequency and reduce the number of phases. The selection of design parameters aim to balance different sources of losses. High switching frequencies increase switching loss (CV^2f), but reduce resistive loss ($I^2_{RMS}R$) caused by inductor current ripple ($\Delta I_{L,PP}$). Assuming a converter operating under CCM with a triangular wave for the inductor current (I_L), the following equation shows that both DC value and peak-to-peak ripple of the inductor current contribute to $I^2_{RMS}R$ loss:

$$I^2_{L,RMS} = I^2_{L,DC} + \frac{\Delta I^2_{L,PP}}{12}. \quad (1)$$

Shown in Fig. 4, $\Delta I_{L,PP}$ of the 3-level converter reaches minimum at 50% duty cycle, increases as duty cycle deviates from 50% and decreases again when duty cycle goes below 25% or over 75%. Taking advantage of small $\Delta I_{L,PP}$ at duty cycles near 50%, the 3-level converter minimizes switching loss by running at low frequencies. As $\Delta I_{L,PP}$ grows at duty cycles away from 50%, the converter runs at higher frequencies to suppress $I^2_{RMS}R$ loss, albeit with larger switching loss. Increasing switching frequency at light loads contradicts the conventional wisdom of using pulse frequency modulation (PFM) in buck converters to reduce frequency at light loads. As duty cycle deviates from 50%, $\Delta I_{L,PP}$ of the 3-level increases while that of the buck converter decreases. This allows the buck to reduce frequency at light loads, while forcing the 3-level converter to increase frequency.

To study how the number of phases affects conversion loss, the following equation expands (1) to a multi-phase 3-level converter, which consists of multiple interleaved copies of a single phase converter:

$$\begin{aligned} I^2_{L,RMS} &= \left(I^2_{L,DC} + \frac{\Delta I^2_{L,PP}}{12}\right) \times N_{PH} \\ &= \frac{I^2_{LOAD}}{N_{PH}} + \frac{\Delta I^2_{L,PP}}{12} \times N_{PH} \\ &(N_{PH}: \text{number of phases}, \\ &I_{L,DC}: \text{DC inductor current per phase}, \\ &\Delta I_{L,PP}: \text{inductor current ripple per phase}). \end{aligned} \quad (2)$$

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

B. Comparison to Buck and SC Converters

Remaining simulation plots (Figs. 5, 6, 7, 11) present efficiencies with optimized design parameters using ranges specified in Fig. 2. Fig. 5 presents a similar efficiency versus output voltage plot of the buck converter for different inductance values, assuming CCM operation across all load conditions. Simulations use a buck converter design similar to one proposed in [7] with series stacks of power FETs using thin-oxide devices. For the same inductor quality ($L/R = 2.5\,nH/\Omega$), larger inductance reduces $\Delta I_{L,PP}$ while increasing inductor series resistance (R_L). At low load currents, 2 nH and 4 nH inductors achieve higher

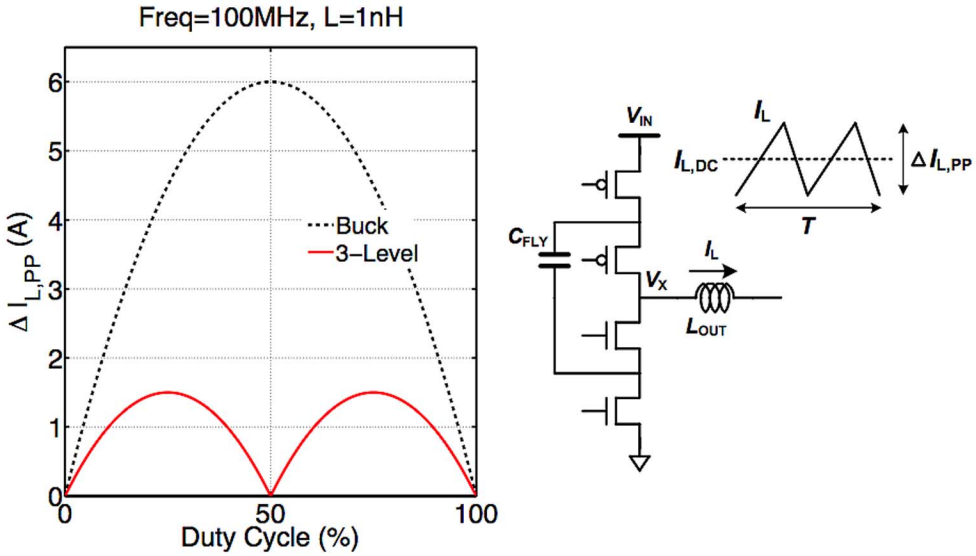

Fig. 4. Simulated peak-to-peak inductor current ripple ($\mathrm{I_{L,PP}}$) of 3-level and buck converters in continuous conduction mode (CCM).

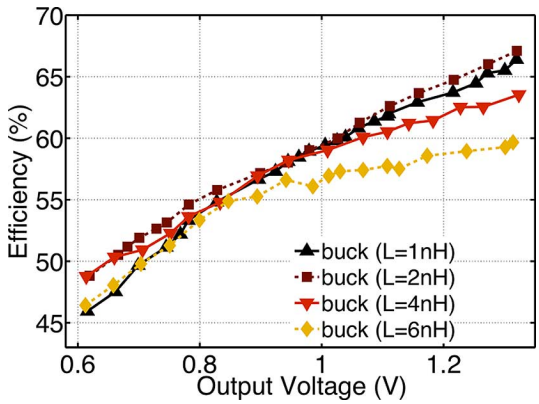

Fig. 5. Simulated conversion efficiencies of buck converters across inductance values ($\mathrm{L/R = 2.5\ nH/\Omega}$).

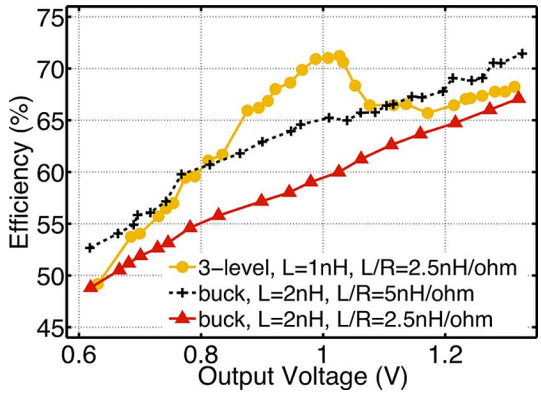

Fig. 6. Simulated conversion efficiencies of 3-level and buck converters across inductor qualities.

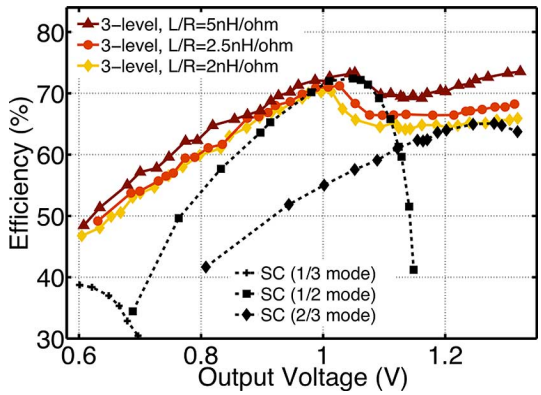

Fig. 7. Simulated conversion efficiencies of 3-level and switched-capacitor converters across inductor qualities.

efficiencies than 1 nH and 6 nH, which suffer from large $\mathrm{\Delta I_{L,PP}}$ and $\mathrm{R_L}$, respectively. At high load currents, $\mathrm{R_L}$ significantly affects conversion loss, allowing 1 nH and 2 nH to achieve higher efficiencies than 4 nH and 6 nH. We choose 2 nH for further comparisons to 3-level converters.

