Power Electronic Systems
Laboratory

IEEE Transactions on Power Electronics, Vol. 32, No. 1, pp. 378-393, January 2017

A 10 W On-Chip Switched Capacitor Voltage Regulator With Feedforward Regulation Capability for Granular Microprocessor Power Delivery

T. M. Anderson,
F. Krismer,
J. W. Kolar,
T. Toifl,
C. Menolfi,
L. Kull,
T. Morf,
M. Kossel,
M. Brändli,
P. A. Francese

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

A 10 W On-Chip Switched Capacitor Voltage Regulator With Feedforward Regulation Capability for Granular Microprocessor Power Delivery

Toke M. Andersen, *Student Member, IEEE*, Florian Krismer, *Member, IEEE*, Johann W. Kolar, *Fellow, IEEE*, Thomas Toifl, *Senior Member, IEEE*, Christian Menolfi, *Member, IEEE*, Lukas Kull, *Member, IEEE*, Thomas Morf, *Senior Member, IEEE*, Marcel Kossel, *Senior Member, IEEE*, Matthias Brändli, and Pier Andrea Francese, *Member, IEEE*

***Abstract*—Granular power delivery with per-core regulation for microprocessor power delivery has the potential to significantly improve the energy efficiency of future data centers. On-chip switched capacitor converters can enable such granular power delivery with per-core regulation given a high efficiency, high power density, fast response time, and high output power converter design. This paper details the implementation of an on-chip switched capacitor voltage regulator in a $32\,\mathrm{nm}$ SOI CMOS technology with deep trench capacitors. A novel feedforward control for reconfigurable switched capacitor converters is presented. The feedforward control reduces the output voltage droop following a transient load step. This leads to a reduced minimum microprocessor supply voltage, thereby reducing the overall power consumption of the microprocessor. The implemented on-chip switched capacitor voltage regulator provides a 0.7–$1.1\,\mathrm{V}$ output voltage from $1.8\,\mathrm{V}$ input. It achieves a 85.1% maximum efficiency at $3.2\,\mathrm{W/mm^2}$ power density, a subnanosecond response time with improved minimum supply voltage capability, and a maximum output power of $10\,\mathrm{W}$. For an output voltage of $850\,\mathrm{mV}$, the feedforward control reduces the required voltage overhead by $60\,\mathrm{mV}$ for a transient load step from 10% to 100% of the nominal load. This can reduce the overall power consumption of the microprocessor by 7%.**

***Index Terms*—Feedforward systems, power integrated circuits, regulators, switched capacitor circuits.**

I. Introduction

ELECTRICITY consumption within information and communication technologies (ICT) now approaches 10% of the world's total electricity consumption. That is, $1500\,\mathrm{TWh}$ annual electricity consumption, and forecasts predict as much as $6000\,\mathrm{TWh}$ annual electricity consumption for ICT in 2035 [1]. This tremendous amount of electricity is used to produce, store, transport, process, and display the zettabytes of data produced and used by: 1) data centers; 2) wired and wireless communication infrastructure; and 3) end-user devices such as personal computers, smart phones, and digital televisions. At the heart of all of these lies the data processing units like microprocessor cores, caches, graphic processors, I/O circuits and networks, etc. From a power conversion perspective, these data processing units act as electronic loads powered by a point of load (POL) converter, which is the last power conversion stage in the entire power delivery chain. For high-performance multicore and manycore microprocessors, which are the main target application of this paper, the POL converter typically consists of an external voltage regulator module (VRM).

Switched capacitor (SC) converters have recently become a popular topology for on-chip voltage regulators due to their ease of monolithic integration [2]–[7], [11], [22], [23], [25], [29], [46]. Using only transistors and capacitors available in the semiconductor processes, SC converters offer the possibility to be implemented on the same die as the microprocessor using no external components, thereby providing a monolithic (2D) converter integration with the load.

This paper presents a fully on-chip switched capacitor voltage regulator (SCVR) designed and implemented in a $32\,\mathrm{nm}$ SOI CMOS technology. This semiconductor technology features the high capacitance density and low loss deep trench capacitor, which is known to result in high efficiency and high power density on-chip SCVR designs [3], [5], [13]. Historically, SC converters are perceived as low efficiency, low power density, difficult to regulate, and low power ($< 1\,\mathrm{W}$) converters. This paper, which extends the work from [47], demonstrates an on-chip SCVR that simultaneously achieves high efficiency, high power density, subnanosecond regulation capability, and high output power. As shown in the state-of-the-art overview in Fig. 1(a), this design is among the highest efficiency and highest power density on-chip SC converters presented to date. Also, as shown in Fig. 1(b), this converter delivers $10\,\mathrm{W}$ maximum output power. This output power is six times or higher than any other on-chip SCVR converter presented to date, and it shows that SC converters can achieve similar performance as their (inductive) buck converter counterparts.

For the application of granular microprocessor power delivery, the on-chip SCVR must always maintain an output voltage above a certain minimum level $V_{\mathrm{out,min}}$ for the microprocessor core to meet setup time requirements. Following a transient load

Manuscript received October 14, 2015; revised December 25, 2015; accepted February 3, 2016. Date of publication February 16, 2016; date of current version September 16, 2016. Recommended for publication by Associate Editor J. A Oliver.

T. M. Andersen, F. Krismer, and J. W. Kolar are with the Power Electronic Systems Laboratory, Swiss Federal Institute of Technology Zurich, 8092 Zurich, Switzerland (e-mail: andersen@lem.ee.ethz.ch; krismer@lem.ee.ethz.ch; kolar@lem.ee.ethz.ch).

T. Toifl, C. Menolfi, L. Kull, T. Morf, M. Kossel, M. Brändli, and P. A. Francese are with the IBM Research Zurich, 8803 Rüschlikon, Switzerland (e-mail: tto@zurich.ibm.com; cme@zurich.ibm.com; lku@zurich.ibm.com; tmr@zurich.ibm.com; mko@zurich.ibm.com; mbr@zurich.ibm.com; pfr@zurich.ibm.com).

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/TPEL.2016.2530745

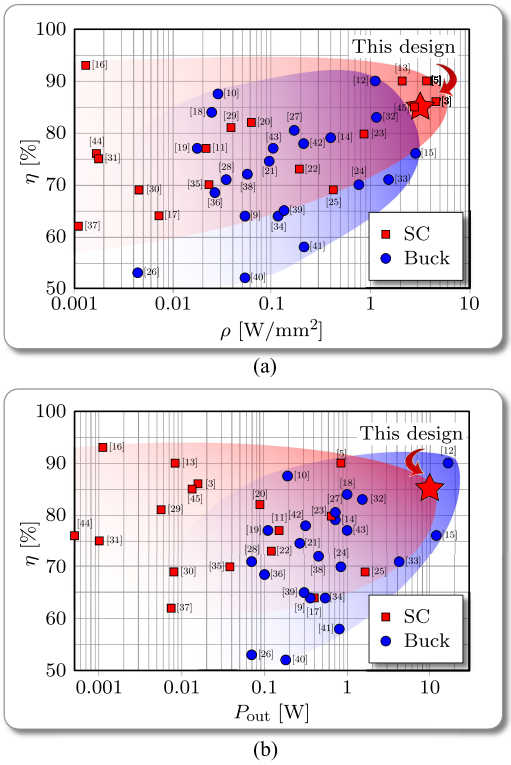

Fig. 1. State-of-the-art overview of (a) quoted efficiency versus power density and (b) quoted efficiency versus maximum output power of published on-chip voltage regulators [3], [5], [9]–[45].

step, the output voltage typically experiences a droop due to the parasitic inductance of the power distribution network (PDN). Therefore, the steady-state output voltage is kept high enough to ensure $V_{\text{out}} > V_{\text{out,min}}$ at all times, thereby introducing a voltage overhead that leads to increased microprocessor energy consumption [12], [48]. The single-bound hysteretic control for interleaved SCVRs provides a subnanosecond response time to a transient load step [5], [8], [22]. However, as reported in [5], the output voltage can still experience a voltage droop despite the subnanosecond response time. As discussed, the droop is neither caused by the regulation loop being too slow nor by the lack of output decoupling capacitance. Instead, a significant droop at the input of the converter is considered to be the root cause of the output voltage droop. A minimum input voltage is required by the SCVR to maintain its output voltage. If the input voltage droops below this minimum voltage, and since the SCVR cannot boost the output voltage, the output voltage will droop. The input voltage droop, and thereby the output voltage droop, can be reduced by implementing sufficient on-chip input decoupling capacitors. However, a large capacitance is needed to reduce the droop sufficiently, and it becomes impractical to implement owing to the large chip area overhead required. This paper presents a novel feedforward control scheme that mitigates the output voltage droop without adding excessive input decoupling capacitors. This feedforward regulation can enable per-core dynamic voltage and frequency scaling (DVFS) with reduced voltage overhead for future energy-efficient microprocessor power delivery.

Section II provides a detailed application-specific motivation for implementing SCVRs for granular microprocessor power delivery with per-core regulation. Furthermore, the converter specifications are defined. The system overview of the implemented on-chip SCVR is presented in Section III, and the SC converter power stage and gate driver are presented in Section IV. In Section V, the implementation of the single-bound hysteretic feedback control is detailed, and in Section VI, the novel feedforward control for reconfigurable SC converters is presented. Section VII presents the experimental results of a 10 W on-chip SCVR implemented in a 32 nm SOI CMOS technology that features the high capacitance density deep trench capacitors.

II. Granular Microprocessor Power Delivery

On-chip SCVRs enable granular power delivery by providing independent voltage domains with various voltage and current specifications. These voltage domains include microprocessor cores, caches, signal I/O's, memory, graphic processors, etc., either all on the same die or on a separate die within the package [11], [12], [49]. The SCVR generates the desired supply voltage from a higher-than-nominal input voltage, thereby becoming an independent POL converter for each voltage domain. Within a voltage domain, the granularity can be further extended by incorporating per-core regulation, where each core is regulated independently of each other. Furthermore, the granularity can be even further extended by considering within-core regulation, which counteracts and stabilizes voltage variations across the core die [50]. Per-core regulation is treated further in the next subsection, whereas within-core regulation is not treated further in this paper.

Using an example of five independent and different voltage domains, Fig. 2 illustrates how SCVRs can be implemented to provide granular power delivery. The example voltages and currents shown are representative for a future high-performance microprocessor system. The typical granular power delivery is shown in Fig. 2(a), where several external VRMs are used to supply each of the independent voltage domains. Each external VRM has different output voltage and current requirements. Additionally, the external VRMs typically take up a significant fraction of the total motherboard area. In this future microprocessor example, a current of approximately 360 A through the combined PDN is imaginable. Most of this current is for V_1 that supplies the microprocessor cores. The 360 A current applies at the highest supply voltage of 1.1 V. At reduced supply voltages, the microprocessor core is expected to draw a lower current due to the reduced clock frequency of the core. This load characteristic matches well to the SC converter efficiency characteristic, which will be evident from the efficiency measurements in Section VII-B where the highest efficiencies are obtained at high output voltages. Furthermore, these high currents are challenging to supply through the PDN due to: 1) wire voltage drops (IR drops) in the parasitic wiring resistance; 2) $L\mathrm{d}i/\mathrm{d}t$ supply instability caused by the parasitic wiring inductance; and 3) C4 solder bumps, which are typically limited to around 200 mA per C4 solder bump. According to the international technology roadmap for semiconductors, a modern high-performance

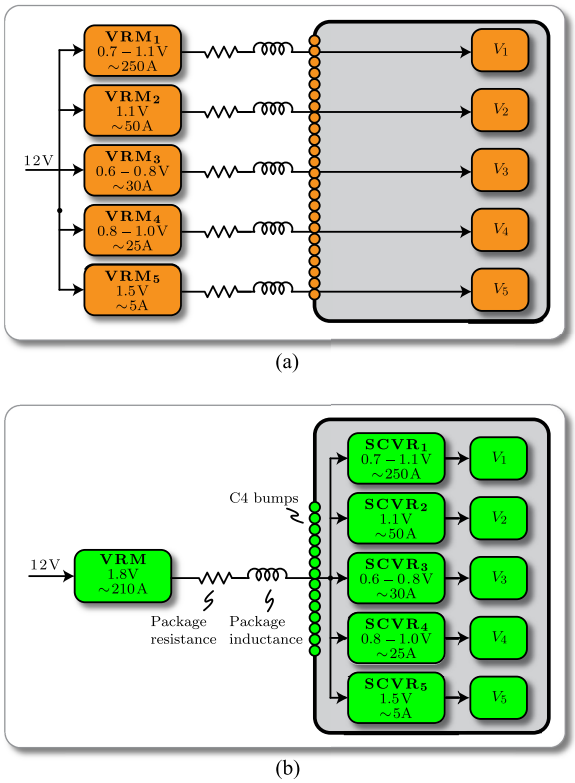

Fig. 2. Example microprocessor system consisting of five independent voltage domains with different current and voltage specifications. (a) Typical power delivery where five external VRMs are designed to supply each independent voltage domain. (b) Granular power delivery where a single external VRM supplies five on-chip SCVRs, each of which supplies an independent voltage domain.

microprocessor package can incorporate more than 3000 C4 solder bumps, and more than two-thirds of these are used to supply power [51]. Hence, supplying approximately 360 A through the PDN would require at least 1800 C4 solder bumps, thereby being feasible.

