

Evaluation of Voltage Stacking for Near-Threshold Multicore Computing

Sae Kyu Lee

David Brooks

Gu-Yeon Wei

Harvard University

33 Oxford St., Cambridge, MA

{saekyu, dbrooks, guyeon}@eecs.harvard.edu

ABSTRACT

This paper evaluates voltage stacking in the context of near-threshold multicore computing. Key attributes of voltage stacking are investigated using results from a test-chip prototype built in 150nm FDSOI CMOS. By “stacking” logic blocks on top of each other, voltage stacking reduces the chip current draw and simplifies off-chip power delivery but within-die voltage noise due to inter-layer current mismatch is an issue. Results show that unlike conventional power delivery schemes, supply rail impedance in voltage stacked systems depend on aggregate power consumption, leading to better noise immunity for high power (low impedance) operation for many-core processors.

Categories and Subject Descriptors

C.5.4 [Computer System Implementation]: VLSI Systems

General Terms

Design

Keywords

Near-threshold computing, power delivery, power management, voltage stacking

1. INTRODUCTION

Power consumption and delivery have emerged as one of the major challenges facing modern computing devices from mobile to server applications. In order to keep power consumption at manageable levels while continuing to offer performance enhancements, design trends have shifted towards multicore computing on a single chip. However, as advances in modern technology pack more and more transistors on a single chip, fixed power budgets due to thermal constraints effectively limit the amount of simultaneously active cores [4]. Near-threshold computing has been proposed as a way

to improve energy efficiency for applications that have plenty of parallelism, the key idea being that lost performance due to reduction in voltage can be reclaimed through increase in aggregate throughput [3]. However, near-threshold operation leads to several challenges that must be overcome. Near-threshold computing suffers from higher susceptibility to V_{TH} and VDD fluctuations due to transistor current's exponential dependence on voltage. Also for a fixed power budget, decrease in supply voltage rapidly increases the amount of current driven through off-chip wire connections. Increase in chip current demand result in higher IR drop and I^2R power loss due to board and package resistance, exacerbated by the fact that off-chip power delivery impedance does not necessarily scale. Increase in IR drop and susceptibility to voltage noise directly conflict with the more stringent voltage margins required in near-threshold operation to maintain performance, which aggravate the problem further.

Dynamic voltage and frequency scaling (DVFS) is another technique that can reduce power consumption and improve energy efficiency. However, as more cores are integrated on chip, traditional DVFS approaches are becoming more difficult to utilize due to the lack of temporal and spatial granularity of existing regulators. Off-chip DC-DC converters suffer from slow voltage transition times, while the high board-level footprint and system cost limit the number of voltage domains that can be implemented. On-chip voltage regulators enable fast voltage transition times and facilitate multiple voltage domains [5], but such regulators suffer from low conversion efficiency, especially for high step-down ratios.

Voltage stacking is an alternative approach to power delivery that powers multiple low-voltage blocks off of a single higher voltage by “stacking” the logic blocks and recycling charge between the layers [8]. By charge recycling, voltage stacking significantly reduces the chip current demands, thereby alleviating the strain on the power delivery network. However, voltage stacking suffers from internal voltage noise issues unique to the voltage stacked power delivery scheme, caused by current mismatches between layers. To evaluate voltage stacking in the context of multicore computing, the underlying key attributes of voltage stacked power delivery system must be fully understood. This paper presents experimental results from a test-chip prototype built in 150nm fully-depleted silicon-on-insulator (FDSOI) process to evaluate voltage stacking for multicore computing. Experimental

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISLPED'12, July 30–August 1, 2012, Redondo Beach, CA, USA
Copyright 2012 ACM 978-1-4503-1249-3/12/07 ...\$10.00.

