

A Programmable Resistive Power Grid for Post-Fabrication Flexibility and Energy Tradeoffs

Kyle Craig^{1,2}, Yousef Shakhsher¹, Sudhanshu Khanna¹, Saad Arrabi¹, John Lach¹,
Benton H. Calhoun¹, and Stephen Kosonocky²

¹Dept. of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA, USA

²Advanced Micro Devices, Fort Collins, CO, USA

ABSTRACT

This paper explores the benefits of splitting a monolithic power gate transistor into parallel, independently controlled, variable weighted power gates to provide programmable post-fabrication power grid resistance. This power gate topology creates energy saving opportunities by providing adjustable localized voltages during active modes and reducing leakage current in idle blocks while retaining data. Measurements show over 30% active energy savings per operation and 90% savings in idle current with retention. A modeling flow for a resistive power grid was also developed that demonstrates the effectiveness of this approach in a Bulldozer processor core.

Categories and Subject Descriptors

1.1 Technologies and Digital Circuits

Keywords

Variable Weighted Headers, Low Power Design, Leakage, Dynamic Voltage Scaling

1. INTRODUCTION

This paper proposes a variable resistance power grid constructed from variable weighted parallel power transistors. This structure provides a dynamic voltage scaling (DVS) approach for active mode at much lower cost than using DC-DC converters and offers a convenient low leakage standby mode with data retention.

Many contemporary and emerging applications such as laptop computers, portable media players, smart phones, and bio-medical devices impose strict constraints on energy consumption due to battery size, yet also demand high performance for short bursts of time to meet timing constraints. Although technology scaling has provided raw performance gains and lower switching capacitances, increasing the battery life by lowering system energy consumption is still an ongoing effort, especially due to increasing energy consumption from leakage current.

Power gating and dynamic voltage scaling are two common solutions to reduce leakage current during standby mode and to tradeoff dynamic energy and delay during active operation, respectively. Power gating uses large transistors in series with the

power supply or ground to cut off leakage during idle mode. One disadvantage of power gating is that data stored in registers is lost. A variety of approaches to deal with this problem include putting registers on a separate supply, using high V_T balloon registers in parallel with core registers, or other alternative dual V_T register circuits. All of these incur overhead and added design complexity. DVS during active mode saves power by lowering the frequency and voltage together when timing slack exists. Applying DVS to multiple blocks requires multiple DC-DC converters that adjust the local voltage levels or alternative schemes that allow local V_{DD} selection from among multiple regulated supplies [1]. Again, the overhead of these approaches can be substantial.

This paper explores an approach for providing dynamic system level flexibility by partitioning large, monolithic power gating transistors into parallel, independently controllable power gates with different widths in a scheme called *variable weighted power gates*, as shown in Figure 1. This structure allows a once monolithic power gate to be partially on and/or to vary its effective size dynamically, providing both boot-time programmability and run-time adaptability that enable each block of a system to independently approach an ideal header resistance configuration for all phases and modes of all applications. We can leverage this functionality to implement fine grained DVS at the block level and to enable block level low leakage standby modes with data retention.

The authors of [2] first introduced separating power gate FETs and controlling them individually as a method to reduce noise. This paper expands on this idea and shows how to apply variable weighted power gates in a flexible fashion to provide additional low power operation modes. A variable weighted power gate can provide a controlled power supply resistance to enable a large number of effective voltages at which a block can operate. The block is therefore not constrained by available voltage rails; it can operate at lower voltages than the rest of the system without changing the entire chip V_{DD} , so it avoids the high overhead of

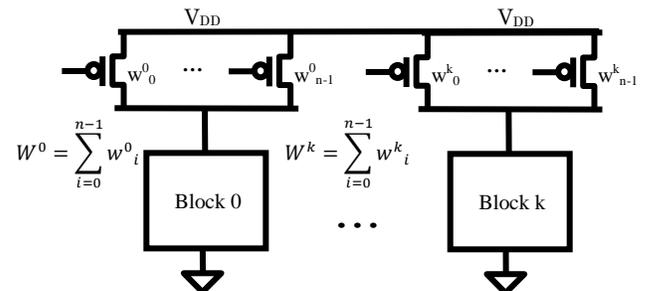


Figure 1. Proposed variable weighted power gate scheme. Multiple header widths are connected in parallel to the same rail with separate gate controls.