The target granular power delivery shown in Fig. 2(b) utilizes SCVRs to supply each independent voltage domain from a single external VRM, which provides a higher-than-nominal supply voltage. Each SCVR is dimensioned and scaled to match the requirements of its respective voltage domain. With 1.8 V input voltage and the same total dissipated power, only around 210 A of current flows through the PDN, thereby directly reducing the issues with IR losses, $L\mathrm{d}i/\mathrm{d}t$ supply instability, and limited C4 solder bumps. Additionally, the design of the external VRM can be simplified when converting down to a fixed and higher voltage instead of a dynamic and lower voltage. Also, the efficiency of the VRM alone can be improved 3–4% [48]. Finally, with only a single external VRM, motherboard area used for power delivery can be significantly reduced, allowing for an overall smaller form factor of the entire microprocessor system.

The 10 W output power delivered by the on-chip SCVR presented in this paper is not enough to supply the entire microprocessor. It is suited to supply a voltage domain or a single core in a per-core regulation scheme requiring 10 W. Besides, the SCVR can generally be scaled to deliver higher or lower output powers depending on the microprocessor load requirements. The 10 W output power in this design is therefore not a specific target in itself; it is merely what could be achieved with the silicon area available for the given semiconductor technology.

By incorporating on-chip SCVR, sensitive analog and I/O circuitry could potentially be disturbed by the switching behavior of the SCVR. However, the entire microprocessor core already generates substrate noise and crosstalk from switching of the logic gates. It is the assumption that if the analog and I/O circuitry can cope with the substrate noise from the microprocessor core (e.g., by proper power plane separation), then those circuits will also cope with an on-chip SCVR using similar techniques. A detailed investigation of potential crosstalk from an on-chip SCVR to sensitive analog and I/O circuitry is beyond the scope of this paper.

A. Per-Core Regulation

DVFS is a popular technique to dynamically adjust the voltage and clock frequency of a microprocessor core to meet, but not exceed, supply voltage demands [48]. Changing the state of a logic circuit electrically translates into charging or discharging the total capacitance C_{tot} of all logic gates that change state. With a clock frequency f_{clk}, the total power dissipation of the logic circuit, e.g., a microprocessor core, can be estimated using

$$P_{\mathrm{logic}} = C_{\mathrm{tot}} V_{\mathrm{sup}}^2 f_{\mathrm{clk}} \tag{1}$$

where V_{sup} is the supply voltage of the logic circuit. DVFS ensures that the supply voltage and clock frequency follow the demand set by the microprocessor core for a given computational workload. From a power delivery perspective, only V_{sup} can be regulated to save system power as both C_{tot} and f_{clk} are determined by the microprocessor core and the workload. From (1), the power dissipation reveals a quadratic dependency of V_{sup} on P_{logic}. Hence, a power management scheme is introduced to meet, but not exceed, the demand in V_{sup}.

Three power management schemes for an example four core microprocessor are shown in Fig. 3. The three example workloads performed by the microprocessor have specific core utilization profiles that dictate the supply voltage requirements for each core [43], [48]. Following a transient load step, the supply voltage typically experiences a droop. Therefore, extra voltage overhead is added on top of the core utilization profile to support the frequency scaling, which is equivalent to a transient load step. However, the voltage overhead directly leads to system energy loss. The function of the power management scheme is to reduce the voltage overhead as much as possible while still meeting the supply voltage demands at all times.

Most power management schemes today support the DVFS shown in Fig. 3(a), where a single supply voltage is delivered to all cores simultaneously, and the DVFS voltage is adjusted to the voltage requirement of the core having the highest utilization profile within a certain workload. As can be seen for workloads 1 and 3, cores with low utilization profiles experience a large voltage overhead, thereby leading to undesired system energy losses. For workload 2 having uniform core utilization, this power management scheme provides no system energy savings. For per-core DVFS shown in Fig. 3(b), the supply voltage

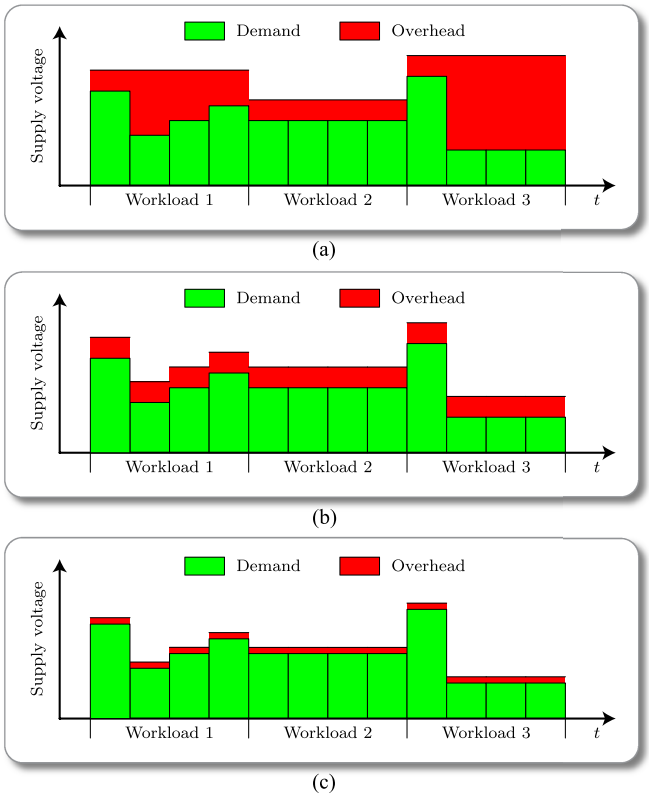

Fig. 3. Power management schemes for multicore and manycore microprocessor systems. (a) DVFS without per-core regulation, which results in an unnecessarily large voltage overhead when only a single core requires a high supply voltage. (b) Per-core regulation reduces this voltage overhead. (c) Further reduction of the supply voltage is achieved if the voltage overhead itself is decreased. Any voltage overhead translates directly to system energy loss.

of each core is independently regulated to match its utilization profile within each workload. This power management scheme significantly reduces the voltage overhead in workload 1 and 3, thereby reducing the system energy losses. However, the voltage overhead to account for transient load steps is still required, and there is again no system energy savings for workload 2. In Fig. 3(c), the improved per-core DVFS with reduced voltage overhead is shown as the ultimate goal where overhead voltage for each core is minimized. To be feasible, the power management scheme must provide a solution to reduce the voltage droop following a transient event, thereby allowing for an overall lower supply voltage with per-core DVFS. This attractive power management scheme leads to minimal system energy losses for all workloads. Although not shown due to clarity of the figure, the overhead voltage could potentially be further reduced at low supply voltage since the transients at lower supply voltages are smaller than for high supply voltage

For a single voltage domain, e.g., V_1 from Fig. 2, per-core DVFS can be supported by incorporating multiple converters on the microprocessor die. According to [48], per-core DVFS can, depending on workload, increase the overall system efficiency by up to 21%. This efficiency improvement stems mainly from reduced power losses in the PDN, and it includes the power losses of the added on-chip converter. Furthermore, if the converter can reduce the voltage overhead while still maintaining the minimum voltage requirement by each core, the per-core DVFS with reduced overhead voltage from Fig. 3(c) can be implemented. These attractive benefits motivate for further investigation and exploration of SCVRs to enable per-core DVFS with reduced overhead voltage.

B. Converter Specifications

Based on the motivation for granular microprocessor power delivery with per-core DVFS given above, the following defines and discusses the main SCVR specifications targeted in this paper:

1) *Monolithic (2D) integration:* requires the SCVR to be implemented onto the microprocessor die, i.e., to be designed using the same semiconductor technology as the microprocessor.
2) *High efficiency:* The 85–90% efficiency range is targeted. Implementing an SCVR with low efficiency would severely limit or even negate the power and energy efficiency gains achieved by reduced IR losses in the PDN and reduced overhead voltages using per-core DVFS.
3) *High power density:* To integrate the SCVR with the microprocessor, the power density of the converter must be high with the respect to the state-of-the-art overview in Fig. 1(a) in order to not take up too much chip area. Therefore, power densities above $1\,\mathrm{W/mm^2}$ are targeted.
4) *High output power:* From the state-of-the-art overview in Fig. 1(b), prior to this paper, only the SC converter design in [25] demonstrates more than $1\,\mathrm{W}$ output power. It is therefore a goal to demonstrate more than $1\,\mathrm{W}$ maximum output power.
5) *Wide output voltage range:* To enable per-core DVFS, the output voltage range is 0.7–1.1 V for a fixed $1.8\,\mathrm{V}$ input supply. For the $32\,\mathrm{nm}$ technology used, the nominal voltage is $0.9\,\mathrm{V}$; hence, the higher-than-nominal $1.8\,\mathrm{V}$ input supply supports the granular power delivery shown in Fig. 2.
6) *Fast transient response:* required to enable the per-core DVFS that supports various workload supply voltage requirements of individual microprocessor cores as shown in Fig. 3(b). Therefore, a response time of the SCVR of $1\,\mathrm{ns}$ is targeted for a step from 10% to 100% of the nominal load current. The nominal load current for this design corresponds to $4.3\,\mathrm{A}$ at $850\,\mathrm{mV}$ output voltage.
7) *Reduced output voltage overhead:* As shown in Fig. 3(c), reducing the output voltage overhead while still maintaining a certain $V_{\mathrm{out,min}}$ at all times is an additional improvement to per-core DVFS. Therefore, reducing the voltage overhead serves as a design goal.

The specifications discussed above are summarized in Table I. Note that the switching frequency is not specified since it is considered a design parameter that follows the design of the converter. The methodology for achieving a Pareto-optimized design of the corresponding SC converter is detailed in [7].

III. SCVR System Overview

Given the converter specifications above, the complete overview of the implemented SCVR is shown in Fig. 4. A

TABLE I
SPECIFICATIONS OF THE ON-CHIP SWITCHED CAPACITOR VOLTAGE REGULATOR

Parameter	Symbol	Specification
Efficiency	η	$> 85\%$
Power density	ρ	$> 1\,\mathrm{W/mm^2}$
Output power	P_{out}	$> 1\,\mathrm{W}$
Output voltage	V_{out}	$0.7 - 1.1\,\mathrm{V}$
Transient response	–	$< 1\,\mathrm{ns}$
Voltage overhead	–	Minimal
Level of integration	–	2D

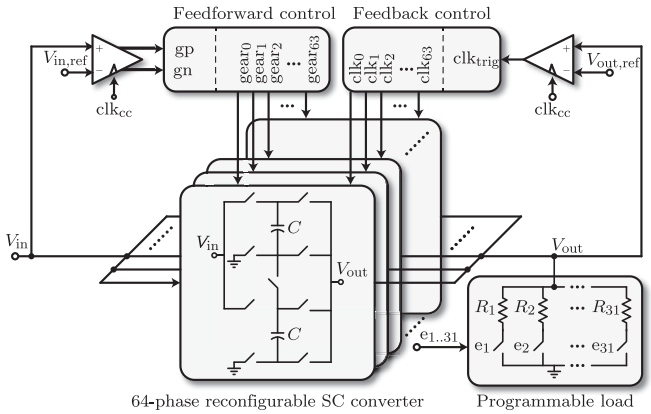

Fig. 4. System overview of the 64-phase 2:1 and 3:2 reconfigurable SCVR. The feedforward control scheme works in conjunction with the single-bound hysteretic control scheme (feedback control).

64-phase interleaving scheme of a 2:1 and 3:2 reconfigurable SC converter is employed. The feedback control is implemented as a single-bound hysteretic control (delta modulation) comprising a clocked comparator and a digital clock interleaver. The feedforward control is implemented using a clocked comparator and a digital gear controller that dynamically changes the configuration (gear) of the interleaved SC converter units. Both the feedback and the feedforward controls use a $4\,\mathrm{GHz}$ clock. The programmable load is configured externally by a digital configuration interface (not shown) by enabling the signals $\mathrm{e}_{1..31}$. For no load, all signals $\mathrm{e}_{1..31}$ are logic 0.

From simulations, 32 interleaved stages are found to be a good match with the $4\,\mathrm{GHz}$ clock for this design and with this semiconductor technology. The 64 interleaved stages are taken by implementing two 32 interleaved converters, but one runs off the inverted clock. However, 128 or more could also be possible, but the maximum switching frequency of each SC converter would be reduced, as discussed in Section V-B. From simulations, the additional ripple reduction for more than 16 interleaved stages compared to 16 is not significant, but even a small improvement can prove useful in an optimized design.

