Chapter 8
Fully Integrated Switched-Capacitor DC-DC Conversion

Elad Alon, Hanh-Phuc Le, John Crossley, and Seth R. Sanders

Abstract This chapter describes techniques to maximize the achievable efficiency and power density of fully integrated switched-capacitor (SC) DC-DC converters. Circuit design methods to support multiple topologies (and hence output voltages) are described. These techniques are verified by a proof-of-concept converter prototype implemented in 0.374 mm^2 of a 32 nm SOI process. The 32-phase interleaved converter can be configured into three topologies to support output voltages of 0.5–1.2 V from a 2 V input supply, and achieves ~80 % efficiency at an output power density of 0.86 W/mm^2.

8.1 Introduction

As parallelism is now the dominant mechanism by which integrated circuit designers improve the computing performance of their chips while remaining within strict power budgets, there is increasing need and potential benefit to utilizing an independent power supply for each processing core. Simply adding off-chip supplies not only incurs significant degradation of supply impedance due to e.g. split package power planes, but also additional cost due to increased motherboard size and package complexity. Therefore, there is strong motivation to fully integrate voltage conversion on the die, as shown in Fig. 8.1.

Although on-die DC-DC converters are currently almost always implemented as linear regulators, achieving high efficiency across a broad range of output voltages necessitates the use of switching converters. Inductor-based switching converters are dominant in off-chip converters, and recent efforts to co-package and reduce the inductor size [1, 2] have brought them closer to complete integration. However, fully integrated DC-DC converters based on CMOS inductors either require costly

E. Alon (✉) • H.-P. Le • J. Crossley • S.R. Sanders
University of California, Berkeley, USA
e-mail: elad@eecs.berkeley.edu

A. Baschirotto et al. (eds.), *Frequency References, Power Management for SoC, and Smart Wireless Interfaces: Advances in Analog Circuit Design 2013*,
DOI 10.1007/978-3-319-01080-9_8,

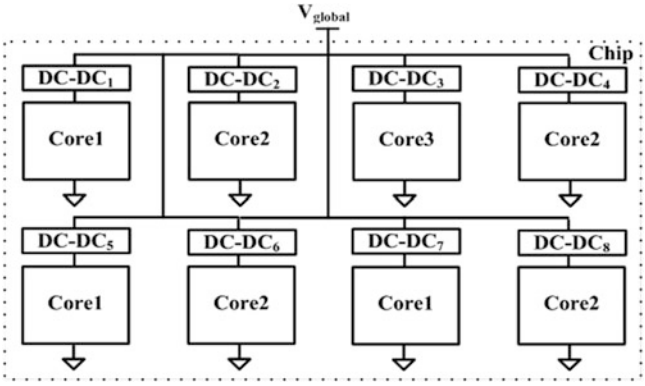

Fig. 8.1 Chip with multiple, local, on-die DC-DC converters

extra fabrication steps [3, 4] (e.g., thick metals or integrated magnetics), or suffer from the high series resistance (and hence low efficiency [5]) of standard on-die inductors. Integrated capacitors on the other hand can achieve low series resistance and high capacitance density, and most importantly, can be used to implement DC-DC converters in completely standard CMOS processes.

Given these advantages, fully integrated switched-capacitor (SC) converters have recently received significant attention from multiple researchers. For example, [6] and [7] both investigated multiphase interleaving for fully integrated SC voltage doublers, with [6] demonstrating high efficiency (82 %) but low power density (0.67 mW/mm^2), and [7] achieving high power density (1.123 W/mm^2), but low efficiency (60 %). The need for high efficiency is self-evident, but high power density is also critical since it sets the area overhead of the converter relative to the on-chip circuitry it is supplying power to.

In order to explore the boundaries of their capabilities, in this chapter we describe a methodology to achieve the optimal tradeoff between efficiency and power density for fully integrated SC converters. Section 8.2 therefore presents an analysis and optimization of SC converter losses as a function of power density, and discusses the use of topology reconfiguration and output impedance control to enable wide output voltage range. Section 8.3 then describes a converter prototype with reconfigurable topology. Measurement results from the prototype converter verifying the predicted performance and proposed techniques are presented in Sect. 8.4, and the chapter is finally concluded in Sect. 8.5.

8.2 SC Converter Analysis and Optimization

To achieve the optimal tradeoff between power density and efficiency, in this section we will analyze the operation and loss mechanisms of SC converters. This analysis will lead to design equations for switching frequency and switch width that

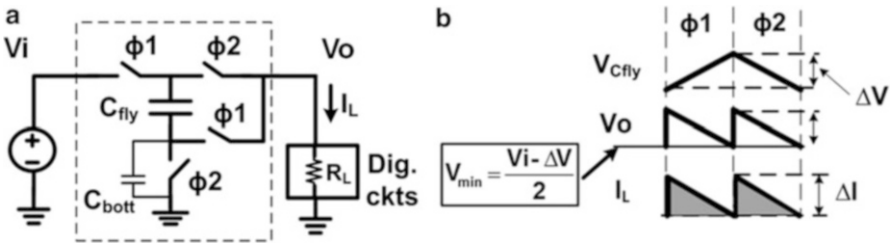

Fig. 8.2 (**a**) 2:1 step-down SC DC-DC converter and (**b**) its operating waveforms

minimize losses in a given technology and power density. Since a single-topology SC converter is only efficient when generating an output voltage within a limited range, this section also describes a simple design strategy for enabling reconfigurable topologies as well as predicting the overall efficiency vs. output voltage.

8.2.1 Operation of a Sample SC Converter

In order to highlight the key loss mechanisms that will set the tradeoff between converter efficiency and area (i.e., power density), we will begin by examining the operation of the 2:1 step-down converter shown in Fig. 8.2a.

Switched-capacitor DC-DC converters typically operate in two phases, each of which ideally has 50 % duty cycle. While it is possible to operate SC DC-DC converters at a fixed switching frequency and use variable duty cycle to adjust the output impedance [8, 9], maximum efficiency can only be achieved by optimizing switching frequency and operating with 50 % duty cycle.