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, or 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, California, USA.

Copyright 2012 ACM 978-1-4503-1249-3/12/07...\$10.00.

extra DC-DC converters. In addition, a low energy standby mode can be provided that reduces leakage current in idle blocks but enables data retention and incurs lower overhead when returning to normal operation than does full power gating. A similar leakage reduction mode was used to give state retention modes in SRAM arrays [3] for reducing idle power.

Extensive simulation results from both a 32nm and 90nm technology demonstrate the effectiveness of variable weighted power gates to create multiple effective voltages through controlled power supply resistance. Measurements from a 90nm bulk test chip and measurements from a 32nm SOI x86 processor show the effectiveness of this technique for dynamic energy savings and leakage reduction respectively. Finally, a modeling methodology was created to enable design and verification of a full commercial processor utilizing a programmable power grid.

2. PROGRAMMABLE POWER GRID RESISTANCE FOR LOCAL DVS

As discussed above, circuits that utilize DC-DC converters for DVS incur significant delay and energy overhead for each voltage transition. Internal regulators require significant area, and external regulators require dedicated V_{DD} pins that increase package and board costs. As a result, DVS schemes usually employ global DVS in which the entire circuit or multiple cores operate at the same voltage.

Variable weighted power gates can be utilized as a controlled resistance to efficiently provide local voltage control without adding new DC-DC regulators, changing the voltage output of the existing regulators, or adding metal routing complexity. These power gates provide the effective voltage by utilizing the voltage drop across the power FET as a controlled resistor. The virtual supply rail (V_{rail}) decreases voltage as the effective power gate width decreases. As seen in Figure 2, there is no feedback loop to adjust the output voltage. Rather, the voltage rail is allowed to settle during operation, thus providing an output voltage on the voltage rail for the circuit that is proportional to the circuit current load and the power gate effective resistance. This virtual rail voltage not only depends on the size of the power gate, but also on the activity factor of the block and extrinsic and intrinsic decoupling capacitance on the virtual rail. This virtual rail voltage sets the delay and energy of the operation.

Assuming frequent operation and no idling, the virtual rail will droop to a voltage below the nominal V_{DD} . Figure 3 shows a 90nm CMOS simulation of the virtual voltage rail of successive operations of a 32b Kogge Stone adder for different fraction of enabled PMOS power gate widths, i.e., the width of enabled power gates enabled divided by the total width of power gates in the

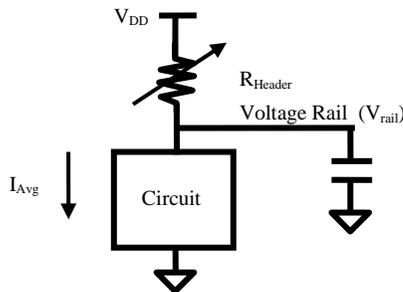


Figure 2. Variable weighted power gate used as a power grid resistance. R_{Header} varies based on header size.

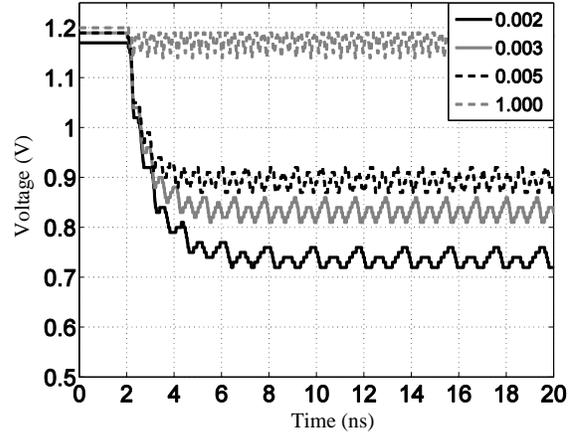


Figure 3. Virtual rail of a 32b Kogge Stone adder during successive adder operations over several header widths.

design, including the extracted virtual rail capacitance and an ideal V_{DD} applied to the circuit. Notice that the rail settles near a certain voltage for each power gate width. In this operation, the energy is reduced by up to 37% in simulation as compared to circuit operation tied directly to V_{DD} with no power gate with a maximum increase of $\sim 2.8x$ in delay for frequent operations. This effectively implements a light weight DVS mechanism with no additional circuits required (except control), which can be used in lieu of or in addition to conventional DVS methods.

