Managing Mismatches in Voltage Stacking with CoreUnfolding

EHSAN K. ARDESTANI and RAFAEL TRAPANI POSSIGNOLO, UC Santa Cruz

JOSE LUIS BRIZ, Universidad Zaragoza

JOSE RENAU, UC Santa Cruz

Five percent to 25% of power could be wasted before it is delivered to the computational resources on a die, due to inefficiencies of voltage regulators and resistive loss. The power delivery could benefit if, at the same power, the delivered voltage increases and the current decreases. This article presents *CoreUnfolding*, a technique that leverages voltage Stacking to improve power delivery efficiency. Our experiments show that about 10% system-wide power can be saved, the voltage regulator area can be reduced by 30%, di/dt improves 49%, and the power pin count is reduced by 40% ($\approx 20\%$ reduction in packaging costs), with negligible performance degradation.

CCS Concepts: • **Hardware** → **Power networks**; *Chip-level power issues*; *Power grid design*;

Additional Key Words and Phrases: Voltage stacking, power delivery network, microarchitecture

ACM Reference Format:

Ehsan K. Ardestani, Rafael Trapani Possignolo, Jose Luis Briz, and Jose Renau. 2015. Managing mismatches in voltage stacking with CoreUnfolding. *ACM Trans. Archit. Code Optim.* 12, 4, Article 43 (November 2015), 26 pages.

DOI: <http://dx.doi.org/10.1145/2835178>

1. INTRODUCTION

The increase in electric current draw poses growing challenges to the design and implementation of the power delivery system in modern computing devices. Among these challenges are the need for more pins and pads in the package dedicated to power delivery, elevated I^2R power loss, higher IR and Ldi/dt voltage noise, and bigger, less efficient voltage regulators (VRs).

On-package and integrated on-die voltage regulators have been proposed [Kim et al. 2012; Ramadass et al. 2010; Chang et al. 2010; Burton et al. 2014] to offer a finer temporal granularity and better response to di/dt , which boosts the benefit of dynamic frequency and voltage scaling (DVFS). Nonetheless, on-chip VRs suffer from less efficiency compared to off-chip VRs [Kim et al. 2008, 2012; Burton et al. 2014].

This work was supported in part by the National Science Foundation under grants CNS-1059442-003, CNS-1318943-001, CCF-1337278, and CCF-1514284. Any opinions, findings, and conclusions or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the NSF.

Briz's contribution was partially supported by grants TIN2013-46957-C2-1-P and T48 research group (Aragón Gov. and European ESF), and HiPEAC-3 NoE (European FET FP7/ICT 287759).

Authors' addresses: E. K. Ardestani, R. T. Possignolo, and J. Renau, Computer Engineering Department, Jack Baskin School of Engineering, University of California Santa Cruz, 1156 High Street, Santa Cruz, CA 95064; emails: {eka, rafaeltp, renau}@soe.ucsc.edu; J. L. Briz, Depto. de Informática e Ingeniería de Sistemas, Escuela de Ingeniería y Arquitectura, University of Zaragoza, Pedro Cerbuna, 12, 50009 Zaragoza - España; email: briz@unizar.es.

Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2015 ACM 1544-3566/2015/11-ART43 \$15.00

DOI: <http://dx.doi.org/10.1145/2835178>

A possible remedy to the power delivery challenges is voltage stacking [Gu and Kim 2005; Rajapandian et al. 2006; Shenoy et al. 2011a; Lee et al. 2012; Candan et al. 2014], in which multiple components are “stacked” in a series configuration on top of each other. The current passing through one component is recycled by the next component, reducing the total current demand. To maintain the same power delivered to each component, the stack is supplied with a higher voltage. In other words, voltage stacking delivers the same power to the logic in a different form: higher voltage but lower current. This alleviates the aforementioned challenges; for example, VRs can be designed for better efficiency delivering the same power at higher voltage and lower current [Shenoy et al. 2010; Lee et al. 2012; Kim et al. 2011].

A main challenge in voltage stacking, which we address in this article, is load mismatch between the stack levels. Consider a configuration with two stacked cores: running different applications on each core can result in cores demanding different power. Because the same current passes through all the components in the stack, one core can restrict power delivered to the other, potentially inducing timing failure. The load mismatch can be managed by adding extra VRs [Shenoy et al. 2011a]. However, this requires VRs with equal size to VRs in conventional power delivery solutions to guarantee handling of worst-case mismatches. Also, such configuration can potentially result in increased system-wide power consumption. These complications limit the applicability of voltage stacking because the only real benefit seems to be reducing the number of power delivery pads in the package [Zhan and Sapatnekar 2008] and voltage noise [Gu and Kim 2005].

This article proposes *CoreUnfolding*, a novel method to use voltage stacking *within* each core. Not only does *CoreUnfolding* have the same advantages as voltage stacking (reduced number of power delivery pads and improved di/dt , voltage drop, and noise) but also it improves the system-wide energy efficiency and reduces the area dedicated to voltage regulators. The novelty resides in guaranteeing a maximum load mismatch between the stacks, achieved by stacking two groups of functional units within a core, clustered based on the correlation and magnitude of their power consumption. This allows an efficient system that uses a much smaller VR. The proposed microarchitecture easily scales to multicore homogeneous and heterogeneous configurations and works with single-threaded, multithreaded applications. It is also compatible with power gating techniques, since the stacking occurs within a core while cores are not stacked and therefore can be independently turned on/off.

CoreUnfolding partitions the functional units in a core into two groups: *Header* and *Footer*. Thus, the data cache could be in the *Header* group, and the reorder buffer in the *Footer* group. The partition aims to balance the load or power consumption in the two groups under different workloads. We study the correlation of power consumption between components among functional units. We observe that the activity and power consumption among some functional units tightly correlate, which makes them good candidates for stacking.

In some infrequent cases, the mismatch between the levels of the stack could create a high load imbalance between the *Header* and *Footer* groups, which requires some sort of safeguard. We resort to dynamic load-balancing techniques to match up the power consumption of *Header* and *Footer* for this uncommon case. In this article, we cap these mismatches by using two mechanisms: “dummy activity” and DVFS. Consequently, the additional voltage regulator (*sVR*) is only required to match as little as 10% of the total power consumption. As a result, the proposed solution allows for a much smaller total voltage regulator size down to 70% of the size in the baseline. The benefit would be higher considering the thermal implications.

Equally important is the power savings of the system. *CoreUnfolding* works with both on-chip and off-chip voltage regulators. In our experiments, the VR needs to support up

to 60% of the current that the baseline VR provides, which results in up to 10% overall power savings.

The main contributions of this article are as follows:

- The idea of applying voltage stacking within a core
- Using power average and transient power correlation to decide which units to stack
- Capping the maximum mismatch to allow smaller and more efficient VRs

2. RELATED WORK

Voltage stacking has been proposed as a paradigm shift to tackle the problem of increased challenges of power delivery as the technology scales further. Voltage stacking in the logic allows for operation at higher voltage and lower current for the same power budget. This does not reduce the energy consumption of the logic. However, it allows for the power delivery to operate in a more efficient way [Shenoy et al. 2010; Kim et al. 2011].