A. 2:1 and 3:2 Reconfigurable Power Stage

The 2:1 and 3:2 reconfigurable power stage is shown in Fig. 5. It consists of two flying capacitors and nine transistors that are operated as switches [5], [6], [23], [29], [47]. The reconfigurable power stage supports a high efficiency over the entire output

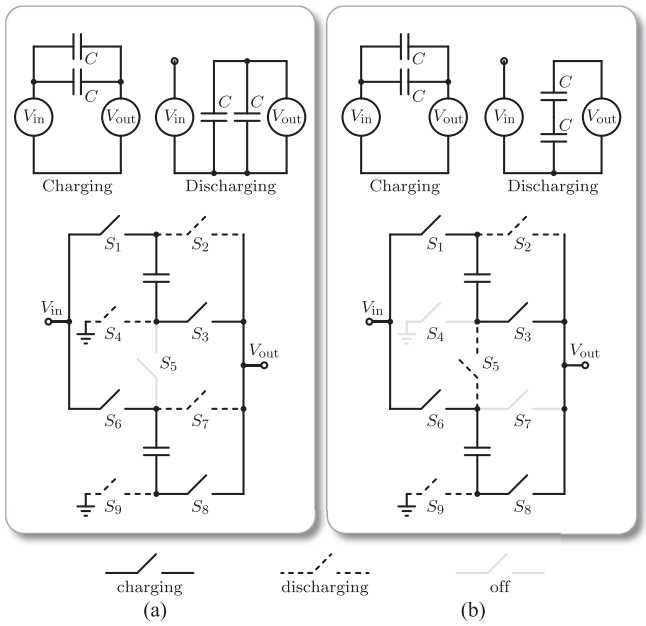

Fig. 5. The (a) 2:1 and (b) 3:2 reconfigurable SC converter power stage including the switch configuration in the charging and the discharging state.

voltage range of 0.7–$1.1\,\mathrm{V}$. Note that the power stage also can be operated at lower output voltages than $0.7\,\mathrm{V}$, but at a lower efficiency.

In each configuration, the two flying capacitors are sequentially switched between a charging and a discharging state at 50% duty cycle. In the 2:1 configuration shown in Fig. 5(a), the switch S_5 connecting the two flying capacitors is always off, and the power stage reduces to two 2:1 converters operated in parallel. In the charging state, the two paralleled flying capacitors are in series with the input and the output nodes. In the discharging state, the two paralleled flying capacitors are in parallel with the output node, and the input node is unconnected. In the 3:2 configuration shown in Fig. 5(b), the parallel connection of the two flying capacitors is in series between the input and the output nodes in the charging state. In the discharging state, the series connection of the two flying capacitors is in parallel with the output node, and the input node is unconnected.

B. Interleaving

Interleaving is a popular technique employed in SCVRs [8], [16], [23], [25], [29], [52], [53]. Instead of implementing only a single SC converter unit, the SC converter is divided into several smaller units having the input clock signals phase shifted with respect to each other. Fig. 6 shows how employing interleaving is a means to simultaneously reduce the input current and the output voltage ripples[1] without excessive input and output decoupling capacitors.

With a single-phase implementation as shown in Fig. 6(a), the output voltage and input current ripples typically require

[1]Although the input current and output voltage ripple decays in Fig. 6 are shown to be linear, they conceptually are exponential waveforms resulting from the charging and discharging of the flying capacitor through the transistor on-state resistances and the load [7].

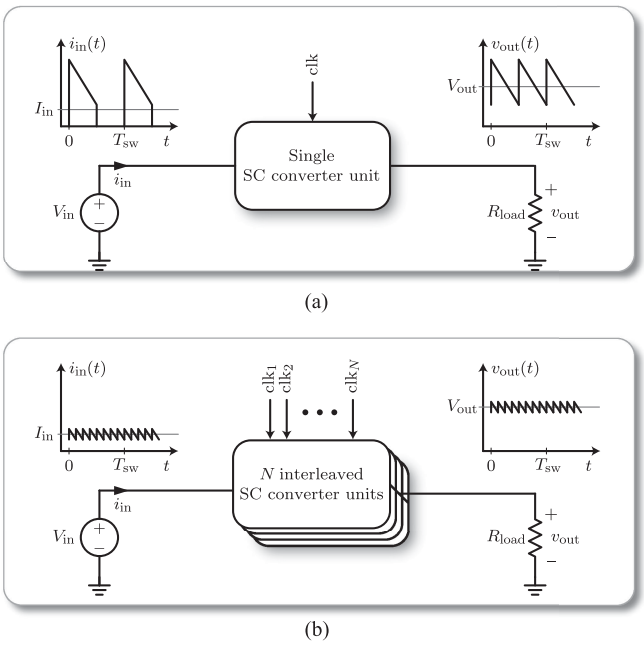

Fig. 6. SC converter implementation both with and without interleaving. Interleaving SC converter units is a means to simultaneously reduce the input current ripple and the output voltage ripple, thereby greatly reducing (or even omitting) the need for input and output decoupling capacitors.

substantial decoupling capacitance to meet ripple specifications. This directly translates into occupying significant chip area with capacitors. Considering the output voltage ripple, the ripple in itself might not be critical for the operation of the microprocessor core as long as the $V_{\rm out,min}$ requirements are met, but the larger ripple can be viewed as voltage overhead, which is sought reduced according to Fig. 3.

With an N-phase interleaving implementation as shown in Fig. 6(b), whenever an SC converter phase changes state from charging to discharging or vice versa, the flying capacitors of the remaining $N-1$ phases effectively act as decoupling to that switching event. Hence, the input and output decoupling capacitors required to reduce the steady-state ripples can be greatly reduced or even omitted, thereby saving precious chip area and at the same time reduce the voltage overhead.

IV. Power Stage and Gate Driver

The complete circuit schematic of the SC converter unit is shown in Fig. 7. It consists of a gate driver that generates the transistor gate signals, multiplexers that distribute the gate driver signals to the transistor gates (based on the configuration), and the 2:1 and 3:2 reconfigurable power stage. This SC converter is implemented using thin-oxide transistors in the $32\,\rm nm$ SOI CMOS technology. For good reliability, the allowed maximum voltage of the transistors is $V_{\rm max}=1.2\,\rm V$, which is lower than the input voltage of $V_{\rm in}=1.8\,\rm V$ for this application [5], [12]. Thick-oxide transistors within the technology would be able to withstand the input voltage, but the resulting efficiency degradation due to worsened $R_{\rm on}Q_{\rm g}$ merit makes thick-oxide transistors for this high-efficiency design impractical to use.

A. Stacked Voltage Domain Gate Driver

The gate driver, which generates the four main gate signals[2] $v_{\rm g,xY}$ for the power transistors $S_{1-9(\rm s)}$, has to be designed ensuring that no pair of transistor terminals is exposed to overvoltage conditions. With a maximum transistor voltage of $1.2\,\rm V$, the $1.8\,\rm V$ input supply requires $2\times$ stacking to avoid overvoltage situations.[3] Therefore, the gate driver in Fig. 7(a) employs a stacked voltage domain where the upper voltage domain for $v_{\rm g,xH}$ is maintained between $V_{\rm in}$ and $V_{\rm out}$ and the lower voltage domain for $v_{\rm g,xL}$ is maintained using $V_{\rm out}$ and $\rm gnd$ [3].

The input clock $\rm clk_{\it n}$ of the *n*'th SC converter unit is fed into the lower voltage domain. A level shifter circuit shifts the input clock to the upper voltage domain. In each voltage domain, the clock signal is passed through a latch with built-in delay (nonoverlapping clock) in order to generate a dead time interval between the signal edges to avoid shoot-through currents in the power transistors.

The stacked voltage domain gate driver ensures a safe startup with no overvoltage of any transistor. However, the load must be disconnected, i.e., no current is drawn from $V_{\rm out}$, during startup. The reason for the safe startup is the symmetrical design of the level shifter and nonoverlapping clock circuits that divide the input supply ramp up close to equally between the two domains. The safe startup operation has been examined thoroughly in simulations under all relevant corner conditions.

B. Gate Multiplexers and Buffers

The four main gate signals $v_{\rm g,xY}$ are distributed to the power stage transistors based on whether the converter is in the 2:1 or the 3:2 configuration. From Fig. 7(b), four multiplexer are required, one for each of the gate signals $v_{\rm g4}$, $v_{\rm g5}$, $v_{\rm g7}$, and $v_{\rm g7s}$ that change clock source based on the configuration. Tapered buffers after the multiplexers provide the drive strength to turn the transistors on and off in the power stage. Tapered buffers are also added for the gate signals that do not require multiplexers.

The gate signals of transistors S_5 and $S_{5\rm s}$ in the 3:2 configuration are unusual as the gates are tied to $V_{\rm out}$. This implementation is feasible since the transistors' source potential V_2 is equal to $V_{\rm out}/2$ when conducting, resulting in a very simple and robust gate drive of S_5. However, the gate–source voltage when conducting is only $V_{\rm out}/2$, which increases the on-state resistances[4] of these transistors compared to the other transistors. The increased on-state resistances result in a slight penalty in the maximum achievable power density in the 3:2 configuration [3].

[2]Regarding $v_{\rm g,xY}$, subscript x is n for NMOS transistors and p for PMOS transistors, and subscript Y is H for high side and L for low side.

[3]Higher input voltages of, e.g., $2.5\rm V$ or $3.3\,\rm V$ is feasible with SC converters having higher conversion ratios, but the $3\times$ or $4\times$ stacking required to withstand the higher input voltage would significantly increase circuit complexity in order to ensure that no transistor encounters an overvoltage situation during operation, including startup.

[4]To a first order, the on-state resistance of a MOS transistor is inversely proportional to the gate–source voltage.

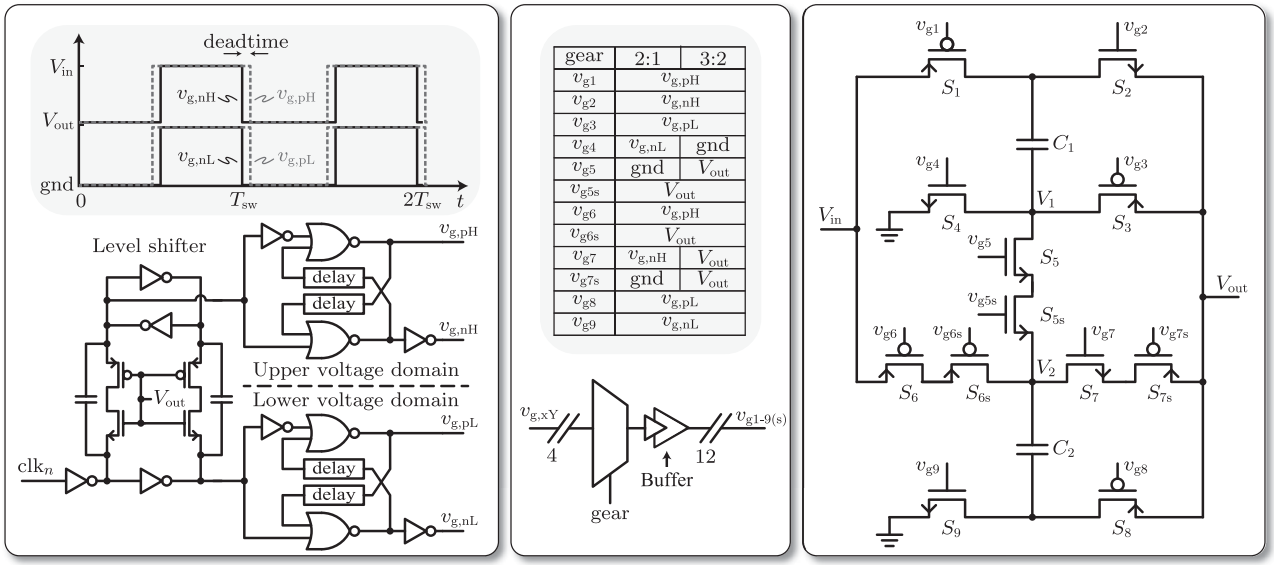

(a) (b) (c)

Fig. 7. Complete SC converter unit. (a) Stacked voltage domain gate driver consisting of a level shifter and, for each domain, a nonoverlapping clock that generates the dead time. (b) Gate multiplexers that select the appropriate gate signals for each transistor based on the configuration. (c) 2:1 and 3:2 reconfigurable SC power stage.

C. Stacked Transistor Implementation

The reconfigurable SC converter power stage in Fig. 5(c) inherently exposes transistors S_5 and S_6 to overvoltage, and transistor S_7 inherently exhibits undesired turn-on behavior. Therefore, the power stage in Fig. 7(c) includes three stacked transistors $S_{5\rm s}$, $S_{6\rm s}$, and $S_{7\rm s}$. The purpose of $S_{5\rm s}$ and $S_{6\rm s}$ is to avoid overvoltage situations inherent in the power stage. The voltage drop across the turned-off transistors get reduced to values smaller than $1.2\,\rm V$ owing to the transistor stacking because the source potential of the stacked transistors cannot get lower than its gate voltage, and hence, the gate voltage clamps the voltage across the turned-off transistor correspondingly. Regarding $S_{7\rm s}$, it is important to note that the transistors in the $32\,\rm nm$ semiconductor technology are laid out symmetrically, meaning that there is no physical difference between the drain and source terminals. This implies that a transistor can be turned on with zero gate–source voltage if the gate–drain voltage is above the threshold voltage. The purpose of $S_{7\rm s}$ is to prevent undesired turn on due to layout symmetry.