Since a 3-level converter adds flying capacitors on-die, it occupies larger die area than a buck using the same inductor. Assuming the buck converter can use additional die area to implement larger, higher quality inductors, Fig. 6 compares conversion efficiencies of 3-level and buck converters, providing similar or higher quality inductors to buck converters. The 3-level converter uses 16 nF of $\mathrm{C_{FLY}}$, and both buck and 3-level converters use 10 nF of $\mathrm{C_{OUT}}$, operating with up to four phases. To make a fair comparison between converters with different $\mathrm{V_{OUT}}$ ripple characteristics, as proposed in [10], we calculate conversion efficiency using the minimum value of $\mathrm{V_{OUT}}$ ripple, instead of the average $\mathrm{V_{OUT}}$ value. For the same inductor quality ($\mathrm{L/R = 2.5\ nH/\Omega}$), the 3-level converter exhibits higher efficiency than the buck converter. Both converters suffer from degrading efficiencies at low voltages, but the slope of 3-level is steeper than that of the buck. This is because $\mathrm{\Delta I_{L,PP}}$ of the 3-level increases as duty cycle deviates from 50%, while that of the buck decreases. Using a higher quality inductor ($\mathrm{L/R = 5\ nH/\Omega}$) allows the buck to achieve higher efficiencies than the 3-level converter at low and high loads.

Fig. 7 compares the conversion efficiency of the 3-level converter to a reconfigurable SC converter that can switch between

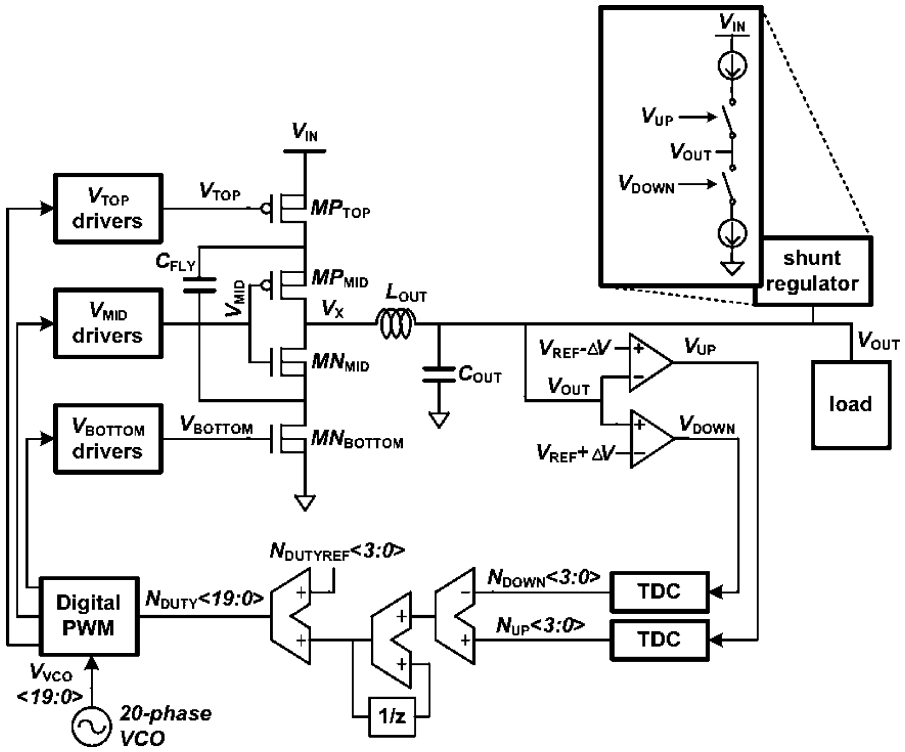

Fig. 8. Block diagram of 3-level converter with slow digital feedback control and fast shunt regulation. Finer duty cycle control is necessary to avoid limit-cycling.

three modes – 1/3, 1/2 and 2/3. Simulations use a SC converter design similar to one in [10] with series stacks of thin-oxide devices to support high input voltage. While the 3-level converter has 16 nF of C_{FLY} and 10 nF of C_{OUT}, SC can use C_{FLY} as an output decoupling capacitor, obviating additional C_{OUT}. Assuming the same die area for the two converters, the SC converter can use 26 nF of C_{FLY} without any C_{OUT}. For the 3-level converter, we assume that C_{FLY} is MOS capacitors placed underneath the inductor to avoid additional area overhead (as explained later in Fig. 13). Since 16 nF of MOS capacitance occupies 1.6 mm^2 in UMC 130 nm technology, four 0.4×0.4 mm inductors occupying 0.64 mm^2 can fit on top of C_{FLY}. In contrast to the 3-level and buck converters, the SC converter does not need a thick metal layer for high quality inductors. For fair comparison, we present conversion efficiencies across inductor qualities that represent different metal characteristics. Again, efficiency is calculated using minimum V_{OUT} instead of average V_{OUT}.

Assuming an inductor built with two metal layers in parallel using the digital logic process in UMC 130 nm ($L/R = 2\ nH/\Omega$), the SC converter in 1/2 mode achieves higher efficiency than the 3-level converter at the center where duty cycles are in the vicinity of 50%, while the 3-level converter exhibits higher efficiencies at light loads than the SC converter in 1/3 mode. The trend is similar assuming an inductor built with two metal layers (one 2 μm thick layer) using the RF process in UMC 130 nm ($L/R = 2.5\ nH/\Omega$). The 3-level has the potential for even higher efficiencies when assuming an ultra-thick metal available in modern process technologies that enables an even higher quality inductor ($L/R = 5\ nH/\Omega$), albeit with higher cost. Although the inductor adds series resistance, the 3-level converter has the following benefits when operating at 50% duty cycle. First, the inductor allows the 3-level to have a lower per-phase peak current than the SC converter, reducing resistive loss [21]. Second, the inductor reduces loss caused by charge redistribution in the 3-level converter. Whenever capacitors switch between series stack and parallel configurations in the SC converter, the resulting charge redistribution between C_{FLY} and C_{OUT} increases conversion loss [22]. In contrast, the inductor in the 3-level converter sits between C_{FLY} and C_{OUT} to store a portion of the charge otherwise lost to charge redistribution.

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

III. Implementation of 3-Level Converter

Fig. 8 presents an overall block diagram comprising a set of thin-oxide transistors used as power FETs for power conversion, drive circuitry for the power FETs, a flying capacitor, an on-die LC filter, and control circuitry for voltage regulation. A relatively slow digital feedback loop sets the signals out of the digital pulse-width modulator (DPWM) that feed drivers to switch the 3-level converter with appropriate duty cycles (D). In parallel, a fast shunt regulator [18] on the output reacts to sudden load current transients to maintain a steady voltage. The overall design target is to minimize conversion loss, on-die area overhead, voltage fluctuations, and dynamic voltage scaling time. This section further studies the components in Fig. 8 and looks at circuit implementations in detail.

The 3-level converter uses four power FETs, a flying capacitor, and an output LC filter to generate a wide range of

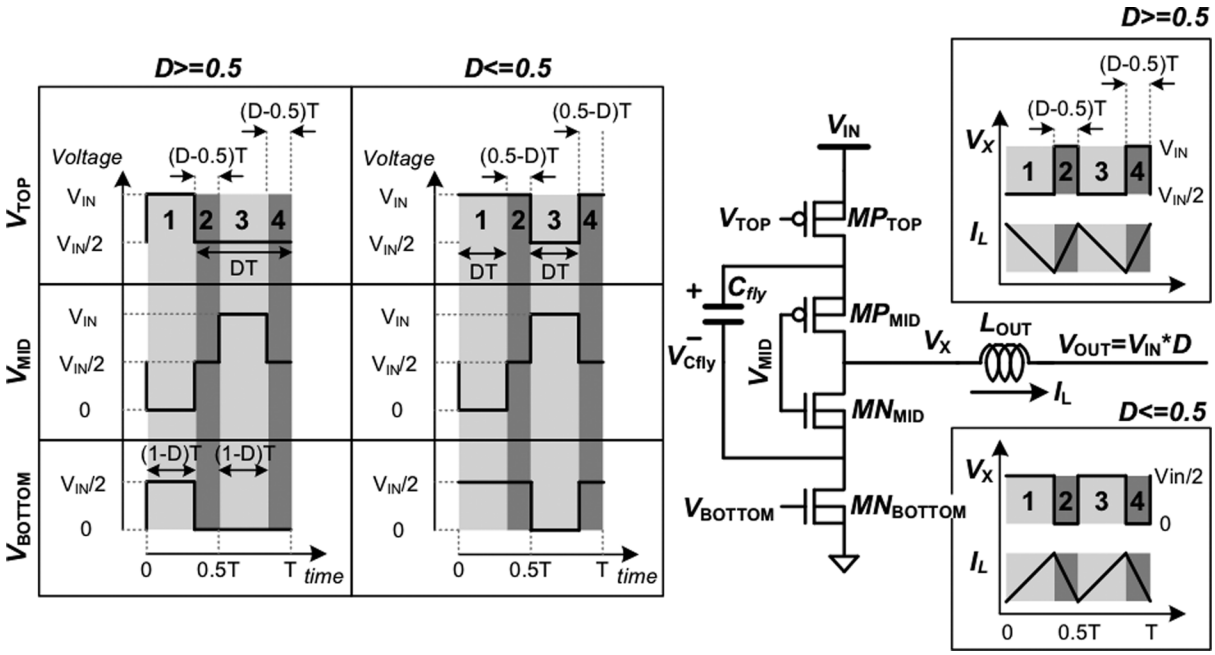

Fig. 9. Schematic of the proposed 3-level power converter. Signal timing diagrams illustrate different operating modes.

