

A RISC-V Vector Processor with Tightly-Integrated Switched-Capacitor DC-DC Converters in 28nm FDSOI

Brian Zimmer¹, Yunsup Lee¹, Alberto Puggelli¹, Jaehwa Kwak¹, Ruzica Jevtic¹, Ben Keller¹, Stevo Bailey¹, Milovan Blagojevic^{1,2}, Pi-Feng Chiu¹, Hanh-Phuc Le¹, Po-Hung Chen¹, Nicholas Sutardja¹, Rimas Avizienis¹, Andrew Waterman¹, Brian Richards¹, Philippe Flatresse², Elad Alon¹, Krste Asanović¹, and Borivoje Nikolić¹

¹Dept. of Electrical Engineering and Computer Sciences, University of California, Berkeley

²STMicroelectronics, Crolles, France

Abstract

This work demonstrates a RISC-V vector microprocessor implemented in 28nm FDSOI with fully-integrated non-interleaved switched-capacitor DCDC (SC-DCDC) converters and adaptive clocking that generates four on-chip voltages between 0.5V and 1V using only 1.0V core and 1.8V IO voltage inputs. The design pushes the capabilities of dynamic voltage scaling by enabling fast transitions (20ns), simple packaging (no off-chip passives), low area overhead (16%), high conversion efficiency (80-86%), and high energy efficiency (26.2 DP GFLOPS/W) for mobile devices.

Introduction

Optimal energy efficiency requires tight integration of the power supply control with the microprocessor. Alternatives to high latency off-chip regulators are integrated low-drop-out (LDO) regulators [1], buck converters with off-chip inductors [2], and SC-DCDC converters [3]. Traditional interleaved SC-DCDC converters stabilize the output voltage to minimize frequency margining for supply variation, but in principle, efficiency could be increased by using a non-interleaved design that avoids charge sharing [4]. In the non-interleaved operation, SC-DCDC unit cells switch simultaneously to avoid charge sharing losses, and an adaptive clock translates a higher instantaneous voltage into a higher frequency to exploit the rippling supply voltage.

Integrated System Implementation

Figure 1 shows the chip architecture. The 64-bit single-issue in-order scalar core implements the open-source RISC-V instruction set [5]. The scalar core has a memory management unit that supports page-based virtual memory, an IEEE 754-2008-compliant floating-point unit, a high-performance 64-bit vector accelerator with vector gather and scatter support, and L1 caches. The processor boots Linux and executes both compiled scalar and vector code with single- and double-precision floating-point operations, including fused multiply-add. Two voltages, a 1.0V core and 1.8V I/O, are supplied to the on-chip converters. The SC-DCDC converter is partitioned into twenty-four $90\mu\text{m} \times 90\mu\text{m}$ unit cells surrounding the core (16% area overhead), and generates four dynamically reconfigurable average output voltages of 1.0V, 0.9V, 0.67V, and 0.5V. An adaptive clock generator adjusts the clock period each cycle based on the instantaneous converter output voltage. Level shifters and asynchronous FIFOs separate the core and uncore voltage domains. Large random variations in SRAM memory cells typically limit voltage scaling, so custom SRAMs were implemented to enable voltage scaling down to 0.45V. Each 4KB SRAM uses 8T cells and has 512 words of 72 bits with 2:1 interleaving.

Figure 2 depicts the on-chip SC-DCDC converter. Four possible discrete SC-DCDC configurations generate voltages between 0.45V and 1V to enable a wide operating range. Hopping between values enables intermediate average voltages. A lower bound (hysteretic) controller switches the cells when the output voltage V_{out} drops below an external reference.

The adaptive clocking scheme, shown in Figure 3, ensures that the system operates at the maximum instantaneous frequency. The rippling supply voltage powers a tunable replica circuit (TRC), adjustable from 4 to 124 FO1 inverter delays, to mimic the critical path delay at the instantaneous voltage level. When the TRC generates a pulse, the controller selects one of the sixteen DLL phases to send to the core as a clock edge. Figure 3 shows that the TRC accurately tracks the 100mV voltage ripple for each SC-DCDC mode; different TRC lengths are used for different SC-DCDC modes.