Under this condition (shown in Fig. 8.2), during phase ϕ1, the flying capacitor C_{fly} is connected between the input node Vi and the output node Vo. The charge drawn from Vi though C_{fly} charges up this capacitor and flows to the load. In phase ϕ2, C_{fly} is connected between Vo and GND, and thus the charge previously stored on the flying capacitor is transferred to the output. Since the switching cycle is often much smaller than the charge/discharge time constant (which is set by $R_L C_{fly}$), the ramp rate of the voltage across the capacitor is relatively constant, and hence the load can be treated as a current source.

As will be described later, in order to maximize efficiency it is desirable to utilize all available capacitance within the converter itself. Therefore, we will assume that there is no explicit output filtering capacitor, which in the case of the simple SC converter described so far, makes the peak-to-peak voltage ripple across the capacitor and the converter's output equal, as shown in Fig. 8.2b. This voltage ripple has a direct implication on the loss – and hence the achievable efficiency – of the converter.

8.2.2 Loss Analysis

The voltage ripple across the capacitors scales with the load current, and will therefore act as a form of series loss similar to the switch conduction losses. In addition, any SC converter will also have shunt losses that are independent of the load current, including gate and bottom plate capacitor switching losses. Note that the control circuitry for an SC converter will also contribute to shunt loss, but will be neglected here since this loss can usually be made relatively small. These losses can be modeled as shown in Fig. 8.6, where the series losses are represented by an equivalent output resistance R_o [10, 11], the shunt losses by the parallel resistor R_p, and the transformer represents the ideal voltage conversion ratio.

In order to show the relationship between voltage ripple across the capacitor and loss, we should recall that most fully integrated switched-capacitor converters will be delivering power to synchronous digital circuitry. The performance of synchronous digital systems is determined by the operating frequency, which in turn is set by the minimum average voltage over a clock period. Since the clock period of most digital circuits will be short in comparison to the converter's switching period, the performance of these circuits is typically simply set by the minimum voltage V_{min} of the supply [12]. In this case, the efficiency of the converter should be calculated relative to the power that would have been consumed by the load if it was constantly operating at exactly V_{min} [12]. In other words, the ideal power consumed by the load is:

$$P_{L-min} = V_{min} I_L \tag{8.1}$$

where $I_L = \frac{V_{min}}{R_L}$. However, due to the voltage ripple from the converter, and assuming that this ripple is relatively small compared to the nominal voltage, the average power dissipated by the load is approximately:

$$P_{L-tot} \approx \left(V_{min} + \frac{\Delta V}{2}\right)\left(I_L + \frac{\Delta I}{2}\right) \tag{8.2}$$

where ΔV is the output voltage ripple (due to the operation of the converter) and $\Delta I = \Delta V / R_L$.

Although P_{L-tot} is indeed dissipated by the load, any power consumed beyond P_{L-min} should be counted as loss since this extra power does not contribute to an increase in performance. In order to quantify this loss, we need to calculate V_{min} and ΔV; as shown in Fig. 8.2b, for the 2:1 converter considered here, V_{min} is lower than the ideal output voltage $Vi/2$ by $\Delta V/2$:

$$V_{min} = \frac{Vi}{2} - \frac{\Delta V}{2}, \tag{8.3}$$

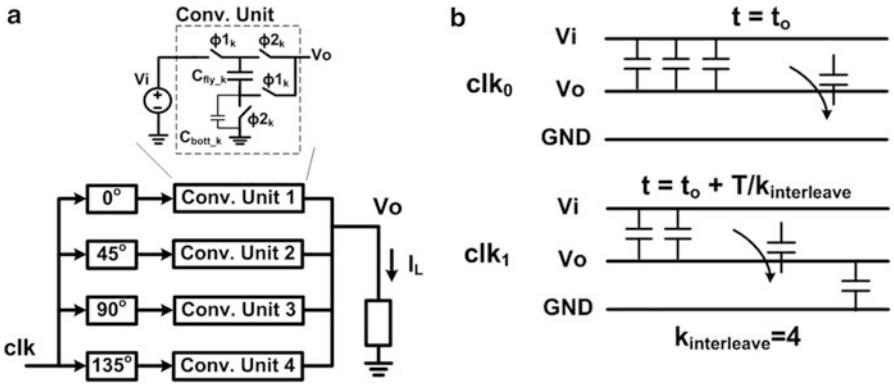

Fig. 8.3 (**a**) Sample four-phase interleaved SC converter and (**b**) operation of its flying capacitors

and the voltage ripple ΔV is set by:

$$\Delta V = \frac{I_L}{C_{fly}} . \frac{T}{2} = \frac{I_L}{2C_{fly} f_{sw}}, \tag{8.4}$$

where T is the switching period and $f_{sw} = 1/T$ is the switching frequency.

As should be clear from Eq. 8.2, the loss caused by the operation of the converter is due to both the voltage ripple ΔV as well as due to the excess current flowing in the load ΔI. The loss due to the voltage ripple ΔV is unavoidable because the voltage drop $\Delta V/2$ in Eq. 8.3 is inherent to the fact that charge (power) is being delivered through a capacitor. However, the current ripple ΔI can be eliminated if the ripple in the output voltage above V_{min} is minimized.

Fortunately, the ripple in the output voltage and hence the load current ripple can be reduced by multiphase interleaving. As described in [6], [7], and [13], multiphase interleaving is implemented by partitioning the converter into sub-units and switching each one of these units on a different clock phase. Figure 8.3 depicts a sample four-phase interleaved design and the operation of the flying capacitors in clock phase 0 (clk_0) and clock phase 1 (clk_1). Each unit in this converter uses 1/4 of the total capacitance and a clock that is 45° phase-shifted from its neighbor.