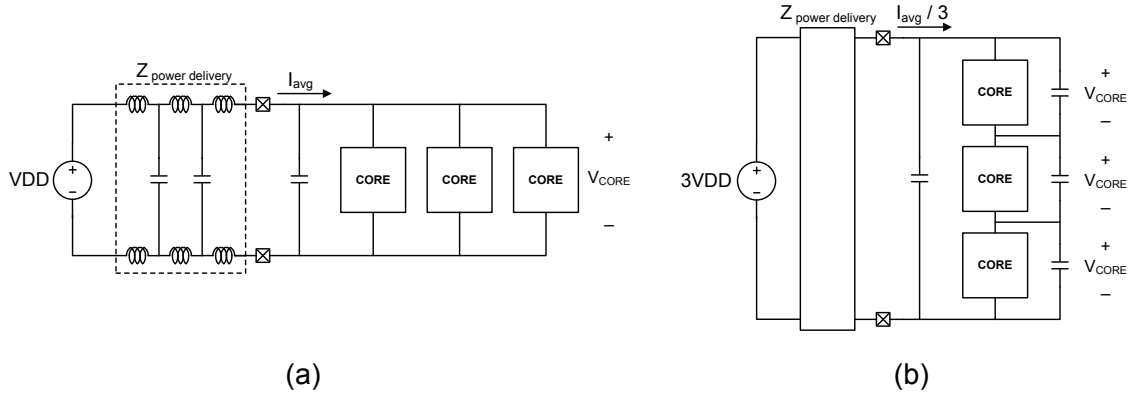


Figure 1: Block diagrams of (a) Conventional Power Delivery Scheme, (b) Voltage Stacking

results show that while susceptibility to internal voltage is a concern, voltage stacking has promising attributes for application in future many-core processors.

2. VOLTAGE STACKING

Voltage stacking is an energy-efficient approach for delivering power that provides high power-supply voltage to the chip and divides up the voltage consumed by the logic blocks via stacking with charge recycling between the layers. Figure 1 presents a conceptual block diagram of voltage stacking for a three core example. Instead of running the cores in parallel off of VDD, the cores are stacked on top of one another and a 3x higher voltage is applied across the stack. If each core consumes same amount of power, the stack voltage ($3V_{\text{DD}}$) will distribute evenly between the cores, with each V_{CORE} settling at VDD and the average current flowing through the power-delivery system is reduced by a factor of 3. For a voltage-stacked design of n -stacks, the average current demand of the chip would decrease by a factor of n .

Voltage stackings addresses some of the key power delivery issues facing modern processors by delivering a higher voltage to the chip and thereby reducing the chip current draw.

- I^2R power loss due to board and package resistance reduces by a factor of n^2 compared to the conventional power delivery scheme of Figure 1(a), significantly reducing the power wasted through the power delivery subsystem.
- IR drop across the power delivery subsystem also reduces by a factor of n . However, in a conventional power delivery scheme, IR drop across the power delivery subsystem affects all cores equally, while in a voltage stacked design, the effect of the IR drop across the voltage stack is distributed across the n -stack layers. Therefore, IR drop across a core operating voltage (V_{CORE}) is ultimately reduced by a factor of n^2 .
- Voltage regulators benefit from lower step-down ratio requirements of voltage stacking, leading to higher regulator efficiency and reduction in design complexity. This results in overall system-level power savings

and potential reduction in board-level components and real estate.

While voltage stacking has the potential to alleviate the inefficiencies related to power delivery, voltage stacking introduces additional concerns that must be addressed. KCL dictates that identical current flows through each of the stack layers. However, current consumption of cores can fluctuate as a function of workload, various power-saving schemes (e.g., clock gating) and data, leading to inter-layer current imbalances. There is an implicit self-regulating loop built into the stack that distributes the internal voltage levels to compensate for any current imbalances between the stack layers. For example, if a stack layer consumes less current compared to other layers, its voltage increases to compensate at the expense of reducing the voltage across the other stack layers. Hence, any aggregate inter-layer core activity mismatch manifests itself as internal voltage noise. Therefore, balancing current consumption across the layers is a key challenge facing voltage stacking. Given the understanding of the potential problems of voltage stacking, a test-chip prototype, presented in the next section, was built to evaluate the key attributes and challenges of voltage stacked multicore systems.