1) Stacked Transistor $S_{5\rm s}$: From Fig. 5, it is clear that S_5 in Fig. 7(c) should always be off in the 2:1 configuration. However, simply grounding its gate to turn it off leads to an overvoltage situation since the voltage at node V_2 is $V_{\rm in}$ in the charging phase of the 2:1 configuration, thereby resulting in an unacceptably high gate–source voltage of $-V_{\rm in}$ for S_5. Stacked transistor $S_{5\rm s}$, which has its gate tied to $V_{\rm out}$, results in a tolerable gate–source voltage of $V_{\rm out}-V_{\rm in}$. However, S_5 is still needed due to layout symmetry: without S_5, the gate–drain of $S_{5\rm s}$ would equal $V_{\rm out}$ in the discharging state of the 2:1 configuration because $V_1=0\,\rm V$. This would lead $S_{5\rm s}$ to undesirably turn on. For these reasons, both S_5 and $S_{5\rm s}$ are needed, and together they hinder the overvoltage situation and ensure the desired turn off in the 2:1 configuration.

2) Stacked Transistor $S_{6\rm s}$: From Fig. 7(c), it is seen that the voltage V_2 can get as low as $V_{\rm out,min,3:2}/2=900\,\rm mV/2=450\,\rm mV$ in the 3:2 configuration. This low voltage imposes an overvoltage situation of $V_{\rm in}-450\,\rm mV=1.35\,\rm V$ on the drain–source and gate–drain terminals of S_6. For the gate–drain overvoltage, the stacked transistor $S_{6\rm s}$ effectively overcomes this overvoltage situation in a similar manner as for $S_{5\rm s}$ discussed above. With the gate of $S_{6\rm s}$ tied to $V_{\rm out}$, the desired switching in both configurations is solely determined by the switching of S_6. For the drain–source overvoltage, the stacked transistor implementation shares the voltage between the two transistors, thereby eliminating the overvoltage situation.

3) $S_{7\rm s}$: Stacked transistor $S_{7\rm s}$ is inserted to prevent an undesired turn on event of S_7 in the 3:2 configuration, where, according to Fig. 5, it should always be turned off. As before, the voltage at V_2 can be as low as $450\,\rm mV$, which, with $v_{\rm g7}$ tied to $V_{\rm out}$, results in a positive gate–drain voltage of S_7, causing it to undesirably turn on due to layout symmetry. Stacked transistor $S_{7\rm s}$ effectively ensures that S_7 does not turn on in the 3:2 configuration, while at the same time ensuring the transistor turns on as desired in the 2:1 configuration.

V. Single-Bound Hysteretic Control Scheme

A single-bound hysteretic control scheme is implemented as the feedback control loop. This regulation scheme is advantageous due to: 1) its simple implementation; 2) its inherently stable operation; and 3) its high control bandwidth [8], [54]. This regulation scheme modulates the switching frequency to regulate the SC converters' equivalent output resistance to achieve the desired output voltage. Switching frequency modulation is a popular regulation technique for on-chip SC converters [8], [16], [22], [23], [25], [52], [53].

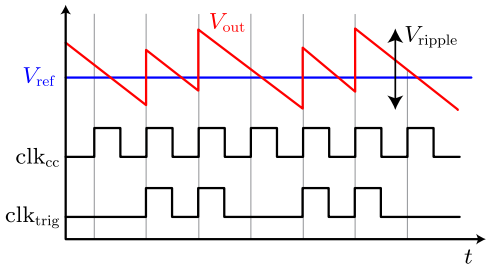

Fig. 8. Single-bound hysteretic control scheme using a clocked comparator to produce a trigger clock $\mathrm{clk_{trig}}$ whenever the output voltage is below the reference voltage at the sampling event.

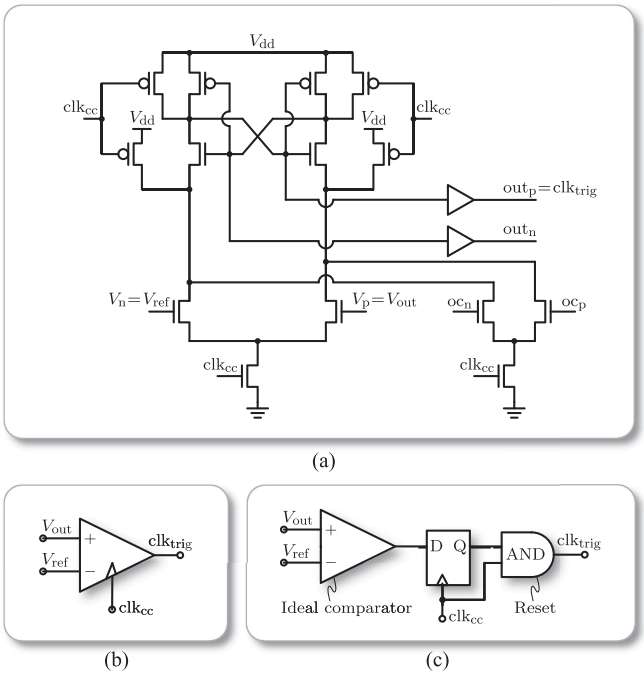

Fig. 9. Clocked comparator with reset used to implement the single-bound hysteretic control scheme. The output $\mathrm{clk_{trig}}$ can be considered as the sampling of an ideal comparator that is reset before the subsequent sampling event.

The concept of the single-bound hysteretic control is shown in Fig. 8. As can be seen, $\mathrm{clk_{trig}}$ transitions to logic high whenever V_{out} is less than V_{ref} at the rising edge of $\mathrm{clk_{cc}}$. A rising edge on $\mathrm{clk_{trig}}$ causes the digital clock interleaver (discussed below) to change the state of the next SC converter unit to deliver more charge to the output, thereby causing a rise in the output voltage. If V_{out} is greater than V_{ref} at the rising edge of $\mathrm{clk_{cc}}$, $\mathrm{clk_{trig}}$ remains logic low and the clock pulse is skipped, and no SC converter unit changes state.

A. Clocked Comparator

The circuit schematic of the clocked comparator with reset is shown in Fig. 9. Based on $\mathrm{clk_{cc}}$, the clocked comparator generates $\mathrm{clk_{trig}}$ used in the single-bound hysteretic control illustrated in Fig. 8. From the circuit schematic in Fig. 9(a), the clocked comparator main inputs are the differential transistor pair with V_{p} and V_{n}. The comparator structure is a sense-amp latch, also known as a Strong-ARM latch. Offset calibration of the comparator is performed by the second differential transistor pair with $\mathrm{oc_p}$ and $\mathrm{oc_n}$. The offset calibration is performed on-chip by incorporating a dual comparator: when one comparator measures the inputs, the other calibrates and vice versa.

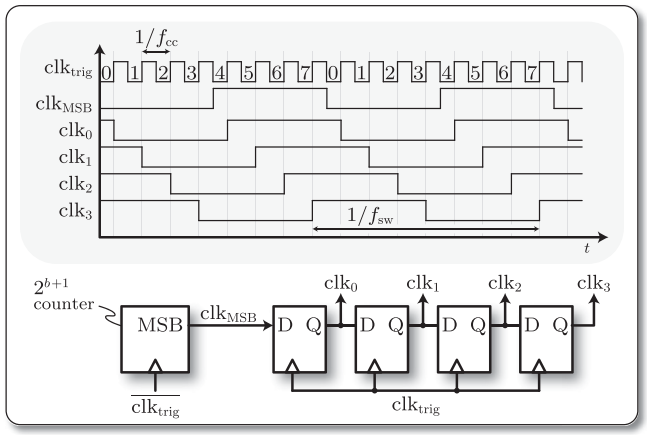

Fig. 10. Digital clock interleaver that divides the high-frequency $\mathrm{clk_{trig}}$ signal into $N = 2^b$ clock phases for the N interleaved SC converter units. Shown for $b = 2$.

The clocked comparator with reset outputs a pulse on $\mathrm{out_p}$ and no pulse on $\mathrm{out_n}$ whenever $V_{\mathrm{p}} < V_{\mathrm{n}}$ following a rising edge of $\mathrm{clk_{cc}}$. It outputs a pulse on $\mathrm{out_n}$ and no pulse on $\mathrm{out_p}$ whenever a $V_{\mathrm{p}} > V_{\mathrm{n}}$ following a rising edge of $\mathrm{clk_{cc}}$. For the single-bound hysteretic control, $\mathrm{out_n}$ is not used. The detailed analysis and design of the clocked comparator, which initially is designed to be used in high-speed analog to digital converters, are treated further in [55].

In Fig. 9(b), the clocked comparator symbol used in the remainder of the thesis is illustrated. The clocked comparator equivalent circuit, which is shown in Fig. 9(c), consists of an ideal comparator followed by a flip-flop that performs the sampling event at each rising edge of the comparator clock $\mathrm{clk_{cc}}$. The AND gate provides a reset of the output trigger signal $\mathrm{clk_{trig}}$ before the subsequent sampling event. Propagation delays are ignored in the equivalent circuit, but it should be ensured that $\mathrm{clk_{trig}}$ contains no glitches.

B. Digital Clock Interleaver

Fig. 10 shows the implementation of the digital clock interleaver, where a shift register performs frequency division of the high-frequency $\mathrm{clk_{trig}}$ signal. The outputs are $N = 2^b$ clock phases, each of which is fed to an SC converter unit as shown in the overview in Fig. 4. In prior art, implementations of this regulation scheme use the inverted output of the last flip-flop and feed it to the first flip-flop of the shift register [8], [53]. Doing this requires the shift register to be properly initialized. As shown in Fig. 10, we use the most significant bit (MSB) of a $(b + 1)$-bit counter as input to the shift register. This solution does not require any initialization of the shift registers since the desired flip-flop states are reached once the counter has completed one full cycle from zero. This still holds if the counter is not zero at start up. Hence, using the counter's MSB as input to the shift register is a very robust implementation.

Assuming that no pulses of $\mathrm{clk_{trig}}$ are skipped (as illustrated in Fig. 10), the maximum switching frequency of each SC converter unit is

$$f_{\mathrm{sw,max}} = \frac{1}{2}\frac{f_{\mathrm{cc}}}{N} = \frac{f_{\mathrm{cc}}}{2^{b+1}} \tag{2}$$

where the factor $1/2$ is due to the fact that it takes two rising edges of $\mathrm{clk_{trig}}$ per SC converter clock period and f_{cc} is the frequency of the clocked comparator. If $f_{\mathrm{sw,max}}$ is reached for a given f_{cc}, the desired output regulation cannot be achieved and the SC converter delivers less power than required. If this occurs, it typically leads to a lower output voltage in steady state or an output voltage droop during a transient load step.

C. Loop Latency

The loop latency, which is the propagation delay from when the sampling event occurs until the corresponding SC converter unit changes state, poses an upper limit to f_{cc}

$$f_{\mathrm{cc}} < \frac{1}{t_{\mathrm{lat}}}. \tag{3}$$

If the above criterion is not met, the subsequent sampling event occurs before the present sampling event has had its effect on the output voltage. If this double sampling event happens, two SC converter units change switching states where only the change of one unit is required. Although a double sampling event in itself is not critical for the stability of the controller, the output voltage will experience a higher ripple, thereby not fully exploiting the ripple reduction of the interleaving scheme [54].

For the test chip, the loop latency is minimized by: 1) optimizing the size (and thereby propagation delay) of logic gates in the stacked voltage domain gate driver; 2) optimizing the digital controller implementation to process without unnecessary latching; and 3) keep buffers and time-critical wiring at a minimum without sacrificing robustness. For this design, the loop latency achieved is around $200\,\mathrm{ps}$.

VI. Novel Feedforward Control

The novel regulation concept, discussed next, is a feedforward control since it regulates based on the input voltage [47]. However, a feedback control scheme, e.g., the single-bound hysteretic control, that regulates based on the output voltage is still required. In Fig. 11, a conceptual overview of the typical feedback control in conjunction with the new feedforward control is illustrated. The SC converter equivalent model (neglecting the switching losses) consists of a transformer with a fixed conversion ratio M and an equivalent output resistance R_{eq} [7], [8]. As shown in Fig. 11(a), the typical feedback control regulates R_{eq} to achieve the desired output voltage as shown in the corresponding flowchart. Typically, as in the single-bound hysteretic control scheme, R_{eq} is modulated by the switching frequency.

The novel feedforward control, which can be implemented with a reconfigurable SC converter, introduces an additional control loop as depicted Fig. 11(b). As shown in the corresponding flowchart, the feedforward control dynamically changes the configuration to a higher voltage conversion ratio $M_2 > M_1$ when an input voltage droop is detected. Once the input voltage has recovered (by the external VRM), the configuration is changed back to M_1, where the converter operation is more efficient. Although shown for the 2:1 and 3:2 reconfigurable SC converter, the principle can be extended to other conversion ratios as well.

