

10.8 A 32nm Fully Integrated Reconfigurable Switched-Capacitor DC-DC Converter Delivering 0.55W/mm² at 81% Efficiency

Hanh-Phuc Le¹, Michael Seeman¹, Seth R. Sanders¹, Visvesh Sathe², Samuel Naffziger², Elad Alon¹

¹University of California, Berkeley, CA

²AMD, Fort Collins, CO

With the rising integration levels used to increase digital processing performance, there is a clear need for multiple independent on-chip supplies in order to support per-IP or block power management. Simply adding multiple off-chip DC-DC converters is not only difficult due to supply impedance concerns, but also adds cost to the platform by increasing motherboard size and package complexity. There is therefore a strong motivation to integrate voltage conversion blocks on the silicon chip.

The key challenge associated with realizing such integrated converters is achieving high efficiency at the high power densities required by high-performance digital logic. In typical CMOS processes, on-die capacitors have significantly higher Q and energy density and lower cost than on-die inductors, leading to several recent efforts in exploring fully integrated switched-capacitor (SC) DC-DC converters. For example, both [1] and [2] investigated interleaving to reduce the output ripple of fully integrated SC voltage doublers, with [1] demonstrating high efficiency (82%), and [2] showing high power density (~0.5-1.1W/mm²). In this paper, we expand upon these previous designs by demonstrating a fully integrated step-down SC converter capable of achieving high efficiency (81%) at high power density (0.55W/mm²) while supporting a wide range of output voltage levels.

Careful analysis of the 4 major loss components of integrated SC converters (Fig. 10.8.1a) is critical to optimizing converter design parameters (particularly switching frequency, f_{sw}) and hence achieving high efficiency and power density. At low power density, efficiency is largely set by stray capacitance (bottom/top plate), while at high power density, efficiency is set by flying capacitor density and switch speed. Since the converter is targeted at digital loads, it is important to recognize that the minimum output voltage V_{min} sets the performance of the digital blocks. Therefore, the increased load current caused by the converter's output ripple should be included as additional loss [3]. Fortunately, interleaving reduces ripple without impacting V_{min} , and is thus effective in mitigating this loss.

As shown in Fig. 10.8.1b, in order to enable variable conversion ratios while maintaining efficiency, the converter is partitioned into multiple standard cells, each consisting of 1 flying capacitor and 5 switches. 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 variable converter topology. In this demonstration, we group 2 standard cells in a converter unit supporting 3 topologies with conversion ratios of 2/3, 1/2, and 1/3 (1.33V, 1V, and 0.66V with a 2V input). Intermediate voltage levels are generated by controlling f_{sw} and thus the output impedance of the converter [4], which is equivalent to linear regulation off of the ideal voltage.

Each converter unit operates in 2 non-overlapping clock phases $c1$ and $c2$ (Fig. 10.8.2) with controllable dead-time. Using versions of $c1$ and $c2$ that either swing between V_o and V_{ss} , V_i and V_o (noted by $_h$), or V_i and V_{ss} (noted by $_fs$), it is fairly straightforward to drive the gates of all of the (thin-oxide) switches except for M4, M5, and M7. In the 1/3 (2/3) mode when switch M4 (M7) should always be off, the source of M4 (M7) is driven above (below) the rails of a standard inverter driver. Dual-rail power gating and voltage clamps (M4a, M7a, D4, and D7) are therefore added to INV4 and INV7 in order to ensure these switches remain off, as illustrated in Fig. 10.8.2. Other than the addition of small parasitic capacitance, these added elements have no effect on the operation of the converter in other modes.

Driving switch M5 is even more challenging since its source and drain voltages can reside anywhere among V_i , $(V_i/2+V_o/2)$, V_o , and $(V_o/2)$ in the 3 modes of operation. Fortunately however, the physical/logical ordering of the standard cells can be leveraged to realize a simplified driver design. Specifically, since C1 is always intended to be stacked above C2, C1_{pos} and C1_{neg} can be used as virtual rails for a "flying" inverter INV5 (Fig. 10.8.2). By connecting the input of INV5 to C2_{pos}, M5 is automatically controlled by the actions of the other switches on the two flying capacitors.

The wide variation in voltage levels also makes level shifting the clock signals from V_o - V_{ss} to V_i - V_o challenging. Many conventional level-shifter designs lead to V_o -dependent unbalanced duty cycles and timing mismatch between $c1$, $c2$ and $c1_h$, $c2_h$, and consequently to short-circuit currents. Although increasing dead-time can mitigate these issues, it comes at the penalty of reduced efficiency due to lower effective power-transfer duty cycle. To minimize level-shifter-induced clock mismatch, the design shown in Fig. 10.8.3 DC biases $lv3$ at its trip point through R1 and AC couples the V_o - V_{ss} input clock through a capacitor C_{ls1} . In order to enable low minimum f_{sw} without excessive area, R1 is implemented by minimum-sized pass-gates in series with two pairs of back-to-back diodes. These diodes reduce the voltage swing across the pass-gate and hence increase their equivalent resistance. To further reinforce matched timing, C_{ls2} is added to couple $lv4$ and $lv5$, and fewer than 3 levels of logic are used after the level shifter.

As previously mentioned, interleaving is critical to SC converter efficiency when driving digital loads, and therefore this design utilizes 32-phase interleaving. The layout is arranged so that 2 adjacent converter units operate 180° out-of-phase and hence can share a single distributed clock phase and level shifter. As also shown in Fig. 10.8.3, at a higher level, the phases are arranged to maximize the symmetry of the switching actions around the center of the array.

In order to validate its performance, the SC converter is integrated into a 32nm SOI test chip. Figure 10.8.7 shows the test-chip micrograph. To measure efficiency, the test chip includes an array of digitally-controlled PMOS load transistors that mimic the I-V characteristics of standard digital blocks. Figures 10.8.4a and 10.8.4b show the converter efficiency and optimal f_{sw} in the 1/2 mode while supplying the load circuits, and Fig. 10.8.5 shows the converter's efficiency across output voltage in the 3 operating modes. The converter achieves a peak efficiency of 81% at 0.55W/mm², with a good correlation between measured and predicted peak efficiencies in the 2/3 and 1/2 modes. The most likely causes of the significant discrepancy in efficiency in the 1/3 mode appears to be un-modeled leakage from V_i and V_o , and that switches M3, M8, and M9 approach sub-threshold operation at low V_o ; further testing is in progress to definitively root cause these issues. Note that since the 32nm process is developmental, we expect converter efficiency to improve and match more closely with analytical predictions in a production process with a higher performance.

Acknowledgement:

The authors would like to thank BWRC sponsors, IFC, AMD, and NSF Infrastructure Grant No. 0403427.

References:

- [1] T. Van Bruessegem and M. Steyaert, "A 82% Efficiency 0.5% Ripple 16-Phase Fully Integrated Capacitive Voltage Doubler," *IEEE Symp. VLSI Circuits*, pp. 198-199, Jun., 2009.
- [2] D. Somasekhar, et al., "Multi-phase 1GHz Voltage Doubler Charge-Pump in 32nm logic process," *IEEE Symp. VLSI Circuits*, pp. 196-197, Jun., 2009.
- [3] E. Alon and M. Horowitz, "Integrated Regulation for Energy-Efficient Digital Circuits," *IEEE J. Solid-State Circuits*, pp. 1795-1807, Aug., 2008.
- [4] M.D. Seeman, S.R. Sanders, "Analysis and Optimization of Switched-Capacitor DC-DC Converters," *IEEE Trans. Power Electronics*, pp. 841-851, Mar., 2008.




Figure 10.8.7: Chip Micrograph.