3. TEST CHIP OVERVIEW

Block diagram of the test-chip prototype comprising a 3×3 array of power-consuming *cores* is presented in Figure 2. The test-chip is organized as three voltage-stack layers with three cores per layer and a total of four voltage rails for the array. Only VDDH and VSSC connect to an external voltage source ($V_{\text{DD,EXT}}$) and the internal rails VDDM and VDDL settle to levels dictated by the aggregate current consumption of the *cores*. *Peripheral support* circuitry, containing level shifters and voltage monitoring circuits, feeds signals into and reads signals out of the array, and operates off of separate rails VDDIO and VSSIO connected to a separate external source. The test-chip operates at VDDH=1.5V with each *core* ideally operating at 0.5V.

Each *core* contains two sets of power-consuming blocks shown in Figure 3. The *logic load* block contains three sets of four independent logic paths, all roughly 15 fanout-of-4 inverter (FO4) delays long. The logic paths operate at one

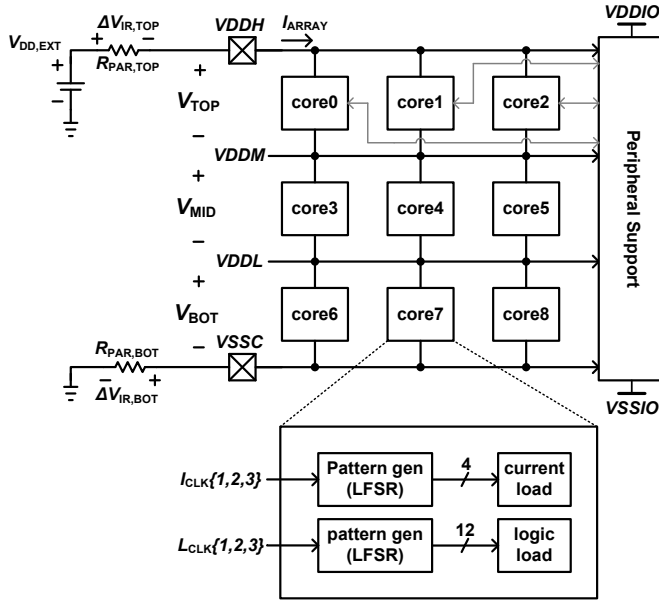


Figure 2: Block diagram of the test-chip prototype with a 3x3 voltage-stacked array of power-consuming blocks

of three clock frequencies ($I_{CLK}\{0,1,2\}$), with input data fed by pattern generators comprising a bank of linear feedback shift registers (LFSRs), programmable via scan chain. This enables configurations (i) to measure logic delay sensitivity to voltage and (ii) to inject fixed or pseudo-random patterns of switching noise into the array at different layers. The *current load* block contains a bank of nMOS switches operating as programmable 4-bit binary-weighted current sources, resulting in a total of 45 *current load* settings per stack layer. This block models average current consumption of processor logic and the switches (operating in saturation) approximate the impedance across the rails of digital CMOS circuits. Separate pattern generators allow the *cores* to mimic larger magnitudes of current fluctuations and coarser-grained switching activity (e.g. clock gating) at one of three clock frequencies ($L_{CLK}\{0,1,2\}$).

Voltage monitor circuit, shown in Figure 4, is used to capture the internal rail voltages in real time. The voltage monitor circuit consists of a source follower (SF) stage that senses the rail voltage followed by a transconductance (Gm) stage that converts the source follower output voltage into I_{OUT} [6]. The test-chip contains four separate voltage monitor circuits for each internal voltage rail (VDDH, VDDM, VDDL, VSSC).

The test-chip was fabricated in a 150nm FDSOI CMOS process from MIT Lincoln Laboratory. An annotated die photo is shown in Figure 5. Not only is voltage stacking easier with SOI process due to the absence of body effect, ultra-thin body SOI technologies are one of the candidates for future aggressively scaled process nodes beyond 32nm. Hence, this test-chip encompasses key attributes to evaluate the potential of using voltage stacking in future designs.




