

Voltage-Stacked Power Delivery Systems: Reliability, Efficiency, and Power Management

An Zou, *Student Member, IEEE*, Jingwen Leng, Xin He, *Member, IEEE*, Yazhou Zu, Christopher D. Gill, *Senior Member, IEEE*, Vijay Janapa Reddi, Xuan Zhang, *Member, IEEE*

Abstract—In today’s manycore processors, energy loss of more than 20% may result from inherent inefficiencies of conventional power delivery system (PDS) design. By stacking multiple voltage domains in series to lower the step-down conversion ratio of the off-chip voltage regulator module (VRM) and reduce energy loss along the path of the power delivery network (PDN), voltage stacking (VS) offers a novel alternative power delivery technique to fundamentally improve power delivery efficiency (PDE). However, voltage stacking suffers from aggravated supply voltage noise from current imbalance, which hinders its adoption. In this paper, we investigate practical voltage stacking implementation in manycore processors to improve power delivery efficiency (PDE) and achieve reliable performance, while maintaining compatibility with advanced power management techniques. We first present the system configuration of a voltage-stacked manycore processor. We then systematically characterize supply voltage noise in voltage stacking, identify global and residual differential currents as its dominant contributors, and calculate the possible worst supply voltage noise. We next propose a hybrid voltage regulation solution, based on a charge-recycling off-chip voltage regulator and distributed integrated voltage regulators, to mitigate supply voltage noise effectively. We also study the compatibility of voltage stacking with higher level power management techniques. Finally, the performance of a voltage-stacked GPU system is comprehensively evaluated. Simulation results show that our approach can achieve 93.5% power delivery efficiency, reducing the power loss by 13.6% compared to conventional single-layer power delivery system.

Index Terms—Power Delivery System, Manycore Architecture, Voltage Stacking, Supply Noise, Integrated Voltage Regulator

I. INTRODUCTION

Computers consume a non-trivial proportion of the total electricity energy both globally and in the U.S [1], [2]. For example, it is predicted that the power consumption of world data centers alone will approach 1,000 TWh within a decade, which is more than the amount now consumed for all purposes by Japan and Germany combined [3], [4]. A closer examination of the power delivery path in modern computing systems reveals a provocative finding: transmitting

Manuscript received July 23, 2019.

An Zou, Xin He, Christopher D. Gill and Xuan Zhang are with the Washington University in St. Louis, St. Louis, MO, 63130 USA, e-mail: {anzou@wustl.edu, hex0102@gmail.com, cdgill@cse.wustl.edu, xuan.zhang@wustl.edu}

Jingwen Leng is with the Shanghai Jiao Tong University, Shanghai, 200240 China, e-mail: {leng-jw@sjtu.edu.cn}

Yazhou Zu is with the University of Texas at Austin, Austin, TX 78712 USA, e-mail: {yazhou.zu@utexas.edu}

Vijay Janapa Reddi is with the University of Texas at Austin, Austin, TX, 78712 USA and the Harvard University, Cambridge, MA, 02138 USA, e-mail: {vj@eecs.harvard.edu}

Xuan Zhang and Jingwen Leng are the corresponding authors of this paper.

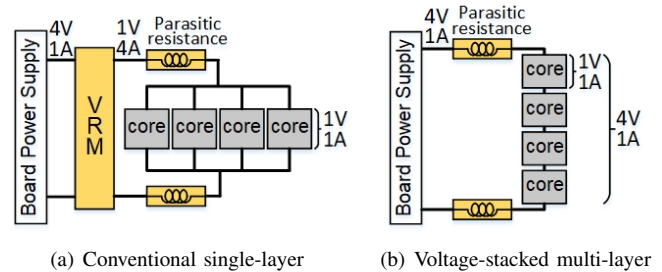


Fig. 1: Conventional single-layer and voltage-stacked multi-layer power delivery system. (PCB board voltage: 4V; each core requires 1V voltage and 1A current)

and distributing electricity across tens or hundreds of miles in the grid to reach the power plug incurs only a 6% power loss [5], whereas delivering the power for “the last centimeter” from the PCB board to the manycore processor chip can waste more than 20% of the power [6]–[8]. Thus, improving the efficiency of the last power delivery stage of today’s manycore processors yields not only very large economic benefits, but also a smaller carbon footprint and environmental benefits.

Despite the importance of improving manycore processor power delivery efficiency (PDE), the energy loss in a conventional single-layer power delivery system (PDS), such as that shown in Fig. 1(a) is difficult to eliminate. Three main inefficiency sources are directly associated with the PDS. The first is the voltage conversion loss incurred in converting a higher supply voltage at the board level to a lower supply voltage required by the microprocessor [9]. The second is the power delivery network (PDN) loss in parasitic resistance in transferring the electron charges from the off-chip power source to the distributed on-chip computing units [10], [11]. The third is the supply voltage margin to accommodate supply voltage noise and process variation. Generally speaking, the three inefficiencies become worse with lower supply voltages, increasing power density, and higher power ratings. Although various techniques have been proposed in prior work to reduce PDN loss by moving the voltage regulation closer to the point-of-load [12], [13], they are not capable of addressing the inefficiencies from the voltage regulator simultaneously, and thus are fundamentally unable to close the efficiency gap.

Voltage stacking shown in Fig. 1(b), also known as charge recycling [14] or multi-story power delivery [15], is a novel technique that allows efficient power delivery through a single high voltage source to multiple serially-stacked voltage domains. Due to the inherent voltage division among the voltage domain in series, it obviates the need for step-

down voltage conversion and reduces the currents flowing through the PDN. Ideally, if the current loads from all the voltage domains are perfectly balanced, then the input voltage is evenly divided with no supply voltage noise fluctuation. Voltage stacking's theoretical peak power delivery efficiency under balanced power activity is close to 100%, making it an attractive solution. However, in real applications, voltage stacking is seriously limited by its exacerbated supply voltage noise caused by the current imbalance between the serially-stacked voltage domains [16]. This limitation prevents wide adoption in practical systems that require consistent and reliable operation. In this paper, we systematically investigate the feasibility and potential benefits of applying voltage stacking to a graphic processing unit (GPU) processor to improve its power delivery efficiency.

II. BACKGROUND AND RELATED WORK

A. Power Delivery System

The power delivery system (PDS) in modern processors consists of a step-down voltage regulation module (VRM) on the motherboard; sockets, off-chip decoupling capacitors and electrical connections at the board, package, and chip levels in the form of PCB traces; and socket bumps and C4 bumps, where undesirable parasitic resistance and inductance reside. The decoupling capacitors (C) and the parasitic resistance (R) and inductance (L) along the connection path form the electrical model of the PDN in a computing system with a conventional PDS. To study the power delivery efficiency and system reliability, it is usually sufficient to assume the output of board level VRM is an ideal voltage source, and we adopt this convention in this work.

In a conventional setting, voltage conversion using a step-down VRM is necessary because the voltage level at the board is higher than the digital supply of a processor. Yet due to inherent inefficiency of step-down VRMs, energy is lost during the voltage conversion. Resistive parasitics along the PDN path also contribute to energy loss and incur voltage drop across the resistance, which is known as IR-drop. These two major efficiency losses can approach 20% or more in advanced technology nodes and under peak power operations.

Moreover, because of the non-ideal effect of the parasitic RLC network, electrons cannot be delivered instantaneously from the VRM output to immediately satisfy the fast changing current loads of various on-chip components. This lag results in on-chip voltage fluctuations and causes supply voltage noise reliability issues during operation.

B. Voltage Stacking

In voltage stacking (VS), the step-down VRM can be eliminated by serially stacking the voltage domains. It can be intuitively understood as allowing electron charges to recycle through the stacking layers in series. In addition to eliminating step-down conversion loss, voltage stacking lowers the PDN loss due to resistive parasitics, because in a N-layer voltage-stacked system, the PDN path current is reduced by $N\times$, which corresponds to $N^2\times$ reduction in power loss. These

efficiency improvements have been demonstrated in prior work [17], [18].

A theoretical peak PDE close to 100% can be achieved using voltage stacking [18] when all the stacking layers have balanced activities, and hence the same transient current demands. In practice, though, applying voltage stacking in real computing systems, where activity mismatches abound both spatially and temporally, proves to be challenging. As has been shown in previous studies, such activity mismatches can cause severe voltage fluctuations in a voltage-stacked system [8], [19]–[23]. The aggravated noise problem remains one of the most obstinate obstacles preventing voltage stacking adoption in the mainstream.

C. Supply Voltage Noise

Due to its impact on system reliability, supply voltage noise has been diligently studied and characterized for conventional single-layer PDS in single-core [24], [25], multi-core [26], [27], and manycore GPU processors [21], [28]–[31]. While circuit techniques such as load line compensation are effective at taming IR-drop induced noise [32], dynamic Ldi/dt noise, and resonance noise in particular, are more dominant and harder to tackle [33], [34], and often demand a cross-layer solution. However, a voltage-stacked manycore processor experiences more serious and complex supply voltage noise behavior due to the interactions between the cores that can lead to constructive or destructive noise composition. This aggravated supply voltage noise prevents the wide adoption of voltage stacking in mainstream computing systems, despite its higher power delivery efficiency. Up until now, it has been only intuitively understood that the supply voltage noise in voltage stacking is from an imbalanced workload, and a systematic supply voltage noise study of manycore processors with multi-layer voltage-stacked PDS is still lacking.

D. Power Delivery Efficiency

The underlying physical mechanism to convert and transfer electron charges from the higher supply voltage on the motherboard to the much lower supply voltage on the microprocessor chip invariably causes energy loss. The energy loss can be broken down into three parts:

First, energy is lost in voltage conversion to step down the supply voltage [9]. We define the conversion efficiency of a voltage regulator (η_{VR}) as the ratio between the power it delivers at the voltage regulator output over the power it consumes at the input. η_{VR} is usually a function of the step-down conversion ratio. A high performance off-chip switching VRM can deliver over 90% conversion efficiency, but the efficiency is degraded at a lower output voltage with a higher step-down ratio [35].