output voltages. Fig. 9 illustrates the converter's operation via signal waveforms associated with the power FETs ($\mathrm{MP_{TOP}}$, $\mathrm{MN_{BOTTOM}}$, $\mathrm{MP_{MID}}$, and $\mathrm{MN_{MID}}$) and the output inductor for two scenarios: $\mathrm{D} \geq 0.5$ and $\mathrm{D} \leq 0.5$. As previously described, node $\mathrm{V_X}$ can swing between three voltage levels by iterating through four steps per switching period (T) that control the power FETs and $\mathrm{C_{FLY}}$.

For $\mathrm{D} \geq 0.5$, step 1 turns on $\mathrm{MN_{BOTTOM}}$ and $\mathrm{MP_{MID}}$, placing $\mathrm{C_{FLY}}$ between $\mathrm{V_X}$ and 0. In step 3, $\mathrm{MP_{TOP}}$ and $\mathrm{MN_{MID}}$ turn on, placing $\mathrm{C_{FLY}}$ between $\mathrm{V_{IN}}$ and $\mathrm{V_X}$. As in a SC converter, where two capacitors alternate between series-stack and parallel configurations, steps 1 and 3 generate $\mathrm{V_{Cfly}}$ and $\mathrm{V_{IN}} - \mathrm{V_{Cfly}}$, respectively, on $\mathrm{V_X}$. Assuming the ideal case where $\mathrm{V_{Cfly}}$ is equal to $\mathrm{V_{IN}}/2$, $\mathrm{V_X}$ stays at $\mathrm{V_{IN}}/2$ in steps 1 and 3. In steps 2 and 4, $\mathrm{V_X}$ connects to $\mathrm{V_{IN}}$ through $\mathrm{MP_{TOP}}$ and $\mathrm{MP_{MID}}$. By adjusting the time spent in each step, the converter can generate any voltage between $\mathrm{V_{IN}}$ and $\mathrm{V_{IN}}/2$ at $\mathrm{V_{OUT}}$.

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

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

A. Power FETs

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

Again referring to Fig. 9, the stacked structure using thin-oxide devices requires voltage stress across each device to be limited to $\mathrm{V_{IN}}/2$. Input signals to the power FETs need to be carefully set in order to meet this requirement in each step. For this purpose, the input signal to $\mathrm{MP_{TOP}}$ ($\mathrm{V_{TOP}}$) swings between $\mathrm{V_{IN}}$ and $\mathrm{V_{IN}}/2$, while $\mathrm{V_{BOTTOM}}$ swings between $\mathrm{V_{IN}}/2$ and 0. To limit voltage stress on the middle FETs ($\mathrm{MP_{MID}}$ and $\mathrm{MN_{MID}}$), their input ($\mathrm{V_{MID}}$) swings across three voltage levels, $\mathrm{V_{IN}}$, $\mathrm{V_{IN}}/2$ and 0. In step 1 for $\mathrm{D} \geq 0.5$, $\mathrm{V_{MID}}$ is set to 0 to simultaneously turn $\mathrm{MP_{MID}}$ on and turn $\mathrm{MN_{MID}}$ off. In step 2, both $\mathrm{MP_{MID}}$ and $\mathrm{MN_{MID}}$ remain in their respective on and off states from step 1. However, as $\mathrm{V_X}$ goes up to $\mathrm{V_{IN}}$, $\mathrm{V_{MID}}$ must increase to $\mathrm{V_{IN}}/2$ to meet voltage stress requirements on $\mathrm{MP_{MID}}$ and $\mathrm{MN_{MID}}$. In step 3, $\mathrm{V_{MID}}$ is set to $\mathrm{V_{IN}}$ to turn $\mathrm{MN_{MID}}$ on and turn $\mathrm{MP_{MID}}$ off. Step 4 sets $\mathrm{V_{MID}}$ to $\mathrm{V_{IN}}/2$ again to alleviate voltage stress as seen in step 2. When the converter operates with $\mathrm{D} \leq 0.5$, similar voltage stress constraints must be observed.

This circuitry requires an additional voltage, $\mathrm{V_{IN}}/2$, to generate inputs for power FETs that switch between two sets of supply rails ($\mathrm{V_{IN}}$ and $\mathrm{V_{IN}}/2$ or $\mathrm{V_{IN}}/2$ and ground). To generate $\mathrm{V_{IN}}/2$, we use an external power source with 20 μF on-board and 660 pF on-chip decoupling capacitance. Since the pFETs switching between the top supply rails ($\mathrm{V_{IN}}$ and $\mathrm{V_{IN}}/2$) are larger than nFETs between the bottom rails ($\mathrm{V_{IN}}/2$ and ground), current usually flows into the power source that provides $\mathrm{V_{IN}}/2$. An integrated linear regulator [16] can replace