Voltage stacking can improve power delivery systems [Rajapandian et al. 2006; Gu and Kim 2005; Zhan and Sapatnekar 2008], but managing the load mismatch between the stacked devices remains a challenge. Lee et al. [2012] observe this issue through real measurement but do not propose a solution. Just stacking the components like Lee et al. [2012] would result in timing failure. Zhan and Sapatnekar [2008] use voltage stacking to reduce the number of pins in 3D ICs. To avoid increased power consumption due to the load mismatch between stacked levels, they partition the modules in the floorplan and assign them to appropriate level, and use extra voltage regulators to manage the load mismatch. However, their method does not guarantee bounded mismatch, and the extra voltage regulators increase the area overhead. Gu and Kim [2005] use voltage stacking for voltage noise reduction. To avoid using an extra voltage regulator, they use a “digital voltage regulator,” which introduces extra activity to balance the load between levels. We call this “dummy activity.” Only relying on dummy activity for the whole processor, if even possible, would incur cycle time penalty.

Voltage stacking has been proposed in the server level for data centers [Candan et al. 2014]. To regulate the voltage for each of the stack levels, differential power converters were employed. Instead of trying to regulate the voltages to the nominal value, this proposal uses acceptable bands of voltages. The stacking was made at the mother board level, that is, at the 12V range. With very regular workloads, they are able to improve the voltage conversion efficiency considerably.

Multilevel ladder voltage converters have been proposed to be used in voltage-stacked circuits [Kesarwani et al. 2013; Schaef and Stauth 2016; Lee et al. 2015]. Such topologies have the advantage of delivering multiple-output voltages from a single converter, with high efficiency. In this article, we leverage the results of a fully integrated switched capacitor ladder DC-DC converter [Lee et al. 2015] that yields great efficiency when regulating multiple levels of voltage.

Through voltage stacking, this work proposes a framework to manage and bound the load mismatch between the stacked levels in a processor. Similar to Zhan and Sapatnekar [2008], we use extra VRs to regulate the mismatches but cap the demand by using dummy activity for rather uncommon high mismatches. As a safeguard, we use DVFS to scale down the power in case of extreme mismatch, so the scaled mismatch can be regulated by the small voltage regulator. Consequently, the whole system can benefit from lower current draw. This leads to system-wide power savings, better voltage margin, and reduced complexity and area of the voltage regulators in the system.

Fig. 1. (a) For the same power output, reducing the current output of a given VR allows more efficient VRs. (b) In general, VRs are more efficient when optimized (at design time) to lower current and higher output voltage (i.e., smaller step down). This is illustrated by VR1 [Linear Technology Inc. 2014] and VR2 [Linear Technology Inc. 2012] (from the same company), which are designed for the same power output but with different efficiencies.

3. BACKGROUND

3.1. Voltage Regulators

An ideal VR delivers power from a power source to the load without any losses. Unfortunately, the regulator itself consumes power with some *conversion efficiency*, defined as the ratio of power delivered to the load by the regulator to the total power into the regulator. Regulator losses are dominated by switching power and resistive losses, which depend on the size of the switching power transistors, switching frequency, and load conditions (e.g., load current and voltage levels) [Kim et al. 2008].

The benefit of higher voltage and lower current output can be considered through two knobs: static and dynamic. At runtime, drawing less current from the VR could reduce the resistive loss and lead to better efficiency, in particular for high loads. In very light loads, the switching power loss can become dominant. Figure 1(a) shows the efficiency against the output current of an Intel-compliant off-chip VR for desktop/server applications. Increasing the voltage from 0.8V to 1.6V, and reducing output current by half, increases the efficiency in that particular operating point by 10%.

Nonetheless, designing VR modules involves thorough optimization between a number of parameters, including the switching and resistive loss. VRs designed to deliver lower current, with higher output voltage, are, in general, more efficient [Wei 2004]. We refer to this as the static knob. Figure 1(b) shows the efficiency for two VRs from the same manufacturer with different maximum output currents (8A and 4A) [Linear Technology Inc. 2014, 2012]. For the same power, it is possible to increase the efficiency by over 10% when the VR with smaller maximum output is used. Lower heat dissipation due to the higher efficiency can have a secondary benefit on the package footprint and cost as well.

Note that those are not fundamental reasons for this gap in efficiency. Those are empirical observations from current commercial voltage regulators. Design time optimizations and new VR designs could change that scenario, but there is no indication of such in the near horizon. Nevertheless, from the observations made in this section, it seems plausible to consider that, by doubling the voltage and reducing the current by half, it would be possible to improve by 10% to 15% the efficiency of VRs in a real design. To be conservative, we consider that 10% improvement in efficiency is possible in the use cases found; actual gains may differ.

Fig. 2. Different power delivery schemes. Both schemes deliver the same amount of power, though in different forms. For the same loads (a) delivers the power in the form of $V \times I$, and (b) in the form of $2V \times \frac{I}{2}$. The latter one could be more efficient.

3.2. Voltage Stacking

Voltage stacking configures components in a series manner instead of the conventional parallel way, as shown in Figure 2. To deliver the same power, the stack is supplied with a higher voltage. In general, a stack of n components, each supplied with V_{dd} in the conventional configuration, is supplied with $n \times V_{dd}$. Charge passing through one component is recycled by the next component. Hence, the total current demand is reduced, ideally to $\frac{1}{n}$ [Shenoy et al. 2011b].

Voltage stacking can increase the efficiency of the power delivery subsystem, as (1) it requires the power delivery to operate at a higher voltage and lower current, (2) it requires a lower step-down ratio, and (3) it results in less IR drop across the power delivery. Note that the logic in the best case would consume the same amount of power as it does in conventional power delivery configuration.

While voltage stacking increases the efficiency of power delivery network, it suffers from load mismatch between stacked levels. This is because current demand (Z in Figure 2) for each component could change at runtime, for example, due to changes in the workload behavior. However, the same current flows through all the levels in the stack. This results in different voltage drops across each level, where the delivery should supply each level with the same voltage. The load mismatch and deviation of V_{mid} from the supposed value ($\frac{2V_{dd}}{2} = V_{dd} = V_{mid}$) can be seen as voltage noise for each component, and could even result in timing failure. Managing the load mismatch and maintaining the overall benefit of the voltage stacking has remained an open problem. This article addresses this challenge. In our experiments, we focus on two-level stacking.

Body biasing of the transistors connected to V_{mid} can introduce issues as those transistors are not connected to V_{dd} or Gnd rails. To support independent body bias for each stack level, triple well or Silicon-on-Isolator (SOI) technology must be used. Those processes provide insulation between different wells and allow the differential body biasing for the transistors connected to V_{mid} . They are commonly offered by all the major fabs.