The total charge (per switching cycle) required by the output is the same as that in the converter without interleaving, but is equally divided among each unit. Thus, the charge flowing through each unit flying capacitor in the interleaved design is the same as it would be in the original design. As illustrated in Fig. 8.4a, the voltage ripple on each unit capacitor required to deliver that charge is therefore essentially identical to the previous ΔV from Eq. 8.4. As a result, V_{min} is unchanged. However, because the charge delivered to the output is divided more finely, the output voltage and current ripple are reduced by the interleaving factor ($k_{interleave} = 4$), as shown in Fig. 8.4b. This leads to a reduction in the loss associated with the current ripple:

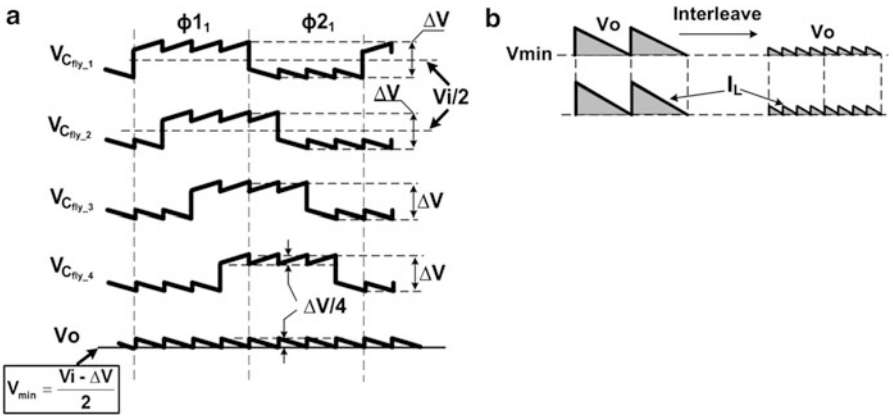

Fig. 8.4 (**a**) Flying capacitor voltages and (**b**) their effect on output voltage and current ripple of the four-phase interleaved converter in Fig. 8.3

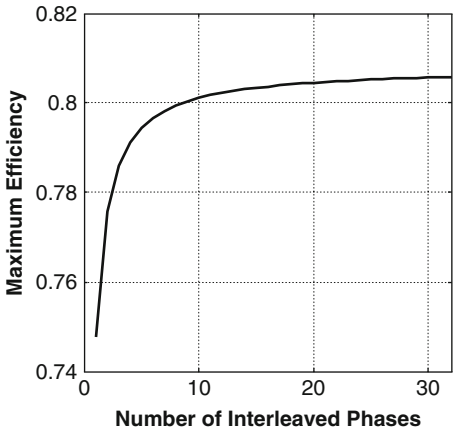

Fig. 8.5 Efficiency of an example 2:1 SC converter as a function of $k_{interleave}$

$$P_{L-tot} \approx \left(V_{min} + \frac{\Delta V}{2}\right)\left(I_L + \frac{1}{k_{interieave}} \cdot \frac{\Delta I}{2}\right) \tag{8.5}$$

As shown in Fig. 8.5, interleaving the converter by roughly by a factor of 10 (which is relatively simple in an integrated design) is sufficient to essentially eliminate the efficiency penalty due to load current ripple. In other words, extreme levels of interleaving are generally not necessary – especially if they would result in significant control overhead. Assuming sufficient interleaving (i.e., $k_{interleave} > \sim 10$) we can generally ignore the loss caused by the current ripple, resulting in the classic SC loss [10] given by:

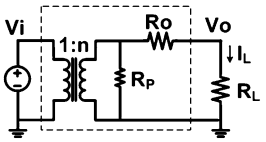

Fig. 8.6 Simplified SC converter model for calculation of losses

$$P_{C_{fly}} = I_L . \frac{\Delta V}{2} = \frac{I_L^2}{M_{cap} C_{fly} f_{sw}} \tag{8.6}$$

where M_{cap} is a constant related to the converter's output resistance and is determined by the converter's topology (e.g., $M_{cap} = 4$ for a 2:1 SC converter).

Beyond the intrinsic SC loss, the finite conductance of the switches leads to another series loss term. To simplify the equations, we will assume here that all of the switches have identical characteristics (regardless of type or gate overdrive), but it is straightforward to extend the analysis to handle differences between each switch. The switch conductance loss $P_{R_{sw}}$ is therefore set by:

$$P_{R_{sw}} = I_L^2 \frac{R_{on}}{W_{sw}} M_{sw} \tag{8.7}$$

where R_{on} is the switch resistance density measured in Ω·m, W_{sw} (m) is the total width of all switches, and M_{sw} is a constant determined by the converter's topology. For the 2:1 converter in Fig. 8.2, there are four switches, and each occupies 1/4 of the total switch area. During each half of a switching period, two of the four switches conduct the current flowing into to the output, resulting in:

$$\begin{aligned} M_{sw} &= N_{switches,tot} \times \left(\frac{T_{ph1}}{T} \times N_{sw,on,ph1} + \frac{T_{ph2}}{T} \times N_{sw,on,ph2} \right) \\ &= 4 \times \left(\frac{1}{2} \times 2 + \frac{1}{2} \times 2 \right) = 8 \end{aligned} \tag{8.8}$$

As shown in Eqs. 8.6 and 8.7, the intrinsic switched-capacitor loss and switch conductance loss are both set by the load current, and can hence be modeled by the equivalent output resistance R_o in Fig. 8.6. The total series loss is therefore approximately set by:

$$P_s = I_L^2 R_o = P_{R_{sw}} + P_{C_{fly}} \tag{8.9}$$

The other key portion of an SC converter's losses stems from shunt losses due to switching the parasitic capacitance of the flying capacitors and power switches. Any flying capacitor – particularly fully integrated ones – will have parasitic capacitance associated with both its top plate and its bottom plate. In steady-state operation, both of these plates experience approximately equal voltage swings.

Therefore, we will group both losses caused by the top-plate capacitor $C_{top\ -\ plate}$ and the bottom-plate capacitor $C_{bottom\ -\ plate}$ into one parasitic capacitor switching loss $P_{bott\ -\ cap}$, given by:

$$P_{bott-cap} = M_{bott} V_o^2 C_{bott} f_{sw} \tag{8.10}$$

where M_{bott} is a constant determined by the converter's topology (e.g., $M_{bott} = 1$ for a 2:1 SC converter) and $C_{bott} = C_{bottom-plate} + C_{top-plate}$. For simplicity, once again assuming that all of the switches have identical characteristics, the gate parasitic capacitance switching loss $P_{gate\ -\ cap}$ is given by:

$$P_{gate-cap} = V_{sw}^2 W_{sw} C_{gate} f_{sw} \tag{8.11}$$

where V_{sw} is the gate voltage swing and C_{gate} is the gate capacitance density (F/m) of the switches.