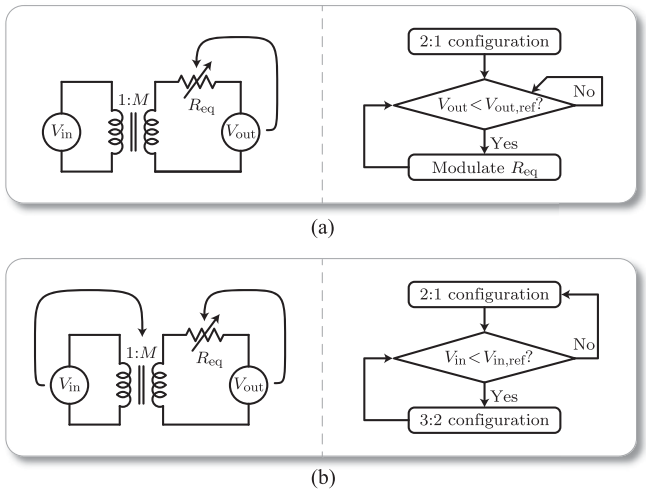

Fig. 11. Typical SC converter feedback control modulates R_{eq} to achieve the desired output voltage, whereas the novel feedforward control dynamically changes the conversion ratio M when an input voltage droop is detected. (a) typical feedback flow chart (b) additional feedforward flowchart.

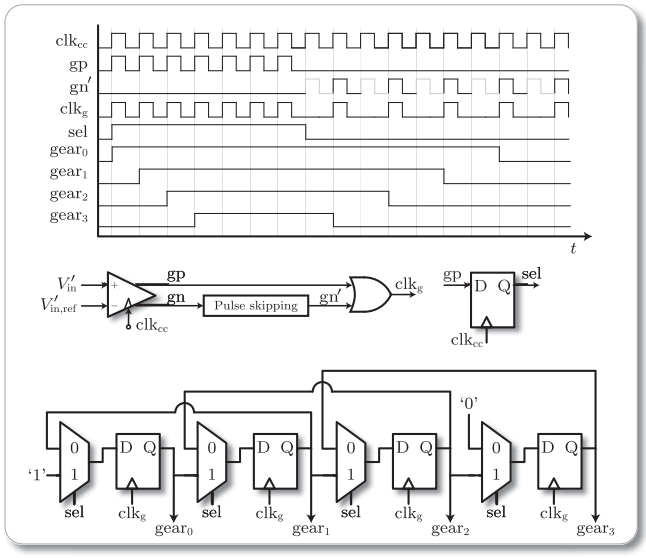

Fig. 12. Implementation of the digital gear controller that dynamically changes the configuration (gear) of the reconfigurable SC converter when an input voltage droop is detected by the clocked comparator. Shown for $N = 4$ ($b = 2$) interleaved phases.

A. Digital Gear Controller

Fig. 12 shows the circuit schematic of the digital gear controller for an example four-phase reconfigurable SCVR. The input voltage V'_{in} is compared with the reference $V'_{\mathrm{in,ref}}$ by a clocked comparator having both positive (gp) and negative (gn)

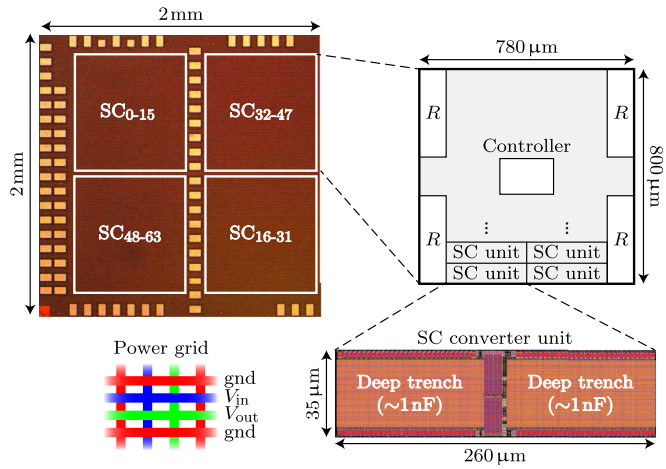

Fig. 13. Chip photo and SC converter unit layout of the 10 W SCVR implemented in the 32 nm technology with deep trench capacitors. The total converter consists of four 16-phase SC converter instances, where R denotes the programmable load, which is distributed among the converter instances. The total active converter area is $1.968\,\mathrm{mm}^2$.

outputs, where V'_{in} and $V'_{\mathrm{in,ref}}$ denote scaled voltages of V_{in} and $V_{\mathrm{in,ref}}$, respectively, to avoid overvoltage situations. For interleaved designs, which are considered here, it is found from simulations that changing the configuration of all converter units simultaneously leads to unnecessarily high ripples at the output node.[5] Therefore, the digital gear controller, which implements the feedforward control, is designed to change the configuration one at a time. The clocked comparator is clocked with the same high-frequency clock $\mathrm{clk_{cc}}$ used in the single-bound hysteretic control.

The gear signals are governed by a bidirectional shift register, where the direction is controlled by the select signal (sel). When $V'_{\mathrm{in}} < V'_{\mathrm{in,ref}}$, a rising edge of gp appears triggering sel to go high, and logic 1 is stored in the first flip-flop, causing gear_0 to go high. Consecutive gp triggers cause the following gear signals to go high, and when all gear signals are high, subsequent gp triggers have no further impact since logic 1 is stored in all flip-flops, i.e., all gear signals are high and all SC converter units operate in the 3:2 configuration. Once $V'_{\mathrm{in}} > V'_{\mathrm{in,ref}}$, a rising edge on gn appears triggering sel to go low, and logic 0 is stored in the last flip-flop, causing gear_3 to go low. Again, consecutive gn triggers cause the gear signals to go low one at a time. From simulations, it is found that pulse skipping of gn (denoted gn′) leads to the smoothest transition back to the original conversion ratio. The pulse skipping in gn′ is shown in Fig. 12 by the gray tone of every second pulse. However, pulse skipping of more pulses is possible as well.

VII. Experimental Results

The chip photo of the implemented 64-phase interleaved SC converter is shown in Fig. 13. The input power is supplied in the middle pad row (in P_{in}/gnd pairs) for a symmetrical power delivery to the chip. Having even more input power and gnd pads and using, e.g., flip chip bonding would further improve the input power delivery to the chip. The double pad rows to the left are for Kelvin probing at various locations on the chip. The layout of the 2:1 and 3:2 reconfigurable SC converter unit is reused from the design presented in [5] and [6] with only minor modifications. A regular top-level power grid consisting of V_{in}, V_{out}, and gnd covers the entire active chip area to minimize power grid resistance and inductance. For the SC converter units, the deep trench capacitors take up 72.1%, the transistors 27.3%, and the gate driver 0.6% of the total converter area.

The entire converter consists of a total of 192 SC converter units. Estimations based on extracted simulations at $850\,\mathrm{mV}$ output voltage result in a total capacitance of the entire converter of $384\,\mathrm{nF}$. Furthermore, the total on-state resistance of each of the power transistors is estimated to $6\,\mathrm{m\Omega}$ (paralleling 192 transistors each having $1.2\,\Omega$ on-state resistance), and the total equivalent series resistance of each deep trench capacitor is estimated to $3\,\mathrm{m\Omega}$ (paralleling 192 capacitors each having $0.6\,\Omega$ on-state resistance).

No additional output decoupling capacitance ($C_{\mathrm{out}} = 0\,\mathrm{nF}$) is implemented due to the high number of interleaved stages. For this design with a resistive load, minimal capacitance from the load is expected. However, in general, the capacitance of the microprocessor load will add to the total decoupling.

A. Thermal Model

Before proceeding with the efficiency and power density measurements, a thermal model is developed to predict the on-chip temperature during operation. Since the on-chip resistance of the programmable load array cannot be measured when the converter is operating, the resistance is measured under "cold" conditions, i.e., when the converter is not operating. For low-power implementations, the temperature increase is typically small and can in most cases be neglected. However, for high-power implementations, the temperature dependence of the on-chip resistance must be taken into account when estimating the converter's output power, from which efficiency and power density are estimated.

A thermal model of the SCVR design is setup using the 3D-ICE thermal model simulator [56], [57]. The entire chiplet measures $3\,\mathrm{mm} \times 3\,\mathrm{mm}$, but the SCVR design takes up $2\,\mathrm{mm} \times 2\,\mathrm{mm}$ and is placed in the lower left corner of the chiplet. Using the 3D-ICE simulator, the expected heat flux per region is mapped to the model. Since the load is integrated on the same chip as the converter, the entire input power is dissipated on the chip. A converter efficiency of 90% is assumed. Using the floor plan in Fig. 13, this means that 90% of the input power is uniformly dissipated in load regions and 10% of the input power is uniformly dissipated in the converter regions. The model furthermore assumes the silicon die thickness to be $780\,\mu\mathrm{m}$, and a thermal interface material of thickness $20\,\mu\mathrm{m}$ glues the chiplet to a $5\,\mathrm{mm} \times 5\,\mathrm{mm}$ water-cooled copper cold plate, where the cooling water is kept at a temperature of $27^{\circ}\mathrm{C}$ using a chiller.

[5]However, the higher ripples when changing all configurations simultaneously are not necessarily an issue regarding $V_{\mathrm{out,min}}$ of the microprocessor core. It is only a less smooth transition.

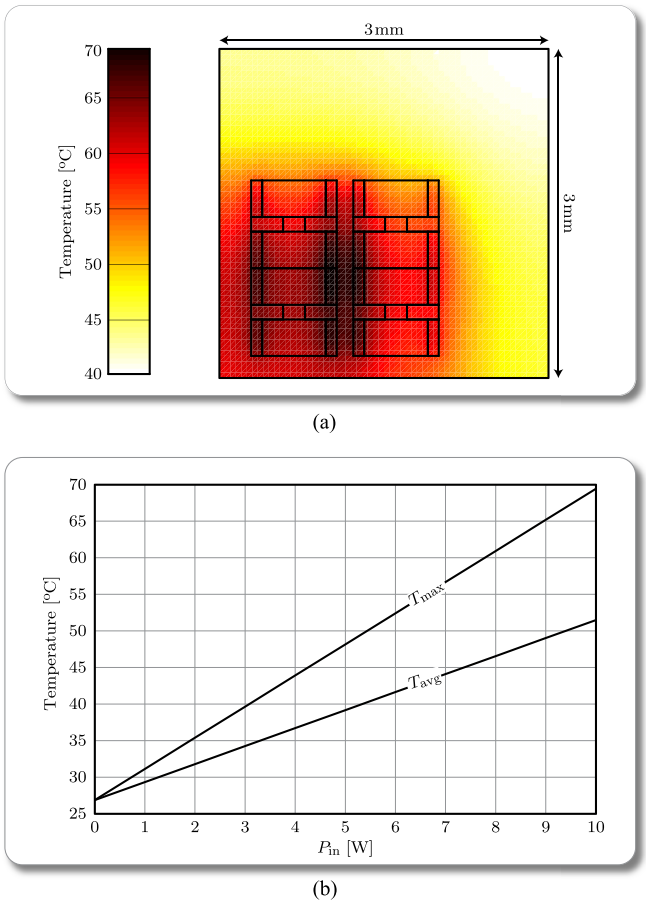

Fig. 14. Thermal model simulation results using 3D-ICE. (a) Simulated heat map shown for $P_{\text{in}} = 10\,\text{W}$. (b) Maximum and average temperatures as a function of the input power.

The results of the 3D-ICE simulations are shown in Fig. 14. The simulated heat map for an input power of 10 W is shown in Fig. 14(a), where the color coding resembles the maximum temperature on the chip. As can be seen, the maximum on-chip temperature for 10 W is 69°C. Using the same simulation setup for other power levels, Fig. 14(b) shows the maximum and average die temperatures as a function of the input power. As seen, the cooling setup manages to keep the on-chip temperature below 70°C, which is considered appropriate for the measurement setup.

The resistance of the on-chip programmable load is measured using a four-point measurement setup. Using the chiller to heat the water flowing through the cold plate to a predefined temperature, the chip is heated to approximately the same temperature which enables the characterization of the on-chip load resistance over temperature. Fig. 15 shows measured load resistances over temperature for 15 of the 31 resistors active and all of the resistors active. A close to linear increase in measured on-chip resistance over temperature is observed. Similar measurements are carried out for all resistance levels provided by the on-chip programmable load, and the measured temperature-correlated load resistances are used to determine the converter's output power, efficiency, and power density discussed next.

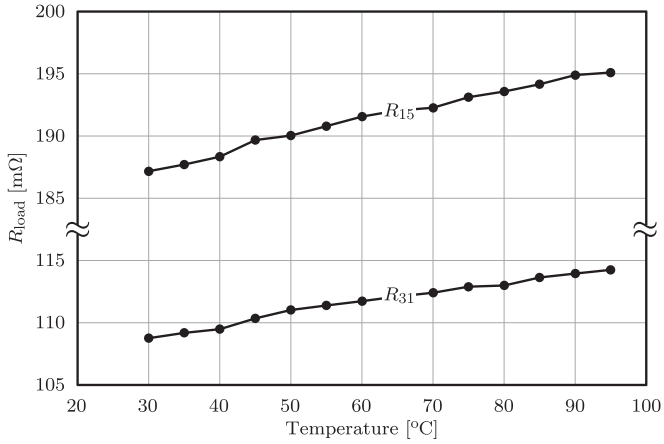

Fig. 15. Measured on-chip load resistance over temperature. Although shown for 15 of the 31 resistors active (R_{15}) and all resistors active (R_{31}), similar measurements are carried out for all other load levels provided by the on-chip programmable load.