The second part of the energy loss occurs in the power delivery network mostly because of heat dissipation when current runs through the parasitic resistance that exists along the path of the PDN, which is related to the IR-drop component of supply voltage noise [10], [11]:

$$\eta_{PDN} = \frac{R_{core}(V_{core})}{(R_{PDN} + R_{core}(V_{core}))} \quad (1)$$




Fig. 2: Illustration of on-chip power routing for conventional and voltage-stacked power delivery configurations

where R_{PDN} represents the total parasitic resistance contributed by the PDN, and R_{core} represents the equivalent resistive impedance of the computational load as a function of V_{core} . The definition of R_{core} suggests that $R_{core} = V_{core}/I_{core}$. For fixed V_{core} value, R_{core} is a measure of the power rating.

The final and often overlooked part is the energy overhead incurred by raising the supply by a non-negligible voltage margin, $\Delta V = V_{core} - V_{min}$, to accommodate and sustain fault-free operation [36], [37]. We can express this component as $\eta_{\Delta V}$:

$$\eta_{\Delta V} = \frac{P_{core}(V_{min})}{P_{core}(V_{core})} = \frac{V_{min}I_{core}(V_{min})}{V_{core}I_{core}(V_{core})} \quad (2)$$

P_{core} and I_{core} represent the power consumption and the current load of the processor core as a function of the core supply voltage (V_{core} and V_{min});

Based on above analysis, the full power delivery system efficiency can be expressed as

$$\eta_{PDS} = \frac{P_{core}(V_{min})}{P_{src}} = \eta_{VR} \cdot \eta_{PDN} \cdot \eta_{\Delta V} \quad (3)$$

where P_{src} is the total power drawn from the source.

Voltage stacking can reduce the power loss in step-down voltage regulator by eliminating supply voltage conversion and power loss in power delivery network parasitic resistance by reducing path current. Besides, previous works [21], [38] further prove that voltage stacking can also diminish the voltage margin to further improve power delivery efficiency. In this paper, for a fair comparison, we assume voltage stacking has a same voltage margin with conventional single-layer power delivery system and mainly focus on its improvements in conversion efficiency (η_{VR}) and PDN efficiency (η_{PDN}).

E. Related Work

Proof-of-concept circuits [15], [39] and silicon prototypes [14], [17], [18], [40], [41] have been presented previously to explore voltage stacking using low-power microcontrollers, along with design methodology for floorplanning and placement [42], [43]. These pioneering works demonstrate the feasibility of voltage stacking, but they are often limited to simple assembly of uncorrelated cores with low power density. Inter-layer current imbalance has been discussed qualitatively [16] as a contributor to the supply voltage noise in voltage-stacked systems, but without rigorous quantitative derivation of worst-case conditions. To overcome supply voltage noise, most voltage stacking prototypes [14], [17], [18], [40], [41] resort to employing charge-recycling integrated voltage regulators (CR-IVR) to actively balance the current mismatches. However,

TABLE I: GPU voltage-stacked system configuration

Configuration	Value	Configuration	Value
PCB supply voltage	4.1V	SM core voltage	1V
No. of SM cores	16	Clock frequency	700M
Voltage-stacked layers	4	SM cores per layer	4
SM core ave power	5W	SM core max power	14W
Threads per SM core	1536	Threads per warp	32
Registers per SM core	128KB	Shared memory	48KB

the overhead and trade-offs from CR-IVR require further discussion and should be reduced.

Built upon these early prototypes, a number of novel approaches have been proposed to take advantage of voltage stacking under different scenarios, such as 3D-IC with varying TSV, on-chip decoupling capacitance, and package parameters [44], [45]; optimal system partitioning to unfold CPU cores [20], [38]; and GPU systems with supercapacitors [46] operating under near-threshold voltages [21], [38]. CoreUnfolding [20] is a novel method that voltage stacking can be used within each core. However, it is highly invasive as it requires separating function units inside the core to balance the groups of units. Voltage-Stacked GPUs [47], [48] models the power grid of the voltage-stacked GPU as a linear dynamic system to derive the power control strategy for supply noise guarantee, but the architecture-level power control scheme sacrifices the system performance for reliability.

III. SYSTEM CONFIGURATION

In this paper, we use the GPU system with a NVIDIA Fermi architecture as a representative manycore processor. Table I lists the configuration details of Fermi architecture. The Fermi architecture GPU has 16 streaming multiprocessor (SM) cores [49]. We use a 4×4 voltage stacking structure: 16 SM cores are stacked in 4 layers and each layer has 4 SM cores, as 4V is generally available on the board and SM cores require 1V. Our analysis and solutions are not limited to this 4×4 voltage stacking configuration and can be applied to other manycore processors with arbitrary voltage stacking configurations.

A. Power Grid Routing and PDN Modeling of VS

Voltage stacking can be implemented in both 2D and 3D-IC chips, but for a fair comparison with conventional power delivery methods, we focus on voltage stacking implementation in a 2D planar technology. To properly isolate the transistors in each voltage layer from the global substrate, voltage stacking often relies on advanced process technology such as triple wells or Silicon on Insulator (SOI) to establish local body




Fig. 3: Power delivery network (PDN) of a 2x4 voltage-stacked manycore processor

biasing voltages [16], [22], [23], [50], [51]. A hierarchical structure is used in 2D power routing, as shown in Fig. 2(a). The top metal layers are for global power grid which connects cores or modules. The next layers are local power grids connecting the function blocks such as ALU and Reg. Finally local power grids in the bottom metal layers connect to the logic gates. As illustrated by the power/ground routing scheme in Fig. 2(b) and 2(c), topologically stacking the voltage domains on a 2D chip can be achieved with minimal modifications by re-routing the top metal layers from parallel connections to series connections, leaving the local power/ground grids in the lower metals and the physical floorplans of the underlying blocks largely intact. Assuming this minimally-invasive routing method, we derive the voltage stacking PDN model shown in Fig. 3 based on the typical RLC circuits and parameters introduced previously to study GPU manycore processors [28], [31]. Note that there is parasitic resistance (R_S) between the vertically-connected cores (modeled by current sources), as depicted in Fig. 3. Our study focuses on the SM core power grid to clearly demonstrate the benefit of voltage stacking and evaluate the proposed hybrid regulation methodology, since its peak and average powers account for 80% and 93% of the total GPU chip power consumption [52]. Similar scheme can also be adapted to other on-chip components like SRAM in a voltage-stacked configuration [39].

B. Communication Across Layers

A voltage-stacked system suffers inherently complex communications across different voltage layers: instructions and data are communicated among memory, cache, and core registers, which are in different voltage layers. In the GPU system, SM cores do not directly communicate with each other, and the cross-layer communication mainly happens between SM cores and Memory Partition Units through an interconnection network. There are two interconnection networks with butterfly topology and 22 nodes: one for traffic from SM cores to Memory Partitions, and one for traffic from Memory Partitions back to SM cores. Cross layer communication requires extra level shifters added to the interconnection network. Several level shifter designs are suitable for a stacked architecture, such as capacitive-coupling-based (conventional) [53], two-stage cross-coupled (TSCC) [54]–[58], Wilson current mirror (WCM) [55], [57], [59], stacked Wilson current mirror (Stacked) [60], switched-capacitor (Tong) [22] and modified switched-cap (Mod-Tong). Tested by Ebrahimi [61] with input signal at 1GHz, Tong has the best energy-delay trade-off.




Fig. 4: Supply voltage noise decomposition

IV. SUPPLY VOLTAGE NOISE ANALYSIS

A. Supply Voltage Noise Characterization

Unlike previous empirical approaches [28], [31], we develop an analytical modeling framework to study and characterize supply voltage noise responses in voltage stacking PDN, especially in the presence of both correlated and uncorrelated core activities. The cornerstone of our analytical approach lies upon the decomposition and superposition principles in the fundamental circuit theory.

1) *Noise Decomposition & Superposition*: Since the basic electrical model of voltage stacking PDN consists of only linear components, including the RLC and ideal voltage and current sources, the superposition principle in linear systems generally holds, allowing us to decompose the core current to different components to reveal their distinctive characteristics. Without loss of generality, let us assume a voltage-stacked system that consists of N_L vertically-stacked layers with N_V cores on each layer. For example, Fig. 3 shows a $N_L = 2$ and $N_V = 4$ voltage-stacked system. The cores that align vertically are defined as a voltage stack. To facilitate later analysis, we adopt the s-domain expressions for current sources and give the following definitions:

$$I_{i,j}^{core}(s) = I^G(s) + I_i^{ST}(s) + I_{i,j}^R(s) \quad (4)$$

$$I^G(s) = \frac{\sum_{i=1}^{N_V} \sum_{j=1}^{N_L} I_{i,j}^{core}(s)}{N_V N_L} \quad (5)$$

$$I_i^{ST}(s) = \frac{\sum_{j=1}^{N_L} I_{i,j}^{core}(s)}{N_L} - I^G(s) \quad (6)$$

$$I_{i,j}^R(s) = \frac{(N_L - 1)I_{i,j}^{core}(s) - \sum_{k=1, k \neq j}^{N_L} I_{i,k}^{core}(s)}{N_L} \quad (7)$$

where $I_{i,j}^{core}(s)$ is the current contributed by the core in the i^{th} stack and the j^{th} layer. It is decomposed into three components: $I^G(s)$, $I_i^{ST}(s)$, and $I_{i,j}^R(s)$, in Eq. (4) - (7). $I^G(s)$ represents the global current component shared by all the cores, $I_i^{ST}(s)$ represents the common current components shared by the cores in the i^{th} stack, and $I_{i,j}^R(s)$ is the residual current components after removing the global and per-stack common terms. Now, the supply voltage noise at the core (in the i^{th} stack and the j^{th} layer) can be expressed by the current components working on their respective effective impedances, $Z_{eff,i}^G$, $Z_{eff,i}^{ST}$, and $Z_{eff,i,j}^R$, and causing superimposed supply voltage noises, $\Delta V_{core,i,j}^G$, $\Delta V_{core,i,j}^{ST}$, and $\Delta V_{core,i,j}^R$, as described in Eq. (8) and Fig. 4.

$$\Delta V_{core,i,j} = \Delta V_{core,i,j}^G + \Delta V_{core,i,j}^{ST} + \Delta V_{core,i,j}^R = I^G Z_{eff,i}^G + I_i^{ST} Z_{eff,i}^{ST} + \sum_{i=1}^{N_V} \sum_{j=i}^{N_L} I_{i,j}^R Z_{eff,i,j}^R \quad (8)$$

To illustrate how the decomposition and superposition in Eq. (4) - Eq. (8) help us analyze and characterize supply voltage noise effects in voltage stacking, we use a simplified RLC network of a 2×2 voltage stacking PDN, as shown in Fig. 5.

