

Scalable series-stacked power delivery architectures for improved efficiency and reduced supply current

Robert Pilawa

Enver Candan, Josiah McClurg, Sai Zhang, Pradeep Shenoy* Phil Krein, Naresh Shanbhag University of Illinois at Urbana-Champaign *Texas Instruments

> PwrSoC 2014 October 8, 2014

Outline

- Background and motivation
- Series-stacked architecture
 - Previous work
 - Advantages and challenges
- Top down:
 - General purpose computing units
 - Extreme efficiency data center power delivery
 - Scalability, experimental results
- Bottom up:
 - Specialized compute cores
 - Power converter and load integration
 - New switched-capacitor architectures
- Conclusion





- Continued voltage reduction
- CPU power has remained fairly constant





VRM limitations





- Large current on I/O pins
- Difficult to realize large step-down voltage conversion on-chip
- Efficiency is limited by power converter efficiency

Can we leverage the increased core count for power delivery purposes?





Other Applications – Low Voltage Sources



Solar 0.5 V cells 30 V modules 600 V strings Battery systems

- 3-12 V cells
- Up to 400 V DC bus





Series-stacking is widely used in low voltage DC sources

Series-stacked Architecture - Promise







Prior Work









Rajanpandian et al. "High Voltage Power Delivery Through Charge Supporting Vertically-Recycling, JSSC 2006

K. Kesarwani et al. "A Multi-Level Ladder Converter Stacked Digital Voltage Domains," APEC 2013

S.K. Lee et al. "Evaluation of Voltage Stacking for Near-**Threshold Multicore** Computing", ISLPED 2012

- We are looking at massively scaled architectures for future multi-core architectures
- Rely on high efficiency power electronics

Data Center Architectures

Before:

- Single computer
- Single processor
- Single core
- Stacking not beneficial
- Today:
 - Warehouses of computers
 - Multi-processor severs
 - Multi-core CPUs
 - Stacking helpful, across the full system





Blue Waters @ Illinois



1867

Data center application





- Motivation here is conversion efficiency
- Shares many constraints with stacked cores
- Proof-of-concept demonstration

Series-stacked Architecture - Challenges

- Voltage regulation
- Grounding
 - Design change
- Communication across voltage domains
 - Ethernet 1500 V isolation
 - Fiber-optic
- Hot-swapping, reliability





P. S. Shenoy, and P. T. Krein, "Differential power processing for dc systems," *Power Electronics, IEEE Transactions on,* vol. 28, no. 6, pp.2980 – 2997, June 2013.

Proposed Solution – Differential Power Processing

- ✓ Voltage regulation by injecting or rejecting current from nodes.
- ✓ Bidirectional DPP converters process only the difference in power
- Bulk power is delivered to the series-stacked servers without being processed.





Differential Power Processing – Low Voltage Ratings





McClurg ECCE 2014

- Non-isolated
- Inefficient power transfer
- Order-dependent



Candan INTELEC 2014

- Isolated converters
- Minimum power transfer
- Order independent

Dual Active Bridge DC-DC Converter

- Four prototype DAB converter are designed as DPP converters.
- Simple phase shift modulation is used.
- Symmetrical design at both sides of the transformer.

DAB CONVERTER SPECIFICATIONS AND KEY COMPONENTS

Rated Power	120 W
Peak Efficiency	95%
Switching Frequency	175 kHz
Modulation Technique	Simple phase shift
Control Mode	Bidirectional Hysteresis
Switch	DrMOS - Vishay SiC780ACD
Digital Isolator	TI - ISO7241C
Microcontroller	TI - C2000 Piccolo



120 W 12V-12V dual active bridge converter



*BoM is around \$30

E. Candan, "A Series-stacked power delivery architecture with isolated converters for energy efficient data centers," 17 Master's Thesis, University of Illinois at Urbana-Champaign, 2014.

System Level Control

Control Objectives:

- Server voltage regulation.
- Virtual Bus voltage regulation.
- Voltage sampling only.
- No communication between converters.
- Highest possible light-load efficiency
- Bi-directional Hysteresis Control





Experimental Setup and Tests



A. Waterland, Stress POSIX workload generator. [Online]. Available: http://people.seas.harvard.edu/~apw/stress/

Experimental Setup – Conventional Architecture





- A best-in-class PSU with 96% peak efficiency.
- Identical web traffic and computational tests.
- The same measurement unit.

*SynQor PQ60120QEx25 - \$200

Experimental Results – Computation Test



Typical waveforms



AVERAGE INPUT AND OUTPUT POWERS DURING COMPUTATION TEST

$< P_{in} > = V_{Bus} \times I_{Bus}$	426.60 W
$\langle P_{out} \rangle = \sum_{i=1}^{4} V_{s,i} \times I_{s,i}$	426.11 W
Efficiency	99.89 %





COMPARISON OF PROPOSED ARCHITECTURE WITH CONVENTIONAL ARCHITECTURE

	Web Traffic Test		Computation Test	
	Proposed	Conventional	Proposed	Conventional
$< P_{in} > [W]$	241.09	252.87	426.60	447.59
$< P_{out} > [W]$	237.98	238.58	426.11	426.51
$< P_{loss} > [W]$	3.11	14.29	0.49	21.08
Efficiency [%]	98.71	94.35	99.89	95.29

With server to virtual bus DPP:

✓ 4.6 times reduction in average power loss for web traffic
✓ 40 times reduction in average power loss for computation

Note that the standard power supply for this system has 80-90% efficiency

Ongoing work

- Core-level emulation1GHz ARM Cortex-A8
- 5V, 1A Power requirements
- On-board PMU
- Resonant SC DPP converters





- Load balancing strategies
 - Modified Hadoop scheduler

Working with software and CPU architecture partners (you all should)

Bottom Up – Stackable Cores

- Today:
 - Tight voltage regulation
 - Transient conditions
 - Error-free computing
 - Digital logic has priority
- Tomorrow?
 - Deeply scaled CMOS (post-CMOS)
 - Increased variations
 - Embrace errors?
 - Communication-inspired computing
 - Power delivery and computing equal partners



Systems On Nanoscale Information fabriCs Center (SONIC)

Compute VRM







Zhang et al.: A 0.79 pJ/k-gate, 83% efficient unified core and voltage regulator architecture, JSSC 2014

2x2mm, IBM 130 nm CMOS

Embrace the ripple







- Demonstrated a scalable series-stacked computer architecture
 - 12 V servers
 - Differential power processing
 - 40x loss reduction compared to state-of-the-art
- Outlined design challenges and opportunities for core-level work
- Requires careful cooperation/co-design of power electronics, CPU, and software
- Emerging compute units compatible with seriesstacking may be best way forward

Post-CMOS devices may be the ultimate driver for this technology

- Texas Instruments
 - Pradeep Shenoy
- Google
 - Google Faculty Research Award
- UIUC Strategic Research Initiative
 - Profs. Phil Krein, Naresh Shanbhag, Yi Lu

Questions?