B. Measured Efficiency and Power Density

Measurements are carried out using GBB PicoProbe needles on the unpackaged chip die mounted on a probe station. The input and output voltages are measured using Kelvin contacts to account for the voltage drops of cable and contact resistances. An Agilent E3633A power supply is used as input supply. The input power P_{in} is estimated using the current displayed on the input supply and the measured Kelvin input voltage. However, P_{in} does not include the power consumption of the digital controller as it is not possible to separate that power consumption from the total digital power consumption, which includes several housekeeping functions for testing that are not part of the digital controller. The output power P_{out} is measured using the Kelvin output voltage and the on-chip resistance, which is measured using the thermal model discussed above. Hence, the measured efficiency taking temperature effects into account is

$$\eta(T_{\max}) = \frac{P_{\text{out}}(T_{\max})}{P_{\text{in}}} = \frac{V_{\text{out}}^2}{R_{\text{load}}(T_{\max}(P_{\text{in}}))}\frac{1}{P_{\text{in}}} \qquad (4)$$

where $T_{\max}$ is the maximum operating temperature in Celsius, and $R_{\text{load}}(T_{\max}(P_{\text{in}}))$ is the measured load resistance evaluated at the maximum die temperature from Fig. 15. The maximum die temperature $T_{\max}(P_{\text{in}})$ is determined from the measured input power using the thermal model results in Fig. 14(b).

The measured efficiency over output power and power density for four different output voltages at $V_{\text{in}} = 1.8\,\text{V}$ is shown in Fig. 16. The efficiency at nominal load in the 2:1 configuration for $V_{\text{out}} = 0.85\,\text{V}$ is 83% at $1.9\,\text{W/mm}^2$ power density. For the same load with $V_{\text{out}} = 1.1\,\text{V}$, the efficiency is 85% at $3.2\,\text{W/mm}^2$ power density. Finally, the 10 W output power is achieved at 84% efficiency and $5\,\text{W/mm}^2$ power density for $V_{\text{out}} = 1100\,\text{mV}$.

Also shown in gray-scale in Fig. 16 are the efficiency and output power measurement results when disregarding the influence of the temperature on the load resistance, i.e., when $R_{\text{load}}(T_{\max} = 30°\text{C})$ in (4). As can be seen, both the efficiency and the output power are overestimated when disregarding the

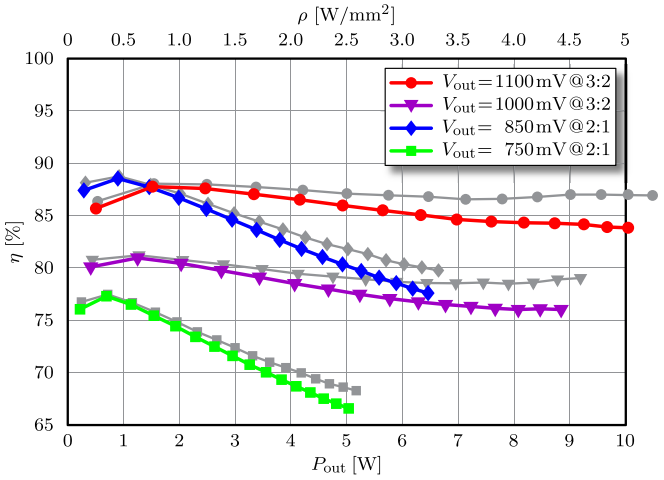

Fig. 16. Measured efficiency over output power and power density for four different output voltages at $V_{\rm in} = 1.8\,{\rm V}$. The maximum output power is 10 W at $V_{\rm out} = 1100\,{\rm mV}$ in the 3:2 configuration. The gray-scale results are when disregarding the influence of the temperature on the load resistance when estimating the converter output power and efficiency.

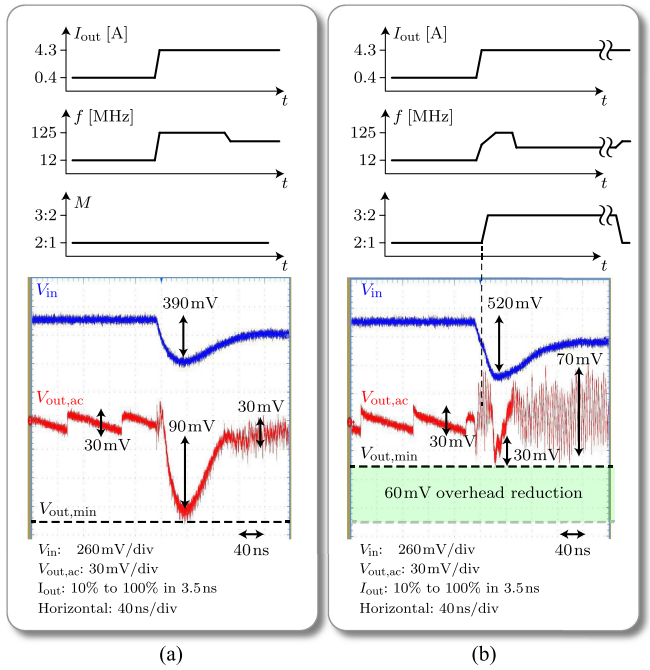

Fig. 17. Measured transient responses for (a) without and (b) with the novel feedforward control shown for $V_{\rm out} = 850\,{\rm mV}$ and $V_{\rm in} = 1.8\,{\rm V}$. The feedforward control effectively reduces the voltage overhead required to maintain a certain $V_{\rm out,min}$. The output voltage droops have been obtained without external decoupling capacitance, i.e., $C_{\rm out} = 0\,{\rm nF}$.

temperature effects, especially for high output powers. The gray-scale points highlight the importance of a thermal model when using on-chip resistive loads, since the error in efficiency due to thermal properties of the load resistances can have a large impact on the measured efficiency points.

C. Measured Transient Response

Transient responses are measured using a 20 GHz, 50 GS/s Textronix DSA72004 oscilloscope. The Kelvin contacts are probed using 40 GHz needles from GGB Industries, Inc., and 30 GHz Sucoflex cables are used to connect the probes to the oscilloscope.

The controller is clocked at $f_{\rm cc} = 4\,{\rm GHz}$ to achieve the sub nanosecond response time. The measured transient responses are shown in Fig. 17. The load change from 0.4 A to 4.3 A load current within 5 ns. The 4.3 A load current is chosen since it applies for the entire output voltage range, whereas the 10 W maximum output power only applies for $V_{\rm out} = 1.1\,{\rm V}$. The switching frequency of each SC converter unit changes from around 12 MHz at 0.4 A load current to a maximum of 125 MHz according to (2).

Without the feedforward control as shown in Fig. 17(a), the sub nanosecond feedback control maintains the output voltage for a short duration following the transient event. However, the significant droop of the input voltage causes the output node to experience a large droop, which leads to a relatively low $V_{\rm out,min}$. These results are in agreement with previously presented transient responses from [5]. It should be noted that the clock frequency $f_{\rm cc}$ can be reduced during light-load conditions to improve light-load efficiency.

With the feedforward control as shown in Fig. 17(b), the reconfigurable SC converter dynamically changes from the 2:1 to the 3:2 configuration when the input voltage droop is detected. As observed, the resulting output voltage droop is significantly reduced, leading to an improved $V_{\rm out,min}$. For this design, the voltage overhead is reduced by 60 mV, which can be used to reduce the steady-state output voltage and still comply with $V_{\rm out,min}$ requirements. In this particular case for $V_{\rm out} = 850\,{\rm mV}$, the voltage overhead reduction leads to $60\,{\rm mV}/850\,{\rm mV} = 7\%$ reduction in overall power consumption of the microprocessor.

As furthermore seen in Fig. 17(b), the input voltage droop is worsened by the feedforward control. However, from an application point of view, ensuring the output voltage droop to always be above $V_{\rm out,min}$ is all that matters. Furthermore, the larger ripple after the transient event is a result of the converter being in the 3:2 configuration, which, for that configuration, is a relatively low output voltage operation with higher ripple. Besides, for a digital load, the combined gate capacitance of all logic gates not switching will act as additional decoupling in a practical application. From an application point of view, this ripple is not considered to be critical since digital loads such as microprocessor cores are inherently insensitive to supply noise. The ripple can be reduced by adding additional output decoupling capacitance, where the amount of capacitance depends on the specific ripple requirements of the microprocessor. Although not shown, the converter transitions back to the more efficient 2:1 configuration once the transient has settled completely using the pulse skipping scheme of the digital gear controller discussed in Section VI. In conclusion, the feedforward control is an enabler for per-core DVFS with improved $V_{\rm out,min}$, which, as shown in Fig. 3(c), has the potential to save significant amounts of compute energy in future multicore and manycore microprocessor systems.

TABLE II
DETAILED COMPARISON WITH PRIOR ART

Design	**This work**	Le [22] ISSCC 2013	El-Damak [16] ISSCC 2013	Jain [53] VLSI 2013	Meyvaert [25] ECCE 2011	Andersen [5] ISSCC 2014
Technology	**32 nm SOI**	65 nm Bulk	32 nm SOI	22 nm Tri-Gate	90 nm Bulk	32 nm SOI
Conversion ratios (M)	**2:1, 3:2**	5:2, 3:1	3:1, 2:1 3:2, 1:1	2:1, 3:2 5:4, 1:1	2:1	2:1, 3:2
Capacitor technology	**Deep trench**	MOS	Ferroelectric	MIM	MOS	Deep trench
Interleaving	**64**	18	4	8	21	16
C_{fly} / C_{out}	**1 nF / 0**	3.88 nF / 0	1 nF / 10 nF	– / 100 pF	0.57 nF / 0	1 nF / 0
V_{in}	**1.8 V**	3 V – 4 V	1.5 V	1.23 V	2.35 V – 2.6 V	1.8 V
V_{out}	**0.7 V – 1.1 V**	1 V	0.4 V – 1.1 V	0.45 V – 1 V	1.03 V	0.7 V – 1.1 V
$P_{\text{out,max}}$	**10 W**	121 mW	1.1 mW	88 mW	1.65 W	840 mW
t_{response}	**< 1 ns**	< 1 ns	< 1 ms	3 ns – 5 ns	< 15 µs	< 1 ns
V_{droop}	**30 mV**	76 mV	–	25 mV	95 mV	94 mV
$V_{\text{ripple,pp}}$	**30 mV**	–	–	43 mV	–	30 mV
η_{max} per M	**82.7%, 85.1%**	71.5%, 73%	90%, 91% 93%, 80%	82%, 71% 73%, 68%	69%	86%, 90%
ρ [W/mm^2] @ η_{max}	**1.9, 3.2**	0.19, 0.19	0.0006, 0.0010 0.0013, 0.0016	0.062, 0.100 0.126, 0.243	0.42	2.2, 3.7

It should be noted that the input droop can be reduced by incorporating sufficient input decoupling capacitance. However, from simulations, the amount of input decoupling required would be impractically large, at least the size of the entire converter. The size penalty is amplified by the fact that the deep trench capacitor can withstand around 1.5 V; hence, two capacitors in series are required to support the 1.8 V input voltage, thereby giving little capacitance for the area. However, input decoupling capacitance within the microprocessor package would also improve the response. The amount required varies from package to package and depends on the layout and resulting parasitic package inductances obtainable.

D. Comparison With Prior Art

This SCVR design is compared in detail with prior art in Table II. The selected designs resemble state of the art achieved with bulk [22], [25], silicon on insulator (SOI) [5], [16], and trigate [53] semiconductor technologies. The reported efficiency and a high power density benchmarks for the various capacitor technologies MOS, MIM, ferroelectric, and deep trench capacitor are compared. Finally, transient responses are compared and the voltage droops, where reported, can be held up against the implemented output decoupling capacitance C_{out}, if any. As a result of the limited designs published, the prior art varies significantly in SC converter topology, semiconductor technology, capacitor technology, and target application.

Having no on-chip decoupling capacitors, the 30 mV voltage droop achieved in this work is comparable with the design of [53] that uses precious chip area for on-chip decoupling capacitors. For the designs in [5], [22], and [25] without on-chip decoupling capacitors, the droop achieved in this work is reduced from more than 76 mV reported in [22] to 30 mV in this design.

The availability of the deep trench capacitor results in an SCVR design that, like [5], achieves both a high efficiency and a high power density. In contrast, [16] achieves high efficiency but low power density using ferroelectric capacitors, and [22], [25] achieve high power density but low efficiency using MOS capacitors. The 10 W maximum output power of this design is more than six times higher than any published prior art.