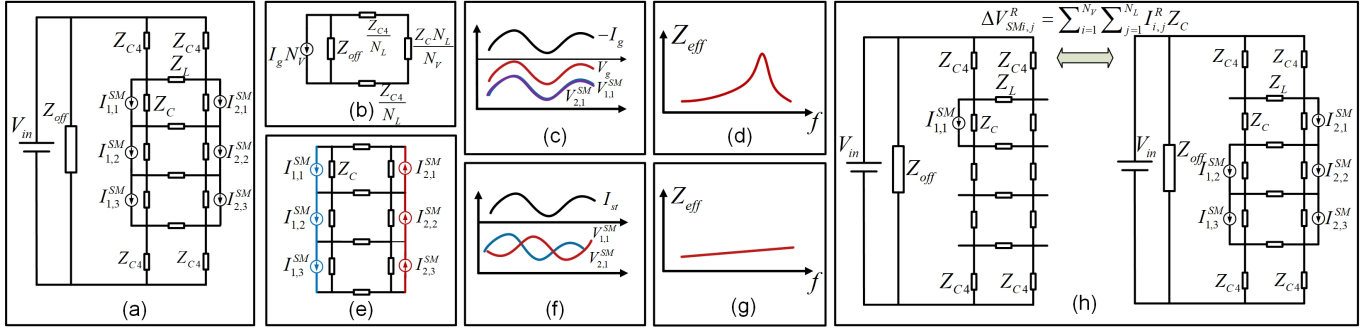


Fig. 5: Illustrative example for noise decomposition using 2×3 voltage stacking network: (a) simplified 2×3 network; (b) equivalent network for I^G ; (c) voltage response with I^G ; (d) equivalent impedance for I^G ; (e) equivalent network for I^{ST} ; (f) voltage response with I^{ST} ; (g) equivalent impedance for I^{ST} ; (h) equivalent network for I^R

2) *Global Uniform Current*: Since $I^G(s)$ is a uniform component across all the cores, the effective network can then be transformed by removing the path between equal-potential nodes and merging the parallel components as in Fig. 5(b) according to our 2×2 example. We can derive the supply voltage noise caused by I^G with an analytical expression for a general $N_L \times N_V$ network:¹

$$\Delta V_{i,j}^G = I_{i,j}^G Z_{eff}^G = I_{i,j}^G \left(\frac{Z_{C4}}{N_L} + \frac{Z_S}{N_L} + \frac{N_V}{N_L} Z_{off} \right) // Z_C \quad (9)$$

Due to the uniform nature of the global current, all cores share the same common mode, Z_{eff}^G , and thus the same $\Delta V_{core,i,j}^G$. Eq. 9 also applies to the case when $N_L = 1$, which is a conventional single-layer PDN. From Eq. 9 and the typical impedance profile of Z_{eff}^G shown in Fig. 5(c), we can see that in a $N_L \times N_V$ voltage stacking PDN, $\Delta V_{core,i,j}^G$ peaks at the dominant resonant frequency of Z_{off} , similar to the conventional single-layer, but its magnitude is reduced by $N_L \times$ when stacked.

3) *Local Uniform Through-stack Current*: Following our definition of $I_i^{ST}(s)$, we can see that since $\sum_{i=1}^{N_V} I_i^{ST}(s) = 0$, there is no current going through Z_{off} according to Kirchhoff's Current Law (KCL) and the entire branch can be eliminated. The linear circuit network is again transformed to a simpler form as in Fig. 5(d). For example, in our 2×2 example, we can derive $\Delta V_{core,i,j}^{ST}$, for $i = 1, 2$ and $j = 1, 2$ respectively, as a function of the unit current stimulus I_i^{ST} and complex impedances in the form of Z_L and Z_C :

$$\Delta V_{core,i,j}^{ST} = I_i^{ST} Z_{eff(i)}^{ST} = I_i^{ST} \frac{1}{N_L} [Z_C // Z_L] \quad (10)$$

where $\Delta V_{core,i,j}^{ST}$ represents the supply voltage noise induced by I_i^{ST} , the common current components shared by all the cores in the i^{th} stack. All cores in the i^{th} stack share the same common-mode $\Delta V_{core,i,j}^{ST}$ disturbance. The resulting expression suggests that on the first order, the combined effect of all the I_i^{ST} exerts differential voltage fluctuations between the vertical stacks, and it is further voltage divided across the cores in the same stack, as illustrated in Fig. 5(d). The dividing ratio depends on the ratio of Z_L/Z_C , and in its high-frequency limit asymptotically approaches Z_L/N_L . The analytical results of the local uniform through-stack current again suggest that by moving from single-layer to multi-layer, the supply voltage noise experienced at each core level and contributed by this current component is reduced by $N_L \times$ on average.

4) *Residual Per-Core Differential Current*: On closer inspection of Eq. 7, $I_{i,j}^R$ can be rearranged as the summation

of differential currents in the form of $I_{i,j}^{core} - I_{i,k}^{core}$, where $k \neq j$. The summation suggests that the remaining voltage noise effect, unaccounted for by the global and the local terms, ΔV^G and ΔV_i^{ST} , are induced by the aggregated differential currents. This differential current represents the mismatched part of current between cores which will not only cause voltage noise at itself but also cause noise at other cores. For example, at $core(i, j)$, the noise from residual current is from its own residual current and other cores' residual current:

$$\Delta V_{core,i,j}^R = I_{i,j}^R Z_{eff,i,j}^R + \sum_{n \neq i} \sum_{m \neq j} I_{n,m}^R Z_{eff,n,m}^R \quad (11)$$

where $I_{i,j}^R Z_{eff,i,j}^R$ is the supply voltage noise caused by its own residual current, and $\sum_{n \neq i} \sum_{m \neq j} I_{n,m}^R Z_{eff,n,m}^R$ is the supply voltage noise caused by residual current from other cores. Most importantly, this type of residual per-core differential current is unique to voltage stacking, since these terms simply vanish when $N_L = 1$.

B. Dominating Supply Voltage Noise

Based on the above system configuration, we characterize the effective impedances, Z_{eff}^G , $Z_{eff,i}^{ST}$, and $Z_{eff,i,j}^R$, of each current component defined in Eq. (8). The effective impedance for $core(1, 1)$ is shown in Fig. 6. Due to location symmetry, the effective impedances of other cores are similar to $core(1, 1)$. We divide the frequency range into low frequency ($< 10MHz$), medium frequency ($10MHz - 50MHz$), and high frequency ($> 50MHz$). From the effective impedance curve, we can see that both $Z_{eff,i,j}^R$ at low frequency and Z_{eff}^G at high frequency (especially at resonance), have relatively large magnitudes. The corresponding low frequency residual current components and high frequency (resonance) global current components that excite these effective impedances can thus cause large supply voltage noise, and we identify them as the dominant causes of voltage noise in voltage stacking.

C. Worst-Case Supply Voltage Noise

Identifying the root cause of noise is not sufficient for rigorous reliability analysis. We must also consider what core activity conditions can result in the worst-case supply voltage noise. Understanding the condition and the magnitude of worst-case would help us determine the necessary and sufficient noise mitigation strategy to guarantee reliable operation in real-world voltage-stacked systems.

¹symbol // is the circuit symbol for parallel connection




Fig. 6: Effective impedance of current components

After characterizing Z_{eff}^G , Z_{eff}^{ST} and Z_{eff}^R , and establishing the relationship between ΔV_{core} as a function of these impedances, searching for the load current conditions that would result in worst-case supply voltage noise can now be performed in the frequency domain. We formulate it as an optimization problem of finding the optimal frequency distribution of each core current $I_{i,j}^{core}$ to maximize their combined effects $\Delta V_{m,n}^{core}$ on $core(m,n)$. This optimization can be solved as a linear programming problem, and the process is described in *Algorithm 1*. The optimization variables are each core current distribution at different frequency range $I_{i,j}^{core}(s)$. The optimization objective function is the supply voltage noise $\Delta V_{m,n}^{core}$ at $core(m,n)$ and the constraints are from voltage noise decomposition Eq. (4) - (7) and peak GPU SM core power, as shown in Table I. This linear optimization formulation with a general constraint of max power/current allows us to search the vast space of arbitrary synthetic core current stimuli from all possible activity combinations, including the effects cause by clock gating and power gating, and therefore can quantitatively represent the worst-case supply voltage noise for rigorous reliability analysis.

Algorithm 1 Maximize supply voltage noise

Optimization Variables:

Each core current frequency distribution $I_{i,j}^{core}(s)$

Objective Function:

$$\Delta V_{corei,j} = \Delta V_{corei,j}^G + \Delta V_{corei,j}^{ST} + \Delta V_{corei,j}^R \text{ in Eq. (8)}$$

Subject to:

- 1: $\forall i, j; 0 \leq I_{i,j}^{core}(s)$
- 2: $\forall i, j; I_{i,j}^{core}(t) = \mathcal{F}^{-1}(I_{i,j}^{core}(s)) \leq \text{peak current (14A)}$
- 3: $\forall i, j; I_{i,j}^{core}(s), 0 \leq s \leq \text{clock frequency (700MHz)}$
- 4: Eq. (4) - (7): current decomposition rules

The numerical solution of the linear programming problem based on the GPU configurations in Table I gives us a glimpse of the core current distribution and combination that act together and cause the largest supply voltage fluctuation

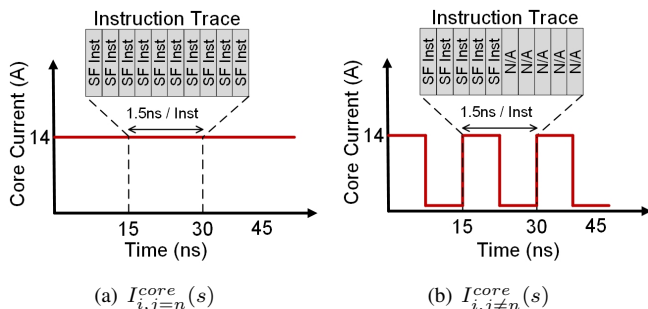


Fig. 7: An example instruction trace contributing to worst-case supply noise

TABLE II: Freq. Distribution of decomposed core current

Core Current	Frequency	Major Component
$I_{i,j=n}^{core}(s)$	low frequency ($< 10MHz$)	residual current
$I_{i,j≠n}^{core}(s)$	high frequency ($> 50MHz$)	global current

at $core(m,n)$, as shown in Table II. The currents, $I_{i,j=n}^{core}(s)$, are distributed at low frequency with major components of residual currents, while the currents, $I_{i,j≠n}^{core}(s)$, are distributed at the resonant frequency of Z_{eff}^G with major components the global currents. This worst-case scenario is plausible in real GPU applications shown in Fig. 7 when the $I_{i,j=n}^{core}(s)$ are alternating between idle and Sine/Cosine special function instructions (SF Inst) at the resonant frequency, while $I_{i,j≠n}^{core}(s)$ are at peak power executing Sine/Cosine special function instructions (SF Inst). We compare the worst-case noise derived by our optimization algorithm with three other scenarios based only on heuristics: (1) all cores have low frequency residual currents, (2) all cores have high frequency global currents, and (3) all cores have randomly distributed currents. From the supply voltage noise histograms in Fig. 8, we can see that the worst-case rigorously derived by our method is more severe than the heuristic ones, and therefore is more representative as a stressmark for supply voltage noise reliability analysis.