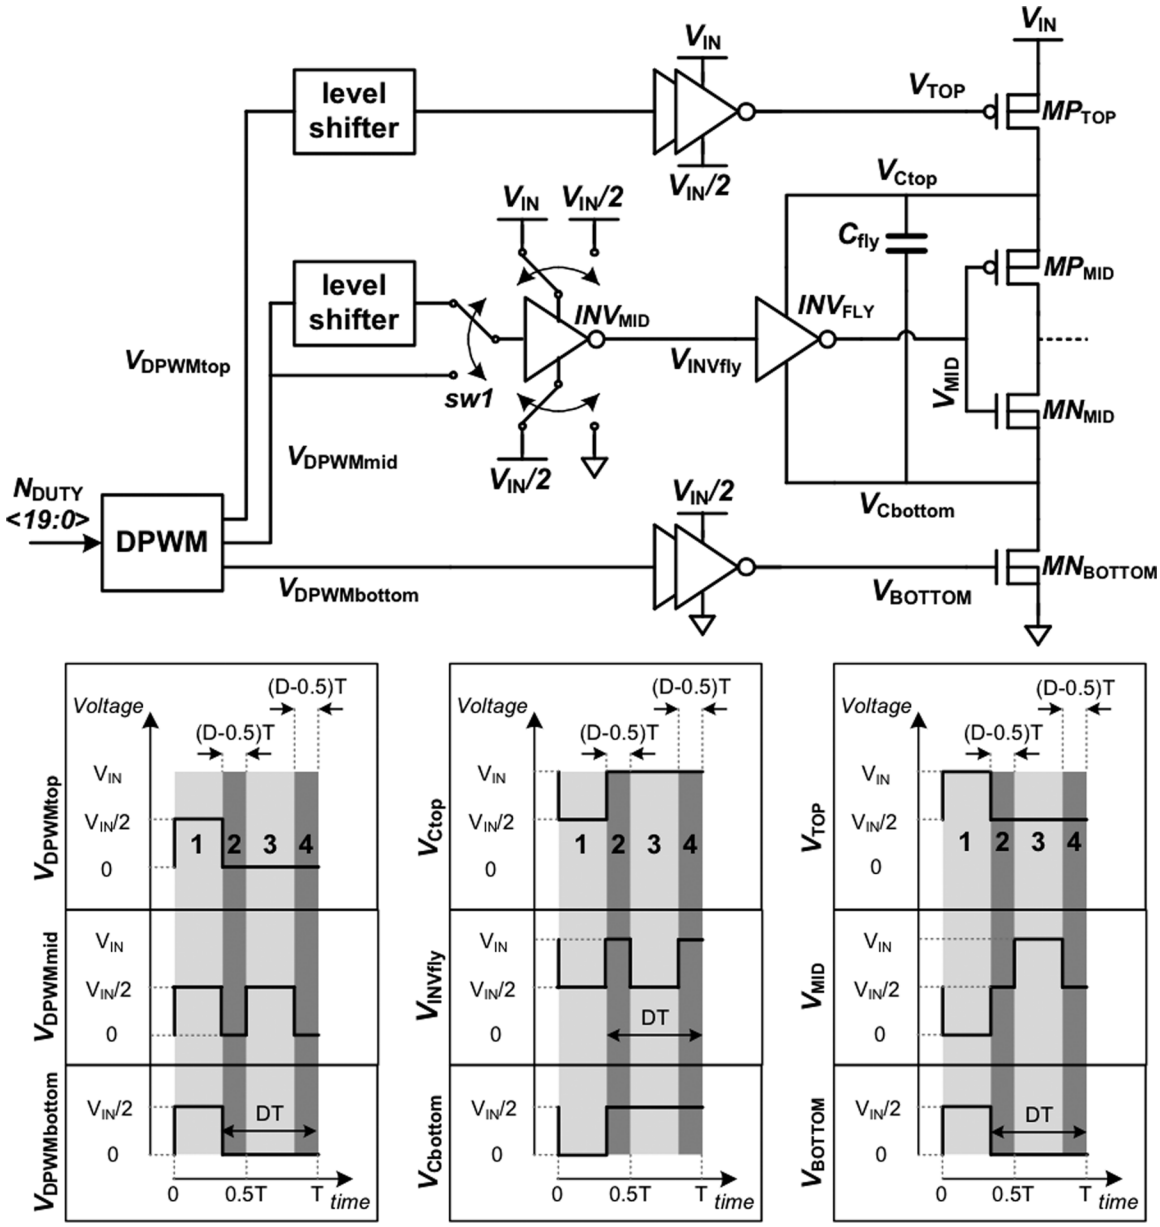

Fig. 10. Schematic and waveforms that drive power FETs when duty cycle is over 50%.

the external source by bleeding in current caused by the imbalance between top and bottom rails without adding significant power loss.

B. Driver Circuits

Creating appropriate signals to limit voltage stress on the power FETs requires careful design of the circuitry that generates V_{TOP}, V_{MID}, and V_{BOTTOM}. Fig. 10 presents schematics of the drivers for the four power FETs and associated signal waveforms for the case when $D \geq 0.5$. A digital pulse-width modulator (DPWM) block generates signals based on a digital, thermometer coded representation of the desired converter duty cycle, $N_{DUTY}\langle 19:0\rangle$, using a 20-phase VCO. The DPWM consists of digitally controlled switches that choose two VCO phases that determine the duty cycle of the output signal. While inverters can generate V_{BOTTOM} from the DPWM output signal ($V_{DPWMbottom}$), V_{TOP} requires a level-shifter [15] to shift the DPWM output ($V_{DPWMtop}$), which swings between $V_{IN}/2$ and 0, up to swing between V_{IN} and $V_{IN}/2$.

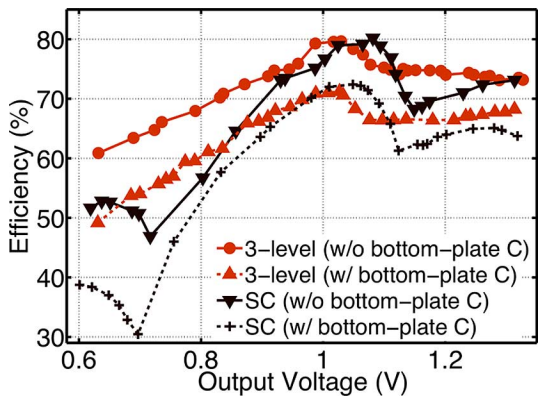

Fig. 11. Simulated conversion efficiencies of 3-level and SC converters with and without bottom-plate parasitic capacitance.

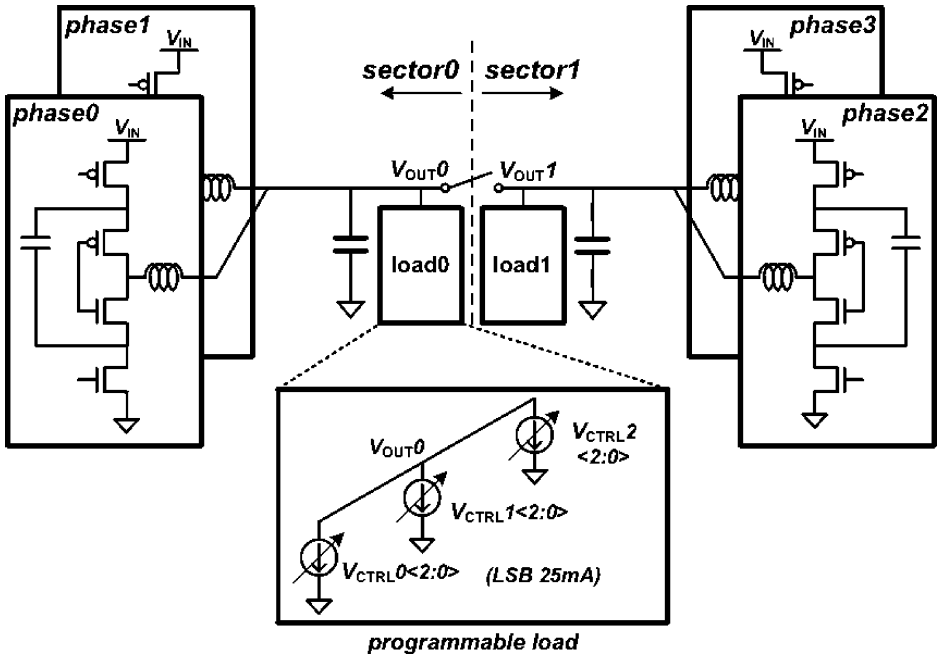

Fig. 12. High-level architecture of the 3-level converter test-chip prototype.

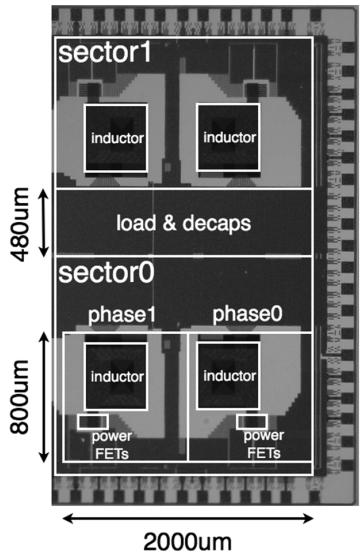

Technology	130nm CMOS
Load Power	0.2-1W
Input Voltage	2.4V
Output Voltage	0.4-1.4V
Inductor per phase	1nH
Total Flying Capacitance	18nF
Total Output Capacitance	10nF
Total Input Capacitance	1nF
Switching Frequency	50-250MHz
Peak Efficiency	77%

Fig. 13. Die micrograph of the converter with dimensions of main blocks. Flying capacitors are placed underneath the inductors to reduce area overhead. The table shows converter specifications.