8.2.3 Loss Optimization

In order for the converter to achieve the highest overall efficiency at a given power density we must minimize the total loss, which is set by the combination of the four previously discussed components:

$$\begin{aligned} P_{loss} &= \left(P_{C_{fly}} + P_{R_{sw}}\right) + \left(P_{bott-cap} + P_{gate-cap}\right) \\ &= \left(\frac{I_L^2}{M_{cap} C_{fly} f_{sw}} + I_L^2 \frac{R_{on}}{W_{sw}} M_{sw}\right) + \left(M_{bott} V_o^2 C_{bott} f_{sw} + V_{sw}^2 W_{sw} C_{gate} f_{sw}\right) \end{aligned} \tag{8.12}$$

For a given technology, R_{on} and C_{gate} are set by the available transistors, and hence are essentially fixed. Similarly, the intrinsic switched-capacitor loss $P_{C_{fly}}$ (and hence the overall loss) will always be minimized by utilizing as large of a flying capacitor C_{fly} as possible given the chip area constraint. Therefore, at a given power density, the only two variables that can be optimized to minimize the total losses are switch width W_{sw} and switching frequency f_{sw}.

Increasing either switch width or switching frequency decreases the series losses at the cost of increasing the shunt loss. Minimizing the converter's total losses therefore boils down to setting the values of W_{sw} and f_{sw} to balance series and shunt losses. As we will detailed next, the power density required of the converter plays an important role in determining the most dominant loss components, and hence how W_{sw} and f_{sw} should be set to minimize loss.

At high power densities (i.e., large I_L or equivalently small R_L, where $R_L = V_o/I_L$), W_{sw} and f_{sw} must both increase with the load current in order to suppress the series losses. Since $P_{gate\ -\ cap}$ is proportional to both width and frequency while $P_{bott\ -\ cap}$ scales only with switching frequency, beyond a certain load current

the bottom plate loss becomes the least significant term. To arrive at simple analytical equations for the optimal f_{sw} and W_{sw} in this regime, we can thus ignore the bottom plate portion of the shunt losses. In this case, the optimal f_{sw} and W_{sw} will be:

$$f_{sw_{opt}} = \frac{1}{\sqrt[3]{M_{cap}^2 M_{sw}}} \cdot \sqrt[3]{\frac{V_o^2}{V_{sw}^2} \times \frac{1}{R_{on} C_{gate} \left(R_L C_{fly}\right)^2}} \tag{8.13}$$

$$W_{sw_{opt}} = \sqrt[3]{M_{sw}^2 M_{cap}} \cdot \sqrt[3]{\frac{V_o^2}{V_{sw}^2} \frac{R_{on}^2 C_{fly}}{R_L^2 C_{gate}}} \tag{8.14}$$

Under these conditions and with the optimal f_{sw} and W_{sw}, the minimum normalized loss (which sets the efficiency $\eta = (1 + P_{loss}/P_L)^{-1}$) is approximately:

$$\frac{P_{loss}}{P_L} = 3 \sqrt[3]{\frac{M_{sw}}{M_{cap}}} \cdot \sqrt[3]{\frac{V_{sw}^2}{V_o^2} \frac{R_{on} C_{gate}}{R_L C_{fly}}} \tag{8.15}$$

This relative loss expression highlights the tradeoff between power density and efficiency (Fig. 8.7). For a given technology and converter topology, increasing the power density by a factor of x at a given output voltage implies that R_L also decreases by a factor of x, leading to an increase in the minimum normalized loss by a factor of $\sqrt[3]{x}$.

This relative loss expression also highlights that the most important technology metric guiding the selection of the switches is the product of gate voltage swing squared and intrinsic time constant (i.e., ${V_{sw}}^2 R_{on} C_{gate}$). Similarly, since it is the ratio of this switch metric to the load voltage squared multiplied by the effective time constant for charging/discharging the flying capacitors (i.e., ${V_o}^2 R_L C_{fly}$), increasing the density of the capacitors also directly improves efficiency at a given power density.

Although the previous analysis provides a clear intuitive picture of the relationship between power density and efficiency, it is only accurate at high power densities where the loss due to switching the "bottom-plate" parasitics of the flying capacitors is negligible compared to the other losses. Both the optimal switching frequency and the switch area scale down as output power drops, and hence at low power densities the losses due to driving the switch parasitic capacitors become much smaller than the all of the other losses. Therefore, in this regime we can approximately find the optimum loss by ignoring the switch gate loss and finding the optimum switching frequency $f_{sw,opt}$:

$$f_{sw,opt} = \frac{1}{\sqrt{M_{cap} M_{bott} k_{bott}}} \frac{1}{C_{fly} R_L} \tag{8.16}$$

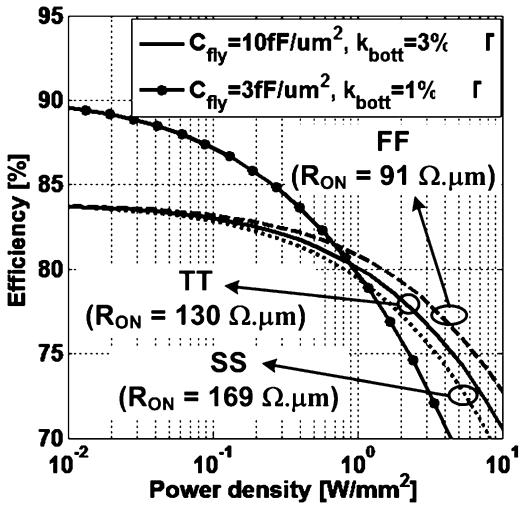

Fig. 8.7 Analytical predictions of optimized power density vs. efficiency for a 2:1 SC converter. The switch characteristics of a 32 nm CMOS technology (i.e., $R_{ON} = 130\ \Omega{\cdot}\mu m$, $C_{gate} = 3$ fF/μm, $V_{sw} = 1$ V) were used to generate these curves, which also highlight the impact of +/−30 % variations in R_{ON}

where $k_{bott} = C_{bott}/C_{fly}$ is the parasitic to flying capacitance ratio. Although the switch gate losses were assumed to be small, we can still size the switches to minimize the total loss in Eq. 8.12 with the frequency found in Eq. 8.16:

$$W_{sw} = \sqrt{\frac{V_o^2}{V_{sw}^2} \cdot \frac{R_{on} C_{fly}}{R_L C_{gats}}} \cdot \sqrt{M_{sw}^2 M_{cap} M_{bott} k_{bott}} \tag{8.17}$$