V. NOISE MITIGATION BY HYBRID REGULATION

A. Hybrid Regulation Framework

To combat elevated and hard-to-predict supply voltage noise and guarantee reliable operation in spite of worst-case conditions in voltage-stacked manycore processors, we explore a hybrid voltage regulation mechanism using both on-chip charge-recycling integrated voltage regulators (CR-IVRs) and an off-chip charge-recycling voltage regulator module (CR-VRM). Fig. 9 shows the framework of the proposed hybrid regulated voltage stacking using either switched-capacitor or low dropout voltage regulators. This hybrid approach takes advantage of the unique merits of on-chip and off-chip voltage regulators and simultaneously avoids their individual defects.

Unlike step-down voltage regulators converting supply voltage, charge-recycling voltage regulators move extra charge between different layers to balance current and maintain a stable voltage of each layer. Because the direction and amplitude of extra charge keeps changing with core workload conditions, charge-recycling voltage regulators should support bidirectional fast switching current. Voltage regulators, such as low drop-out voltage regulators and switched capacitor voltage regulators, can be used as charge-recycling voltage regulators, while inductor based voltage regulators, such as buck converters, do not support bidirectional fast switching of current movement and incur extra Ldi/dt noise. Multi output switched capacitor (SC) voltage regulators are the most widely used charge-recycling voltage regulators, because they have higher power efficiency, but they require each layer to have the same voltage. Previous work has demonstrated a multi-output switched-capacitor integrated voltage regulator [40] that balances the layer currents in voltage-stacked systems.




Fig. 8: Histogram comparison between analytically derived worst case and other heuristic core activation patterns

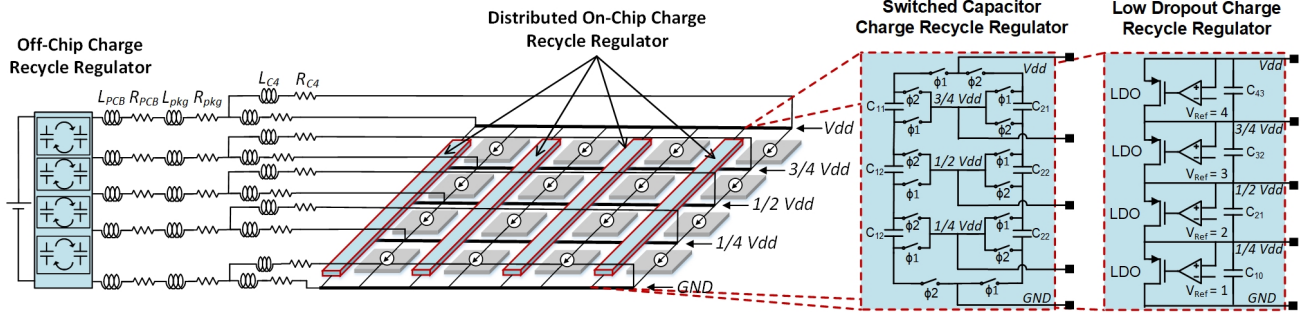




Fig. 9: Hybrid voltage regulation based on distributed on-chip CR-IVRs and off-chip CR-VRM



(a) Voltage distribution among the 16 SMs during execution of the backp benchmark



(b) Voltage distribution among the 16 SMs during execution of the blackscholes benchmark

Fig. 10: Voltage distribution among the 16 SMs

Although low drop-out voltage regulators have lower power efficiency, they do not force each layer to have exactly the same voltage, and hence are more suitable to support dynamic voltage and frequency scaling in voltage stacking.

B. Centralized and Distributed Integrated Voltage Regulator

Located closer to the point-of-load, on-chip integrated voltage regulators enjoy fast regulation response, but have limited on-die area and capacity, making them suitable for reducing high-frequency noise of smaller magnitude. According to the analysis in Section IV, one of the dominant causes of worst-

case supply voltage noise is high frequency global currents. This noise can be mitigated by on-chip CR-IVRs.

By moving charges across the stacking layers, the CR-IVR effectively behaves as an additional parallel impedance connected to the original effective impedance Z_{eff}^G . It thus reduces the supply voltage noise caused by global current:

$$\Delta V_{corei,j}^G = I^G [Z_{eff}^G / (Z^{CR-IVR} + Z^{CR-IVR-path})] \quad (12)$$

Here, Z^{CR-IVR} is the impedance of the on-chip charge-recycling voltage regulator, and $Z^{CR-IVR-path}$ is the parasitic impedance of the on-chip power grid between the core and voltage regulator. By deploying CR-IVR with the desired impedance, $\Delta V_{corei,j}^G$ from global current I^G can be effectively mitigated. The effective impedance of a multi-output switched-




Fig. 11: Supply voltage noise distribution

capacitor CR-IVR can be expressed:

$$Z^{CR-IVR} = \sqrt{Z_{SSL}^2 + Z_{FSL}^2} \quad (13)$$

$$Z_{SSL} = \frac{1}{C_{total} f_{SW}} \left(\sum_1^n |a_{c,i}| \right)^2 \quad Z_{FSL} = \frac{G_{total}}{D_{cycle}} \left(\sum_1^n |a_{r,i}| \right)^2 \quad (14)$$

where, C_{total} is the fly capacitance, G_{total} is the total switch conductance, f_{SW} is the switching frequency, and D_{cycle} is the duty cycle. Further, $a_{c,i}$ and $a_{r,i}$ are charge multiplier vectors [44], [62]. $Z^{CR-IVR-path}$ is the other important factor that determines the supply voltage noise mitigation. It is related to the distance between the core and the voltage regulator. As the regulator is located far from the load, the noise mitigation effect will be reduced because the parasitic impedance between the core and the voltage regulator contributes to a larger $Z^{CR-IVR-path}$. One effective way to enhance the noise mitigation is by distributing a large centralized voltage regulator to smaller distributed ones, because the distributed voltage regulators can be located closer to each core.

We next will demonstrate the effectiveness of hardware regulation by on-chip charge-recycling integrated voltage regulators (CR-IVRs) and compare the regulation effects of centralized and distributed CR-IVRs. We first simulate the transient voltage waveforms of all the SMs with one centralized CR-IVR physically located in the middle of each layer and plot their voltage distribution using box plots. The statistics presented in Fig. 10 were collected from the benchmark *backp* and benchmark *blackscholes*, but similar results are observed for all the benchmarks from both NVIDIA CUDA SDK and Rodinia 2.0 benchmark suites. Comparing the standard deviations and peak-to-peak values of all the SM core voltages in the proposed voltage-stacked GPU, with centralized CR-IVR and without CR-IVR, reveals that the regulation effect is uneven among the SMs. This phenomenon is highlighted in the histograms in Fig. 11(a). We have the histogram of the voltage distribution across SM1 and SM2, collected with 500,000 samples over a typical $10\mu s$ period from the benchmark *backp*. SM2 exhibits the smallest supply voltage noise spread, yet noise worsens at SM1, because SM2 is closer to the centralized CR-IVR and has a smaller $Z^{CR-IVR-path}$ than SM1.

Now, we leverage the scalability of CR-IVR in a distributed design. The distributed CR-IVR divides the original centralized design into four equal sub-IVRs and connects each sub-IVR directly to the SMs in each layer, with each sub-IVR consisting of 1/4 of the total switched capacitance. The extra implementation overhead of the distributed design is mainly due to the duplication of control logic, which accounts for negligible area and power consumption compared to the rest

TABLE III: Switched Cap. Regulator Parameters

Design Parameters	CR-IVR	CR-VRM
Topology of VR	Multi-output SC	Multi-output SC
Number of VR	4	1
Switch frequency	50MHz	1MHz
Total capacitor / VR	1.24uF	624uF
Capacitor density	50nF/mm ²	0.2uF/mm ²
Switch on resistance	130Ω · um	37600Ω · um
Area per VR	24.8mm ² (Die)	3.12cm ² (Board)




Fig. 12: Effective impedance after employing CR-VRM

of the CR-IVR circuitry. The resulting SM voltage distribution using the distributed regulation is presented in Fig. 11(b). The location dependence is now completely removed and the same regulation effect is achieved across the board. A optimal design parameters [63] of distributed CR-IVR are shown in Table IV.