TABLE I
SPECIFICATIONS OF ON-CHIP SPIRAL INDUCTORS MODELED USING ASITIC [23] AND MOS CAPACITORS

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

The middle FETs, $\mathrm{MP_{MID}}$ and $\mathrm{MN_{MID}}$, need a special driver to generate $\mathrm{V_{MID}}$ that swings across three different voltages, $\mathrm{V_{IN}}$, $\mathrm{V_{IN}}/2$ and 0. The buffer, $\mathrm{INV_{FLY}}$, needs to dynamically switch between two configurations—sitting between $\mathrm{V_{IN}}$ and $\mathrm{V_{IN}}/2$ and sitting between $\mathrm{V_{IN}}/2$ and 0. Since $\mathrm{C_{FLY}}$ alternates between the same two configurations, one way to implement $\mathrm{INV_{FLY}}$ is to place it between the top ($\mathrm{V_{Ctop}}$) and bottom plate ($\mathrm{V_{Cbottom}}$) of $\mathrm{C_{FLY}}$, creating a "flying inverter" [10]. In step 1, $\mathrm{INV_{FLY}}$ follows $\mathrm{C_{FLY}}$ to sit between $\mathrm{V_{IN}}/2$ and 0. Input to $\mathrm{INV_{FLY}}$ ($\mathrm{V_{INVfly}}$) is set to $\mathrm{V_{IN}}/2$ to generate 0 at $\mathrm{V_{MID}}$. In steps 2 to 4, $\mathrm{INV_{FLY}}$ sits between $\mathrm{V_{IN}}$ and $\mathrm{V_{IN}}/2$ with $\mathrm{V_{INVfly}}$ swinging between $\mathrm{V_{IN}}$ and $\mathrm{V_{IN}}/2$ to generate $\mathrm{V_{MID}}$. While this is the case for $\mathrm{D} > 0.5$, $\mathrm{V_{INVfly}}$ needs to swing between $\mathrm{V_{IN}}/2$ and 0 for $\mathrm{D} < 0.5$ ($\mathrm{V_{INVfly}}$ is fixed at $\mathrm{V_{IN}}/2$ for $\mathrm{D} = 0.5$). To accommodate both cases, $\mathrm{D} > 0.5$ and $\mathrm{D} < 0.5$, the buffer, $\mathrm{INV_{MID}}$, that generates $\mathrm{V_{INVfly}}$ sits between $\mathrm{V_{IN}}$ and $\mathrm{V_{IN}}/2$ for $\mathrm{D} \geq 0.5$, while it sits between $\mathrm{V_{IN}}/2$ and 0 for $\mathrm{D} \leq 0.5$. $\mathrm{INV_{MID}}$ switches between these two configurations using power switches. The switch (SW1) that connects the input to $\mathrm{INV_{MID}}$ is an analog 2:1 mux built with thick-oxide devices to accommodate input signals ranging from 0 to $\mathrm{V_{IN}}$.

C. Passive Elements

For high efficiency, it is crucial to design high quality passive elements while not incurring excessive on-die area overhead. Table I shows specifications for the spiral inductor implemented using top two metal layers in parallel to reduce series resistance. To save on-die area, the flying capacitor resides under the inductor. Since the flying capacitors can potentially

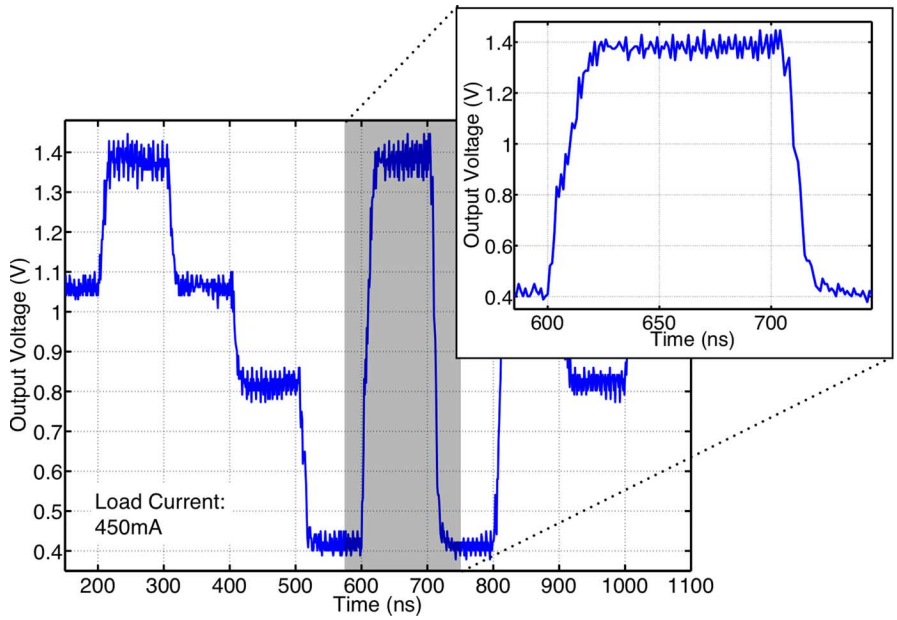

Fig. 14. Measured snapshot of fast dynamic voltage scaling of the converter operating in open-loop. Voltage scales from 1.4 V to 0.4 V and vice versa within 20 ns.

inject noise into the inductor, a patterned ground shield protects the inductor from noise coupling [17]. The flying capacitor is implemented with a MOS gate capacitor, because of its higher density compared to metal wire capacitors. However, both sides of the flying capacitor swing by $V_{IN}/2$, which impacts the design of the MOS capacitor.

While a triple-well nFET offers slightly higher capacitor density, a pFET incurs less area overhead associated with the surrounding wells. Hence, we opted to implement the MOS cap using a pFET with drain, source, and body all tied together. A major overhead of this choice comes from the junction capacitance between the P-substrate and N-well, which adds large bottom-plate parasitic capacitance that exacerbates switching loss. Fig. 11 presents simulated conversion efficiencies of SC and 3-level converters including and excluding bottom-plate parasitic capacitance. Both converters benefit from a $\sim$10% efficiency gain across a wide range of loads when bottom-plate parasitic capacitance is eliminated. This motivates using a process technology with high density capacitors with less bottom-plate parasitic capacitance.

D. Feedback Loop and Shunt Regulator

Building on the previous blocks that generate an output voltage with respect to different duty cycles, we now turn our attention to the relatively slow digital feedback loop and shunt regulator loop that regulate V_{OUT} to a desired level, especially under load fluctuations. Revisiting Fig. 8, both loops share a pair of fast voltage comparators with hysteresis to sense whether the output voltage is above or below a desired reference level, V_{REF}. In the digital loop, a pair of simple time-to-digital converters (TDC) generates 4-bit thermometer codes, $N_{UP}\langle 3{:}0\rangle$ and $N_{DOWN}\langle 3{:}0\rangle$, whose difference corresponds to the V_{OUT}-V_{REF} error within each switching cycle. Accumulating the difference between $N_{UP}\langle 3{:}0\rangle$ and $N_{DOWN}\langle 3{:}0\rangle$, and adding it to a reference duty cycle, $N_{DUTYREF}\langle 3{:}0\rangle$, results in a digital code, $N_{DUTY}\langle 19{:}0\rangle$, that feeds the DPWM described above. $N_{DUTYREF}\langle 3{:}0\rangle$ can be programmed externally and changes only when the converter needs to dynamically scale the output voltage. Simultaneously changing V_{REF} and $N_{DUTYREF}\langle 3{:}0\rangle$ together enables nanosecond-scale voltage scaling, as opposed to only adjusting V_{REF} and slowly accumulating error through the digital loop. $N_{DUTY}\langle 19{:}0\rangle$ can generate a range of duty cycles between 25% and 75%, in 5% steps, which leads to 120 mV output voltage resolution for a 2.4 V V_{IN}. The coarse resolution hinders the feedback from providing tight regulation, often resulting in steady-state limit-cycling [24]. Finer-grain duty cycle control, possible using a VCO with a larger number of phases, is necessary to achieve tighter regulation.

Since the digital loop cannot easily track sudden load current transients, there is a supplemental shunt regulator that suppresses output voltage fluctuations by detecting when V_{OUT} crosses low or high thresholds and injecting or extracting current [18]. Based on the V_{UP} and V_{DOWN} signals from the two comparators, the shunt regulator can either turn on pFETs sitting between V_{IN} and V_{OUT} to inject current to V_{OUT}, or turn on nFETs between V_{OUT} and 0 V to extract current from V_{OUT}. Since V_{OUT} varies widely, the shunt regulator uses thick-oxide devices for pFETs sitting between V_{IN} and V_{OUT}. In contrast, maximum voltage stress is 1.4 V for nFETs sitting between V_{OUT} and 0, allowing for thin-oxide devices.