4. COREUNFOLDING

We apply voltage stacking within cores. This is based on the observation that current draws, that is, power consumption, of different functional units (FUs) in a core correlate with each other. We exploit this correlation to guide the stacking. FUs with well-correlating power consumption are good candidates for stacking at design time, minimizing the chance of a mismatch between the two stacked levels. We partition the FUs according to their power magnitude and correlation, and assign each partition

Fig. 3. Power consumption correlation matrix of functional units can guide stacking among them. In this example, *fiw* and *frf* can be stacked.

to a level in the stack. To ensure correct functionality at runtime, we use additional dynamic techniques with minimal overhead.

4.1. Design-Time Load Balancing Across Stacked Levels

Figure 3 shows the correlation between power consumption of four functional units in an out-of-order processor running the *applu* benchmark. Each line and column represent a FU. The intersection of FUs shows the correlation of the two FUs; points distributed closer to the diagonal indicate better correlation. It shows the correlations for data TLB (*dtlb*), data cache (*dcache*), floating-point instruction window (*fiw*), and floating-point register file (*frf*). In this set, *dtlb* and *dcache* correlate well with each other ($\text{cor} = 0.981$). *fiw* and *frf* also correlate very well with each other ($\text{cor} = 0.999$). However, there is no good correlation between *fiw* and *dcache* ($\text{cor} = 0.352$) or *fiw* and *dtlb* ($\text{cor} = 0.362$). This correlation among the units, in this case, suggests stacking of *dtlb* with *dcache*, and *fiw* with *frf*. As will become clear later in the article, this one-on-one correlation should be seen as a presentation of concept; *CoreUnfolding* will aggregate the power consumption of different FUs instead of comparing pairs of FUs isolated.

The magnitude of power consumption of the units should be considered in the partitioning as well. For example, even though the power consumption of *dtlb* and *dcache* correlates well, a stack comprising only these two units would still have load mismatch, as the average power consumption of *dcache* is three to four times the power consumption of *dtlb*.

For the remainder of this article, we refer to the top level in the stack as *Header* and the lower level as *Footer*. Figure 4 shows an example of a two-level stacked configuration. The two partitions will be stacked one on the top of the other, while the middle layer between them (V_{mid}) is shared across the partitions. This simplifies stacking of the units, as more units can contribute in matching the average and transient power consumption of the two levels. It is also simple to implement on the chip, as the middle layer between the stacked units will be treated in a similar way as other power delivery nets (e.g., V_{dd}) in terms of decoupling capacitance (decap) specification and implementation.

Fig. 4. The functional units will be stacked based on their power consumption and the correlation of power among units. The middle layer for all the units will form a V_{mid} power net.

Fig. 5. Simulation results of the level converter model in SPICE for 40nm CMOS technology. The converter delay is comparable with a *not* gate with fan-out of 4 (FO4).

4.2. Interlevel Communication

Communication between functional units located at different partitions, and hence different levels, requires a level converter to adjust the logic's electrical level to that of the destination. Level converters have been studied in several works (e.g., see Ishihara et al. [2004]). We use a level converter similar to the one presented in Gu and Kim [2005] (see Figure 5).

Level converters should be utilized judiciously, as they consume power, add area, and delay overhead. This actually presents the following tradeoff: finer FU granularity would provide more FUs or components per design, allowing a better matching partition. However, increasing the number of FUs would result in an increase in the number of level converters.

To obtain the characteristics of the level converters, we implement them in SPICE using a 45nm technology [Zhao and Cao 2006]. To fairly consider leakage, we assume a

clock width of 23 fan-out4 (FO4) gates [Choudhary et al. 2011]¹ and scale the leakage of each gate accordingly. Figure 5(c) shows the transient simulation result for the *Header* to *Footer* converter using SPICE. The simulation also adds 10% voltage noise on the V_{mid} voltage to ensure robust functionality. The proposed level converter consumes 54uW of power at 1V, which is about 10 times the power consumption of a *not* gate at the same technology. While the power consumption overhead is negligible (less than 0.1% of the total power in our experiments), the level converter adds a delay as big as one FO4 *not* gate. It might not be possible to absorb such a delay in the critical path of performance. For simplicity, we assume a worst-case situation with a full-cycle overhead when talking between voltage domains. Section 6 will discuss the exact cost.

To minimize the overhead of interlevel communication, particularly for the critical path of performance, we consider the communication cost among blocks for any partitioning solution in the partitioning process. Our experiments show that splitting the core into 20 to 30 FUs would result in a partitioning allowing good correlation and small power overhead due to a level converter.

4.3. Systematic Partitioning

We define the stacking problem as partitioning of the FUs in a core into two *Header* and *Footer* partitions. The partitioning has to satisfy three goals: (1) both partitions have matching average power consumption, (2) total power consumption of both partitions correlates well with each other, and (3) the communication of blocks between partitions is minimized. Note that goals 1 and 2 refer to power, but goal 1 is related to the average power (i.e., average over the execution time), while goal 2 is related to instantaneous power (i.e., average power over a delta time; we use 100 cycles).

To estimate the cost of the communication between blocks (Σcut), we have verified which FUs communicate with each other and what is the weight of such communication (i.e., how many wires are used) considering a 64-bit architecture, namely, $w_{i,j}$ for FUs i and j . For each solution, if i and j are in the same cluster, we add $w_{i,j}$ to the total cost; if they are in the same cluster, we add 0.

Any classical partitioning algorithm can be adopted to partition the functional units into two *Header* and *Footer* groups. We choose the Genetic Algorithm (GA). Each individual, representing a valid solution, comprises a string of binary genes; “0” in location i in an individual means that the FU_i belongs to the *Footer*, while “1” means it belongs to the *Header*.

The cost function to be minimized is a superposition of the mismatch in power consumption, correlation between the two partitions, and the cost of interlevel communication:

$$CostFunc = \alpha \times \Delta Pow + \beta \times \Delta Cor + \gamma \times Comm_{critical} \quad (1)$$

$$\Delta Pow = (Mean(Power_{Header}(t)) - Mean(Power_{Footer}(t))) / Mean(Power_{total}(t)) \quad (2)$$

$$\Delta Cor = 1 - Cor(Power_{Header}(t), Power_{Footer}(t)) \quad (3)$$

$$Comm_{critical} = \Sigma cut / (\#FUs)^2. \quad (4)$$

4.4. Runtime Load Balancing

Ideally, partitioning the FUs, based on the magnitude (average power) and correlation of their instantaneous power consumptions, into *Header* and *Footer* groups and

¹The actual range in the reference is from 23 to 33 FO4. We consider the lowest width to account for the worst-case impact.

stacking them would result in no load mismatch between the *Header* and *Footer* during the execution. However, in reality, the utilization of the FUs differs across workloads and varies over time. This results in transient mismatches (6% on average in our experiments) that have to be managed.