C. Off-Chip Charge-Recycling VR

Compared with CR-IVR, off-chip CR-VRMs have slower response time, but they offer better efficiency [64], [65] and do not consume expensive die area. It is important to note that although on-chip CR-IVRs can be designed to provide similar regulating capacity as an off-chip counterpart, they incur large area overhead, sometimes exceeding the total area of the logic cores, making them impractical in real systems. Therefore, off-chip CR-VRM is a better and more economical choice for regulating supply voltage noise at low frequency. Similarly, the addition of the CR-VRM results in an effective parallel impedance connected with the original $Z_{eff}^R(i,j)$ through the C4 pad, package, and PCB. In this case, the supply voltage noise caused by residual current becomes

$$\Delta V_{corei,j}^R = \sum_i^{N_V} \sum_j^{N_L} I_{i,j}^R [Z_{eff}^R(i,j) // (Z^{CR-VRM} + Z^{CR-VRM-path})] \quad (15)$$

where Z^{CR-VRM} is the impedance of the off-chip charge-recycling voltage regulator module; $Z^{CR-VRM-path}$ includes the parasitic impedances of not only the on-chip power grid but also the C4 pads, package, and PCB board between the CR-VRM and the cores. A design optimization similar to that for CR-IVR is applied to arrive at an optimal set of design parameters, as summarized in Table IV. The new effective impedance of the residual current after employing on-chip CR-IVR and off-chip CR-VRM is shown in Fig. 12. With reduced effective impedance, the supply voltage noise, $\Delta V_{corei,j}^G$ and $\Delta V_{corei,j}^R$, are significantly mitigated.

D. Charge-Recycling VR Power Loss

In voltage stacking, most of the current goes through the stacked layers, the occasional residual current is absorbed by decoupling capacitors, and only the accumulated residual current components goes through the CR-IVR or CR-VRM.

For the accumulated residual current, we call it imbalanced current. When imbalanced current is recycled by charge-recycling VR, power losses in these VRs are unavoidable.

1) *Switched-capacitor charge-recycling VR*: A switched capacitor voltage regulator suffers mainly from the following four types of power losses:

Intrinsic switched-capacitor loss: A switched capacitor voltage regulator delivers current to a synchronous digital system, whose frequency is determined by the clock frequency, set by the minimum voltage over a clock period. Power loss in the voltage ripple over the minimum voltage is the intrinsic switched-capacitor loss [66], which is

$$P_{intrinsic} = I_{imbalance} \frac{\Delta V}{2} = \frac{I_{imbalance}^2}{M_{cap} C_{fly} f_{sw}} \quad (16)$$

Switching conductance Loss: Also, the finite conductance of the transistor switch has a series power loss:

$$P_{R_{sw}} = N I_{imbalance}^2 R_{sw} D \quad (17)$$

Plate Parasitic Capacitance Loss: In steady-state operation, both the top and the bottom plates experience approximately equal voltage swings, and parasitic capacitance causes extra power loss:

$$P_{plate-cap} = M_{bott} V^2 C_{plate} f_{sw} \quad (18)$$

Switching Parasitic Capacitance Loss: The loss in voltage swings at the switch transistor parasitic capacitance, which can be expressed as

$$P_{sw-cap} = N C_{sw} V^2 f_{sw} \quad (19)$$

Among these losses, intrinsic switched-capacitor loss and switch conduct loss are the main components [63].

2) *Low Drop-out Charge-Recycling VR*: Low drop-out voltage regulators suffer from following three power losses [67]:

Switch Conduct Loss: The main cause of loss in low drop-out voltage regulators is the power dissipated as heat on the transistor switch resistance. It is highly dependent on the difference between the input and output voltage:

$$P_{R_{sw}} = I_{imbalance}^2 R_{sw} \quad (20)$$

Switching Parasitic Capacitance Loss: The gate parasitic capacitance switching loss is similar to the loss happens in switched capacitor charge recycle voltage regulator.

Control Logic Loss: In LDO, a feedback control logic circuitry is used to control the voltage at the reference value. The loss is due to the current flowing through the operational amplifier, the resistive voltage divider, and the voltage reference generator in the control logic circuitry; the sum of these currents is called quiescent current. The quiescent current can be reduced by optimizing the components, making it negligible when compared to the load current consumption.

E. Hybrid Regulated VS Power Delivery Efficiency

The power delivery efficiency of hybrid regulated manycore voltage stacking can be described as

$$\eta_{PDS} = \frac{P_{core}}{P_{core} + P_{PDN} + P_{CR-IVR} + P_{CR-VRM}} \quad (21)$$

$$= \frac{I_{core} V_{core}}{I_{core} V_{core} + \left(\frac{I_{core}}{N}\right)^2 R_{PDN} + P_{CR-IVR} + P_{CR-VRM}}$$

where, P_{core} is the power consumed by cores, and P_{PDN} is the power loss in the parasitic resistance along the power delivery network. P_{CR-IVR} and P_{CR-VRM} are the power loss in CR-IVR and CR-VRM, derived in Eq. (16) - (20).

In the ideal case, if there is no residual current, the power losses in CR-IVR and CR-VRM are negligible and the power delivery efficiency can approach nearly 100%. In the normal case, most of the current goes through the stacked layers, and only the accumulated residual current components goes through the CR-IVR or CR-VRM where it introduces a small amount of power loss. In the worst-case, when cores in one layer are powered off and cores in the other layers are powered on, all the current consumed by cores is imbalanced current and goes through the CR-IVR and CR-VRM, causing significant power loss. However, this worst case seldom happens, and the system on average achieves high power delivery efficiency. Consequently, the hybrid regulation scheme not only mitigates the supply voltage noise but also maintains the high power delivery efficiency.

VI. ADVANCED POWER MANAGEMENT

We consider a well-designed power delivery system should be compatible with advanced high-level power management techniques. Among them, the most common techniques are dynamic voltage and frequency scaling (DVFS) and power gating. In this section, we will discuss applying DVFS and power gating together with the proposed hybrid regulation in voltage-stacked GPU systems.

A. Dynamic Voltage and Frequency Scaling

Dynamic voltage and frequency scaling (DVFS) adjusts the supply voltage and frequency of a voltage domain to boost performance or save power. In voltage stacking each layer can be divided into different voltage domains and the division should be consistent across different layers to maintain the stacked power delivery. In this paper, we assume the SMs in each layer share one voltage domain in the proposed GPU manycore voltage-stacked system. When DVFS is applied in voltage stacking, each layer (voltage domain) may have a different voltage. The low drop-out voltage regulator will be used to recycle the imbalance current because LDO can support that each layer has its own voltage. We will use LDO hybrid regulations in following voltage stacking DVFS analysis and evaluations. When one voltage domain needs to change to a different supply voltage, the low drop-out voltage regulator can change the reference voltage and conversion ratio to adjust the voltage of each layer [17], [19], [68].

Compared to the original voltage stacking, DVFS maintains the stacked power delivery but may brings more frequent current imbalance. As the current imbalance introduced by DVFS does not go beyond the worst cases studied in Section IV where one layer is totally powered off, the proposed hybrid regulation can effectively guarantee the system's stability under DVFS operation. Although the amplitude of the imbalanced current does not exceed the worst case, the extra current imbalance introduced by DVFS will cause more power loss in CR-IVR and CR-VRM. Compared with the original voltage stacking, part of efficiency benefit will be sacrificed under DVFS.

TABLE IV: LDO Regulator Parameters

Design Parameters	CR-IVR	CR-VRM
Number of VR	4	1
Switch frequency	50MHz	2MHz
Total capacitor per VR	1.1uF	600uF
Capacitor density	50nF/mm ²	0.2uF/mm ²
Switch on resistance	130Ω · um	37600Ω · um
Area per VR	22.0mm ² (Die)	3.1cm ² (Board)

Algorithm 2 Power Saving from DVFS and Power Gating

Input Variables:

DVFS / power gating command: $f_{i,j}^{core} / P_{i,j}^{core-gate}$

Output Variables:

Power estimation of each Core: $P_{i,j}^{core}$

Steps:

1: Replace $I_{i,j}^{core}$ with $f_{i,j}^{core} / P_{i,j}^{core-gate}$ in Eq. (4) - (7).

2: Calculate residual frequency / gated power: $f_{i,j}^R / P_{i,j}^{R-gate}$.

3: Residual current can be known as:

$$I_{i,j}^R = \alpha CV f_{i,j}^R / \frac{P_{i,j}^{R-gate}}{V}$$

4: Calculate VR loss P_{CR-IVR} / P_{CR-VRM} in Eq. (16)-(20).

5: Return power estimation:

$$P_{i,j}^{core} = \alpha CV^2 f_{i,j} / P_{i,j}^{core-gate} - P_{CR-IVR} - P_{CR-VRM}$$

B. Power Gating

Power gating turns off the circuitry inside a core or the core itself for a while when not in use. Power gating introduces current imbalance and also causes supply voltage noise. The most severe imbalance happens when one layer is totally powered off while other layers are working. This scenario is already captured by supply voltage noise worst case analysis in Section IV and supply voltage can be also guaranteed by the proposed hybrid regulation as described in Section V. Similar to DVFS, the extra imbalanced current introduced by power gating will cause more power loss from CR-IVR and CR-VRM thus degrading efficiency gains.

C. Power Management Hypervisor in Voltage Stacking

The DVFS, power gating and other power management techniques optimize the power and performance tradeoffs based on the commands from software operating system. At the software commands level, power management techniques should taken voltage stacking into consideration and many techniques such as fast thread migration [69] can balance the workload before current imbalance happens. To make the correct decision at the software level, the power management techniques first need to know the potential power benefit and performance loss and then find the proper tradeoff point. To estimate the potential power benefit that each core can earned, we introduce a power management technique estimator for the software level power management as described in Algorithm 2. The estimator can evaluate the potential net power consumption of each core considering the extra power loss in power delivery system.

At the hardware power delivery system, we provide a power delivery efficiency guaranteed power management hypervisor of DVFS or power gating instructions for voltage stacking. According to Section IV and V, the power loss in voltage stacking comes from the accumulated residual current

Algorithm 3 Power Management Hypervisor in VS

Input Variables:

Commands in conventional system: $f_{i,j}^{core} / P_{i,j}^{core}$

Output Variables:

Commands for Voltage Stacking: $f'_{i,j}{}^{core} / P'_{i,j}{}^{core}$

Steps:

1: Replace $I_{i,j}^{core}$ with $\alpha CV f_{i,j}^{core} / \frac{P_{i,j}^{core}}{V}$ in Eq. (4) - (7).