Figure 3: Detailed diagrams of power-consuming *current load* and *logic load* blocks within a core




Figure 4: Block diagram of voltage monitor circuits for real-time voltage measurements




Figure 5: Annotated die photo




Figure 6: (a) *Current load* block's measured current and (b) *logic load* block's F_{MAX} vs voltage

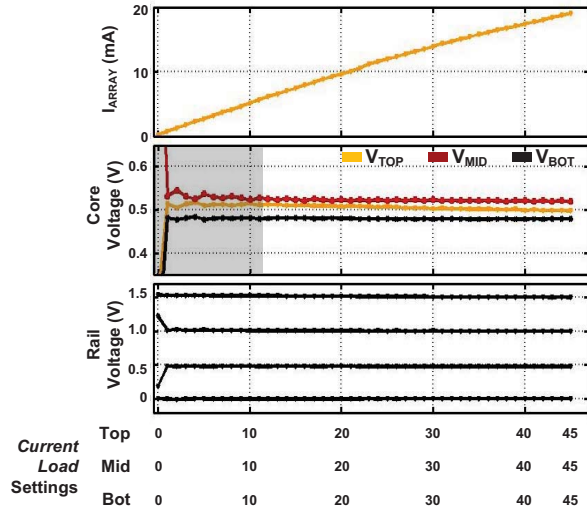


Figure 7: Measured static array current and voltage levels of voltage-stacked array for balanced *current load* setting

4. EXPERIMENTAL RESULTS

Measured current versus voltage relationship for the 8x nMOS switch of the *current load* is presented in Figure 6(a). The the maximum operation frequency (F_{MAX}) of a logic path in the *logic load* versus voltage across the stacks is shown in Figure 6(b). Exponential dependence of F_{MAX} on voltage (Figure 6(b)) clearly illustrate that the *cores* operate in the near-threshold regime at the intended *core* operating voltage of 0.5V.

Under static conditions, balanced inter-layer current consumption evenly distribute stack voltage across the stack layers as dictated by KCL. Figure 7 plots the measured static array current (I_{ARRAY}) and the voltage levels that the internal stack voltages settle to, all versus identical inter-layer static *current load* settings, as the *current load* setting is increased from 0 to 45. For identical current load block settings, the array evenly splits VDDH (at 1.5V) into $\sim 0.5V$ levels. Inherent mismatches between the layers cause small disparities between core voltages. Furthermore, close examination of core voltages at low *current load* settings (shaded in gray) show that voltage deviations are larger at low *current load* setting; the largest difference seen when all

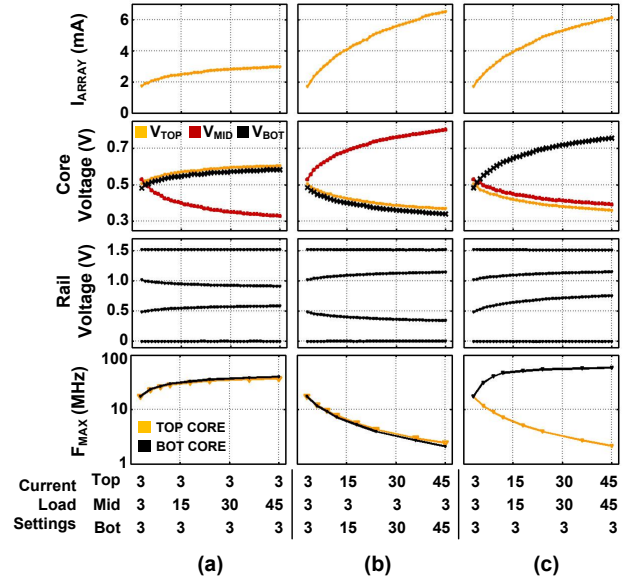


Figure 8: Measured static array current, voltage levels and F_{MAX} of voltage-stacked array for skewed *current load* settings with current increased only in (a) Middle Layer (b) Top and Bottom Layers (c) Top and Middle Layers

blocks only leak (settings of 0). This is due to internal rail impedance, which is further explained later on in the paper.

Figure 8 plots the measured static array current, the voltage levels and F_{MAX} of logic paths in the top and bottom layers, all versus various static skewed *current load* settings. A bug in the level-shifter circuit prevented measurement of F_{MAX} for logic paths in the middle layer. Figure 8 shows that skewed *current load* settings between layers leads to uneven distribution of voltages. For example, as *current load* settings in the middle layer increase relative to the top and the bottom, held at a setting of 3 (Figure 8(a)), voltage across the middle layer decreases. The results show that to maintain balanced voltage distribution across the stack layer as shown in Figure 7, inter-layer current difference must be absorbed using an active regulation scheme, such as push-pull regulators [1], linear regulators [8] or switched-capacitor circuits [2]. However, using voltage regulators to absorb large inter-layer current difference can be costly for sustained current mismatches in terms of energy efficiency.