Measurement Results

The chip is implemented in 28nm ultra-thin body and BOX fully-depleted silicon-on-insulator (UTTB FDSOI) technology [6]. Figure 4 shows measured traces of the rippling core voltage domain for all four possible configurations as well as fast voltage transitions between topologies, which enables very fine-grained DVFS. For all possible converter topologies with adaptive clocking, the processor successfully boots Linux and runs sophisticated user applications such as Python, demonstrating that complex digital logic operates reliably with an intentionally rippling supply voltage.

Figure 5 shows the measured voltage conversion efficiency. The system efficiency measurement accounts for all sources of loss, including losses caused by non-ideal tracking of the voltage ripple. To measure system efficiency, an ideal off-chip regulator directly supplies V_{out} using the bypass DC-DC mode, and total energy versus elapsed time is measured for long running benchmarks. To compare, the SC-DCDC converter and adaptive clock are enabled and the same benchmarks are run while measuring power input to the converter. The system energy efficiency is defined as the ratio in energy for the SC-DCDC-supplied processor to finish the same workload in the same runtime as in the bypass case. The measured efficiency of 80–86% is higher than an ideal LDO, the only other published regulation technique integrated with a functional digital load that does not require off-chip passives. Loss in efficiency can be caused by either converter losses or imperfect tracking of the voltage ripple by the adaptive clock. To gauge the relative contribution of these effects, the efficiency of the converter alone is estimated by characterizing the power at each voltage using a repetitive micro-benchmark and numerically integrating the waveform at V_{out} . The converter alone achieves a maximum efficiency above 90%.

Figure 6 shows the energy efficiency of the system for different workloads and body bias modes enabled by FDSOI technology. The combination of the RISC-V architecture, resilient SRAM, and wide operating range enabled by on-chip voltage conversion and adaptive clocking achieves 26.2 GFLOPS/W with the 1V 1/2 DC-DC configuration when computing double-precision matrix-multiplication using the vector accelerator. Maximum efficiency in the bypass mode is 34 GFLOPS/W at 0.54V and 1W of body bias.

Acknowledgements

The authors would like to thank Tom Burd, James Dunn, Olivier Thomas, and Andrei Vladimirescu for their contributions. This work was funded in part by BWRC, ASPIRE, DARPA PERFECT Award Number HR0011-12-2-0016, Intel ARO, AMD, GRC, Marie Curie FP7, NSF GSFP, NVIDIA Fellowship, and fabrication donation by STMicroelectronics.

References

- [1] Z. Toprak-Deniz *et al.*, in *ISSCC*, 2014, pp. 98–99.
- [2] N. Kurd *et al.*, in *ISSCC*, 2014, pp. 112–113.
- [3] H.-P. Le *et al.*, *JSSC*, vol. 46, no. 9, pp. 2120–2131, 2011.
- [4] R. Jevtic *et al.*, *TVLSI*, vol. PP, no. 99, pp. 1–1, 2014.
- [5] Y. Lee *et al.*, in *ESSCIRC*, 2014, pp. 199–202.
- [6] P. Flatresse *et al.*, in *ISSCC*, 2013, pp. 424–425.

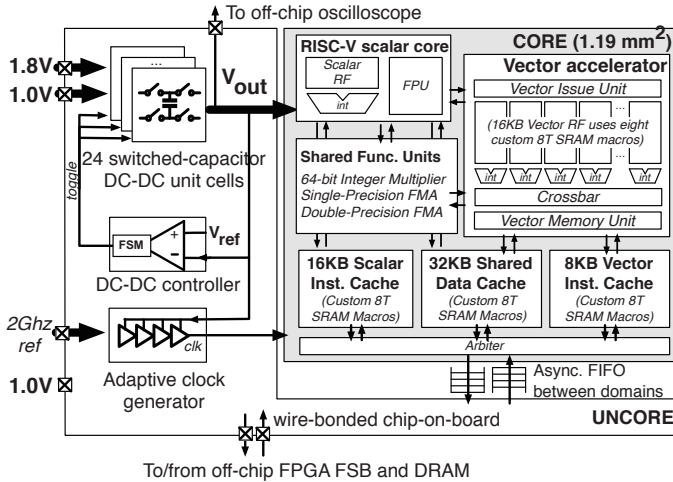


Fig. 1: System block diagram.

