11.1 A Scalable Heterogeneous Integrated Two-Stage Vertical Power-Delivery Architecture for High-Performance Computing

Casey Hardy[1], Hieu Pham[1], Mohamed Mehdi Jatlaoui[2], Frederic Voiron[2], Tianshi Xie[1], Po-Han Chen[1], Saket Jha[1], Patrick Mercier[1], Hanh-Phuc Le[1]

[1]University of California, San Diego, CA
[2]Murata, Caen, France

Emerging high-performance computing needs in data center, autonomous vehicle, and mobile device processors demand increasingly large peak currents at scaled-CMOS-compatible voltages (<1V). To ease otherwise high I^2R losses in power delivery (PD), most applications now target high system-level voltage busses (e.g., 20V) prior to arriving at the processor load. However, voltage incompatibility with scaled CMOS means that conventional approaches require a voltage regulator module (VRM) located laterally off-chip [1] for conversion down to a lower rail (e.g., 1.8V), followed by a fully integrated voltage regulator (FIVR) using external LC filters with inductors embedded into the package substrate [2] (Fig. 11.1.1). The large conversion ratio required by the VRM and the large lateral interconnect losses from the VRM to the processor, followed by low-voltage/high-currents (LV/HC) that flow out of the FIVR to the substrate conductors/passives and then back into the processor at even lower voltages/higher currents, limit the overall efficiency to ~80%, which is not sufficient to meet the needs of emerging energy and thermally-constrained systems. Additionally, LV/HC processor PD requires a large number of pins (e.g., ~50% of all pins for power/gnd), which is problematic for emerging AI/ML applications that require pin-intensive high-throughput memory access. These issues are exacerbated in future systems where more heterogeneous integration (HI) is required.

To address these challenges, the proposed solution shown in Fig. 11.1.1 utilizes a 2-stage approach that gradually tapers the current/voltage delivery while leveraging HI to confine the highest currents near the point-of-load while simultaneously supporting larger input voltages. The 1st stage consists of a hybrid VRM (HVRM), implemented on a 0.18µm HV BCD process, that converts the high voltage (HV) input to an intermediate voltage near 4V rather than directly converting to the LV/HC domain. A 65nm CMOS, 4:1 switched-capacitor voltage regulator (SCVR) co-packaged on a high-density deep-trench-capacitor interposer and mounted under the substrate provides the remaining conversion to the LV domain. This confines the highest system currents with direct vertical delivery at the point-of-load, which reduces associated conduction losses and routing complexity over the traditional solutions. Additionally, as shown in Fig. 11.1.1, the SCVR modules can be readily distributed as an "active" decoupling component in place of passive decoupling capacitors, providing a readily scalable PD solution for larger/higher power in HI systems.

Figure 11.1.2 provides the topology implementations for the converters. The HVRM consists of a 3-level buck that provides increased inductor switching frequency and reduced switching loss benefits over a conventional buck. Rather than utilizing a cascade of two 2:1 SC stages which would require an additional inter-stage hold capacitor (C_H) and 8 power switches, the SCVR utilizes a merged two-stage topology that eliminates the need for C_H [3] and one power switch while also doubling the switching frequency of the 2nd SC stage. The measured steady-state operating waveforms and corresponding states are also provided in Fig. 11.1.2.

The HVRM's flying capacitor, C_1, requires active voltage balancing, since it is fully soft charged/discharged by its inductor current and would be susceptible to V_{C1} drift due to power stage timing and impedance mismatches. The balancer and V_M regulation implementation is shown in Fig. 11.1.3. It utilizes a modified approach from [6] where the V_M regulation loop sets the common mode of the error signals, BAL_P and BAL_M, which determine the duty cycles of the PWM1 and PWM2 signals used to derive the switch control signals. The balancer loop then produces a small differential voltage between BAL_P and BAL_M to finely adjust the switch timing to regulate V_{C1} to $V_{IN}/2$. This allows the balancer and V_M loops to be independently compensated with Z_{REG} since the balancer compensation (Z_{BAL}) does not appear in the common mode path of the V_M loop. Figure 11.1.3 shows the measured performance of the balancer where there is a clear imbalance in V_{LX} when disabled, which is then corrected for when enabled. The closed-loop load step response for the HVRM is also provided in Fig. 11.1.3 that shows the V_M droop and overshoot limited to 0.4V and 0.5V, respectively.