VIII. CONCLUSION

A novel feedforward control for reconfigurable SC converters is presented. The feedforward control dynamically changes the configuration of the SC converter to a higher voltage conversion ratio when an input voltage droop is detected. The feedforward control reduces the output voltage droop from 90 mV to 30 mV, thereby improving $V_{\text{out,min}}$ by 60 mV without the use of dedicated input or output decoupling capacitors. The feedforward control for reconfigurable SC converters reduces the voltage overhead required to meet microprocessor $V_{\text{out,min}}$ requirements. For 850 mV output voltage, the 60 mV voltage overhead reduction can lead to a 7% reduction in overall microprocessor power consumption.

To account for the change of on-chip resistance with temperature, a thermal model is developed to predict the on-chip temperature. Correlating the measured on-chip load resistances with the operating temperature allows for a more accurate efficiency estimation at high output powers. Measurement results of the on-chip SCVR achieve: 1) maximum efficiencies above 85%; 2) power densities above 2.5 W/mm^2; 3) transient responses faster than 1 ns with reduced $V_{\text{out,min}}$ overhead, and 4) output powers up to 10 W. Hence, the feasibility of high-power on-chip SCVR designs is demonstrated experimentally by achieving 10 W maximum output power.

This paper further concludes that on-chip SCVRs, which historically have been perceived as being inefficient, low power, and difficult to regulate, are a viable candidate to enable granular microprocessor power delivery and per-core regulation. The measured performance of the presented converter ranks among

the highest efficiency, highest power density, highest output power, and fastest transient response time on-chip voltage regulators published to date.

REFERENCES

[1] M. P. Mills. (2013, Aug.). The cloud begins with coal. [Online]. Available: www.tech-pundit.com

[2] S. R. Sanders, E. Alon, H.-P. Le, M. D. Seeman, M. John, and V. W. Ng, "The road to fully integrated DC-DC conversion via the switched-capacitor approach," *IEEE Trans. Power Electron.*, vol. 28, no. 9, pp. 4146–4155, Sep. 2013.

[3] T. M. Andersen, F. Krismer, J. W. Kolar, T. Toifl, C. Menolfi, L. Kull, T. Morf, M. Kossel, M. Brändli, P. Buchmann, and P. A. Francese, "A $4.6\,\mathrm{W/mm^2}$ power density 86% efficiency on-chip switched capacitor DC-DC converter in $32\,\mathrm{nm}$ SOI CMOS," in *Proc. IEEE Appl. Power Electron. Conf. Expo.*, Long Beach, CA, USA, Mar. 2013, pp. 692–699.

[4] G. Villar-Piqué, H. Bergveld, and E. Alarcon, "Survey and benchmark of fully integrated switching power converters: Switched-capacitor versus inductive approach," *IEEE Trans. Power Electron.*, vol. 28, no. 9, pp. 4156–4167, Sep 2013.

[5] T. M. Andersen, F. Krismer, J. W. Kolar, T. Toifl, C. Menolfi, L. Kull, T. Morf, M. Kossel, M. Brändli, P. Buchmann, and P. A. Francese, "A sub-ns response on-chip switched-capacitor DC-DC voltage regulator delivering $3.7\,\mathrm{W/mm^2}$ at 90% efficiency using deep-trench capacitors in $32\,\mathrm{nm}$ SOI CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, USA, Feb. 2014, pp. 90–91.

[6] T. M. Andersen, F. Krismer, J. W. Kolar, T. Toifl, C. Menolfi, L. Kull, T. Morf, M. Kossel, M. Brändli, P. Buchmann, and P. A. Francese, "A deep trench capacitor based 2:1 and 3:2 reconfigurable on-chip switched capacitor DC-DC converter in $32\,\mathrm{nm}$ SOI CMOS," in *Proc. IEEE Appl. Power Electron. Conf. Expo.*, Fort Worth, TX, USA, Mar. 2014, pp. 1448–1455.

[7] T. M. Andersen, "On-chip switched capacitor voltage regulators for granular microprocessor power delivery," Ph.D. dissertation, Swiss Federal Inst. Technol., Zurich, Switzerland, 2015.

[8] M. D. Seeman, "A design methodology for switched-capacitor DC-DC converters," Ph.D. dissertation, Univ. California, Berkeley, CA, USA, 2009.

[9] S. Abedinpour, B. Bakkaloglu, and S. Kiaei, "A multistage interleaved synchronous buck converter with integrated output filter in $0.18\,\mu\mathrm{m}$ SiGe process," *IEEE Trans. Power Electron.*, vol. 22, no. 6, pp. 2164–2175, Nov. 2007.

[10] H. J. Bergveld, K. Nowak, R. Karadi, S. Iochem, J. Ferreira, S. Ledain, E. Pieraerts, and M. Pommier, "A 65-nm-CMOS $100\,\mathrm{MHz}$ 87%-efficient DC-DC down converter based on dual-die system-in-package integration," in *Proc. IEEE Energy Convers. Cong. Expo.*, San Jose, CA, USA, Sep. 2009, pp. 3698–3705.

[11] T. M. Van Breussegem and M. S. J. Steyaert, "Monolithic capacitive DC-DC converter with single boundary–multiphase control and voltage domain stacking in $90\,\mathrm{nm}$ CMOS," *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1715–1727, Jul. 2011.

[12] E. A. Burton, G. Schrom, F. Paillet, J. Douglas, W. J. Lambert, K. Radhakrishnan, and M. J. Hill, "FIVR—Fully integrated voltage regulators on 4^{th} generation Intel core SoCs," in *Proc. IEEE Appl. Power Electron. Conf. Expo.*, Fort Worth, TX, USA, Mar. 2014, pp. 432–439.

[13] L. Chang, R. K. Montoye, B. L. Ji, A. J. Weger, K. G. Stawiasz, and R. H. Dennard, "A fully-integrated switched-capacitor 2:1 voltage converter with regulation capability and 90% efficiency at $2.3\,\mathrm{A/mm^2}$," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, USA, Jun. 2010, pp. 55–56.

[14] C. Chia, C.-h. Chang, P.-S. Lei, and H.-M. Chen, "A two-phase fully-integrated DC-DC converter with self-adaptive DCM control and GIPD passive components," *IEEE Trans. Power Electron.*, vol. 30, no. 6, pp. 3252–3261, Jun. 2015.

[15] T. J. DiBene, "$400\,\mathrm{A}$ fully integrated silicon voltage regulator with in-die magnetically coupled embedded inductors," presented at the IEEE Appl. Power Electron. Conf. Expo., Palm Springs, CA, USA, Feb. 2010.

[16] D. El-Damak, S. Bandyopadhyay, and A. P. Chandrakasan, "A 93% efficiency reconfigurable switched-capacitor DC-DC converter using on-chip ferroelectric capacitors," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, USA, Feb. 2013, pp. 374–375.

[17] W. Godycki, B. Sun, and A. Apsel, "Part-time resonant switching for light load efficiency improvement of a 3-level fully integrated buck converter," in *Proc. IEEE Eur. Solid State Circuits Conf.*, Venice, Italy, Sep. 2014, pp. 163–166.

[18] Z. Hayashi, Y. Katayama, M. Edo, and H. Nishio, "High-efficiency DC-DC converter chip size module with integrated soft ferrite," *IEEE Trans. Magn.*, vol. 39, no. 5, pp. 3068–3072, Sep. 2003.

[19] K. Ishida, K. Takemura, K. Baba, M. Takamiya, and T. Sakurai, "3D stacked buck converter with $15\,\mu\mathrm{m}$ thick spiral inductor on silicon interposer for fine-grain power-supply voltage control in SiP's," in *Proc. IEEE 3D Syst. Integr. Conf.*, Munich, Germanic, Nov. 2010, pp. 1–4.

[20] R. Jain, B. M. Geuskens, S. T. Kim, M. M. Khellah, J. Kulkarni, J. W. Tschanz, and V. De, "A 0.45-$1\,\mathrm{V}$ fully-integrated distributed switched capacitor DC-DC converter with high density MIM capacitor in $22\,\mathrm{nm}$ tri-gate CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 4, pp. 917–927, Apr. 2014.

[21] S. S. Kudva and R. Harjani, "Fully-integrated on-chip DC-DC converter with a 450X output range," *IEEE J. Solid-State Circuits*, vol. 46, no. 8, pp. 1940–1951, Aug. 2011.

[22] H.-P. Le, J. Crossley, S. R. Sanders, and E. Alon, "A sub-ns response fully integrated battery-connected switched-capacitor voltage regulator delivering $0.19\,\mathrm{W/mm^2}$ at 73% efficiency," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, USA, Feb. 2013, pp. 372–373.

[23] H.-P. Le, S. R. 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, Sep. 2011.

[24] M. Lee, Y. Choi, and J. Kim, "A $0.76\,\mathrm{W/mm^2}$ on-chip fully-integrated buck converter with negatively-coupled, stacked-LC filter in $65\,\mathrm{nm}$ CMOS," in *Proc. IEEE Energy Convers. Cong. Expo.*, Pittsburgh, PA, USA, Sep. 2014, pp. 2208–2212.

[25] H. Meyvaert, T. Van Breussegem, and M. Steyaert, "A $1.65\,\mathrm{W}$ fully integrated $90\,\mathrm{nm}$ bulk CMOS intrinsic charge recycling capacitive DC-DC converter: Design & techniques for high power density," in *Proc. IEEE Energy Convers. Cong. Expo.*, Phoenix, AZ, USA, Sep. 2011, pp. 3234–3241.

[26] S. Musunuri and P. L. Chapman, "Design of low power monolithic DC-DC buck converter with integrated inductor," in *Proc. IEEE Power Electron. Spec. Conf.*, Recife, Brazil, Jun. 2005, pp. 1773–1779.

[27] J. Ni, Z. Hong, and B. Y. Liu, "Improved on-chip components for integrated DC-DC converters in $0.13\,\mu\mathrm{m}$ CMOS," in *Proc. ESSCIRC*, Athens, Greece, Sep. 2009, pp. 448–451.

[28] K. Onizuka, K. Inagaki, H. Kawaguchi, M. Takamiya, and T. Sakurai, "Stacked-chip implementation of on-chip buck converter for distributed power supply system in SiPs," *IEEE J. Solid-State Circuits*, vol. 42, no. 11, pp. 2404–2410, Nov. 2007.

[29] G. Villar-Piqué, "A 41-phase switched-capacitor power converter with $3.8\,\mathrm{mV}$ output ripple and 81% efficiency in baseline $90\,\mathrm{nm}$ CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, USA, Feb. 2012, pp. 98–100.

[30] Y. Ramadass, A. Fayed, B. Haroun, and A. Chandrakasan, "A $0.16\,\mathrm{mm^2}$ completely on-chip switched-capacitor DC-DC converter using digital capacitance modulation for LDO replacement in $45\,\mathrm{nm}$ CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, USA, Feb. 2010, pp. 208–209.

[31] Y. K. Ramadass and A. P. Chandrakasan, "Voltage scalable switched capacitor DC-DC converter for ultra-low-power on-chip applications," in *Proc. IEEE Power Electron. Spec. Conf.*, Orlando, FL, USA, Jun. 2007, pp. 2353–2359.

[32] N. Sturcken, M. Petracca, S. Warren, P. Mantovani, L. Carloni, A. Peterchev, and K. L. Shepard, "A switched-inductor integrated voltage regulator with nonlinear feedback and network-on-chip load in $45\,\mathrm{nm}$ SOI," *IEEE J. Solid-State Circuits*, vol. 47, no. 8, pp. 1935–1945, Aug. 2012.

[33] N. Sturcken, E. O'Sullivan, N. Wang, P. Herget, B. Webb, L. Romankiw, M. Petracca, R. Davies, R. Fontana, G. Decad, I. Kymissis, A. Peterchev, L. Carloni, W. Gallagher, and K. Shepard, "A 2.5D integrated voltage regulator using coupled-magnetic-core inductors on silicon interposer," *IEEE J. Solid-State Circuits*, vol. 48, no. 1, pp. 244–254, Jan. 2013.

[34] J. Sun, D. Giuliano, S. Devarajan, J.-Q. Lu, T. P. Chow, and R. J. Gutmann, "Fully monolithic cellular buck converter design for 3-D power delivery," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 17, no. 3, pp. 447–451, Mar. 2009.

[35] T. Tong, X. Zhang, W. Kim, D. Brooks, and G.-Y. Wei, "A fully integrated battery-connected switched-capacitor 4:1 voltage regulator with 70% peak efficiency using bottom-plate charge recycling," in *Proc. IEEE Custom Integr. Circuits Conf.*, San Jose, CA, USA, Sep. 2013, pp. 1–4.