Combining the results from Eqs. 8.12, 8.16 and 8.17, the normalized loss in the low power density regime is:

$$\frac{P_{loss,opt}}{P_L} = 2\sqrt{\frac{M_{bott}}{M_{cap}} k_{bott}} + 2\sqrt{\frac{M_{sw}}{\sqrt{M_{cap} M_{bott}}} \cdot \frac{1}{\sqrt{k_{bott}}}} \cdot \sqrt{\frac{V_{sw}^2}{V_o^2} \cdot \frac{R_{on} C_{gate}}{R_L C_{fly}}} \tag{8.18}$$

This result highlights a key intrinsic limit on the efficiency of a switched-capacitor DC-DC converter. Even in very light load conditions (i.e., $R_L = \infty$), the maximum efficiency of the converter is limited by the bottom-plate capacitance ratio k_{bott} and the converter's topology – i.e., by the first term in Eq. 8.18. For example, with a bottom-plate capacitor ratio of 1 %, the efficiency of a 2:1 converter is limited to 90.9 %. Of course, any non-zero load will decrease the efficiency of the converter, but for sufficiently light loads the efficiency will still be dominated by bottom-plate losses.

To illustrate these effects, Fig. 8.7 shows the efficiency vs. power density curves of two optimized converter designs with different flying capacitor characteristics. One converter employs capacitors with high capacitance density but also a higher k_{bott} (e.g., MOS capacitor), while the other employs capacitors with lower density but also lower parasitics (e.g., a MIM or MOM capacitor). At high power densities

(where Eq. 8.15 accurately predicts the minimum loss), high capacitance density directly translates into higher efficiency. However, at low power densities (where Eq. 8.18 is more accurate) the flying capacitors should have as low parasitic capacitance as possible in order to maximize peak efficiency.

Beyond illustrating the importance of selecting an appropriate capacitor given the target power density, Fig. 8.7 also predicts that a 2:1 SC converter using currently available CMOS technology can achieve ~80 % efficiency at a power density of ~1 W/mm^2. While this performance is substantially better than previous predictions or demonstrations of fully integrated DC-DC converters [3, 5–7, 14], it is only achievable at a single output voltage.

8.2.4 Output Voltage Range Considerations

Unlike in inductor-based converters where charge is saved and transferred in the form of current in the inductors – which enables efficient control of the output voltage by modulating the DC voltage applied to one side of the inductor – SC converters save and transfer charge as a voltage on the flying capacitors. The output voltage of a SC converter is thus determined by its topology.

To efficiently achieve a wider output voltage range, SC converters require reconfigurable topologies that can support multiple conversion ratios [15, 16]. By using a given number of reconfigurable topologies, an SC converter can support the same number of discrete open-circuit voltage levels. Intermediate voltages between these discrete levels can then be obtained by controlling R_o, which is equivalent to linear regulation off of the open-circuit voltages.

As discussed in [10] and shown here in Eqs. 8.9 and 8.12, the converter's output impedance R_o can be adjusted by controlling one or the combination of switching frequency f_{sw} [16], switch sizing W_{sw}, and effective flying capacitance C_{fly} [17]. Figure 8.8 shows the resulting efficiency vs. output voltage for a converter operating off a 2 V input and allowing reconfiguration into one of three possible topologies with conversion ratios of 1/3, 1/2, and 2/3. Even with linear regulation performed only by adjusting f_{sw} (which is slightly sub-optimal), Fig. 8.8 predicts that such a converter could achieve above 70 % efficiency for most output voltages spanning from ~0.5 up to ~1.2 V.

8.3 SC Converter Circuit Design

In order to verify the previously described optimization strategy as well as the predicted performance, a prototype SC converter was designed and fabricated in a 32 nm SOI test-chip in collaboration with AMD [16]. Although proper selection of the flying capacitor, switch width, and switching frequency (as outlined in the previous section) are critical to achieving a converter with high efficiency and

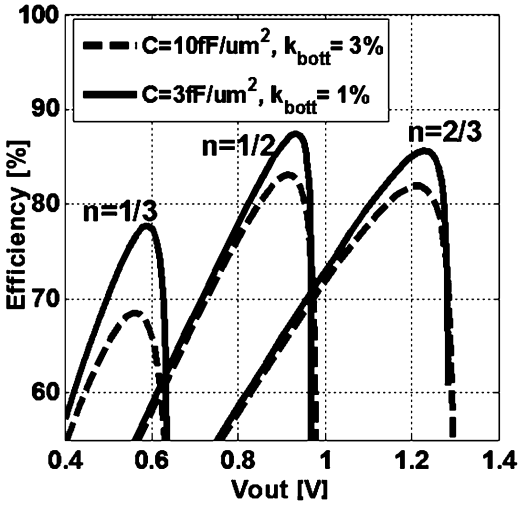

Fig. 8.8 Predicted efficiency vs. Vo with three reconfigurable topologies for two capacitor implementations. For both types of capacitors, the load is adjusted so that the converter is supplying 0.1 W/mm^2 with a Vo of 0.95 V. These curves assume that R_L varies along with the output voltage in the same manner as a CMOS ring oscillator

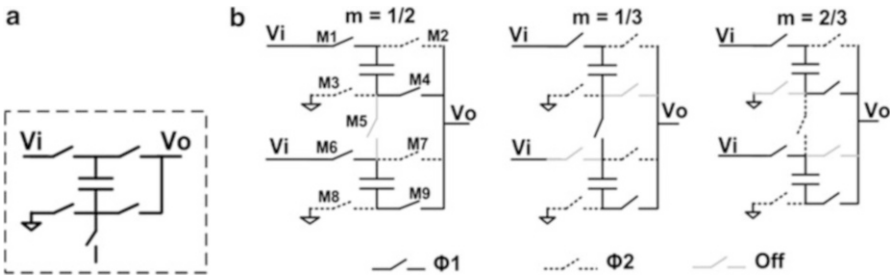

Fig. 8.9 (**a**) Standard cell and (**b**) reconfigurable converter unit

minimal area overhead, the need to support reconfigurable topologies (in this design, 2/3, 1/2, and 1/3) and multiple output voltage results in several circuit design challenges which must be overcome as well.