We categorize mismatches by their power magnitude. A mismatch is said to occur when there is a difference in power of more than 10% (that is considered to be an acceptable fluctuation during the operation of a chip):

Small Magnitude: Mismatches up to 20% of the power consumption are considered small.

Moderate Magnitude: Mismatches in the 20% to 50% range are considered moderate.

High Magnitude: Mismatches of more than 50% of the power consumption are considered high.

To manage these transient mismatches, we propose the use of circuit and microarchitectural mechanisms. Mismatches that are very short (i.e., $< \approx 100$ clock cycles) are too fast for microarchitectural mechanisms; therefore, we rely on decaps to filter out those high-frequency mismatches. Note that the total available on-die decoupling capacitance now has to be split between V_{dd} , Gnd , and V_{mid} ; thus, we propose the use of on-package decaps as well.

The **Small Magnitude** mismatches are dealt with by a secondary VR (sVR) for regulation (details in Section 4.8). sVR is a small and fast voltage regulator that can efficiently balance the V_{mid} voltage. The reason to use the sVR only for small-magnitude mismatches is to avoid the need for a full-sized VR (as occurred in previous voltage stacking proposals).

To limit the maximum current output that sVR has to supply, we use “dummy activity” (DA) for **Moderate Magnitude** mismatches. DA introduces extra activity in a few selected FUs in the level that has less power to match up the power consumption of the group.

Relying too much on DA might require extra logic, which can sacrifice the cycle time. In case of a **High Magnitude** mismatch, DVFS brings down the overall power consumption and the magnitude of mismatch within the range that sVR can regulate, since it is now relatively stronger. This means that managing high-magnitude mismatches implies performance degradation as well. This is an additional use of DVFS, as an addition to current uses of DVFS that include thermal and power management.

High Magnitude mismatches did not occur in our simulations, and **Moderate Magnitude** mismatches occurred a mere 0.5% of time in the execution of all the applications. In any case, voltage scaling is proposed as a safety backup to guarantee correctness in hypothetical extreme cases, like a malicious program running. The sVR is present and operating at all times.

To estimate the mismatch, a controller based on the voltage of V_{mid} is used to trigger DA and DVFS. If the voltage is above/below the set thresholds, custom circuitry in the core activates DA on specific FUs, or DVFS. The on-package decaps are sized to filter very short transients, giving time for DA and DVFS to actuate. The choice of on-package decaps provides large enough capacitors for the task while avoiding the need of using pins for V_{mid} and increase in area.

4.5. Multicore

CoreUnfolding applies the stacking within a core, and the interface of the core to the outside remains the same. The cores in a multicore configuration will be connected in a parallel way, similar to the traditional power delivery system but with higher voltage and lower current. Therefore, multithreaded applications would run on a multicore *CoreUnfolding* system with no extra consideration. To implement multicore

Fig. 6. In a multicore configuration, cores are connected in parallel. The stacking is applied within each core, and the two partitions are flipped across cores so that the minimal possible bias toward one partition in a core can be canceled out by the next core.

CoreUnfolding, each *type* of core in the system goes through the partitioning framework once, and cores with the same type replicate the same partitioning. Note that we distinguish between the core (including private caches) and the shared memory levels and structures such as memory controller.

We notice that swapping *Header* and *Footer* across the cores would provide better results, as the minor bias toward one partition in one core can be canceled out by the next core of the same type (Figure 6). As we observe in our experiments, this yields a very good balance between *Header* and *Footer*.

Another possibility is to design each core with its own V_{mid} , and thus each core acts as an isolated core. The small bias will remain but will not accumulate. This option is required when per-core DVFS is needed.

4.6. Other Voltage Domains

The LLC and memory controller in many processors reside on a different voltage domain than the core. These structures consume a small fraction of the processor power. Nonetheless, the structures within each of these units can be stacked. IOs and PLLs also are among the structures with different voltage domains. However, we do not evaluate stacking for these voltage domains.

4.7. Floorplanning

To avoid adding wire delays between FUs, *CoreUnfolding* does not affect floorplanning. Each FU will either tap V_{dd} and V_{mid} or V_{mid} and Gnd , through a global versus local approach. A good analogy is a scheme with power gating, where a block of logic has its own local power rail. The connection scheme is depicted in Figure 7(b); the dashed lines represent V_{mid} . The analogy with the power-gating technique is shown in Figure 7(a). We keep the power-gating transistors, although that is not strictly necessary for *CoreUnfolding* alone. Also, level converters require both domains, so we also pull wires directly from the global level power rails when needed. This scheme does not incur extra overhead to add a new local power rail. Finally, the addition of the global V_{mid} can be mitigated by the reduction of metal dedicated to V_{dd} and Gnd , since the current through those rails is decreased.

CoreUnfolding requires either a silicon-on-insulator or triple-well technology to isolate wells across the chip [Lee et al. 2012], since silicon in each cluster will be biased with different voltage levels. Figure 7(c) shows the floorplan that was used by ESESC [Ardestani and Renau 2013] in our simulations. The shaded FUs are the FUs in the header group.

The choice of multiple power rails avoids the need for routing the three crossing domain wires of the level converter (Figure 5), avoiding overheads in routability and

Fig. 7. Instead of changing the floorplan, *CoreUnfolding* uses an approach inspired in power gating to avoid overheads.

Fig. 8. Power delivery in *CoreUnfolding* (c and d) versus conventional architecture (a and b). The lower maximum current demand in the *CoreUnfolding* configurations, which is around 60% of the conventional configurations, results in smaller and more efficient voltage regulators, reduced di/dt , and better power delivery. The colored voltage regulators in the figures specify the VR with DVFS functionality.

reducing the costs of wires. The drawback is that it adds cost in area for the extra power rail. Our approach adds one extra rail per cell row.

4.8. Power Delivery Architecture

Figure 8 shows the *CoreUnfolding* configurations against conventional ones. We categorize the voltage regulators regarding their functionality. Step-down regulators perform a fixed step-down conversion from the input to output; they are specified in Figure 8 by dashed lines. Voltage regulator modules with a DVFS controller provide multiple voltage output, in addition to step-down. They are specified by a solid fill in Figure 8. Besides, either of the voltage regulator types can be implemented off the chip or

integrated on the chip. Next we explain each power delivery architecture. The supply voltage to the system is usually 12V, or 3.7V for battery-operated devices.

Conventional off-chip: As shown in Figure 8(a), the off-chip VR (located on the motherboard) performs step-down, usually 12-to- V_{core} , or 3.7V-to- V_{core} for battery-operated devices. As an example, we assume 12V input voltage and maximum $V_{core} = 1.1V$. The VR can output different voltages as requested by the processor. Most processors use a similar configuration.