2.1 Activity Factor

Although the effective V_{DD} of the block is set by the unregulated virtual rail voltage, analysis of the block under maximum current conditions can allow us to select the header width to set the worst case circuit performance. When the circuit consumes lower amounts of current, the virtual voltage, V_{rail} , will not droop as far. This will cut into the active energy savings, but the scheme still saves energy compared to having the full header on. Since the overhead for implementing this scheme is so low, the savings essentially come for free, since the power gates are already utilized in these designs to reduce leakage.

To model the effect of activity factor, a similar methodology as presented in [4] was used. Sixty-four ring five stage oscillators (RO) were simulated in parallel with the ability to disable different ROs with enable signals. With this setup, 64 enabled ROs corresponds to the highest activity factor of 1.0, 32 enabled ROs corresponds to an activity factor of 0.5 and so on. Each of the different activity factors were simulated with a varied amount of enabled header width. Figure 4 shows the impact of activity factor and header width on V_{rail} . The x-axis is the number of enabled parallel ROs while each bar represents a different amount of total header width enabled. The headers partitions were sized for the 1-RO case, and the same sizes were used for all other RO cases. For the 1-RO case, we are able to see a large potential range of V_{rail} values ($\sim 0.6-1.15V$), however as the activity factor increases the V_{rail} range shrinks. For the high RO cases, larger headers widths need to be designed and enabled to regain a large V_{rail} range. In Figure 5, RO performance is shown with varying activity factor and header width. The RO performance is normalized to a RO without any power gates at the nominal V_{DD} . As expected the performance is highly dependent on activity factor and enabled header width and shows similar trends as V_{rail} .

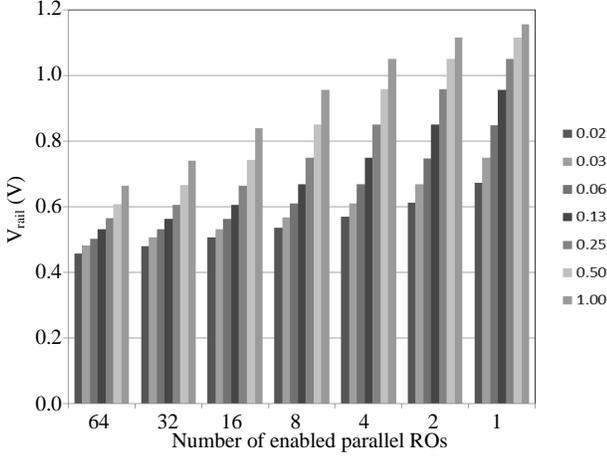


Figure 4. V_{rail} values for varying number of enabled ring oscillators for different normalized header widths

Traditional header sizing methodology still applies, meaning that the total header size needs to be based off the worst case activity factor and target performance requirement. The header width partitioning is highly dependent on activity factor. In order to achieve a flexible power grid resistance, the header width partitioning should be sized for the range of activity factors expected and the performance range desired based on application. This would involve running simulations varying header widths for different characteristic activity factors to characterize the V_{rail} drop, performance degradation and energy savings. This could be done with low level circuit simulations or high level model techniques as discussed later.

2.2 Energy Savings

This section quantifies the energy savings that the resistive grid scheme can provide when used to implement DVS. The following equation for energy per operation is composed of both the dynamic and leakage energy as a function of V_{DD} (the short circuit component is ignored for simplicity):

$$E_{op}(V_{DD}) = C_{eff}(V_{DD}) * V_{DD}^2 + V_{DD} * I_L(V_{DD}) * t_{op}, \quad (1)$$

where C_{eff} is the effective switching capacitance as a function of V_{DD} , and t_{op} is the delay of the component. I_L is the subthreshold leakage current as a function of V_{DD} . Simple modifications can be made to (1) to include the savings of variable weighted power gates by making the equation a function of both V_{DD} and V_{rail} :

$$E_{op}(V_{DD}, V_{rail}) = C_{eff}(V_{rail}) * V_{DD} * V_{rail} + V_{DD} * I_L(V_{rail}) * t_{op}. \quad (2)$$