Conversely, Figure 8 also demonstrates how processor cores within a layer could operate at different voltage/frequency settings from cores at different layers by intentionally skewing inter-layer current consumption. This could be achieved by enabling/disabling different number of cores per layer, essentially trading off core count for higher performance in a particular layer. Voltage stacking shows the potential to achieve workload heterogeneity while obviating voltage regulators.

To fully understand how internal voltages change with current mismatch, we investigate voltage deviation vs current mismatch relationship using a simple two-stack layer

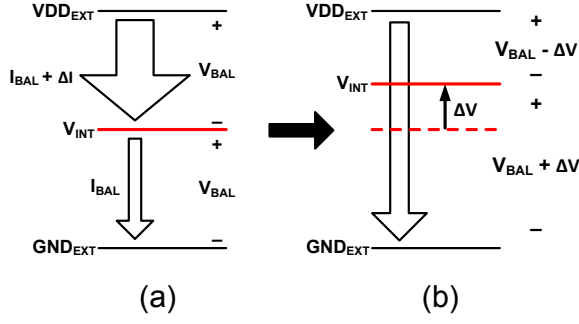


Figure 9: Diagram modeling voltage deviation for a simple two stack layer system

case illustrated in Figure 9, where external voltages VDD_{EXT} and GND_{EXT} are fixed and only V_{INT} change with current mismatch. Starting from a balanced condition ($V_{TOP} = V_{BOT} = V_{BAL}$ and $I_{TOP} = I_{BOT} = I_{BAL}$), if current in the top layer increase as shown in Figure 9(a), internal voltage V_{INT} increases by ΔV to balance the current flow through the stack. Assuming power consuming blocks in each layer are nMOS switches like in the test-chip prototype, the current vs voltage relationship can be modeled using generalized alpha-power law

$$I = K \cdot (V)^\alpha \quad (1)$$

where K represents aggregate width of the transistor and we have ignored V_{TH} for simplicity. Increase in current consumption in the top layer as shown in Figure 9(a) can be modeled as

$$I_{BAL} + \Delta I = (K + \Delta K) \cdot (V_{BAL})^\alpha \quad (2)$$

where ΔK represents the increase in aggregate nMOS device width in the top layer. ΔV due to current mismatch ΔI can then be calculated by solving the KCL equation

$$\begin{aligned} (K + \Delta K) \cdot (V_{BAL} - \Delta V)^\alpha \\ = (K) \cdot (V_{BAL} + \Delta V)^\alpha \end{aligned} \quad (3)$$

and substituting ΔI expression for V_{BAL} , which results in

$$\Delta V = \frac{\sqrt[\alpha]{K + \Delta K} - \sqrt[\alpha]{K}}{\sqrt[\alpha]{K + \Delta K} + \sqrt[\alpha]{K}} \cdot \sqrt[\alpha]{\frac{\Delta I}{\Delta K}} \quad (4)$$

Equations (1)–(4) can easily be expanded to other voltage stacked cases with more layers. Equation (4) illustrates how internal voltage deviations in voltage stacked systems relate to the mismatch in current. For nMOS switches of with quadratic current vs voltage relationship ($\alpha \approx 2$) (Figure 6(a)), the shape of VDDM and VDDL curves result in square root shape as seen in Figure 8. The equations also apply to synchronous CMOS logic blocks ($I \propto CVF$), where K in equation (1) corresponds to aggregate switched capacitance of the logic block, rather than transistor width. The value of α , which determines ΔV vs ΔI relationship, depends on clocking strategy for synchronous logic blocks. Constant frequency operation, corresponding to $\alpha \approx 1$ for equation (1), ignoring capacitance dependence on voltage, yields a close to linear ΔV vs ΔI relationship. On the other hand, current in self-clocked circuits[7], where operating frequency varies with voltage, varies more strongly with voltage, which would result in ΔV vs ΔI relationship more akin to the square root




Figure 10: Measured V_{TOP} values for various initial power consumption conditions for inter-layer current skew (current increased only in top layer)

relationship of the test-chip prototype.