Conventional on-chip: In this configuration, part of the voltage conversion functionality is moved into the chip. The motherboard's VR only performs step-down to 2 to 3V [Burton et al. 2014]. The integrated voltage regulator (iVR) provides the final step-down to the range requested by the processor. Figure 8(b) shows this architecture, an example of which is Intel's Haswell [Burton et al. 2014]; despite being called integrated, depending on the actual implementation, they may use an on-package inductor.

4.8.1. CoreUnfolding. By increasing the supply voltage, *CoreUnfolding* improves power delivery for both off-chip and on-chip voltage regulation configurations. The benefit could be gained through static (design time) and dynamic (runtime) knobs. The lower maximum current demand allows for design time optimization that leads to **smaller area** and **higher efficiency**, and the lower average current demand at runtime allows for less resistive loss in the delivery network. Note that *CoreUnfolding* is not dependent on a particular implementation of VRs but, naturally, can benefit from better VR implementations. The regulation of the V_{mid} is performed by *sVR* and can also be into the package or fully integrated into the die, since it only performs a low-power output step-down from the main input voltage to the core.

CoreUnfolding off-chip: In the case of an off-chip VR, *CoreUnfolding* enables the VR to benefit from a lower step-down ratio [Piqu and Bergveld 2012]. This is shown in Figure 8(c). Note that the VR output voltage is now twice.

CoreUnfolding on-chip: Figure 8(d) shows *CoreUnfolding* with an integrated voltage regulator. In this configuration, the VR on the motherboard only performs the step-down; still, a smaller step-down can be used, since the final supply voltage is higher. In this case, it makes sense to merge iVR and *sVR* into a ladder voltage converter (i.e., a converter capable of supplying more than one output voltage) [Lee et al. 2015].

VR has to be designed for the worst case, which is observed when there is a 20% load imbalance between the stack levels (our cap); in particular, the primary VR (either the iVR in the on-chip configuration or the VR in the off-chip configuration) will be most loaded when the *Header* group consumes 20% more power. The following formulates the ratio of **maximum** output current that the voltage regulators have to provide. Note that for simplicity, we ignore the efficiency of the voltage regulators in computation of the current ratios. The nomenclature is in accordance with Figure 8(a) and Figure 8(c).

In a perfectly balanced condition, the current going through the logic in conventional and *CoreUnfolding* is

$$\gamma = I_H = I_F = 0.5\alpha. \quad (5)$$

I_H and I_F stand for current passing through the *Header* and *Footer* levels, respectively. Perfect balance is not always the case; the resulting current in the case of mismatch is

$$\gamma = \max(I_H, I_F). \quad (6)$$

We cap the demand from *sVR* to 20% using DA technique:

$$\gamma'' = |I_H - I_F| = 0.2 \times 0.5\alpha. \quad (7)$$

Table I. Architectural Parameters

Parameter	Value
#Cores	1 and 4
Freq	3.0GHz single core, 2.5GHz multicore
I\$	32KB 2w (2c hit) private
D\$	32KB 8w (3c hit) private
L2	256KB 16w (12c hit) private
L3	2MB×#cores 16w (12c hit) shared
Coherence	MESI
Mem.	180 cyc best case from LLC
BPred.	10 tab. ogehl 76Kb
Issue/ROB/IWin	4/256/32
Load/StoreQ	48/32

Therefore, the maximum current ratio between the *CoreUnfolding* and conventional configuration is as follows:

$$\gamma = \max(0.5\alpha + \gamma'', 0.5\alpha - \gamma'') = 0.5\alpha + \gamma'' \quad (8)$$

$$\frac{\gamma}{\alpha} = \frac{0.5\alpha + \gamma''}{\alpha} = \frac{0.5\alpha + 0.2 \times 0.5\alpha}{\alpha} = 0.6. \quad (9)$$

For the same power, the primary VR in *CoreUnfolding* is designed for 60% current output of VR in the conventional architecture. The average ratio is closer to 50%, as will be discussed in Section 6. As a result, it is smaller in area and can be optimized for more efficiency in the regulation.

The current demand from *sVR*, γ'' , would be much less than the VR in the conventional configuration. In fact, most of the time there would be no demand from *sVR*, and if there is a demand, it is capped to 20% of the power consumption of the group, or 10% of the power consumption of the chip (which represents $10\times$ less than previous voltage stacking approaches). As a result, *sVR* would be much simpler, smaller, and more efficient than the primary VR.

Recently, industry solutions, such as IBM Power8 [Toprak-Deniz et al. 2014] and Intel Haswell [Burton et al. 2014], are opting for integrated regulators, with power switches distributed across the die, which are more suited for quick responses to load demands and allow more fine-grained control of the voltage in different parts of the chip. For our experiments, we assume a primary VR such as Intel's Haswell regulator. For the secondary VR, we proposed the use of a switched-capacitor DC-DC stacked converter such as the one presented by Lee et al. [2015], which allows for very efficient stacked conversion.

5. SETUP AND EVALUATION METHODOLOGY

5.1. Architectural Simulation Setup

We use ESESC [Ardestani and Renau 2013] for our architectural evaluation, including performance, power, and temperature. The temperature dependency of leakage power is modeled as well. Table I lists the architectural parameter of the core and processor we model.

5.2. Power Delivery Network Setup

ESESC power consumption trace is used to generate a time-varying impedance model for each FU, which is then fed to a transient SPICE simulation that models the power delivery network; this methodology was proposed in Leng et al. [2014]. The off-chip power delivery network is similar to the model presented [Leng et al. 2014]; the on-chip power network is more detailed using the IBM Power Grid Benchmark [Nassif

Fig. 9. A detailed power distribution model is simulated to evaluate the voltage fluctuation in the middle rail.

Table II. Training Set

SPEC	bwaves, lbm, vpr, applu, equake, gcc, mgrid, mesa, art, astar, soplex, twolf, perlbench, swim, crafty, povray, milc, sphinx, dealII, wupwise, vortex, libquantum, mcf, gap, leslied
------	---

Table III. Evaluation Set

SPEC	bwaves, lbm, vpr, applu, equake, gcc, mgrid, mesa, art, astar, soplex, twolf, perlbench, swim, crafty, povray, milc, sphinx, dealII, wupwise, vortex, libquantum, mcf, gap, leslied
PARSEC	blackscholes, bodytrack, canneal, facesim, ferret, fluidanimate, swaptions, x264
SPLASH-2	ocean, fft, fmm, radix

2008] and includes on-package decoupling capacitors. Figure 9 shows the complete power delivery network used, with the simulation parameters. The on-chip sVR is also represented. We model our sVR as a fully integrated, two-level symmetric ladder switched-capacitor DC-DC converter [Lee et al. 2015], with efficiency of around 95% (depending on the load mismatch).