As V_{rail} decreases, we expect a greater than linear dynamic energy reduction due to the linear reduction of $V_{DD} * V_{rail}$, plus a less than linear reduction in C_{eff} , due to device source and drain parasitic junction capacitance being dependent on V_{rail} . Finally, as V_{rail} between V_{DD} decreases, I_L will decrease due to the exponential and the amount of DIBL (drain-induced barrier lowering). To highlight the potential benefits of using this droop to save energy, we

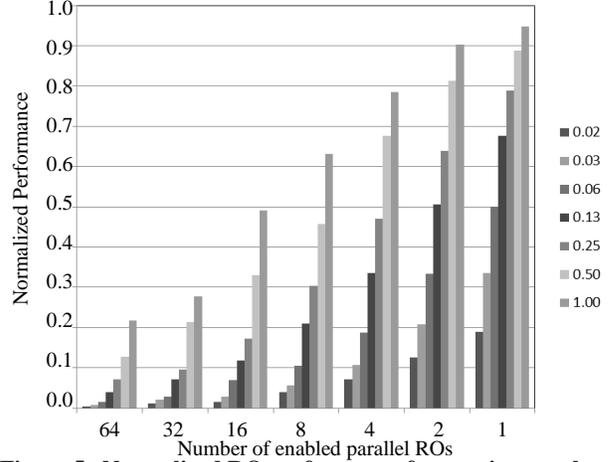


Figure 5. Normalized RO performance for varying number of enabled ROs for different normalized header widths

connected a ring oscillator to variable weighted power gates and allowed the rail to settle for each power gate width.

Figure 6 shows the normalized energy and delay versus the normalized power gate width. As expected, reducing the power gate width decreases the energy and increases the delay. These RO results were confirmed through silicon measurements using a 90 nm commercial bulk technology with four 97 stage ROs in parallel consisting of inverters and delay cells to simulate a high current and activity load. The normalized measured values match the simulated values and show an energy savings of over 30% in silicon.

Additionally, as with any DVS scheme, there is an overhead associated with this variable weighted power gate technique. The virtual rail must eventually recover from the droop during operations when high performance is required. This recovery energy can be defined as:

$$E_{rec} = C_{rail} * V_{DD} * (V_{DD} - V_{rail}). \quad (3)$$

Rail recovery during successive operations should be minimized to save the maximum amount of energy. To minimize this wasteful energy, the clock frequency should be increased, the amount of the

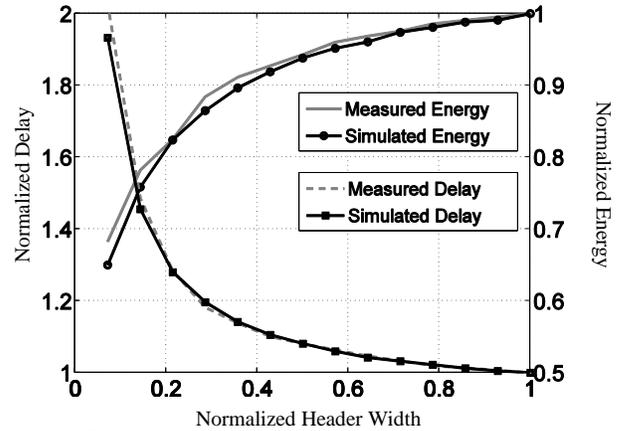


Figure 6. Simulated and measured energy and delay for a ring oscillator with sweeping header size in a 90 nm test chip.

power gate that is enabled should be decreased, or the source voltage should be decreased to reduce the rail's ability to recover. This will allow for most energy efficient operation. We can amortize this recovery energy over multiple cycles by maximizing the number of operations run with this voltage droop, thus reducing energy per operation when the virtual rail is relatively constant.