2: Calculate residual current: $I_{i,j}^R$

3: Dynamic Voltage Frequency Scaling:

for i, j = 1, 2, 3, 4 do

if $|I_{i,j}^R| > |\Delta I_{threshold}|$ then

$$I_{i,j}^{core} = I_{i,j}^{core} - (I_{i,j}^R - \Delta I_{threshold})$$

$$f'_{i,j}{}^{core} = \frac{I_{i,j}^{core}}{\alpha CV}$$

else then

$$f'_{i,j}{}^{core} = f_{i,j}^{core}$$

Return DVFS commands: $f'_{i,j}{}^{core}$

4: Power Gating:

for i, j = 1, 2, 3, 4 do

if $|I_{i,j}^R| > |\Delta I_{threshold}|$ then

$$I_{i,j}^{core} = I_{i,j}^{core} - (I_{i,j}^R - \Delta I_{threshold})$$

$$P'_{i,j}{}^{core} = V I_{i,j}^{core}$$

else then

$$P'_{i,j}{}^{core} = P_{i,j}^{core}$$

Return power gating commands: $P'_{i,j}{}^{core}$


component going through charge-recycling voltage regulators. The hypervisor guarantees the power delivery efficiency by limiting the maximum allowed residual current, described in Algorithm 3. In the hypervisor, the residual current of each core under DVFS and power gating is calculated with Eq. (4) - (7). The residual current threshold $\Delta I_{threshold}$ is given to limited the residual current $I_{i,j}^R$ and guarantee power loss in power delivery system. Then each core whose residual current exceeds the threshold $\Delta I_{threshold}$ or $P_{threshold}$ will be compensated by $I_{i,j}^{core} - \Delta I_{threshold}$ to make sure that the residual current $I_{i,j}^{core}$ and power loss in power delivery system are limited within desired range.

VII. SYSTEM EVALUATION

In this section, we evaluate the hybrid regulated GPU manycore voltage-stacked system in terms of supply voltage noise, power delivery efficiency, advanced power management compatibility, and finally compare it with other power delivery systems. We develop an *hybrid simulation infrastructure* that combines SPICE3 [70] and GPGPU-Sim 3.1.1 (with GPUWattch) [71], [72]. SPICE3 simulates the circuit transient response of the full voltage-stacked power delivery system and the charge-recycling voltage regulators as illustrated in Fig. 9, and GPGPU-Sim 3.1.1 simulates the GPU architecture level system specified in Table I. We use ten representative benchmarks that cover a wide range of scientific and computational domains from two benchmark suites, five from Rodinia 2.0 [73] and five from NVIDIA CUDA SDK [74].

A. Supply Voltage Noise Evaluation

We first evaluate the supply voltage noise across real GPU benchmarks and the worst case derived by Algorithm 1. As



(a) Supply voltage noise comparison between SC / LDO hybrid regulated and default voltage stacking (b) Worst supply noise distribution
 Fig. 13: Evaluation of the supply voltage noise in hybrid regulated voltage stacking system




Fig. 14: Power delivery efficiency comparison between voltage-stacked system with SC/LDO hybrid regulation and conventional single-layer system across ten benchmarks

shown in Fig. 13(a), in default voltage stacking without any voltage regulation, the supply voltage suffers huge noise, especially under the worst case. As demonstrated by the noise histograms in Fig. 13(a) and 13(b), after deploying hybrid regulation in the voltage-stacked GPU system, the supply voltage noise across both the benchmarks and the worse case is limited to a range of 0.2V, comparable to conventional single-layer power delivery system². One of the key strengths of our hybrid approach is its use of the more expensive on-chip regulator for high frequency noise mitigation and the more economical off-chip regulator for low frequency noise mitigation. This choice avoids over design of the on-chip CR-IVR, saves significant on-die area, and provides worst-case guaranteed reliability.

B. Efficiency in Real Applications

We evaluate the system level power delivery efficiency (PDE) by running a wide range of real GPU benchmarks on our integrated hybrid simulation infrastructure. We compare our hybrid regulated voltage-stacked system in Fig. 9 with the conventional single-layer power delivery system with a board-level voltage regulator module (VRM), which is the default GPU power delivery system [28], [75].

The normalized breakdown of the full system power delivery efficiency across benchmarks is shown in Fig. 14. On average, voltage-stacked power delivery system configurations (with hybrid regulation) can deliver power at close to 93.5% efficiency with switched capacitor charge-recycling voltage regulators and 92.3% efficiency with LDO charge-recycling voltage regulators, as compared to 79% for the single-layer VRM (conventional baseline). The charge-recycling voltage

²0.2V is the voltage margin used in commercial GPU systems for tolerable supply noise [28].

TABLE V: SM Core DVFS Frequency and Voltage Pairs

Core freq. (MHz)	700	650	600	550	300
Core voltage (V)	1	0.95	0.91	0.87	0.46

regulator in voltage stacking outperforms the step-down voltage regulator in the single-layer PDS because the former only needs to shuffle the accumulated imbalanced part, usually within 20% of the layer power, whereas the latter delivers the total power. For example, in benchmark Transpose, only 11.8% and 2.9% of current are imbalanced current that goes through CR-IVR and CR-VRM respectively, and causes 3.7% and 1.1% of power loss in switched capacitor CR-IVR and CR-VRM respectively.

C. Compatibility with Advanced Power Management

First we leverage the common and classic DVFS algorithm proposed in [77] to explore per-core DVFS on a voltage-stacked GPU system, which monitors and predicts the application status (compute bound/memory bound) to adjust each core and memory frequency. The SM core frequency and voltage pairs are shown in Table V. In conventional single-layer power delivery system, each cores has its own frequency and voltage. In voltage stacking, the cores in each layer share a voltage domain and the highest voltage and frequency from the cores in one layer is used as the voltage and frequency for this layer. We evaluate DVFS on the voltage stacking and compare with DVFS on conventional single-layer power delivery system in Fig. 15. Although DVFS on voltage stacking causes more power loss than normal execution on voltage stacking, but it still has a higher power delivery efficiency than on conventional power delivery system at most benchmarks except Transpose. This is because GPU benefits the single instruction multiple thread (SIMT) architecture causing a synchronized activity and synchronized DVFS commands for the cores during most of time. Besides, shown in the right bars in Fig. 15, the power delivery efficiency guided hypervisor in Algorithm 3 can further prevent the aggravated power loss in CR-IVR and CR-VRM by limiting the occasional current imbalance from DVFS. Power delivery efficiency guided hypervisor can help voltage stacking achieve a near 90% power delivery efficiency under DVFS operations. The normalized energy consumption across benchmarks of conventional single-layer system, DVFS on conventional single-layer system and DVFS on power delivery efficiency guided voltage stacking is shown in Fig. 16. On conventional single-layer system, DVFS can reduce the energy consumption of cores across most benchmarks. On

TABLE VI: Power delivery system comparison

Power Delivery System	Efficiency	Die Area	Reliable	Compatibility
Single-layer VRM [46]	79.9%	N/A	✓	✓
Single-layer IVR [76]	85.8%	172.3 mm ²	✓	✓
VS IVR [18]	92%	88.3 mm ²	×	×
VS IVR (worst) [18]	92%	912 mm ²	✓	×
VS Hybrid (this paper)	93.5%	99.2 mm ²	✓	✓




Fig. 15: DVFS power saving comparison between conventional single-layer system and voltage-stacked system with hybrid regulation across benchmarks




Fig. 16: Normalized system energy consumption under DVFS power delivery efficiency guided voltage stacking, the energy consumption of cores are also partly reduced compared to conventional single-layer system without DVFS, but cannot reach the same amount as DVFS on single-layer system. This is because power delivery efficiency guided hypervisor modifies the aggressive DVFS commands which cause current imbalance and low power delivery efficiency. Although the energy consumption of cores is higher than DVFS on conventional single-layer system, when the energy loss in power delivery system is taken into consideration, DVFS on voltage stacking achieves the best overall energy consumption.

For power gating, we manually power off the cores in one layer and leave the cores in the other layers under normal execution which will cause the most current imbalance, power loss and the worst power delivery efficiency. The residual current threshold $\Delta I_{threshold}$ in power delivery efficiency guided hypervisor is set to 25%, 50%, and 75% of the core current respectively to protect power delivery efficiency. Fig. 17 describes the full system power delivery efficiency across benchmarks. Compared with voltage stacking without power gating in Fig. 14, continuous imbalanced current from power gating causing more power loss in CR-IVR and CR-VRM. When $\Delta I_{threshold}$ is set to 25% and 50% the core current, the full system power delivery efficiency can still maintain 80%. When $\Delta I_{threshold}$ is set to 75% the core current, the full system power delivery efficiency is lower than 70%. It means that the power benefits from gating the cores in one layer, which is about 1/4 of system power, are all wasted in the power delivery system. Since that when power gating




Fig. 17: Power delivery efficiency under PDE guided power gating and original power gating on voltage stacking

is applied in voltage stacking, at the software level power gating should prefer powering off the cores in the same stack with the help of thread migration. When it is inevitable to power off the cores in the same layer, hardware based power delivery efficiency guided power management hypervisor will be deployed to prevent gating over 50% of one layer from happening frequently to protect the voltage stacking power delivery efficiency.

D. Comparison with Other Power Delivery Systems

In Table VI, we compare the proposed hybrid regulated voltage-stacked power delivery system with other existing and emerging power delivery system configurations. Although charge-recycling voltage regulators are employed, the voltage-stacked system does not suffer a large efficiency penalty, because most currents go through the vertically-stacked grid without incurring energy loss at the regulators. Validated by benchmarks, the proposed voltage-stacked system with hybrid regulation can achieve 93.5% power delivery efficiency on average and can guarantee that the supply voltage noise remains within the reliable region. Besides efficient power delivery and supply voltage noise mitigation, hybrid regulated voltage-stacked systems are also compatible with other advanced high level power management techniques, such as DVFS and power gating. Although when advanced power managements are applied in voltage-stacked system, huge imbalance current may lead to power delivery efficiency loss, power delivery efficiency guided hypervisor are able to limit the magnitude and frequency of imbalance and guarantee the improved power delivery efficiency. Furthermore, many other techniques, like high efficiency charge-recycling circuit, can be explored to further improve the voltage-stacked power delivery efficiency.