ESEC dumps the power trace every 100 cycles; we verified that this is enough of a resolution for the modeled power delivery network. Even using a maximum power density of $1W/mm^2$ is enough to sustain the transient mismatches that appear on the order of 30ns (100 cycles) because most of the capacitance needed comes from the on-package decoupling capacitors [Popovich et al. 2007]. On-package decaps are cheaper than dedicating more on-chip area for on-chip decaps, but due to pad parasitics, they present smaller decoupling power. To size these decaps, some preliminary transient SPICE simulations were used. The decap size was chosen so that V_{mid} fluctuation in a 100-cycle interval was smaller than 10%. Those simulations show that just one tantalum SMD capacitor [Vishay Sprague 2013] is enough to keep V_{mid} with less than 10% drop for 300 cycles for the whole multicore simulation.

5.3. Benchmarks

We use applications from SPEC CPU 2000 [Henning 2000], CPU 2006 [Henning 2006], PARSEC [Bienia et al. 2008], and SPLASH-2 [Woo et al. 1995] benchmark suites. Our training set consists of applications from the SPEC benchmarks (Table II), and the evaluation is done with all three benchmarks (Table III). This allows us to evaluate the impacts of partitioning in unknown workloads and show the robustness of the method.

Table IV. A Partitioning Solution

Partition	Functional Units
<i>Header</i>	FP RATs, Branch Predictor, Branch Target Buf., itlb, icache, dcache, FP inst. window (FPIW), int Register File (iRF)
<i>Footer</i>	Int RATs, Free Pools, Ret. Addr. Stack (RAS), Instruction Buffer, icache cntrl, dtlb, dcache Cntrl, Load Store Queue, Store Sets, Int Inst. Win., Exe. Units, FPU, FP Reg. File, Reorder Buf.

5.4. Evaluation Methodology

We go through two phases in our evaluation methodology. The first phase, **Training and Partitioning**, gathers power consumption information across benchmarks. Then it uses that information to extract the correlations and generates a partitioning solution. The next phase, **Validation**, performs full simulation to measure the mismatches with a fine-grained resolution. For all the simulations, we skip the first 2B instructions and simulate up to 10B instructions.

5.4.1. Training and Partitioning. We use sampling to gather power consumption information. The sampling simulates 50K instructions every 1M instructions and dumps power consumption information at the end of every sample. We only apply training for the SPEC applications. The power traces from all the benchmarks are aggregated and use the result based on which we perform the GA partitioning.

It is also possible to perform the training per benchmark for better results. Implementation of this, however, would require support for reconfiguration of the stacking. We do not consider reconfiguration in this work, but we will report the result to show how good the global partitioning performs (evaluated as *base* and *peak* partitioning in Section 6). For all the applications, the GA converges within the first 100 generations. The execution time on average was 3 hours. One can perform full simulation for training as well, but we noticed that sampling-based information was good enough to guide the training. Table IV shows the resulting partition, used through our evaluation. If we go back to Figure 3, we note that dtlb and dcache have very good correlation and indeed ended up in different clusters, as one would expect. Note, however, that this is not necessarily the case for all the FUs, since the clustering is done considering the combination of FUs rather than individual FUs. This partitioning solution resulted in around 1K level converters (considering a 64-bit architecture), which we consider to be a small overhead.

6. EVALUATION AND RESULTS

We first compare *CoreUnfolding* against the conventional configuration for area, power, performance, and di/dt . Then we evaluate the quality of the partitioning algorithm on which *CoreUnfolding* is based.

6.1. Area

Higher current draw requires larger inductance and capacitance components, as well as wider drivers. Multiphase VRs and shared current bus have been proposed to enable delivering higher current [Zhou et al. 2000; Huang et al. 2003]. Despite the considerable advances in increasing VR current density recently, especially with integrated VRs such as the Intel Haswell FIVR [Burton et al. 2014] and the IBM Power8 [Toprak-Deniz et al. 2014], VR still takes a lot of area. Considering an Intel Haswell-like processor with

Fig. 10. *CoreUnfolding* reduces the maximum as well as the average current drawn by each core. At peak, *CoreUnfolding* (c) needs 40% less current than baseline (b).

maximum current draw of roughly 100A and $31\text{A}/\text{mm}^2$ [Burton et al. 2014; Intel 2014], the *iVR* takes $\frac{100\text{A}}{31\text{A}/\text{mm}^2} = 4.12\text{mm}^2$.

Figure 10 shows the distribution of current in the baseline (b) and *CoreUnfolding* (c). For multicore applications, the results are reported per core. *CoreUnfolding* reduces the maximum and average current across all the applications. Based on the formulations in Section 4.8.1, the maximum current should be reduced by 40%; the smaller reductions in Figure 10 are due to the fact that these applications do not necessarily draw maximum current. Given the maximum current draw reduction, and considering 10% area overhead of *sVR* compared to the conventional VR, *CoreUnfolding* would allow for a reduction of 30% in VR die area (or 1.24mm^2 in the Haswell-like solution). This does not take into account on-package inductors needed, which take up to $20\text{mm} \times 8\text{mm}$ area in the FIVR solution [Burton et al. 2014].

At high current densities, thermal is a major limiting factor. A nonideal VR with 70% to 80% efficiency [Burton et al. 2014] with current density of $8\text{A}/\text{mm}^2$ would have power densities around $1.6\text{W}/\text{mm}^2$ ($8\text{A}/\text{mm}^2 \times 1\text{V} \times (100-80)\% = 1.6\text{W}/\text{mm}^2$), which is higher than the core itself. This could incur extra cooling. Another possibility is to increase the area of the voltage regulator for power density of $1\text{W}/\text{mm}^2$, to keep the power density of the VR below that of the core. Hence, for high-density voltage regulators, the power density is proportional to the inefficiency as well. Better efficiency in this case could translate to denser and smaller modules. From the observations made from currently available commercial voltage regulators (Figure 1), we estimate that *CoreUnfolding* would be able to improve by around 10% the VR efficiency (by doubling the output voltage and reducing by 40% peak current). The power density in the *CoreUnfolding* configuration would be ($8\text{W}/\text{mm}^2 \times (100-88)\% = 0.96\text{W}/\text{mm}^2$). Since this is below the $1\text{W}/\text{mm}^2$ that conventional cooling solutions can handle, there is no need to increase the area to reduce the density. Hence, at both fronts (current and power densities), *CoreUnfolding* results in a smaller voltage regulator.

6.1.1. Sizing of *sVR*. The size of *sVR* is dictated by the dynamic load-balancing engagement thresholds, that is, thresholds for using *sVR*, DA, and DVFS. The value for this threshold should be decided based on the common case, that is, expected load mismatch ratio. The average load imbalance was about 6% of the total power. Less than 1% of the time, there is a mismatch greater than 20% of the total power. Therefore, we conservatively pick the threshold of 20% for *sVR*. As a result, the *sVR* will be very small, up to 10% of the baseline voltage regulator. Hence, *sVR* can be fabricated in a distributed way across the chip using stacked power transistors [Rajapandian et al. 2005]. On-chip implementation of *sVR* is crucial for *CoreUnfolding* as such an implementation provides a quick response to mismatches, even easing the pressure on the amount of

decoupling capacitance required for very short transient mismatches. Overall, up to 30% of the area dedicated to on-chip VR can be saved.