2.3 Leakage Current Reduction

Many systems and blocks within systems spend large amounts of time idling. Additionally, blocks such as register files or memory may need to retain their data, which is not supported by most power gating schemes. Variable weighted power gates provide a low energy solution that enables data retention with reduced idle leakage current. Leakage current can be reduced through reducing power gate size by dropping the voltage across the active devices which reduces the amount of DIBL, thus reducing device leakage. In Figure 7, the leakage current is measured for a 32nm SOI four-core x86 processor SOC chip which has a variable weighted power gate ring around the core [5] on actual silicon hardware. By changing the power gate width by disabling distributed sections of the footer ring via configuration bits, the idle current can be gradually reduced from 100% to a lowest bound of 10%. Tester characterization can determine the lowest setting to allow core state retention. These measurements confirm that the variable width header scheme supports very low leakage standby modes at low overhead.

3. OPPORTUNITY FOR REGULATION

We investigated the opportunity for dynamic grid voltage control in a commercial x86 four core processor SOC using typical P-state (power state) occupancy data. A P-state defines a voltage/frequency pair independently for each core in the processor, where P0 is the fastest state and P3 is the slowest. To analyze the potential savings of our approach deployed on a larger scale, we applied a model of our variable weighted power gate resistive grid technique (based on equation 2) to our processor simulations. Since all cores share a common V_{DD} , the lowest core P-state sets the operating V_{DD} for all cores, leading to non-optimal V_{DD} for cores with higher P-states. The cores running at a lower P-state than required are only able to use frequency scaling in the absence of a resistive grid, since there is no other mechanism to lower the local V_{DD} . However, Figure 8 shows the opportunity for power savings using variable-weighted power gates. In the figure, the label P1@P0 indicates the total power of core(s) that are running at the P1 frequency while another different core in the SOC is running at a P0 state. During the SysMark trace, the core-wise P-state occupancy is determined by the operating system. By including variable weighted power gates at each core, the cores running at a higher P-state will run at their near optimal V_{DD} during periods of high activity.

The figure shows up to ~15% power savings opportunity by using the variable weighted power gate resistive grid technique. Power/performance results can vary depending on the P-state frequency, voltage settings, and the profile of system activity. By allowing individual core-wise voltage settings, the system has more flexibility to differentiate high performance modes from lower performance modes, which can allow opportunity for additional performance boosting when a single core is running at a low P-state. Definition of the P-state voltage-frequency pair can be determined at SOC characterization time while running a thermal

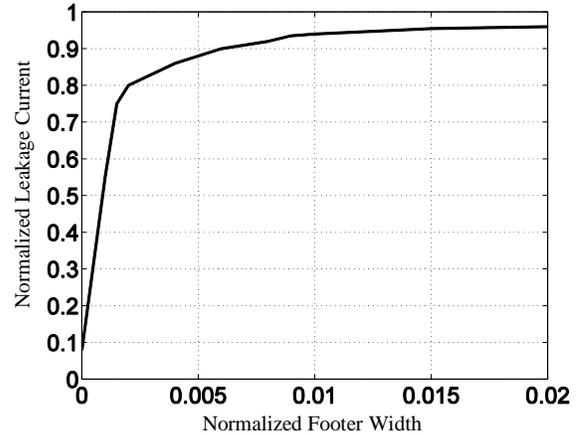


Figure 7. Measured idle current reduction using variable weighted power gates in a 32nm 4-core x86 processor at 1.2V.

design point (TDP) workload in each core to ensure enough voltage margin to accommodate the maximum P-state frequency, at the expense of some power savings.

4. MODELING POWER SUPPLY RESISTANCE

Since the application of this technique requires characterization of the power supply resistance, we developed a design flow, using a commercial power integrity tool, for applying the approach to arbitrary digital designs. We use this design flow to model a full commercial processor using the proposed method for implementing a programmable resistive power grid.

4.1 Route Level Macro Model

A commercial power integrity tool, Apache Redhawk [6], was used to model the effectiveness of variable weighted power gates as a controlled power supply resistance in a large system. In order to demonstrate the feasibility of a controlled power supply resistance to create local effective voltages, a Redhawk set up was created for a single power gated Route Level Macro (RLM) used in a commercial 32nm core. An ideal V_{DD} and V_{SS} were applied to the power grids, which were modeled from metal 11 (M11) to M9 and M11 to M1 respectively, while the virtual V_{DD} grid, modeled from M8 to M2, was observed as the power gate width was varied. The