Unlike traditional power delivery schemes where circuitry has low-impedance connection to fixed supply rails, voltage stacking suffers from relatively high impedance paths for one or both rails, depending on the layer. Hence, the impedance of the internal supply rails of a voltage stack layer is highly dependent on the aggregate power consumption of the cores. Figure 10 plots the voltages that V_{TOP} settle to vs normalized inter-layer current difference when the *current load* settings are skewed by increasing only the top layer current, for three different values of balanced current I_{BAL} . Higher initial I_{BAL} values (higher chip power consumption), result in smaller voltage deviation for the same inter-layer current mismatch, due to lower impedance of the internal supply rails. This can also be seen from equation (4), where higher current consumption (for the same voltage conditions) translate to higher K value, leading to a smaller ΔV for the same change in current $\Delta I(\Delta K)$. This suggests that voltage stacked systems become more immune to inter-layer current mismatches as overall power consumption is increased.

This dependence is illustrated in Figure 11 by measuring the internal voltage variation in the test-chip when pseudo-random noise is injected into the array for different balanced current values. Noise is injected by toggling *current load* bits. Real-time internal voltage trace of $VDDM$ and $VDDL$, along with histogram plots of V_{TOP} , V_{MID} and V_{BOT} are shown. Pseudo-random noise of fixed magnitude was injected into all three stack layers for four different balanced *current load* settings. It can be clearly seen that as power consumption is increased, voltage stack becomes more robust to noise, demonstrating that the voltage variation depends on current mismatch to current consumption ratio, not the absolute value of current mismatch. This attribute suggests that for voltage stacked multicore systems, as the number of cores are increased, power fluctuation of a single core due to architectural events such as cache miss will have lesser effect on internal voltage. For many-core systems, the worst case scenario, when all cores in a stack layer experience power fluctuations at once can be easily avoided by workload scheduling via software. This suggests that for many-core voltage stacked systems, coordinated workload scheduling could provide efficient voltage regulation during