[36] G. Villar and E. Alarcon, "Monolithic integration of a 3-level DCM-operated low-floating-capacitor buck converter for DC-DC step-down

donversion in standard CMOS," in *Proc. IEEE Power Electron. Spec. Conf.*, Rhodes, Greece, Jun. 2008, pp. 4229–4235.
[37] A. K. P. Viraj and G. A. J. Amaratunga, "A monolithic CMOS 5V/1V switched capacitor DC-DC step-down converter," in *Proc. IEEE Power Electron. Spec. Conf.*, Orlando, FL, USA, Jun. 2007, pp. 2510–2514.
[38] N. Wang, J. Hannon, R. Foley, K. McCarthy, T. O'Donnell, K. Rodgers, F. Waldron, and S. C. Ó Mathúna, "Integrated magnetics on silicon for power supply in package (PSiP) and power supply on chip (PwrSoC)," in *Proc. IEEE Electron. Syst. Integr. Technol. Conf.*, Berlin, Germany, Sep. 2010, pp. 1–6.
[39] M. Wens and M. Steyaert, "A fully-integrated $0.18\,\mu\mathrm{m}$ CMOS DC-DC step-down converter, using a bondwire spiral inductor," in *Proc. IEEE Custom Integr. Circuits Conf.*, San Jose, CA, USA, Sep. 2008, pp. 17–20.
[40] M. Wens and M. Steyaert, "A fully-integrated $130\,\mathrm{nm}$ CMOS DC-DC step-down converter, regulated by a constant on/off-time control system," in *Proc. IEEE Eur. Solid State Circuits Conf.*, Edinburgh, U.K., Sep. 2008, pp. 62–65.
[41] M. Wens and M. Steyaert, "A fully integrated CMOS $800\,\mathrm{mW}$ four-phase semiconstant ON/OFF-time step-down converter," *IEEE Trans. Power Electron.*, vol. 26, no. 2, pp. 326–333, Feb. 2011.
[42] J. Wibben and R. Harjani, "A high-efficiency DC-DC converter using $2\,\mathrm{nH}$ integrated inductors," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 844–854, Apr. 2008.
[43] W. Kim, D. Brooks, and 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, Jan. 2012.
[44] J. Kwong, Y. Ramadass, N. Verma, M. Koesler, K. Huber, H. Moormann, and A. Chandrakasan, "A $65\,\mathrm{nm}$ sub-V_t microcontroller with integrated SRAM and switched-capacitor DC-DC converter," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, USA, Feb. 2008, pp. 318–616.
[45] A. Paul, D. Jiao, S. Sapatnekar, and C. H. Kim, "Deep trench capacitor based step-up and step-down DC/DC converters in 32 nm SOI with opportunistic current borrowing and fast DVFS capabilities," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Singapore, Jun. 2013, pp. 49–52.
[46] Y. K. Ramadass and A. P. Chandrakasan, "Voltage scalable switched capacitor DC-DC converter for ultra-low-power on-chip applications," in *Proc. IEEE Power Electron. Spec. Conf.*, Orlando, FL, USA, Jun. 2007, pp. 2353–2359.
[47] T. M. Andersen, F. Krismer, J. W. Kolar, T. Toifl, C. Menolfi, L. Kull, T. Morf, M. Kossel, M. Brändli, and P. A. Francese, "A feedforward controlled on-chip switched-capacitor voltage regulator delivering $10\,\mathrm{W}$ in $32\,\mathrm{nm}$ SOI CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, USA, Feb. 2015, pp. 1–3.
[48] W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," in *Proc. IEEE Symp. High Perform. Comput. Archit.*, Salt Lake City, UT, USA, Feb. 2008, pp. 123–134.
[49] L. Chang, D. J. Frank, R. K. Montoye, S. J. Koester, B. L. Ji, P. W. Coteus, R. H. Dennard, and W. Haensch, "Practical strategies for power-efficient computing technologies," *Proc. IEEE*, vol. 98, no. 2, pp. 215–236, Feb. 2010.
[50] P. Zhou, D. Jiao, C. H. Kim, and S. S. Sapatnekar, "Exploration of on-chip switched-capacitor DC-DC converter for multicore processors using a distributed power delivery network," in *Proc. IEEE Custom Integr. Circuits Conf.*, San Jose, CA, Sep. 2011, pp. 1–4.
[51] (2013). International technology roadmap for semiconductors. [Online]. Available: www.itrs.net
[52] T. V. Breussegem and M. Steyaert, "A 82% efficiency 0.5% ripple 16-phase fully integrated capacitive voltage doubler," in *Proc. IEEE Symp. VLSI Circuits*, Kyoto, Japan, Jun. 2009, pp. 198–199.
[53] R. Jain, B. Geuskens, M. Khellah, S. Kim, J. Kulkarni, J. Tschanz, and V. De, "A 0.45-$1\,\mathrm{V}$ fully integrated reconfigurable switched capacitor step-down DC-DC converter with high density MIM capacitor in $22\,\mathrm{nm}$ tri-gate CMOS," in *Proc. IEEE Symp. VLSI Circuits*, Kyoto, Japan, Jun. 2013, pp. 174–175.
[54] R. Jain and S. Sanders, "A $200\,\mathrm{mA}$ switched capacitor voltage regulator on $32\,\mathrm{nm}$ CMOS and regulation schemes to enable DVFS," in *Proc. the IEEE Conf. Power Electron. Appl.*, Birmingham, U.K., Aug. 2011, pp. 1–10.
[55] L. Kull, T. Toifl, M. Schmatz, P. A. Francese, C. Menolfi, M. Brändli, M. Kossel, T. Morf, T. M. Andersen, and Y. Leblebici, "A 3.1 mW 8b 1.2 GS/s single-channel asynchronous SAR ADC with alternate comparators for enhanced speed in $32\,\mathrm{nm}$ digital SOI CMOS," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3049–3058, Dec. 2013.
[56] A. Sridhar, A. Vincenzi, D. Atienza, and T. Brunschwiler, "3D-ICE: A compact thermal model for early-stage design of liquid-cooled ICs," *IEEE Trans. Comput.*, vol. 63, no. 10, pp. 2576–2589, Oct. 2014.
[57] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunschwiler, and D. Atienza, "3D-ICE: Fast compact transient thermal modeling for 3D ICs with inter-tier liquid cooling," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design*, San Jose, CA, USA, Nov. 2010, pp. 463–470.

Toke M. Andersen (S'10) received the B.Sc. and M.Sc. degrees from the Technical University of Denmark, Kongens Lyngby, Denmark, in 2008 and 2010, respectively, and the Ph.D. degree from ETH Zurich, Switzerland, Zurich, in 2015, where he worked on on-chip power conversion for microprocessor power delivery in collaboration IBM Research Zurich.

He is co-Founder and Senior R&D Engineer at Nordic Power Converters, where he develops innovative VHF power supplies for LED lighting and charger applications.

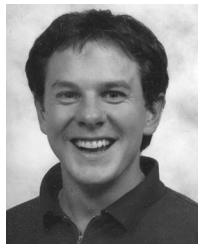

Florian Krismer (S'05–M'12) received the M.Sc. degree from the University of Technology Vienna, Vienna, Austria, in 2004, and the Ph.D. degree from the Power Electronic Systems Laboratory (PES), Swiss Federal Institute of Technology (ETH) Zürich, Zürich, Switzerland, in 2011.

He is currently a Postdoctoral Fellow at PES, ETH Zürich. His research interests include the analysis, design, and optimization of high-current and high-frequency power converters.

Johann W. Kolar (S'89–M'91–SM'04–F'10) received the M.Sc. and Ph.D. degrees (*summa cum laude / promotio sub auspiciis praesidentis rei publicae*) from the University of Technology Vienna, Austria.

Since 1984, he has been working as an independent international consultant in close collaboration with the University of Technology Vienna, in the fields of power electronics, industrial electronics and high performance drives. He has proposed numerous novel PWM converter topologies, and modulation and control concepts and has supervised over 60 Ph.D. students. He has published over 650 scientific papers in international journals and conference proceedings, 3 book chapters, and has filed more than 120 patents. The focus of his current research is on ultra-compact and ultra-efficient SiC and GaN converter systems, wireless power transfer, Solid-State Transformers, Power Supplies on Chip, and ultra-high speed and bearingless motors.

Dr. Kolar has received 21 IEEE TRANSACTIONS and Conference Prize Paper Awards, the 2014 IEEE Middlebrook Award, and the ETH Zurich Golden Owl Award for excellence in teaching. He initiated and/or is the founder of four ETH Spin-off companies. He is a Member of the Steering Committees of several leading international conferences in the field and has served from 2001 through 2013 as an Associate Editor of the IEEE TRANSACTIONS ON POWER ELECTRONICS. Since 2002 he also is an Associate Editor of the *Journal of Power Electronics* of the Korean Institute of Power Electronics and a Member of the Editorial Advisory Board of the *IEEJ Transactions on Electrical and Electronic Engineering.*

Thomas Toifl (S'97–M'99–SM'09) received the Dipl.-Ing. (M.S.) degree and the Ph.D. degree (with highest honors) from the Vienna University of Technology, Vienna, Austria, in 1995 and 1999, respectively.

In 1996, he joined the Microelectronics Group of the European Research Center for Particle Physics (CERN), Geneva, Switzerland, where he developed radiation-hard circuits for detector synchronization and data transmission, which were integrated in the four particle detector systems of the new large hadron collider. In 2001, he joined IBM Research Zurich, Rüschlikon, Switzerland, where he has been working on multigigabit per second, low-power communication circuits in advanced CMOS technologies. In that area, he authored or coauthored 14 patents and numerous technical publications. Since July 2008, he has been managing the I/O Link technology group at IBM Research Zurich.

Dr. Toifl received the Beatrice Winner Award for Editorial Excellence at the 2005 IEEE International Solid-State Circuits Conference.

Christian Menolfi (S'97–M'99) received the Dipl. Ing. and Ph.D. degrees in electrical engineering from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, in 1993 and 2000, respectively.

From 1993 to 2000, he was with the Integrated Systems Laboratory, ETH Zurich, as a Research Assistant, where he worked on highly sensitive CMOS VLSI data-acquisition circuits for silicon-based microsensors. Since 2000, he has been with IBM Research Zurich, Rüschlikon, Switzerland, where he is involved in the design of multigigabit low-power communication circuits in advanced CMOS technologies.

Lukas Kull (S'10–M'14) received the M.Sc. degree in electrical engineering from the Swiss Federal Institute of Technology, Zurich, Switzerland, in 2007, and the Ph.D. degree from the Swiss Federal Institute of Technology, Lausanne, Lausanne, Switzerland, in 2014.

He joined IBM Research Zurich, Rüschlikon, Switzerland, in 2010. He has authored or coauthored more than ten patents and 20 technical publications. His research interests include analog circuit design for high-speed low-power analog-to-digital converters, analog circuit design, IR, and THz imaging.

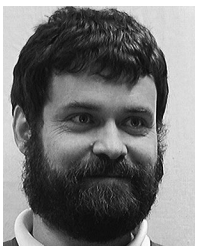

Thomas Morf (S'89–M'96–SM'09) received the B.S. degree from the Zurich University of Applied Science, Zurich, Switzerland, the M.S. degree from the University of California at Santa Barbara, Santa Barbara, CA, USA, and the Ph.D. degree from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland.

From 1996 to 1999, he led a research group in the area of InP-HBT circuit design and technology also at ETH. In 1999, he joined IBM Research Zurich, Rüschlikon, Switzerland. His current research interests include electrostatic-discharge circuit protection, electrical and optical high-speed high-density interconnects, and THz antennas and detectors. He has coauthored more than 100 papers and is a coinventor of more than 20 issued patents.

Marcel Kossel (S'99–M'02–SM'09) received the Dipl.Ing. and Ph.D. degrees in electrical engineering from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, in 1997 and 2000, respectively.

He joined IBM Research Zurich, Rüschlikon, Switzerland, in 2001, where he is involved in analog circuit design for high-speed serial links. His research interests include analog circuit design and RF measurement techniques. He has also worked in the field of microwave tagging systems and radio-frequency identification systems.

Matthias Brändli received the Dipl.Ing. (M.Sc.) degree in electrical engineering from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, in 1997.

From 1998 to 2001, he was with the Integrated Systems Laboratory, Swiss Federal Institute of Technology, working on deep-submicrometer technology VLSI design challenges, digital video image processing for biomedical applications, and testability of CMOS circuits. In 2001, he joined the Microelectronics Design Center, ETH Zurich, where he was involved in numerous digital and mixed-signal application-specific integrated circuit design projects, worked on electronic design automation, and contributed to teaching. In 2008, he joined IBM Research Zurich, Rüschlikon, Switzerland, where he works on multigigabit low-power communication circuits in advanced CMOS technologies.

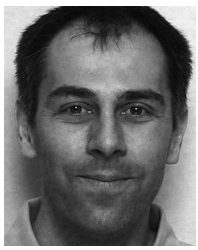

Pier Andrea Francese (M'01) received the Degree in electrical engineering (*cum laude*) from the Politecnico di Milano, Milano, Italy, and the Ph.D. degree from the Swiss Federal Institute of Technology, Zurich, Switzerland, in 1993 and 2005, respectively.

He worked in the field of IC product development for Teradyne, Philips Semiconductors, and National Semiconductor. In June 2010, he joined IBM Research Zurich, Rüschlikon, Switzerland, where he is currently designing circuits for high-speed I/O links in advanced CMOS technologies.