To provide a fast transient response in the 2nd stage, the SCVR is regulated using a lower-bound control, shown in Fig. 11.1.4. The control loop employs a clocked comparator, triggered by an external high-frequency clock, CLK_EXT, that modulates the switching frequency, f_{SC}, of the SCVR according to the load to regulate V_O to V_{REF}. The efficiency of the SCVR strongly depends on the capacitance density, operating voltage, ESR, and SRF of the flying capacitors. Multi-layer ceramic capacitors (MLCCs) do not offer the best trade-off between these parameters and are only used for initial demonstration purposes. Instead, silicon capacitors shown in Fig. 11.1.4, particularly those based on nano-porous instead of micro-porus structures, can offer much better performance. Here, custom 3D nano-porus structures with a depth of only a few micrometers, which achieve a capacitance density ~100× larger than 2D counterparts with no trench, are built into the substrate of a custom silicon interposer (Fig. 11.1.7). Thanks to the high-k materials and manufacturing process, these capacitors have a state-of-the-art density of 1.3µF/mm², feature low ESR/ESL, and are stable over voltage, temperature, and frequency with negligible capacitance derating with DC bias. The SCVR's integrated flying capacitors, C_2, C_3, and decoupling capacitors for V_M, V_O, and V_{DD}, are isolated from each other using 3D native lateral isolation structures. The top metal layer on the interposer is used to connect to the SCVR and allow arraying multiple SCVRs to support a larger load. The capacitors are placed very close to the SCVR I/Os to minimize interconnect ESR/ESL. The landing pattern and pads were designed for the heterogeneous assembly of the SCVR onto the silicon interposer. This interposer prototype is intended to be connected to a test PCB using multiple thick wire bonds. During full system-level integration, the SCVR interposer can be placed directly between the bumps of the system substrate, enabling efficient vertical PD with a potential 2× power pin reduction.

Figure 11.1.5 shows the performance of the complete PD system with the HVRM and SCVR. In the load step transients, the two-stage PD system is measured with 1 HVRM powering 3 SCVR chips with separated outputs, V_{O1-3}. While V_{O3} provides 0.65A at 0.9V, V_{O1} (0.92V) and V_{O2} (0.95V) receive 5A/µs current steps of 0.45A and 0.6A, respectively. The corresponding inductor current I_L and V_{C1P} illustrate the HVRM's stable and balanced operation, while V_{C2P} and V_{C3P} for V_{O2} demonstrate the lower-bound regulation of the SCVRs with their respective switching frequencies modulated accordingly by the loads. These waveforms are captured simultaneously with two synchronized oscilloscopes. The outputs exhibit less than 40mV voltage droop and no overshoot across load transients, demonstrating the load-line characteristics of SCVR. Figure 11.1.5 also shows measured efficiencies of the two stages individually and in combination, assuming 1 HVRM is connected to 6 identical SCVRs to power one single output. The SCVRs are currently assembled with MLCCs for early system evaluation. With a peak efficiency of 93.1% (94.2%) for the HVRM and 89.9% (92.5%) for the SCVR, the 2-stage heterogeneous PD achieves 86% (87.4%) when converting a 20V (12V) input to a 0.9V (0.95V) output.

Figure 11.1.6 summarizes the heterogeneous integrated PD system prototype performance and compares it with the state-of-the-art academic and industry works with similar voltage conversion ratios. This work achieves ~10% higher peak system efficiency, equivalent to ~44% maximum loss reduction, as compared to other 2-stage candidates in [4] and [2], and demonstrates scalability for future heterogeneously integrated systems that is not achievable with the single stage architecture in [5]. The die micrographs of the converters and interposer are shown in Fig. 11.1.7. The HVRM, SCVR, and interposer are fabricated in a 0.18µm BCD, 65nm CMOS, and 3D nano-porus deep trench processes with areas of 3.67mm², 2.6mm², and 12.1mm², respectively.

Acknowledgement:
This work was supported by the Qualcomm Innovation Fellowship, Semiconductor Research Corporation Task No. 2810.061 through UT Dallas' TxACE, and the NSF CAREER Award No. 2042525. The authors would like to thank Joe Duncan and Marko Koski of Qualcomm.