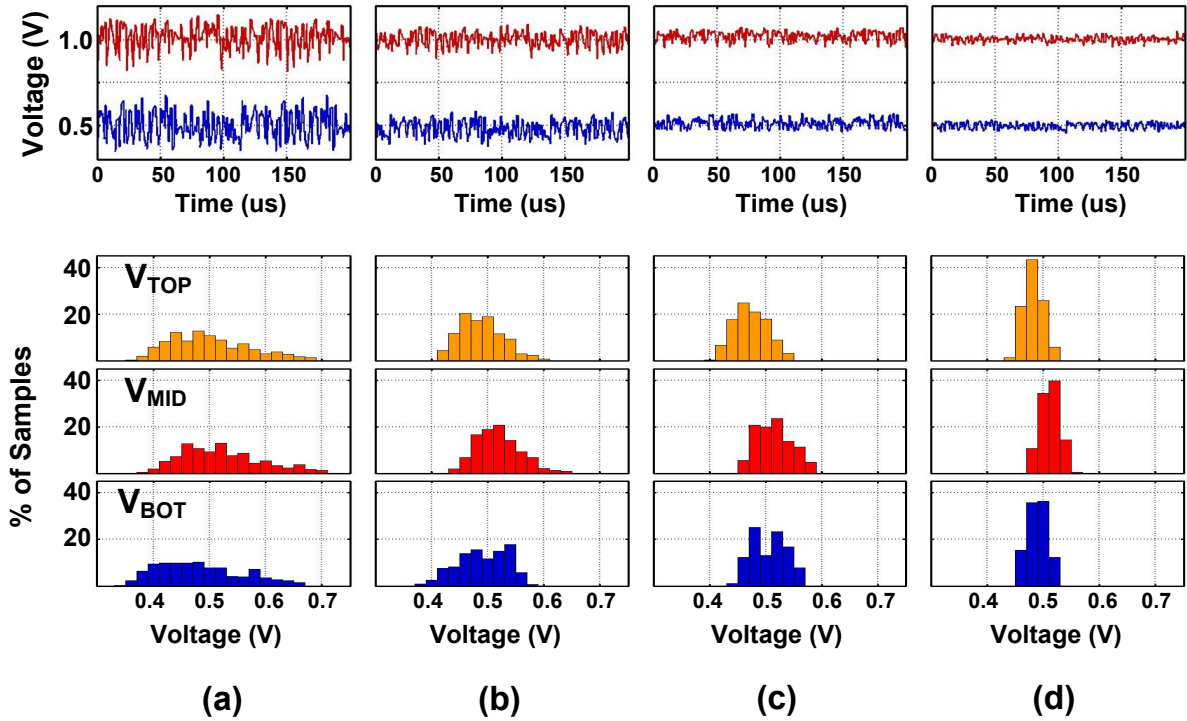


Figure 11: Measured VDDM and VDDL voltage traces (top) and histogram of V_{TOP} , V_{MID} and V_{BOT} when pseudo-random fixed-magnitude noise is injected in all three stack layers for balanced *current* load setting of (a) 3, (b) 8, (c) 15, and (d) 31

high-power states while on-die regulators may only be relied on during low-power states.

5. CONCLUSION

In this paper, an evaluation of voltage stacking is presented in the context of near-threshold multi-core computing. Key attributes of voltage stacking are shown using results from a test-chip prototype built in 150nm FDSOI CMOS. While voltage stacking simplifies off-chip power delivery, it adds extra complexity to combat within-die voltage noise caused by inter-layer current imbalance. Current mismatch between layers must be tightly regulated, especially for near-threshold systems more susceptible to voltage variation. The internal supply rail impedance in voltage stacked systems depends on aggregate power consumption of the stack layers unlike conventional power delivery schemes. Thus voltage noise depends on the current mismatch to chip current consumption ratio rather than the absolute magnitude of the mismatch, suggesting that voltage stacking will scale well with future many-core systems, where power fluctuations in a single core has less impact on voltage noise.

Acknowledgements

This work was partially supported under SRC contract with task ID 1973, as well as partially supported by National Science Foundation grants CCF-0903437, CSR-0720566 and CCF-0702344. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the SRC

or NSF. We also thank MIT Lincoln Labs for fabrication support.

6. REFERENCES

- [1] E. Alon and M. Horowitz. Integrated regulation for energy-efficient digital circuits. *IEEE J. Solid-State Circuits*, 43(8):1795–1807, August 2008.
- [2] L. Chang, R. K. Montoye, B. L. Ji, A. J. Weger, K. G. Stawiasz, and R. H. Dennard. A fully-integrated switched-capacitor 2:1 voltage converter with regulation capability and 90% efficiency at 2.3A/mm². *IEEE Symp. VLSI Circuits*, pages 55–56, June 2010.
- [3] R. G. Dreslinski, M. Wiecekowsi, D. Blaauw, D. Sylvester, and T. Mudge. Near-threshold computing: reclaiming Moore’s law through energy efficient integrated circuits. *Proceedings of the IEEE*, 98(2):253–266, February 2010.
- [4] H. Esmailzadeh, E. Blem, R. St.Amant, K. Sankaralingam, and D. Burger. Dark silicon and the end of multicore scaling. *Proc. 38th Int’l Symp. Computer Architecture*, ACM Press, pages 365–376, June 2011.
- [5] W. Kim, D. Brooks, and G.-Y. Wei. A fully-integrated 3-level dc/dc converter for nanosecond-scale dvs with fast shunt regulation. *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pages 268–270, February 2011.
- [6] M. Nagata, T. Okumoto, and K. Taki. A built-in technique for probing power supply and ground noise distribution within large-scale digital integrated circuits. *IEEE J. Solid-State Circuits*, 40(4):813–819, April 2005.
- [7] L. S. Nielsen, C. Niessen, J. Sparso, and K. van Berkel. Low-power operation using self-timed circuits and adaptive scaling of the supply voltage. *IEEE Transactions on VLSI Systems*, 2(4):391–397, 1994.
- [8] S. Rajapandian, K. L. Shepard, P. Hazucha, and T. Karnik. High-voltage power delivery through charge recycling. *IEEE J. Solid-State Circuits*, 41(6):1400–1410, June 2006.