VIII. CONCLUSION

Voltage stacking fundamentally improves manycore processors power delivery efficiency but suffers aggravated supply voltage noise. According to the analysis using circuit decomposition and superposition, the contributors to supply voltage noise are high frequency global current and low frequency residual current. Then the current configuration leading to the

worst supply voltage is derived as an optimization problem. Based on the characteristics of supply voltage noise, a hybrid regulation, with distributed on-chip and a off-chip charge recycle voltage regulators, is proposed to effectively mitigate supply voltage noise. The supply voltage noise is guaranteed within a safe range even under the worst case. Also, the proposed hybrid regulated voltage-stacked system can not only be compatible with other power management techniques like DVFS and power gating but also maintains a high power delivery efficiency. Compared with conventional power delivery system, the proposed hybrid regulated voltage-stacked system achieves a 13.6% improvement of power delivery efficiency.

ACKNOWLEDGMENT

The research described in this paper was partly supported by NSF CPS grant CNS-1739643, NSF award CCF-1528045, Semiconductor Research Corporation (SRC) task 2810.003 through the University of Texas at Dallas Texas Analog Center of Excellence (TxACE), and the National Natural Science Foundation of China (NSFC) 61702328. We are also grateful to the reviewers for their constructive feedback.

REFERENCES

- [1] L. M. Platchkov and M. G. Pollitt, "The economics of energy (and electricity) demand," *The Future of Electricity Demand: Customers, Citizens and Loads*, vol. 69, p. 17, 2011.
- [2] U.S. Energy Information Administration (EIA), "Annual Energy Outlook 2016 with Projections to 2040," <http://www.eia.gov/forecasts/aeof>.
- [3] L. Liu, C. Li, H. Sun, Y. Hu, J. Gu, T. Li, J. Xin, and N. Zheng, "Heb: deploying and managing hybrid energy buffers for improving datacenter efficiency and economy," in *ACM SIGARCH Computer Architecture News*, vol. 43, no. 3. ACM, 2015, pp. 463–475.
- [4] M. P. Mills, "The cloud begins with coal: Big data, big networks, big infrastructure, and big power," *Digital Power Group*, 2013.
- [5] U.S. Energy Information Administration (EIA), "FAQ:How much electricity is lost in transmission and distribution in the United States?" <http://www.eia.gov/tools/faqs/faq.cfm?id=105&t=3>.
- [6] X. Zhang, T. Tong, S. Kanev, S. K. Lee, G.-Y. Wei, and D. Brooks, "Characterizing and evaluating voltage noise in multi-core near-threshold processors," in *Low Power Electronics and Design (ISLPED), 2013 IEEE International Symposium on*, 2013, pp. 82–87.
- [7] W. Kim *et al.*, "System level analysis of fast, per-core dvfs using on-chip switching regulators," in *HPCA*, 2008.
- [8] S. K. Lee, D. Brooks, and G.-Y. Wei, "Evaluation of voltage stacking for near-threshold multicore computing," in *Proceedings of the 2012 ACM/IEEE international symposium on Low power electronics and design*. ACM, 2012, pp. 373–378.
- [9] R. Jain, B. M. Geuskens, S. T. Kim, M. M. Khellah, J. Kulkarni, J. W. Tschanz, and V. De, "A 0.45–1 v fully-integrated distributed switched capacitor dc-dc converter with high density mim capacitor in 22 nm tri-gate cmos," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 4, pp. 917–927, 2014.
- [10] M. S. Gupta, J. L. Oatley, R. Joseph, G.-Y. Wei, and D. M. Brooks, "Understanding voltage variations in chip multiprocessors using a distributed power-delivery network," in *Design, Automation & Test in Europe Conference & Exhibition, 2007. DATE'07*. IEEE, 2007.
- [11] Intel Corp., "Voltage Regulator Module, Enterprise Voltage Regulator-Down 10.0," <http://www.intel.com/content/www/us/en/power-management/voltage-regulator-module-enterprise-voltage-regulator-down-10-0-guidelines.html>.
- [12] K. Ueda, F. Morishita, S. Okura, L. Okamura, T. Yoshihara, and K. Arimoto, "Low-power on-chip charge-recycling dc-dc conversion circuit and system," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 11, pp. 2608–2617, 2013.
- [13] P. C. Lisboa, P. Pérez-Nicoli, F. Veirano, and F. Silveira, "General top/bottom-plate charge recycling technique for integrated switched capacitor dc-dc converters," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 63, no. 4, pp. 470–481, 2016.
- [14] S. Rajapandian, Z. Xu, and K. L. Shepard, "Implicit dc-dc downconversion through charge-recycling," *IEEE journal of solid-state circuits*, vol. 40, no. 4, pp. 846–852, 2005.
- [15] P. Jain, T.-H. Kim, J. Keane, and C. H. Kim, "A multi-story power delivery technique for 3d integrated circuits," in *Low Power Electronics and Design (ISLPED), 2008 ACM/IEEE International Symposium on*. IEEE, 2008, pp. 57–62.
- [16] S. K. Lee, D. Brooks, and G.-Y. Wei, "Evaluation of voltage stacking for near-threshold multicore computing," in *Proceedings of the 2012 ACM/IEEE international symposium on Low power electronics and design*. ACM, 2012, pp. 373–378.
- [17] K. Blutman, A. Kapoor, A. Majumdar, J. G. Martínez, J. Echeverri, L. Sevati, A. P. van der Wel, H. Fatemi, K. A. Makinwa, and J. P. de Gyvez, "A low-power microcontroller in a 40-nm cmos using charge recycling," *IEEE Journal of Solid-State Circuits*, vol. 52, no. 4, 2017.
- [18] S. K. Lee, T. Tong, X. Zhang, D. Brooks, and G.-Y. Wei, "A 16-core voltage-stacked system with adaptive clocking and an integrated switched-capacitor dc-dc converter," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, no. 4, pp. 1271–1284, 2017.
- [19] K. Blutman, A. Kapoor, J. G. Martínez, H. Fatemi, and J. P. de Gyvez, "Lower power by voltage stacking: A fine-grained system design approach," in *Design Automation Conference (DAC), 2016 53rd ACM/EDAC/IEEE*. IEEE, 2016, pp. 1–5.
- [20] E. K. Ardestani, R. T. Pognolo, J. L. Briz, and J. Renau, "Managing mismatches in voltage stacking with coreunfolding," *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 12, no. 4, 2016.
- [21] R. T. Pognolo, E. Ebrahimi, E. K. Ardestani, A. Sankaranarayanan, J. L. Briz, and J. Renau, "Gpu ntc process variation compensation with voltage stacking," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, no. 99, pp. 1–14, 2018.
- [22] S. K. Lee, T. Tong, X. Zhang, D. Brooks, and G.-Y. Wei, "A 16-core voltage-stacked system with an integrated switched-capacitor dc-dc converter," in *VLSI Circuits (VLSI Circuits), 2015 Symposium on*. IEEE, 2015, pp. C318–C319.
- [23] T. Tong, S. K. Lee, X. Zhang, D. Brooks, and G.-Y. Wei, "A fully integrated reconfigurable switched-capacitor dc-dc converter with four stacked output channels for voltage stacking applications," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 9, pp. 2142–2152, 2016.
- [24] E. Grochowski, D. Ayers, and V. Tiwari, "Microarchitectural simulation and control of di/dt-induced power supply voltage variation," in *High-Performance Computer Architecture, 2002. Proceedings. Eighth International Symposium on*. IEEE, 2002, pp. 7–16.
- [25] M. S. Gupta, V. J. Reddi, G. Holloway, G.-Y. Wei, and D. M. Brooks, "An event-guided approach to reducing voltage noise in processors," in *Design, Automation & Test in Europe Conference & Exhibition, 2009. DATE'09*. IEEE, 2009, pp. 160–165.
- [26] V. J. Reddi, S. Kanev, W. Kim, S. Campanoni, M. D. Smith, G.-Y. Wei, and D. Brooks, "Voltage smoothing: Characterizing and mitigating voltage noise in production processors via software-guided thread scheduling," in *Microarchitecture (MICRO), 2010 43rd Annual IEEE/ACM International Symposium on*. IEEE, 2010, pp. 77–88.
- [27] V. J. Reddi, M. S. Gupta, G. Holloway, G.-Y. Wei, M. D. Smith, and D. Brooks, "Voltage emergency prediction: Using signatures to reduce operating margins," in *High Performance Computer Architecture, 2009. HPCA 2009. IEEE 15th International Symposium on*. IEEE, 2009.
- [28] J. Leng, Y. Zu, M. Rhu, M. Gupta, and V. J. Reddi, "Gpuvolt: Modeling and characterizing voltage noise in gpu architectures," in *Proceedings of the 2014 international symposium on Low power electronics and design*. ACM, 2014, pp. 141–146.
- [29] J. Leng, Y. Zu, and V. J. Reddi, "Gpu voltage noise: Characterization and hierarchical smoothing of spatial and temporal voltage noise interference in gpu architectures," in *High Performance Computer Architecture (HPCA), 2015 IEEE 21st International Symposium on*. IEEE, 2015, pp. 161–173.
- [30] R. Thomas, K. Barber, N. Sedaghati, L. Zhou, and R. Teodorescu, "Core tunneling: Variation-aware voltage noise mitigation in gpus," in *High Performance Computer Architecture (HPCA), 2016 IEEE International Symposium on*. IEEE, 2016, pp. 151–162.
- [31] R. Thomas, N. Sedaghati, and R. Teodorescu, "Emergpu: Understanding and mitigating resonance-induced voltage noise in gpu architectures," in *Performance Analysis of Systems and Software (ISPASS), 2016 IEEE International Symposium on*. IEEE, 2016, pp. 79–89.
- [32] J.-P. Lee, H.-S. Jeon, D.-S. Moon, and B. S. Bae, "Threshold voltage and ir drop compensation of an amoled pixel circuit without a v_{DD} line," *IEEE Electron Device Letters*, vol. 35, no. 1, pp. 72–74, 2014.
- [33] X. Zhang, T. Tong, D. Brooks, and G.-Y. Wei, "Supply-noise resilient adaptive clocking for battery-powered aerial microrobotic system-on-