As with any custom designed VLSI structure, a physical design strategy that enables one to construct larger SC converters blocks by arraying identical sub-converter unit cells is highly desirable. In order to achieve this goal while supporting topology reconfiguration, we therefore propose to partition the converter into a unit cell consisting of one flying capacitor and five switches, as shown in Fig. 8.9a. Conceptually, each standard cell can be configured to operate in series or in parallel with the rest of the cells, leading to a simple physical design strategy that supports multiple conversion ratios. As shown in Fig. 8.9b, for this prototype

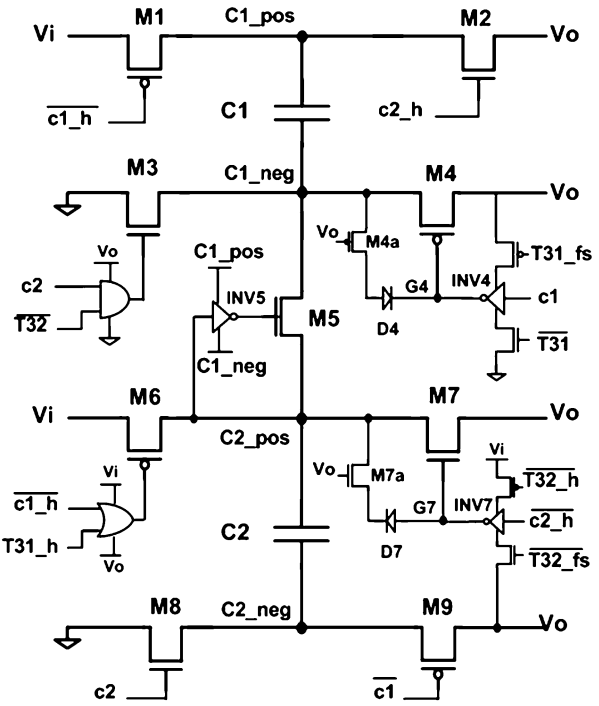

Fig. 8.10 Transistor-level implementation of the converter unit cell. The converter operates off of two non-overlapping clocks c1 and c2

converter we have grouped together two standard cells to form a converter unit supporting three topologies with conversion ratios of 1/3, 1/2, and 2/3 ($V_o = 0.66$, 1, or 1.33 V with a 2 V input). For simplicity, the intermediate voltage levels are generated by controlling f_{sw}. Figure 8.10 highlights the complete transistor-level implementation of this converter unit.

8.4 Experimental Verification

A die photo of the implemented SC converter employing the previously described design techniques is shown in Fig. 8.11. To maximize efficiency at high power densities and mitigate the current ripple losses, this design utilizes standard thin-oxide MOS transistors to implement C_{fly} as well as 32-way interleaving. This level of interleaving was chosen because even at high power densities, the converter's optimal switching frequency is relatively low compared to the intrinsic speed of the transistors.

Since this converter is intended to be co-integrated with the load, measuring the converter's performance requires a careful testing strategy. We will therefore first describe the load structure and its characterization, followed by measured results verifying the design methodology and proposed design techniques.

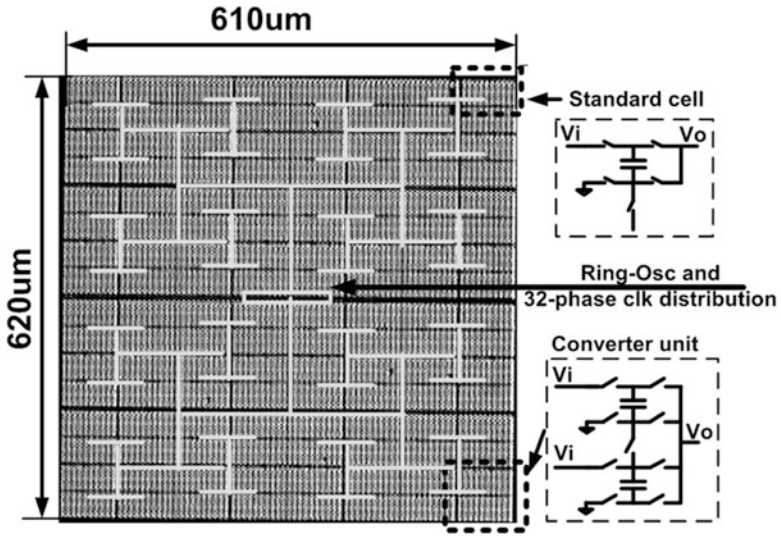

Fig. 8.11 Die photo of the 32 nm SOI SC converter prototype

8.4.1 Test Structure

In order to obtain correct I-V measurements of the on-die loading circuits and thus the efficiency of the converter, the on-die load – which was implemented with a variable-width PMOS device – must be pre-characterized. Four-wire sensing was used to measure the power consumption across the load in order to avoid inaccuracies due to drops in solder bumps, package pins, and PCB traces.

Load characterization was carried out by gating the clock of the converter (disabling it) and then driving the output node Vo of the converter from an off-chip power supply. For each load current (i.e., PMOS transistor width) setting, the voltage supply Vo is swept and the current consumption is measured. Utilizing this data, the power consumed by the load circuit while the converter is in its normal operation can then be extracted simply by measuring Vo.

8.4.2 Measurement Results and Discussion

Figure 8.12 shows the converter's measured efficiency and optimal switching frequencies in the 1/2 mode while supplying the on-die load circuits. For simplicity and in order to obtain optimal efficiency in this demonstration, the switching frequency was adjusted by externally controlling the supply of an on-chip ring oscillator. However, any one of a broad variety of techniques to control switching frequency [18] could be utilized.