IV. Measurement Results

To demonstrate the benefits of the 3-level converter, we designed a test-chip prototype in a 130 nm Mixed-Mode/RF CMOS process from UMC with a 2 μm thick top metal layer. Fig. 12 shows the high-level architecture of the test-chip prototype that consists of a pair of 2-phase, 3-level converters arranged as two identical sectors. The two phases share a single output capacitor to reduce ripple on V_{OUT}. Low-impedance,

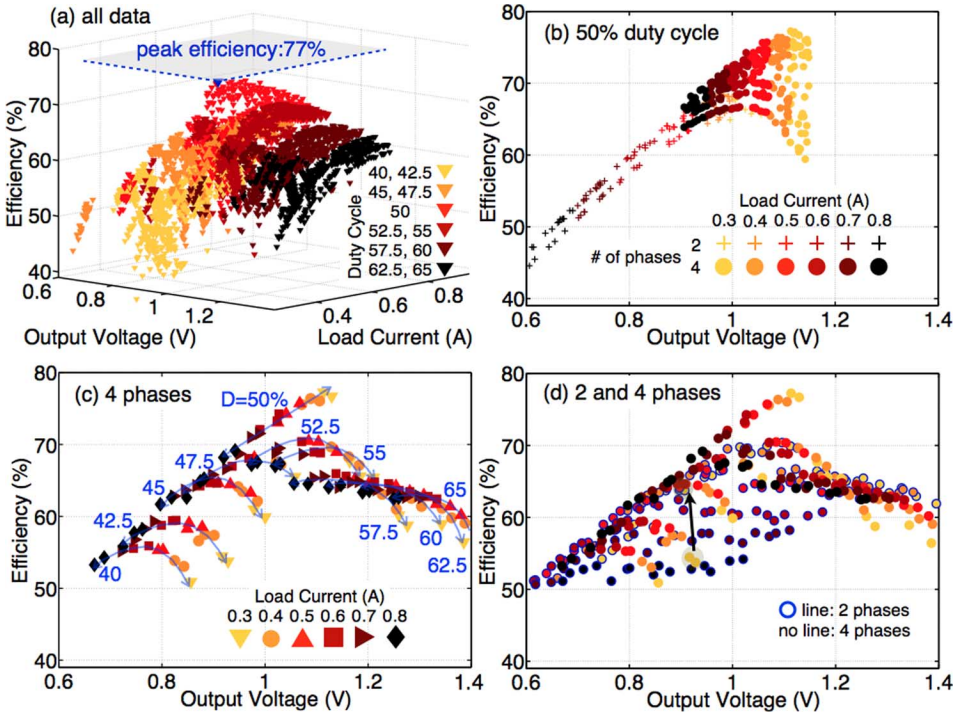

Fig. 15. Measured efficiency of converter operating in open-loop.

TABLE II
BREAKDOWN OF CONVERSION LOSS OF THE 3-LEVEL CONVERTER FOR THREE DESIGN POINTS

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

on-chip switches can connect the two sectors together to create a single 4-phase converter with each phase offset by 90 degrees. Otherwise, the test chip implements two independent 2-phase converters. An ability to disable power FETs further enables multiple 3-level converter configurations consisting of one to four phases. A programmable load in each sector facilitates experimental measurements by sinking up to 0.5 A in 25 mA steps as steady or pseudorandom patterns of current.

Measurement results demonstrate that the 3-level converter can generate a wide range of output voltages using 1 nH integrated inductors. The converter presents nanosecond-scale voltage transition times and peak conversion efficiency of 77%. Fig. 13 presents a die micrograph and a list of specifications for the test chip.

Data captured from a real-time oscilloscope (plotted in Fig. 14) demonstrates the converter can generate output voltages across a wide range—from 0.4 to 1.4 V when the input voltage is 2.4 V — and rapidly scale V_{OUT} by 1 V within 20 ns. Such high-speed voltage transitions at nanosecond time scales enable complex digital systems to leverage temporally fine-grained DVFS and improve system-wide energy efficiency [19].

Fig. 15 summarizes the conversion efficiency measurements made on the test chip in CCM mode. The converter operates in open-loop with fixed duty cycles ranging from 40% to 65% in 5% steps to facilitate measurements across a wide range of conditions. Two converter sectors can also operate with duty

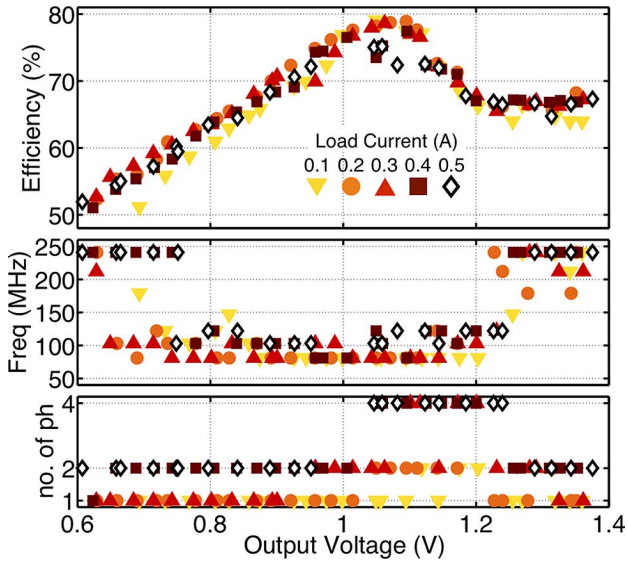

Fig. 16. Measured conversion efficiency with optimal switching frequencies and number of phases.

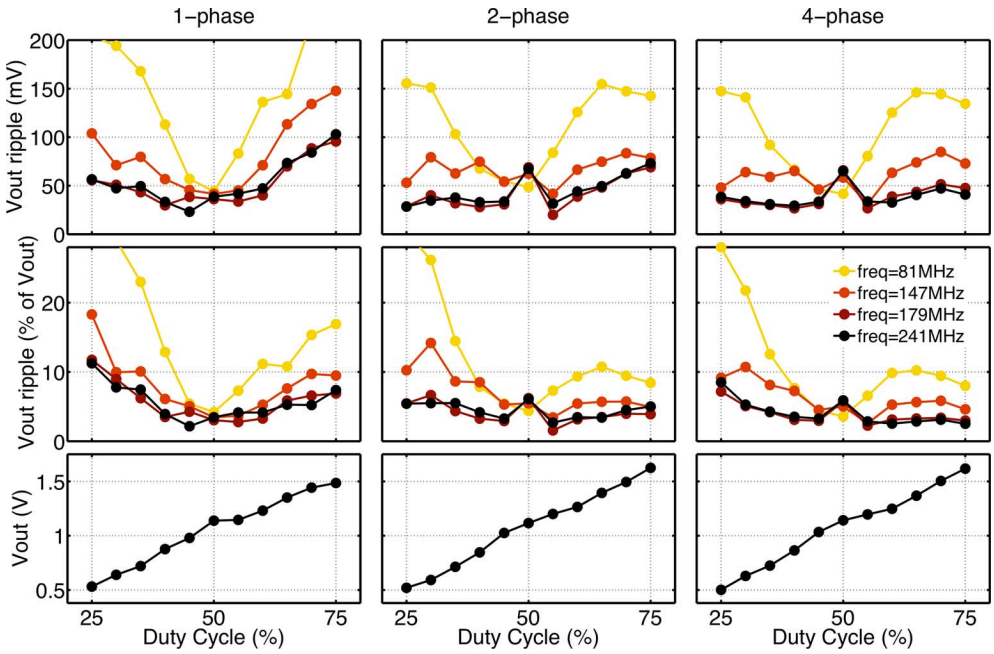

Fig. 17. Measured peak-to-peak output voltage ripple of the 3-level converter with DC load current. Ripple changes across duty cycles, switching frequencies and number of phases.

cycles that differ by 5% to implement finer steps. Since duty cycle is fixed during open-loop measurements, IR drop due to parasitic resistance causes a spread in output voltages with respect to load currents for the same duty cycle. IR drop is larger than expected due to parasitic resistance on the external power supply, bond-wires, and metal trace. Fig. 15(a) aggregates all of the measured efficiencies collected across a range of static load current conditions (0.3 to 0.8 A), duty cycles (40 to 65%), switching frequencies (50 to 160 MHz), and number of phases (1 to 4). Efficiency peaks at 77% for low load current conditions (0.1 W/mm^2) at 50% duty cycle. Fig. 15(b) compares measured data for 50% duty cycle operation using 2 and 4 phases. IR losses increase as load current increases, increasing further for the 2-phase configuration. Higher switching frequency can also degrade efficiency at low load currents due to higher switching losses. Fig. 15(c) plots the upper range of efficiency measurements for the 4-phase configuration by picking the best efficiency data across different duty cycle settings. Trend line overlays again illustrate the spread in output voltages due to IR drop. Efficiency peaks for 50% duty cycle owing to small inductor current ripple as explained in Section II. As duty cycle deviates from 50%, inductor current ripple grows and the corresponding increase in resistive losses degrades conversion efficiency. Fig. 15(d) adds results for the 2-phase configuration (symbols with outlines) to show that fewer phases can improve efficiency at duty cycles away from 50%.