References:

[1] Analog Devices, "LTC3636 Datasheet: Dual Channel 6A, 20V Monolithic Synchronous Step-Down Regulator," Rev. D, June 2021. Accessed on Dec. 1, 2022 <http://www.analog.com/media/en/technical-documentation/data-sheets/LTC3636-3636-1.pdf>

[2] E.A. Burton *et al.*, "FIVR — Fully integrated voltage regulators on 4th generation Intel® Core™ SoCs," *IEEE APEC 2014*, pp. 432–439, Mar. 2014.

[3] L.G. Salem and P.P. Mercier, "A battery-connected 24-ratio switched capacitor PMIC achieving 95.5%-efficiency," IEEE Symp. VLSI Circuits, pp. C340–C341, June 2015.

[4] P. Vivet *et al.*, "2.3 A 220GOPS 96-Core Processor with 6 Chiplets 3D-Stacked on an Active Interposer Offering 0.6ns/mm Latency, 3Tb/s/mm² Inter-Chiplet Interconnects and 156mW/mm² @ 82%-Peak-Efficiency DC-DC Converters," *ISSCC*, pp. 46–47, Feb. 2020.

[5] K. Wei et al., "A Direct 12V/24V-to-1V 3W 91.2%-Efficiency Tri-State DSD Power Converter with Online V_{CF} Rebalancing and In-Situ Precharge Rate Regulation," *ISSCC*, pp. 190–191, Feb. 2020.

[6] X. Liu et al., "A High-Frequency Three-Level Buck Converter with Real-Time Calibration and Wide Output Range for Fast-DVS," *IEEE JSSC*, vol. 53, no. 2, pp. 582–595, Feb. 2018.

2023 IEEE International Solid-State Circuits Conference (ISSCC) | 978-1-6654-9016-0/23/$31.00 ©2023 IEEE | DOI: 10.1109/ISSCC42615.2023.10067315

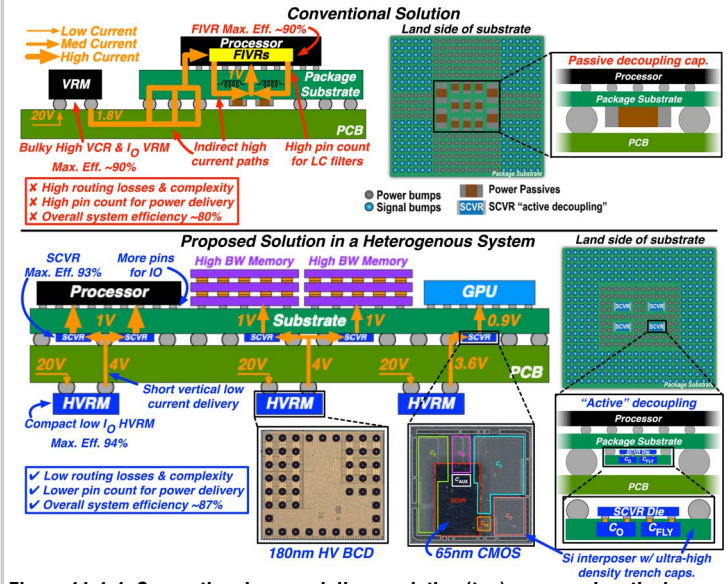

Figure 11.1.1: Conventional power-delivery solution (top), proposed vertical power-delivery solution in a heterogeneous system application example (bottom).