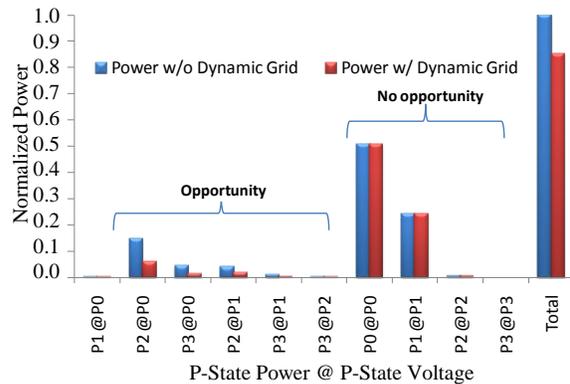


Figure 8. SysMark Trace Segment for a four-core x86 SOC

internal net activity was generated from benchmark simulations of a thermal design point (TDP) benchmark. Figure 9 shows the average virtual V_{DD} and worst case virtual V_{DD} observed for the TDP benchmark. The relatively flat part of the curves is due to the very low impedance sizing required for the maximum header size to enable high frequency operation at the maximum V_{DD} . To achieve variable power gate regulation, fine grain power gate partitioning is needed in the sub 0.2% range of the total power gate width.

4.2 Bulldozer Core Model

Finally, Redhawk was used to model the AMD Bulldozer core [7], which has a similar power gate ring structure seen in [5]. For this simulation, to prevent the simulation time from being prohibitively large, each Route Level Macro (RLM) (roughly 50 in total) in the core was modeled as a time dependent current source and capacitance model, with the exception of the L1 cache and two RLMs without available data. These current profiles were generated from simulations of the Double-precision General Matrix Multiply (DGEMM) benchmark.

A simplified package model was included to capture the real RLC effects seen on hardware. In these simulations, we used power gating footers instead of headers, so the Bulldozer model included real V_{DD} and V_{SS} grids from the C4 bumps down to M10, and virtual V_{SS} (using footer power gates) in M11 and M10. Figure 10 shows a simplified diagram of the setup. This test setup allowed flexibility to vary the power gate width and to observe V_{DD} , V_{SS} , and the virtual V_{SS} during a 25ns window of the DGEMM benchmark. Figure 11 shows the average V_{DD} response to the applied current models in time, showing that our current model profile was functioning correctly.

Figure 12 shows the V_{DD} and virtual V_{SS} voltage profile over time for different normalized power gate widths. The trend in this figure is the same as in Figure 9 the largest change in the virtual V_{SS} happens at very low footer widths. Notice that at 5.07%, the virtual V_{SS} is only slightly above the 100% case, which is expected for a power gate ring system designed for a high performance core.

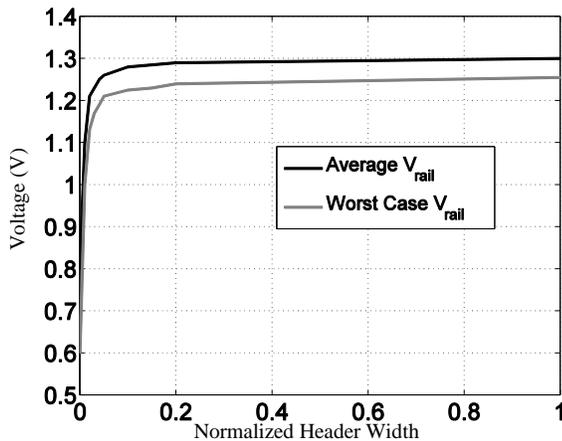


Figure 9. Average and Worst Case V_{rail} of the single RLM during a TDP benchmark.

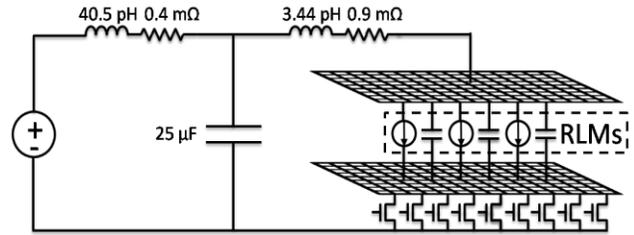


Figure 10. Bulldozer test setup, include RLC package model, RLM current and capacitance models

In this example, dynamic control is required for only a small percentage of the total power gate width. This figure also includes every RLM's V_{DD} and virtual V_{SS} superimposed into a single graph.