The number of pins dedicated to power delivery is also proportional to the maximum current required by the design. The total number of pins in the package defines the minimum footprint size of the package, which could be an important factor nowadays with the emergence of mobile devices. While the maximum current per pin increases over the years, the number of power pins is proportional to the maximum current. Given that *CoreUnfolding* reduces the maximum current by 40%, the number of pins in a processor with *CoreUnfolding* can be reduced by the same amount. Given the number of power pins in current technologies, this reduction would account for about 20% of the total pin count of the package. According to Dong et al. [2010], a 20% reduction in pin count reduces the package cost by about 20%.

This is the first work that caps the demand of the extra voltage regulator used by voltage stacking. Without capping the demand, stacking can reduce the total number of pins, but it cannot reduce the VR area. Note that none of the related works cap the demand from the extra VR. For example, in Zhan and Sapatnekar [2008], the secondary VRs overall need to source the same maximum current as the primary VR. So in the best case, all the area saving in the primary VR will be lost in the secondary VRs. Gu and Kim [2005] do not require *sVR*, but the solution is deemed impractical due to timing overhead of merely relying on a “digital voltage regulator” or dummy activity. So we do not compare it against this work.

6.2. Power

Figure 11 shows the total (dynamic plus leakage) power consumption breakdown for “logic” (split between *Header* and *Footer* for *CoreUnfolding*), voltage regulator (VR), and dummy activity (DA). Both single and multicore configurations are reported. For each benchmark, three bars are shown:

Baseline (*b*) reports the power consumed for the nonstacked baseline core. We consider 80% VR efficiency [Linear Technology Inc. 2012; Infineon Technologies 2013; Linear Technology Inc. 2013].

CoreUnfolding with DA (*d*) reports the power consumed when only DA is used for mismatch management. We report the breakdown of power for *Header* and *Footer* partitions and assume a moderate 10% increase in the efficiency due to lower current draw and voltage drop [Linear Technology Inc. 2012; Infineon Technologies 2013; Linear Technology Inc. 2013]. This configuration does not rely on voltage regulator for load balancing between the stacks. DA effectively manages the load mismatch; however, all the power consumed by DA is wasted power, which accounts for 3.6% of the power consumption. The DA power consumption is relatively more for benchmarks that exhibit more mismatch (see Figure 14 for detailed mismatch distribution during the execution of each benchmark). In reality, merely relying on DA for load management could have cycle time implications.

CoreUnfolding with *sVR* and DA (*p*) reports the power consumed when both DA and *sVR* are used as mismatch management. The power wasted by *sVR* is inversely proportional to its efficiency, which we consider to be around 95%. This extra power is not distinguishable in the plots, compared to the *d* bars, because it adds up to 0.6% of the total power consumption. As mentioned in Section 4, to limit the size of the *sVR*, we use DA to cap the maximum mismatch, hence the maximum current demand.

In our experiments, **High Magnitude** mismatches did not occur, and thus DVFS was not necessary. As one can note in Figure 11, the power consumption of the load management (VR and DA) is a small portion of the total power. For the benchmarks evaluated, the utilization of DA is rare (about 0.5% of the total execution on average).

Fig. 11. Power consumption of different techniques: baseline (b), *CoreUnfolding* with DA (d), and *CoreUnfolding* with DA and VR (p). *CoreUnfolding* utilizing DA and VR for runtime load management saves about 7.5% of the total power.

CoreUnfolding configurations (d and p) also have a component of power consumed by the level converters. However, this component accounts for less than 0.1% of the total power, and hence is not shown. We also do not include the resistive loss in the delivery network, that is, pins and pads for off-chip solutions, and power grid on the chip, in the report of power savings. Considering about $10m\Omega$ resistance for pads and rails, an extra 2% to 3% of power savings would be observed.

The efficiency of voltage regulators varies over the program execution. However, to simplify the evaluation, and without loss of generality, for both base and *CoreUnfolding* configurations, we consider a flat efficiency rate to estimate the possible power savings. *CoreUnfolding* configurations (d and p) have roughly the same power consumption (p around 0.5% higher than d). When the *sVR* is used, it takes some power, but DA is less active, and thus takes less power. DA only provides a very bad management of mismatch, since DA is considerably slow.

6.3. Performance

While power consumption and area overhead of the level converters used for communication between the stack levels is negligible, the performance overhead might not be possible to absorb. The SPICE simulation shows that delay of the level converter is about the same as an FO4 gate. This could account for about 4.3% of the clock width in a modern processor [Choudhary et al. 2011].

Fig. 12. *CoreUnfolding* has less than a 2.2% decrease in performance ignoring improved power, noise, and area provided by *CoreUnfolding*.

The partitioning algorithm tries to minimize utilization of the level shifter for the critical components. However, it is hard to guarantee no additional penalty in all the critical paths. Instead of widening the clock width, we add one cycle to the access time of the functional units that require level shifters. In a real design, adding one cycle of delay in every boundary cross is pessimistic, since most of the crossing will not be in critical paths. Thus, for a real implementation, the latency addition should be more carefully evaluated. Figure 12 shows the results in terms of uIPC for single-threaded applications, already considering the extra latency. uIPC is the retiring rate of micro-operations (result of instruction decode). The simulations show a 2.21% decrease in the performance as a result.

Better voltage supply in terms of IR drop and Ldi/dt results in better performance. Our experiments, presented in more detail in Section 6.4, show 49% less di/dt and 45% less IR drop. Therefore, the peak performance loss can be compensated by the reduced voltage noise and better voltage margin as a result of lower current draw. Also, multithreaded or multiprogrammed applications' performance could even increase due to the power savings and a better thermal profile in case an on-chip voltage regulator is used (similar to Haswell). In such a case, the power savings can allow for a higher turbo multiplier when more than one core is active.

6.4. di/dt

Reduced current improves both resistive and inductive (di/dt) voltage drops. Both types of voltage drops reduce the voltage margin, which is already constrained in modern high-performance processors. Inductance in *CoreUnfolding* should be more or less the same compared to the baseline for on-chip VR solutions. For off-chip VR solutions, the number of pins is reduced in half, but the number of pads does not necessarily change. Hence, the inductance close to the load should not change much. Overall, smaller di/dt relaxes the quick response pressure on the voltage regulator and improves the noise. Figure 13 shows di/dt of the benchmark categories for the conventional and proposed architecture. di/dt was calculated using power traces from ESESC, which were used to calculate the current in each interval. For the multicore benchmarks, the metric is reported per core. We compare the gradients for every 100 processor cycles, which is the highest temporal resolution in our setup, and scale it to the desired unit of $A/\mu s$. b shows the results for the baseline. For *CoreUnfolding*, the metric is shown for both V_{dd} and V_{mid} rails, labeled respectively h and f in Figure 13.