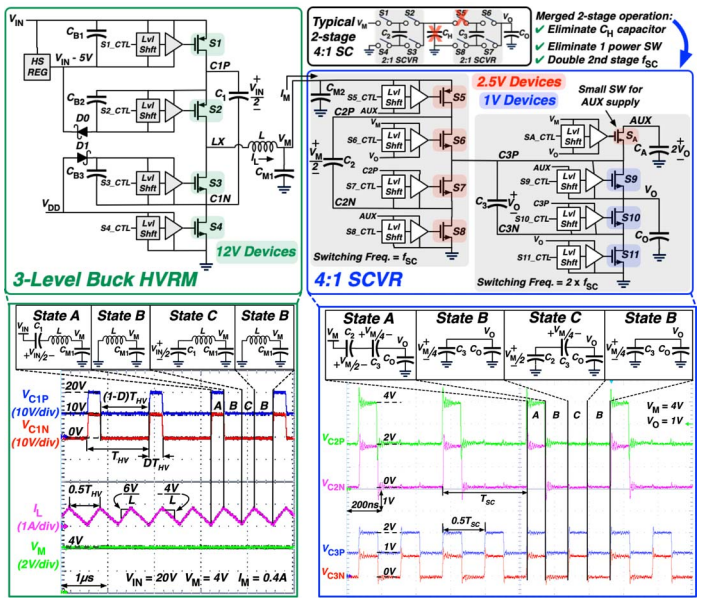

Figure 11.1.2: HVRM/SCVR power stages (top) and steady-state waveforms (bottom).

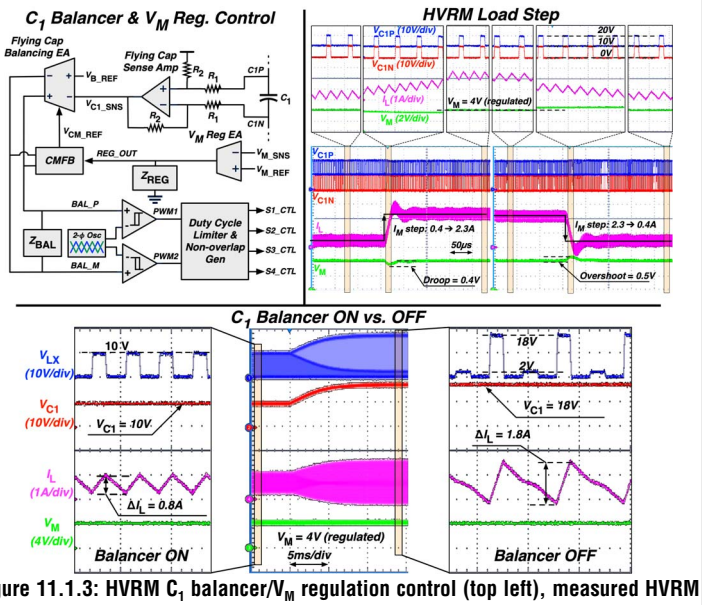

Figure 11.1.3: HVRM C_1 balancer/V_M regulation control (top left), measured HVRM load step (top right), and measured C_1 balancer performance (bottom).

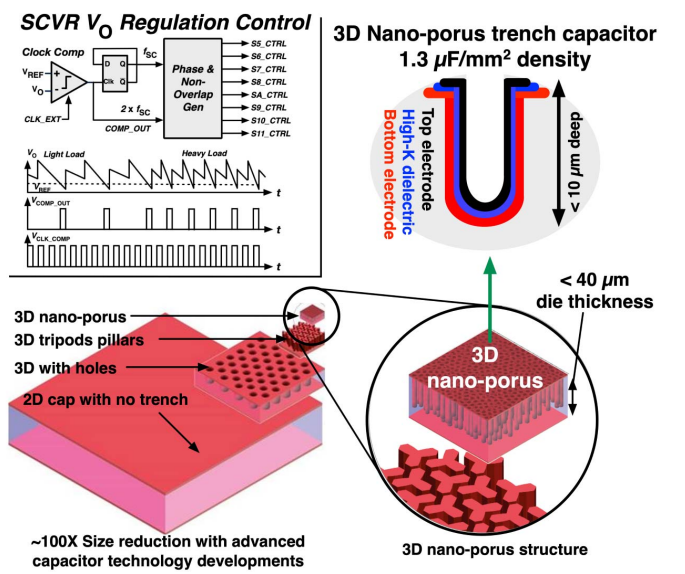

Figure 11.1.4: SCVR output-voltage regulation control (top left) and IPD high density trench capacitor structures used in interposer.