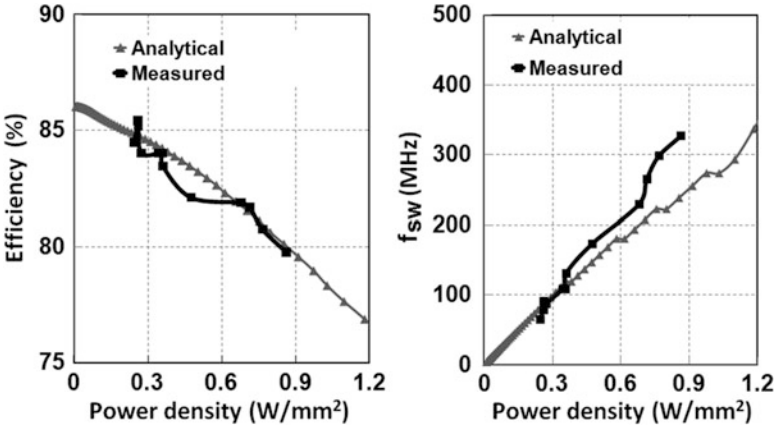

Fig. 8.12 Measured efficiency and optimal switching frequency versus power density in the 1/2 mode with Vi = 2 V and Vo ≈ 0.88 V

The measured converter achieved an efficiency of 79.76 % at 0.86 W/mm^2. The experimental data matches the analytical predictions to within 1.3 % across the range of measured power density (0.24–0.86 W/mm^2). Note that the performance quoted here is better than that reported in [16] due to the availability of new test-chips fabricated in a nearly production version of the process (rather than the developmental process used to obtain the original results).

Figure 8.13 shows the converter's efficiency vs. output voltage in the three operating modes, verifying that the converter functions correctly in all three of the reconfigurable topologies. The measurements in the 2/3 and 1/2 modes match very well with the analytical predictions. The measured efficiency in the 1/3 mode is however much lower. The cause of this discrepancy appears to be un-modeled leakage from Vi and Vo due to over-voltage-stress (~1.4 V) on switches M1, M2, M4, M6, M7 and M9 powered off of the Vi-Vo rails. Therefore, a practical implementation that uses the 1/3 conversion ratio would likely require a lower input voltage (~1.8 V), higher voltage devices, and/or switch cascoding. Despite this issue with the 1/3 conversion ratio, the two reconfigurable topologies enable the converter to maintain an efficiency of over 70 % for most of the output voltage range from 0.7 to ~1.15 V.

The converter's performance is summarized and compared with other work in Table 8.1. This prototype experimentally verifies that by following the design methodology and techniques proposed in this work, both boundaries in efficiency and power density of the previous works in [6] and [7] can be achieved with an implementation in a current commercial process. At 79.76 % efficiency and 0.86 W/mm^2, the proposed design could potentially be integrated into the same space as that already required for decoupling capacitors (as well as serve the same function) in a processor or SoC targeting mobile applications where the load operates at ~100 mW/mm^2. In fact, our recent work [20] demonstrated the efficient

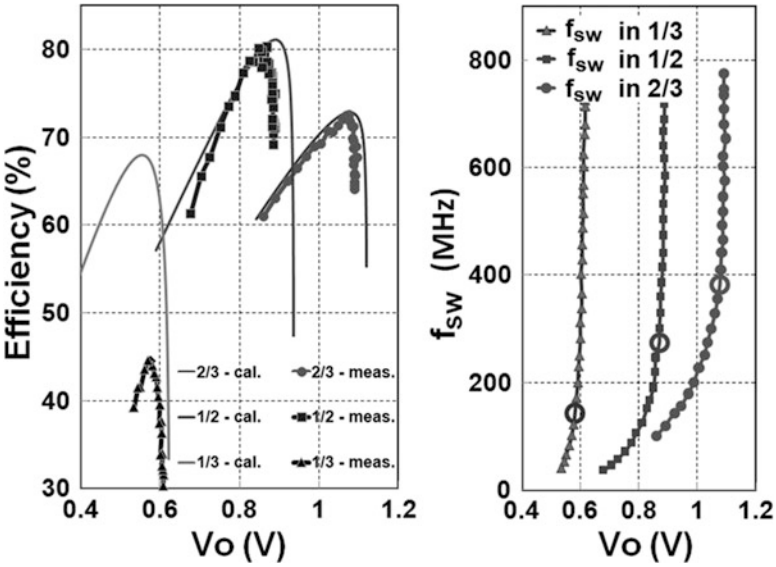

Fig. 8.13 Measured converter efficiency and switching frequency across Vo and topologies with Vi = 2 V and the load circuits set to $R_L \approx 0.9\ \Omega$ at Vo = 0.88 V

Table 8.1 Comparison of recently published fully integrated SC converters

Work	[6]	[7]	[19]	This work
Technology	130 nm bulk	32 nm bulk	45 nm SOI	32 nm SOI
Topology	2/1 step-up	2/1 step-up	1/2 step-down	2/3, 1/2, 1/3 step-down
Capacitor technology	MIM	Metal finger	Deep trench	CMOS oxide
Interleaved phases	16	32	1	32
C_{out}	400 pF (= C_{fly})	0	Yes	0
Converter area	2.25 mm^2	6,678 μm^2	1,200 μm^2	0.378 mm^2 (1.4 % used for load)
Quoted efficiency (η)	82 %	60 %	90 %	79.76 % (in 1/2 step-down)
Power density @ η	0.67 mW/mm^2	1.123 W/mm^2	2.185 W/mm^2	0.86 W/mm^2

integration of a converter interfacing directly to a Li-ion battery and generating a ~1 V output in a standard 65 nm CMOS process. This further opens up the opportunity for mobile SoC's to eliminate the need for external DC-DC converters completely.

In order to expand the applicability of SC converters to even higher performance processors operating at ~1 W/mm^2, the work reported in [19] utilizes ~200 fF/um^2 deep trench capacitors and achieves 90 % efficiency at a power density of 2.185 W/mm^2. This further experimentally verifies the benefit of high density

capacitors in increasing efficiency and power density – as also predicted in Eq. 8.16 and Fig. 8.7. In fact, the analysis from Sect. 8.2 predicts that with 200 fF/μm^2 deep trench capacitors and modern CMOS switches, an optimized SC design may achieve over 88 % efficiency for power densities up to 10 W/mm^2. Thus, the application of the techniques outlined in this chapter along with existing high-density capacitor technologies appears promising in enabling the broad adoption of fully integrated SC converters for on-die power distribution and management.