Using data from Fig. 15, Table II presents the breakdown of conversion loss for three different design points. At low loads (point 1), the 3-level converter runs at 152 MHz with a single phase to reduce loss due to inductor current ripple. At mid-loads (point 2) where duty cycle is 50%, the number of phases increases to 4 and switching frequency decreases, matching the analysis in Section II. However, at high loads, the number of phases does not decrease to 2, but stays at 4. Contrary to the analysis in Section II, using 4 phases exhibits higher efficiency than using 2 phases at high loads because a 2-phase converter suffers larger parasitic resistance in the power delivery wires due to floor-plan issues in the test-chip. The die micrograph in Fig. 13 shows that pads are placed close to each phase of the converter, allowing all phases to have low-impedance connections to power/ground pads. When the converter operates with 2 phases, it has low-impedance connection to about half of the pads that are close to the 2 phases that are turned on. The rest of the pads that are farther away from the 2 phases provide a higher impedance connection with larger parasitic resistance. Compared to a 4-phase converter with short distance to most of the pads, a 2-phase converter suffers from loss due to larger parasitic resistance on the power delivery path.

To further study the effect of frequency and number of phases on efficiency, we measured a second chip across a wider range of switching frequencies. Fig. 16 presents maximum efficiencies for each load current from 0.1 A to 0.5 A plotted across output voltages. As shown in the second subplot, frequency reaches a minimum at the center and increases as duty cycle deviates from 50%, following a U-shaped curve. The optimal number of phases, presented in the bottom subplot, also matches the aforementioned trend of 1 phase at low load, 4 phases at the center and 2 phases at high loads. Since the maximum load current is 0.5 A, lower than 0.8 A in Fig. 15, larger parasitic resistance on the power delivery path has less impact on conversion efficiency, favoring 2 phases over 4 phases at high loads.

Fig. 17 presents peak-to-peak output voltage ripple across duty cycles for 1, 2 and 4-phase configurations and different switching frequencies. In this measurement, the converter operates with DC load current ranging from 0.1 A to 0.7 A that scales linearly with output voltage. As seen in the top row, voltage ripple reaches a minimum at 50% duty cycle for all cases, and increases symmetrically as duty cycle deviates

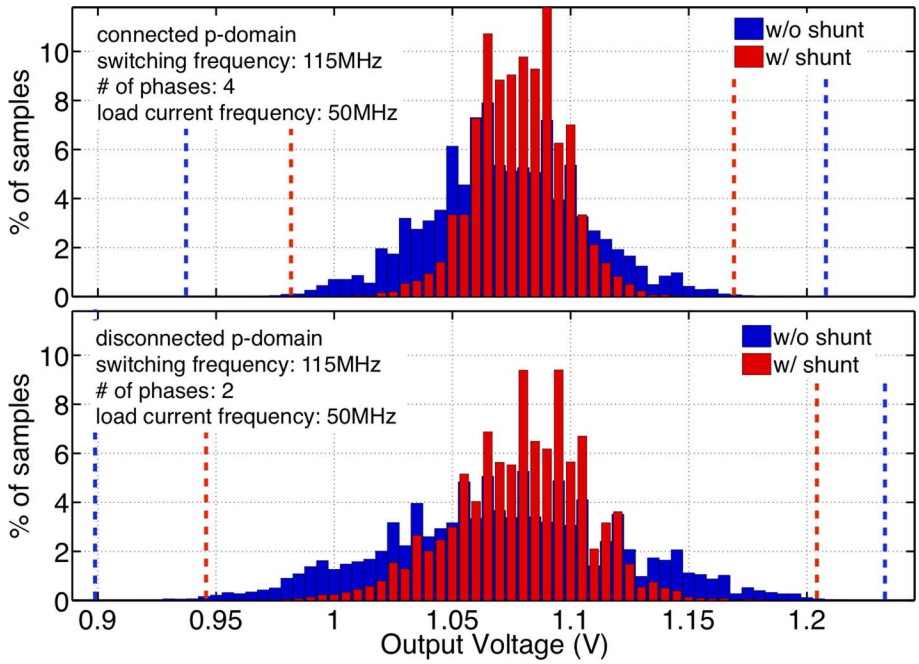

Fig. 18. Histogram of measured voltage noise with and without shunt regulator for connected and disconnected power domains of two sectors.

from 50%, matching the trends of $\Delta I_{L,PP}$ in Fig. 4. Although the absolute magnitude of ripple is roughly symmetric, ripple grows larger as a percentage of V_{OUT} at low output voltages (second row). Interleaving larger numbers of phases helps reduce voltage ripple, especially at extreme duty cycles far from 50%. By increasing the frequency as duty cycle deviates from 50%, and operating with two or four phases, the converter can maintain 5% peak-to-peak ($\pm 2.5\%$) ripple at duty cycles ranging from 30% to 75%, which covers a wide 0.6–1.5 V output voltage range.

Compared to steady-state voltage ripple, rapidly changing load current further increases voltage fluctuation. Fig. 18 presents histogram plots created by sampling the output voltage of the converter. We measure voltage noise due to pseudo-random current patterns generated by the programmable loads, with and without the supplemental shunt regulator turned on. The simulated ramp time of load current is 1.5 mA/ps. With connected sectors (top plot), the shunt regulator is able to reduce peak-to-peak voltage noise from 0.27 V to 0.19 V. These results verify that the shunt regulator can appreciably squeeze the noise distribution together and reduce peak-to-peak voltage excursions, shown in dotted lines. Moreover, connecting the power domains reduces voltage noise as a result of larger output capacitance and some canceling of the pseudorandom load currents.

While the shunt regulator—reacting to threshold crossings—reduces voltage fluctuations, it has two drawbacks. First, internal circuit delays limit how quickly this feedback loop can sense and react. Second, simply relying on thresholds provides limited information as to the magnitude of voltage noise and the appropriate response needed to suppress it. One solution is to use a prediction-based shunt regulator that leverages microarchitecture-level information to reliably predict upcoming voltage droops [20]. The processor can track the history of microarchitecture events using a memory structure to predict events that lead to a surge in load current.

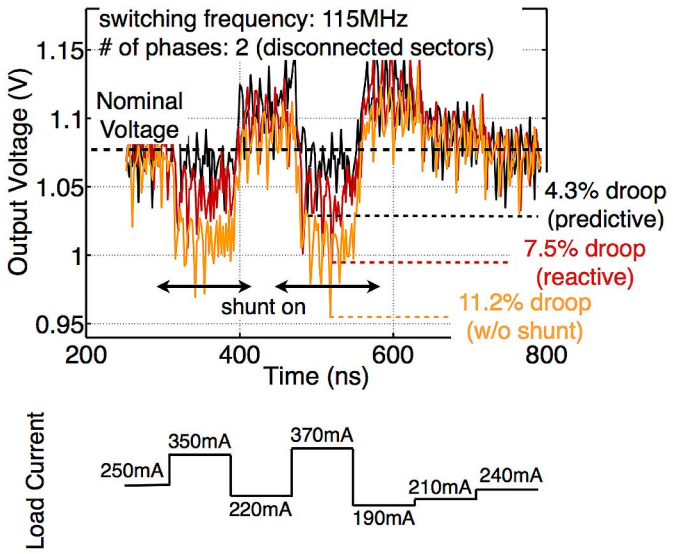

Fig. 19. Comparison of measured on-die voltage noise without shunt regulator, with reactive shunt, and with predictive shunt.

To demonstrate the potential of predictive shunt regulation, we use pulse signals generated externally to turn on the shunt regulator, mimicking signals provided by a processor that predict upcoming voltage droops. Fig. 19 presents snapshots of measured voltage droops due to two consecutive 80 ns wide current pulses of 100 mA and 150 mA. Predictive current shunting reduces the maximum voltage droop by over 40% compared to simply reacting to threshold crossings.

TABLE III
COMPARISON WITH PRIOR DC-DC CONVERTER DESIGNS.

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

* power density including output decoupling capacitance

Lastly, Table III compares recently published DC-DC converters using chip-integrated or package-integrated passive elements. Since the published test-chips use different process technologies, input/output voltage ranges and inductor technologies, it is difficult to make a fair comparison across all of them. The test-chip that is most similar to this work is a buck converter built in 130 nm using on-chip spiral inductors with 2–2.6 V input and 1.1–1.5 V output voltage ranges [8]. Compared to this buck converter, our 3-level converter uses a 4× smaller inductor and exhibits 15 percentage points higher efficiency at comparable power densities.