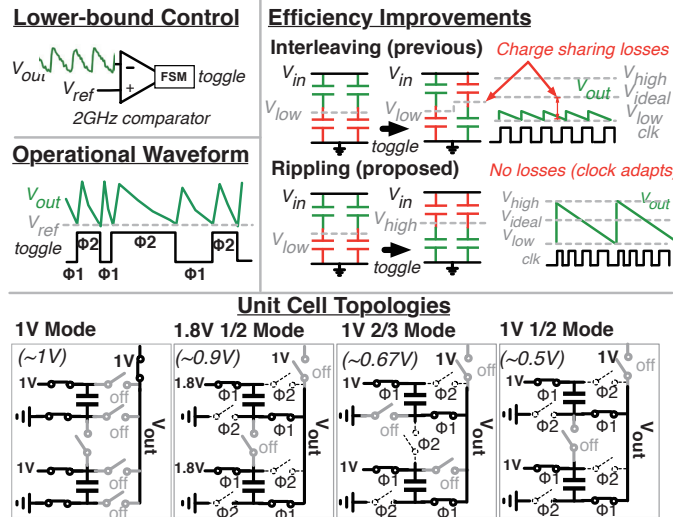


Fig. 2: Reconfigurable SC DC-DC design with four topologies using non-interleaved switching to minimize charge-sharing losses.




Fig. 3: Adaptive clock system with a tunable replica path to track the critical path for constantly changing output voltage.

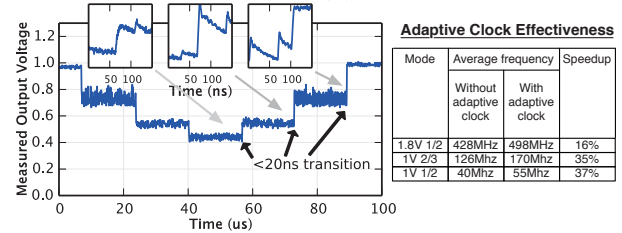



Fig. 4: Oscilloscope measurements of the core voltage V_{out} through a sense pad for all four on-chip regulation modes.

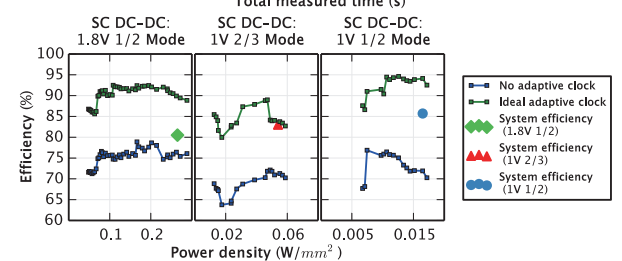



Fig. 5: Measured system efficiency (including overhead of non-ideal adaptive clocking) compared to estimated converter efficiency.

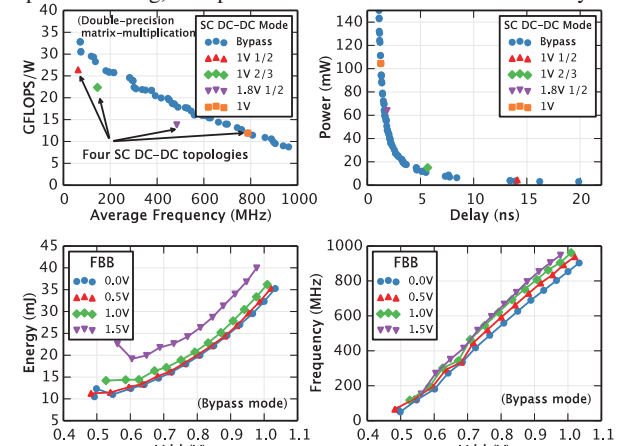


Fig. 6: Measurement of energy efficiency, and the ability of forward body bias (FBB) to trade off energy efficiency with delay.




Fig. 7: Die micrograph and chip summary.