- chip in 40nm cmos,” in *Proceedings of the IEEE 2013 Custom Integrated Circuits Conference*. IEEE, 2013, pp. 1–4.
- [34] X. Zhang, T. Tong, D. Brooks, and G. Y. Wei, “Evaluating adaptive clocking for supply-noise resilience in battery-powered aerial micro-robotic system-on-chip,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 61, no. 8, pp. 2309–2317, 2014.
- [35] Texas Instruments, “LMZ10501 1-A SIMPLE SWITCHER® Nano Module With 5.5-V Maximum Input Voltage,” <http://www.ti.com/product/LMZ10501>.
- [36] Reddi, V.J. and Kanev, S. and Campanoni, S. and Smith, M.D. and Wei, G.Y. and Brooks, D., “Voltage Smoothing: Characterizing and Mitigating Voltage Noise in Production Processors Using Software-Guided Thread Scheduling,” in *Proc. Annual IEEE/ACM Int. Symp. on Microarchitecture*, 2010.
- [37] Y. Kim and L. K. John, “Automated di/dt stressmark generation for microprocessor power delivery networks,” in *Low Power Electronics and Design (ISLPED) 2011 International Symposium on*. IEEE, 2011.
- [38] R. T. Poggio, “Gpu ntc process variation compensation with voltage stacking,” in *Parallel Architectures and Compilation Techniques (PACT), International Conference on*, 2015.
- [39] E. Ebrahimi, R. T. Poggio, and J. Renau, “Sram voltage stacking,” in *Circuits and Systems (ISCAS), 2016 IEEE International Symposium on*. IEEE, 2016, pp. 1634–1637.
- [40] T. Tong, S. K. Lee, X. Zhang, D. Brooks, and G.-Y. Wei, “A fully integrated reconfigurable switched-capacitor dc-dc converter with four stacked output channels for voltage stacking applications,” *IEEE Journal of Solid-State Circuits*, vol. 51, no. 9, pp. 2142–2152, 2016.
- [41] K. Blutman, A. Kapoor, A. Majumdar, J. G. Martinez, J. Echeverri, L. Sevati, A. Van Der Wel, H. Fatemi, J. P. de Gyvez, and K. Makinwa, “A microcontroller with 96% power-conversion efficiency using stacked voltage domains,” in *VLSI Circuits (VLSI-Circuits), 2016 IEEE Symposium on*. IEEE, 2016, pp. 1–2.
- [42] K. Blutman, H. Fatemi, A. B. Kahng, A. Kapoor, J. Li, and J. P. de Gyvez, “Floorplan and placement methodology for improved energy reduction in stacked power-domain design,” in *Design Automation Conference (ASP-DAC), 2017 22nd Asia and South Pacific*. IEEE, 2017, pp. 444–449.
- [43] K. Blutman, H. Fatemi, A. Kapoor, A. B. Kahng, J. Li, and J. P. de Gyvez, “Logic design partitioning for stacked power domains,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2017.
- [44] R. Zhang, K. Mazumdar, B. H. Meyer, K. Wang, K. Skadron, and M. Stan, “A cross-layer design exploration of charge-recycled power-delivery in many-layer 3d-ic,” in *Design Automation Conference (DAC), 2015 52nd ACM/EDAC/IEEE*. IEEE, 2015, pp. 1–6.
- [45] K. Mazumdar and M. Stan, “Breaking the power delivery wall using voltage stacking,” in *Proceedings of the great lakes symposium on VLSI*. ACM, 2012, pp. 51–54.
- [46] Q. Zhang, L. Lai, M. Gottscho, and P. Gupta, “Multi-story power distribution networks for gpus,” in *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2016*. IEEE, 2016, pp. 451–456.
- [47] A. Zou, J. Leng, X. He, Y. Zu, V. J. Reddi, and X. Zhang, “Efficient and reliable power delivery in voltage-stacked manycore system with hybrid charge-recycling regulators,” in *2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*. IEEE, 2018, pp. 1–6.
- [48] A. Zou, J. Leng, X. He, Y. Zu, C. D. Gill, V. J. Reddi, and X. Zhang, “Voltage-stacked gpus: A control theory driven cross-layer solution for practical voltage stacking in gpus,” in *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2018, pp. 390–402.
- [49] NVIDIA, “Whitepaper nvidia’s next generation cudatm compute architecture: Fermi.”
- [50] J. Ervin, A. Balijepalli, P. Joshi, V. Kushner, J. Yang, and T. J. Thornton, “Cmos-compatible soi mesfets with high breakdown voltage,” *IEEE Transactions on Electron Devices*, vol. 53, no. 12, pp. 3129–3135, 2006.
- [51] A. Suchanek, Z. Chen, and J. Di, *Asynchronous circuit stacking for simplified power management*. IEEE, 2018.
- [52] J. Leng, T. Hetherington, A. ElTantawy, S. Gilani, N. S. Kim, T. M. Aamodt, and V. J. Reddi, “Gpuwatch: enabling energy optimizations in gpgpus,” in *ACM SIGARCH Computer Architecture News*, vol. 41, no. 3. ACM, 2013, pp. 487–498.
- [53] J. Gu and C. H. Kim, “Multi-story power delivery for supply noise reduction and low voltage operation,” in *Proceedings of the 2005 international symposium on Low power electronics and design*. ACM, 2005, pp. 192–197.
- [54] J. Zhou, C. Wang, X. Liu, and M. Je, “Fast and energy-efficient low-voltage level shifters,” *Microelectronics Journal*, vol. 46, no. 1, pp. 75–80, 2015.
- [55] S.-C. Luo, C.-J. Huang, and Y.-H. Chu, “A wide-range level shifter using a modified wilson current mirror hybrid buffer,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 61, no. 6, pp. 1656–1665, 2014.
- [56] K.-H. Koo, J.-H. Seo, M.-L. Ko, and J.-W. Kim, “A new level-up shifter for high speed and wide range interface in ultra deep sub-micron,” in *Circuits and Systems, 2005. ISCAS 2005. IEEE International Symposium on*. IEEE, 2005, pp. 1063–1065.
- [57] T. S. Joshi and P. M. R. Nerkar, “A wide range level shifter using a self biased cascode current mirror with ptl based buffer,” in *IJCA Proceedings on National Conference on Emerging Trends in Advanced Communication Technologies*, 2015, pp. 8–12.
- [58] A. Hasanbegovic and S. Aunet, “Low-power subthreshold to above threshold level shifters in 90nm and 65nm process,” *Microprocessors and Microsystems*, vol. 35, no. 1, pp. 1–9, 2011.
- [59] B. Aggarwal, M. Gupta, and A. K. Gupta, “A comparative study of various current mirror configurations: Topologies and characteristics,” *Microelectronics Journal*, vol. 53, pp. 134–155, 2016.
- [60] M. Kumar, S. K. Arya, and S. Pandey, “Level shifter design for low power applications,” *arXiv preprint arXiv:1011.0507*, 2010.
- [61] E. Ebrahimi, R. T. Poggio, and J. Renau, “Level shifter design for voltage stacking.”
- [62] M. D. Seeman, “Analytical and practical analysis of switched-capacitor dc-dc converters,” Master’s thesis, EECS Department, University of California, Berkeley, Sep 2006. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-111.html>
- [63] —, *A design methodology for switched-capacitor DC-DC converters*. University of California, Berkeley, 2009.
- [64] X. Wang, J. Xu, Z. Wang, K. J. Chen, X. Wu, Z. Wang, P. Yang, and L. H. Duong, “An analytical study of power delivery systems for many-core processors using on-chip and off-chip voltage regulators,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 9, pp. 1401–1414, 2015.
- [65] H. Li, J. Xu, Z. Wang, R. K. Maeda, P. Yang, and Z. Tian, “Workload-aware adaptive power delivery system management for many-core processors,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2017.
- [66] H.-P. Le, S. R. Sanders, and E. Alon, “Design techniques for fully integrated switched-capacitor dc-dc converters,” *IEEE Journal of Solid-State Circuits*, vol. 46, no. 9, pp. 2120–2131, 2011.
- [67] P. M. Ponce, D. Schröder, and W. H. Krautschneider, “Trade-off study on switched capacitor regulators for implantable medical devices.”
- [68] E. Candan, “A series-stacked power delivery architecture with isolated converters for energy efficient data centers,” 2014.
- [69] M. Rodrigues, N. Roma, and P. Tomás, “Fast and scalable thread migration for multi-core architectures,” in *2015 IEEE 13th International Conference on Embedded and Ubiquitous Computing*. IEEE, 2015.
- [70] “Ngspice, howpublished = <http://ngspice.sourceforge.net/>, note = Accessed: 2018-12-31.”
- [71] A. Bakhoda, G. L. Yuan, W. W. Fung, H. Wong, and T. M. Aamodt, “Analyzing cuda workloads using a detailed gpu simulator,” in *2009 IEEE International Symposium on Performance Analysis of Systems and Software*. IEEE, 2009, pp. 163–174.
- [72] J. Leng, T. Hetherington, A. ElTantawy, S. Gilani, N. S. Kim, T. M. Aamodt, and V. J. Reddi, “Gpuwatch: Enabling energy optimizations in gpgpus,” in *Proceedings of the 40th Annual International Symposium on Computer Architecture*, ser. ISCA ’13. New York, NY, USA: ACM, 2013, pp. 487–498. [Online]. Available: <http://doi.acm.org/10.1145/2485922.2485964>
- [73] S. Che, M. Boyer, J. Meng, D. Tarjan, J. W. Sheaffer, S.-H. Lee, and K. Skadron, “Rodinia: A benchmark suite for heterogeneous computing,” in *International Symposium on Workload Characterization*, 2009.
- [74] “NVIDIA, howpublished = <https://developer.nvidia.com/>, note = Accessed: 2018-12-31.”
- [75] “Graphics cards voltage regulator modules (vrm) explained,” <https://www.geeks3d.com/20100504/tutorial-graphics-cards-voltage-regulator-modules-vrm-explained>.
- [76] A. Zou, J. Leng, Y. Zu, T. Tong, V. J. Reddi, D. Brooks, G.-Y. Wei, and X. Zhang, “Ivory: Early-stage design space exploration tool for integrated voltage regulators,” in *Proceedings of the 54th Annual Design Automation Conference 2017*. ACM, 2017, p. 1.
- [77] R. Ge, R. Vogt, J. Majumder, A. Alam, M. Burtcher, and Z. Zong, “Effects of dynamic voltage and frequency scaling on a k20 gpu,” in *Parallel Processing (ICPP), 2013 42nd International Conference on*. IEEE, 2013, pp. 826–833.