V. CONCLUSION

Measurement and analysis from a 130 nm test-chip prototype demonstrate the benefits of a fully-integrated 3-level DC-DC converter. Merging the characteristics of the buck and SC converters, the 3-level converter offers a wide output voltage range using a small 1 nH inductor that is suitable for on-chip integration.

For a 2.4 V to 0.6–1.4 V conversion, the converter achieves 77% peak efficiency and voltage scaling across 1 V within 20 ns, which is 100× faster than previously published converters using on-package inductors [3]. Process technologies with smaller bottom-plate parasitic capacitance and thick metal layers offer the potential to further increase the conversion efficiency of future 3-level converter designs.

REFERENCES

[1] K. K. Rangan, G. -Y. Wei, and D. Brooks, "Thread motion: Fine-grained power management for multi-core systems," in *Proc. 36th Annu. Int. Symp. Computer Architecture (ISCA)*, Jun. 2009.

[2] P. Hazucha, G. Schrom, J. Hahn, B. A. Bloechel, P. Hack, G. E. Dermer, S. Narendra, D. Gardner, T. Karnik, V. De, and S. Borkar, "A 233 MHz 80%–87% efficient four-phase DC-DC converter utilizing air-core inductors on package," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 838–845, Apr. 2005.

[3] G. Schrom, F. Faillet, and J. Hahn, "A 60 MHz 50 W fine-grain package integrated VR powering a CPU from 3.3 V," in *Applied Power Electronics Conf.*, 2010.
[4] J. Wibben and R. Harjani, "A high-efficiency DC-DC converter using 2 nH integrated inductors," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 844–854, Apr. 2008.
[5] M. Alimadadi, S. Sheikhaei, G. Lemieux, S. Mirabbasi, and P. Palmer, "A 3 GHz switching DC-DC converter using clock-tree charge-recycling in 90 nm CMOS with integrated output filter," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2007.
[6] S. Abedinpour, B. Bakkaloglu, and S. Kiaei, "A multi-stage interleaved synchronous buck converter with integrated output filter in a 0.18 μm SiGe process," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2006.
[7] G. Schrom, P. Hazucha, F. Paillet, D. J. Rennie', S. T. Moon, D. S. Gardner, T. Kamik, P. Sun, T. T. Nguyen, M. J. Hill, K. Radhakrishnan, and T. Memioglu, "A 100 MHz eight-phase buck converter delivering 12 A in 25 mm^2 using air-core inductors," in *Applied Power Electronics Conf.*, 2007.
[8] M. Wens and M. Steyaert, "An 800 mW fully-integrated 130 nm CMOS DC-DC step-down multi-phase converter, with on-chip spiral inductors and capacitors," in *IEEE Energy Conversion Congress and Expo (ECCE)*, 2009.
[9] Y. K. Ramadass, A. Fayed, 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. 2557–2567, Dec. 2010.
[10] H.-P. Le, M. Seeman, S. Sanders, V. Sathe, S. Naffziger, and E. Alon, "A 32 nm fully integrated reconfigurable switched-capacitor DC-DC converter delivering 0.55 W/mm^2 at 81% efficiency," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2010.
[11] L. Chang, R. K. Montoye, B. L. Ji, A. J. Weger, K. G. Stawiasz, and R. H. Dennard, "Fully-integrated switched-capacitor 2:1 voltage converter with regulation capability and 90% efficiency at 2.3 $\mathrm{A/mm^2}$," in *IEEE Symp. VLSI Circuits*, Jun. 2010.
[12] G. Villar and E. Alarcón, "Monolithic integration of a 3-level DCM-operated low-floating-capacitor buck converter for DC-DC step-down conversion in standard CMOS," in *IEEE Power Electronics Specialist Conf. (PESC)*, 2008.
[13] W. Kim, G. -Y. Wei, and D. Brooks, "A fully-integrated 3-level DC/DC converter for nanosecond-scale DVS with fast shunt regulation," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2011.
[14] V. Yousefzadeh, E. Alarćon, and D. Maksimovic, "Three-level buck converter for envelope tracking applications," *IEEE Trans. Power Electronics*, vol. 21, no. 2, pp. 549–552, Mar. 2006.
[15] S. Rajapandian, K. L. Shepard, P. Hazucha, and T. Karnik, "High-voltage power delivery through charge recycling," *IEEE J. Solid-State Circuits*, vol. 41, no. 6, pp. 1400–1410, Jun. 2006.
[16] P. Hazucha, S. T. Moon, G. Schrom, F. Paillet, D. Gardner, S. Rajapandian, and T. Karnik, "High-voltage power delivery through charge recycling," *IEEE J. Solid-State Circuits*, vol. 42, no. 1, pp. 66–73, Jan. 2007.
[17] F. Zhang and P. R. Kinget, "Design of components and circuits underneath integrated inductors," *IEEE J. Solid-State Circuits*, vol. 41, no. 10, pp. 2265–2271, Oct. 2006.
[18] E. Alon and M. Horowitz, "Integrated regulation for energy-efficient digital circuits," *IEEE J. Solid-State Circuits*, vol. 43, no. 8, pp. 1795–1807, Aug. 2008.
[19] W. Kim, M. Gupta, G. -Y. Wei, and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," in *IEEE Int. Symp. High Performance Computer Architecture (HPCA)*, Feb. 2008.
[20] V. J. Reddi, M. S. Gupta, G. Holloway, G.-Y. Wei, M. D. Smith, and D. Brooks, "Voltage emergency prediction: Using signatures to reduce operating margins," in *IEEE Int. Symp. High Performance Computer Architecture (HPCA)*, Feb. 2009.
[21] R. C. N. Pilawa-Podgurski, D. M. Giuliano, and D. J. Perreault, "Merged two-stage power converter architecture with soft charging switched-capacitor energy transfer," in *IEEE Power Electronics Specialists Conf. (PESC)*, Jun. 2008.
[22] W. –H. Ki, F. Su, and C. –Y. Tsui, "Charge redistribution loss consideration in optimal charge pump design," in *IEEE Int. Symp. Circuits and Systems (ISCAS)*, May 2005.
[23] *ASITIC – Analysis and Simulation of Inductors and Transformers for Integrated Circuits*, [Online]. Available: http://www.eecs.berkeley.edu/~niknejad/asitic.html
[24] A. V. Peterchev and S. R. Sanders, "Quantization resolution and limit cycling in digitally controlled PWM converters," *IEEE Trans. Power Electronics*, vol. 18, no. 1, pp. 301–308, Jan. 2003.

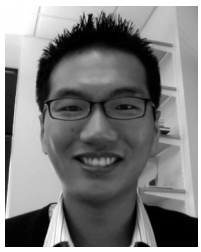

Wonyoung Kim (S'08) is a Ph.D. candidate in the School of Engineering and Applied Sciences at Harvard University, Cambridge, MA. He received the B.S. degree in electrical engineering from KAIST, Korea, in 2006.

He spent 6 months at IBM Austin Research Labs in 2008 as a co-op intern. His research interests are in integrated voltage regulator design and in co-optimizing power conversion circuits and processor architecture for fine-grain power management in chip-multiprocessors.

David Brooks (M'02) received the B.S. degree in electrical engineering from the University of Southern California, Los Angeles, and the M.A. and Ph.D. degrees in electrical engineering from Princeton University, Princeton, NJ.

He is a Gordon McKay Professor of Computer Science in the School of Engineering and Applied Sciences at Harvard University, Cambridge, MA. He joined Harvard in 2002 after spending one year as a research staff member at IBM T. J. Watson Research Center. His research interests include resilient and power-efficient computer hardware and software design for high-performance and embedded systems.

Prof. Brooks is a member of the IEEE and ACM.

Gu-Yeon Wei (M'00) received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1994, 1997, and 2001, respectively.

He is Gordon McKay Professor of Electrical Engineering and Computer Science and Associate Dean for Academic Programs in the School of Engineering and Applied Sciences at Harvard University, Cambridge, MA. His research interests span several areas, including mixed-signal circuit design, computer architecture, and compilers, to address variability and power consumption issues in computing systems that use nanoscale CMOS technologies.