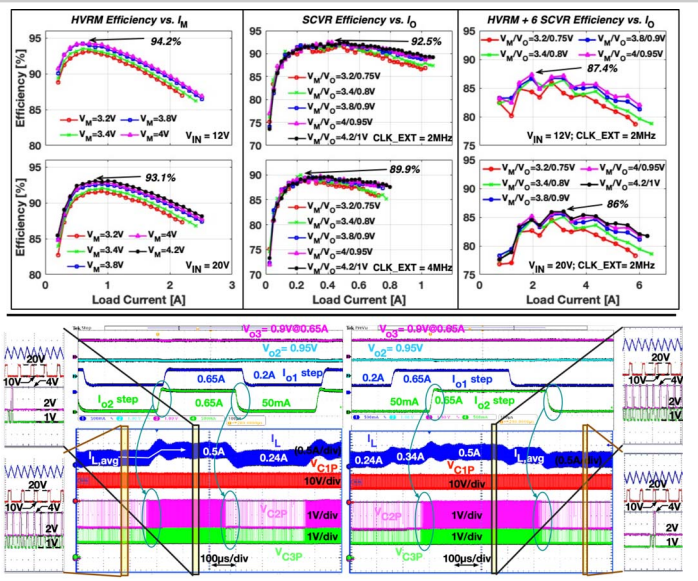

Figure 11.1.5: Efficiency versus output current for the HVRM/SCVR/HVRM + 6 SCVRs (top) and measured load step response when an HVRM is powering 3 SCVRs.

	Parameter	ISSCC '20 [5]	ISSCC '20 [4]	APEC '14 [2]	This work
	Architecture	Single-Stage	Heterogeneous Two-Stage		
Stage 1	Topology	Tri-State DSD	NR	Buck (LTC3636 [1])	3-level Buck
	Process	180nm HV BCD	NR	NR	180nm HV BCD
	V_{IN} / V_O Range	12–24V / 1V	NR / 1.5–2.5V	3.1–20V / 0.6–5V	12–20V / 3.2–4.2V
	Max. Eff.	91.2% @ V_{IN}=12V, P_{OUT}=0.5W^	NR	86.5% @ V_{IN}=20V, P_{OUT}=7.56W^	94.2% @ V_{IN}=12V, P_{OUT}=2.84W
Stage 2	Topology	-	SCVR (4:1–4:3)	Multi-Phase Buck	SCVR (4:1)
	Process	-	CMOS 65nm	CMOS 22nm	CMOS 65nm
	V_{IN} / V_O Range	-	1.5–2.5V / 0.5–1.1V	1.8V / 0.6–1.1V	3.2–4.2V / 0.75–1V
	Max. Eff.	-	82% @ 1.7W	90% @ 1.05–15.75W	92.5% @ 0.4W **
	Integration Level	-	Active Interposer	Integrated on CPU die	Die + IPD Interposer
	Energy Storage Element	-	MOS + MOM + MIM Capacitors	Air Core Package Inductor	IPD Capacitors (Tested w/ 0201 MLCCs)
System	Max. Eff.	91.2%	77.9% *	77.85%^	87.4% ***

^ Estimated from reported data and datasheet
NR: Not reported
* assumed 95% efficiency for 1st stage
** output power of each unit of the SCVR
*** system of 1 HVRM + 6 SCVR

Figure 11.1.6: Performance summary and comparison.

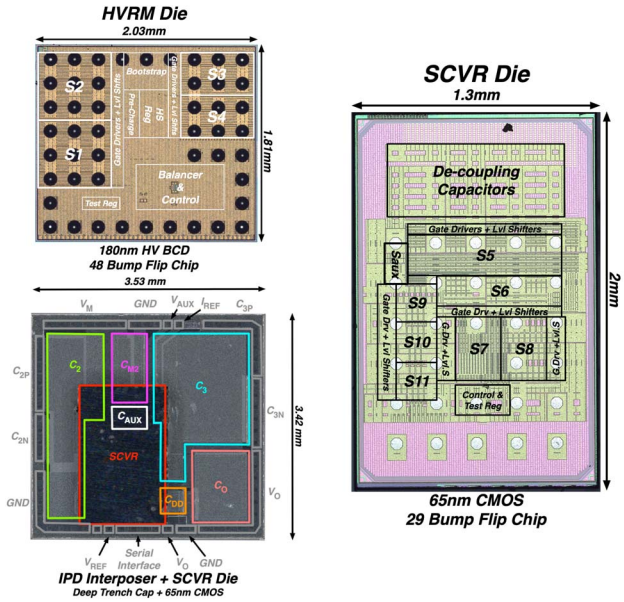

Figure 11.1.7: HVRM, SCVR, IPD interposer die micrographs.

978-1-6654-9016-0/23/$31.00 ©2023 IEEE