The negligible variance in V_{DD} and V_{SS} between RLMs across the chip is due to the robust power grid with low resistance showing that the dominant factor in the virtual V_{SS} droop is caused by the controlled power supply resistance of the power gates. Through variable weighted power gates, we are able to achieve a wide range of virtual V_{SS} supplied to the core.

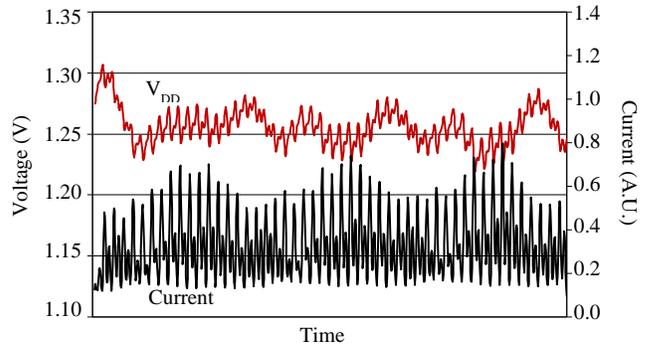


Figure 11. Power grid simulation of V_{DD} and Current profiles over time for AMD Bulldozer core.

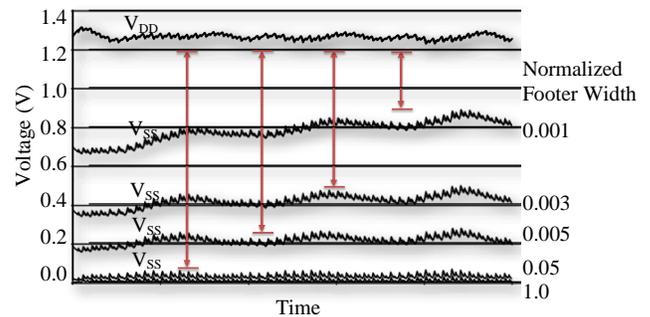


Figure 12. V_{DD} and virtual- V_{SS} across the Bulldozer core during DGEMM benchmark with different power gat widths.

5. CONCLUSIONS

This work demonstrated the use of variable weighted power gates to provide a controlled power supply resistance for designs using power gating and DVS. We showed how a controlled power grid resistance through variable power gate resistance can be used to provide a large number of voltages to a block during active operation. Through extensive simulation and modeling, we can select the correct amount of width to turn on for various desired operating modes. We discussed the opportunity for using the variable weighted power gate in a commercial x86 four core SOC. Finally, through measurements from a 90nm test chip and a 32nm x86 processor, we showed how variable weighted power gates can be used to reduce energy during active mode and to limit leakage current, respectively. Using variable weighted power gates is a low cost solution for providing both boot-time programmability and run-time adaptability that enables each block of a system to approach an ideal configuration for all phases and modes of all applications.

6. REFERENCES

- [1] Shakhsheer, Y., et al., A 90nm Data Flow Processor Demonstrating Fine Grained DVS for Energy Efficient Operation from 0.25V to 1.2V. *Custom Integrated Circuits Conference*, Sept. 2011
- [2] Truong, D.N., et al., A 167-processor computational platform in 65 nm CMOS, *IEEE Journal of Solid-State Circuits*, vol. 44, no. 4, pp. 1130-1144, April 2009
- [3] Zhang, K., et al., SRAM design on 65-nm CMOS technology with dynamic sleep transistor for leakage reduction, *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 895-901, Apr. 2005.
- [4] Wang, A, et al., *Sub-threshold Voltage Circuit Design for Ultra-Low Power Systems*. Springer, New York, NY
- [5] Jotwani, R., et al., An x86-64 Core Implemented in 32nm SOI CMOS *International Solid-State Circuit Conference*, pp. 106-107, 2010.
- [6] Apache Redhawk, “ <http://www.apache-da.com/products/redhawk>”
- [7] Fischer, T., et al., Design Solutions for the Bulldozer 32nm SOI 2-Core Processor Module in an 8-Core CPU, *International Solid-State Circuit Conference*, pp. 78-79, 2011.