8.5 Conclusions

As parallelism increases the number of cores integrated onto a chip, there is a clear need for fully integrated DC-DC converters to enable efficient on-die power management. With the availability of high density and high quality capacitors in existing CMOS processes, switched-capacitor DC-DC converters have gained significant interest as a cost-effective means of enabling such power management functionality.

The area required by a fully integrated SC DC-DC converter to deliver a certain amount of power to the load has direct implications on both cost and efficiency. This chapter therefore describes a methodology to predict and minimize the losses of such a converter operating at a given power density. The chapter further introduced a circuit and physical design strategy to enable topology reconfiguration and hence efficient generation of a wider range of output voltages.

Measured results from a 32 nm SOI prototype confirm the methodology's predictions of ~80 % efficiency at a power density of ~0.5–1 W/mm^2 for a 2:1 step-down converter operating from a 2 V input and utilizing only MOS capacitors. Topology reconfiguration enables the converter to maintain >70 % efficiency for most of the output voltage range from 0.7 to ~1.15 V. Given that this performance was achieved in a standard CMOS process with no modifications or additions, these results illustrate that fully integrated switched-capacitor converters are indeed a promising candidate for low-cost but efficient power management on a per-core or per-functional unit basis.

References

1. Q. Li, Y. Dong, F.C. Lee, High density low profile coupled inductor design for integrated point-of-load converter, in *IEEE Applied Power Electronics Conference (APEC),* 2010, pp. 79–85
2. G. Schrom, et al., A 100MHz eight-phase buck converter delivering 12A in 25mm2 using air-core inductors, in *IEEE Applied Power Electronics Conference* (*APEC*). 2007, pp. 727–730
3. J. Wibben, R. Harjanai, A high-efficiency DC–DC converter using 2 nH integrated inductors. IEEE J. Solid-State Circuit **43**(4), 844–854 (2008)

4. D.S. Gardner, G. Schrom, P. Hazucha, F. Paillet, T. Karnik, S. Borkar, Integrated on-chip inductors with magnetic film. IEEE Trans. Magn. **43**(6) (2007)
5. J. Lee, G. Hatcher, L. Vandenberghe, C.K. Yang, Evaluation of fully integrated switching regulators for CMOS process technologies. IEEE Trans. VLSI **15**, 1017–1117 (2007)
6. T. Van Breussegem, M. Steyaert, A 82% efficiency 0.5% ripple 16-phase fully integrated capacitive voltage doubler, in *IEEE Symposium on VLSI Circuit*, June 2009, pp. 198–199
7. D. Somasekhar, B. Srinivasan, G. Pandya, F. Hamzaoglu, M. Khellah, T. Karnik, K. Zhang, Multiphase 1GHz voltage doubler charge-pump in 32nm logic process. IEEE J. Solid-State Circuit **45**(4), 751–758 (2010)
8. K.D.T. Ngo, R. Webster, Steady-state analysis and design of a switched-capacitor DC-DC. IEEE Power Electron. Spec. Conf. **1**, 378–385 (1992)
9. B.R. Gregoire, A compact switched-capacitor regulated charge pump power supply. IEEE J. Solid-State Circuit **41**(8), 1944–1953 (2006)
10. M.D. Seeman, S.R. Sanders, Analysis and optimization of switched-capacitor DC–DC converters. IEEE Trans. Power Electron. 841–851 (2008)
11. D. Maksimovic, S. Dhar, Switched-capacitor DC-DC converters for low-power on-chip applications. IEEE PESC **1**, 54–59 (1999)
12. E. Alon, M. Horowitz, Integrated regulation for energy-efficient digital circuits. IEEE J. of Solid-State Circuit **43**, 1795–1807 (2008)
13. D. Ma, F. Luo, Robust multiple-phase switched-capacitor DC-DC power converter with digital interleaving regulation. IEEE Trans. VLSI Syst. **16**(6), 611–619 (2008)
14. G. Patounakis, Y. Li, K.L. Shepard, A fully integrated on-chip DC-DC conversion and power management system. IEEE J. Solid-State Circuit **39**(3), 443–451 (2004)
15. Y.K. Ramadass, A.P. Chandrakasan, Voltage scalable switched capacitor DC-DC converter for ultra-low-power on-chip applications, in *IEEE Power Electronics Specialists Conference (PESC)*, 2007, pp. 2353–2359
16. H-P Le, M.D. Seeman, S.R. Sanders, V. Sathe, S. Naffziger, E. Alon, A 32nm fully integrated reconfigurable switched-capacitor DC-DC converter delivering 0.55W/mm^2 at 81% efficiency. IEEE ISSCC Dig. Tech. Pap. 210–211 (2010)
17. Y. Ramadass, A. Fayed, B. Haroun, A. Chandrakasan, A 0.16mm^2 completely on-chip switched-capacitor DC-DC converter using digital capacitance modulation for LDO replacement in 45nm CMOS. IEEE ISSCC Dig. Tech. Pap. 208–209 (2010)
18. S.K. Enam, A.A. Abidi, A 300-MHz CMOS voltage-controlled ring oscillator. IEEE J. Solid-State Circuit **25**(1), 312–315 (1990)
19. L. Chang, R. Montoye, B. Ji, A. Weger, K. Stawiasz, R. Dennard, A fully integrated switched-capacitor 2:1 voltage converter with regulation capability and 90% efficiency at 2.3A/mm^2, in *IEEE Symposium VLSI Circuit*. 2010
20. H.-P. Le, J. Crossley, S.R. Sanders, E. Alon, A Sub-ns response fully integrated battery-connected switched-capacitor voltage regulator delivering 0.19W/mm^2 at 72% efficiency. IEEE ISSCC Dig. Tech. Pap. 372–373 (2013)
21. M. Seeman, S.R. Sanders, Analysis and optimization of switched-capacitor DC-DC converters, in *10th IEEE Workshop on Computers in Power Electronics (COMPEL)*, 2006, pp. 216–224.
22. M. Seeman, *A Design Methodology for Switched-Capacitor DC-DC Converters*, University of California, Berkeley, Technical Report No. UCB/EECS-2009–78, 2009