Fluctuations in power consumption affect di/dt of the V_{dd} rail. In *CoreUnfolding*, the possible load mismatch between the *Header* and *Footer* groups would have a minimal impact on the di/dt of the V_{dd} rail as well. This is because the *sVR* could draw

Fig. 13. *CoreUnfolding* reduces the maximum as well as the average di/dt .

current from the V_{dd} rail for regulation. The V_{mid} rail, on the other hand, is only affected by the capped mismatch rather than the total current draw. Hence, it demonstrates less di/dt . The experiments are performed without dynamic scaling of voltage or frequency. As Figure 13 shows, by reducing the current demand, *CoreUnfolding* reduces the maximum di/dt by 49% on average across the benchmarks.

Decoupling capacitance of the supply nets is reduced in *CoreUnfolding*. Nonetheless, with the decrease in the current draw, the impact of decreased decoupling capacitance on voltage noise is compensated.

Multithreaded applications, that is, PARSEC and SPLASH benchmarks, exhibit a smaller average di/dt but higher maximum. The reason is that threads dynamically spawn or terminate on a core during the execution, which results in sporadic higher di/dt points. Nonetheless, these points are not frequent enough to shift the average.

6.5. Transient Analysis of V_{mid} Voltage and Quality of Partitioning

This section evaluates the quality of the partitioning algorithm in terms of reducing the load mismatch between *Header* and *Footer* groups. This is before the dynamic, runtime load-balancing techniques (DA and VR) get engaged to remove all the mismatches. The partitioning is critical in *CoreUnfolding* to minimize the overhead of runtime load balancing.

Figure 14 shows the results. The base method in this figure partitions the functional units to *Header* and *Footer* groups based on the aggregated data from running multiple benchmarks. The base approach limits the worst-case scenario at the cost of losing the best scenario. The best scenario allows for adjustment per benchmark, that is, to partition the functional units per benchmark. Figure 14(a) presents the results for this under *peak* label. The *peak* is only presented to show how well the actual *base* partitioning performs. We do not use multicore applications to train the partitioning; hence, *peak* does not apply to those applications. Instantaneous power mismatch stays within 10% for 75% of the time for all applications. Also, it is below 20% for 99% of the time for the combination of applications (not shown in the figure); the maximum observed mismatch was around 25%.

To better illustrate how this impacts the voltage in each level of the stack, we run transient SPICE simulations with the power traces for the partitioned core using the model from Section 5.4. Figure 17 shows the trace of the voltage for the *Header* and *Footer* groups for two applications. More than 99% of the time, decaps and *sVR* are able to keep V_{mid} stable within 10% of the nominal voltage. A full transient SPICE simulation of architectural benchmarks is unfeasible; the proposed approach is compatible with industry standards, with the difference of using a default power grid instead of extracting parasitics from GDS.

Another important issue is IR drop across the chip area. Two factors play a role in this; the first is the decreased current flowing through V_{dd} , which tends to reduce the

Fig. 14. Distribution (Q1/Median/Q3) of instantaneous load mismatch for the benchmarks before runtime load balancing is engaged. The *base* method, using one partition for all the applications, performs comparably to *peak*, which adjusts the partitioning per application.

Fig. 15. *CoreUnfolding* reduces IR drop across the chip area.

drop. The second is, given a fixed budget for power rails and the introduction of V_{mid} , the total amount of metal used by V_{dd} will decrease. Overall, we see a decreased IR drop, as shown in Figure 15 (the voltage was scaled for the baseline for better comparison), since the reduction in current is larger than the reduction on V_{dd} metal. The floorplan is partially shown over the figure, just as a reference.

Fig. 16. In multicore configurations, swapping the *Header* and *Footer* clusters helps to bring V_{mid} to balance across the chip area; when there is no swap, the small partitioning bias accumulates.

In the case of multicore, we also evaluate the effects of swapping the cores, as proposed in Section 4.5, as opposed to not swapping them. As discussed, there is a small residual bias toward one partition due to the impossibility of exact balance. To evaluate it, we measure V_{mid} across a dual-core chip, with one core running *mgrid* and one running the *astar* benchmark. The *sVR* is disabled to allow for the analysis of the impact of the swapping only. Our experiments show that swapping brings V_{mid} closer to the midpoint voltage across the chip, even for cores running different benchmarks, from an average of 0.7V to an average of around 0.95V. The results are shown in the heatmap in Figure 16.

6.6. Worst Case and Voltage Noise Analysis

One concern with *CoreUnfolding* is the behavior of the power delivery network during power-up and power-down, in particular the magnitude of inrush current during this phase. To evaluate this, we use the same setup from the previous section but consider start and stop of cores. This also simulates the case where the core is power gated and then turned back on. Figure 17 shows the transient analysis results for current and voltage for a core running *mgrid* benchmark, which is then powered down and back up (see marks on the figure) for both a stacked and a not-stacked configuration.

Fig. 17. *CoreUnfolding* allows for smaller inrush current: (1) mgrid execution, (2) power gated, and (3) execution continued.

The spike after restart is more than two times bigger in the conventional not-stacked configuration than with *CoreUnfolding*, due to the fact that the total current itself is smaller in *CoreUnfolding*.

There is a small difference between the V_{foot} and V_{head} , which is due to a small residual difference in the partitioning. This observation is compatible with the previous finding in the context of voltage stacking [Lee et al. 2012]. $V_{no-stacking}$ is lower due to a higher IR drop, caused by the higher current, even considering fewer resources to V_{dd} and Gnd grids in the stacked configuration (due to the use of resources for V_{mid}).

Also note that counterintuitive voltage noise (and droop) is smaller in the case of *CoreUnfolding*. This is, again, due to the smaller current in V_{dd} , which makes it more stable, but also due to the balancing achieved by the partitioning in *CoreUnfolding* and the presence of mechanisms to cap the maximum mismatch (DA and DVFS).

One interesting point is that the reduction in voltage noise could allow the reduction of voltage margins. Considering an initial voltage margin of 10% and 40% reduction in noise, and thus in voltage margin, an additional 8% of total power could be possible. This would add to around 15% to 18% of total power savings. In our analysis, however, we keep the conservative analysis, not considering this extra gain.

7. CONCLUSION

This article presents *CoreUnfolding*, a microarchitecture-level technique to limit the load mismatch wanted by voltage stacking power delivery systems. To have a balanced load among stacks, (1) *CoreUnfolding* stacks the components inside a core instead of stacking the cores. Components within a core are grouped and stacked based on their instantaneous power correlation and average power magnitude. (2) The minimal mismatches that could still happen are managed by a combination of DA, small secondary VR, and DVFS, depending on the magnitude of the mismatch. Even without DA, the proposed solution reduces the mismatch to about 6% of the total execution on average.

A key contribution of this article is the capacity to reduce overall current in the power supply, thus enabling smaller (estimated around 30%) and more efficient (estimated around 10%) VRs. This is possible by capping its usage through novel static mismatch reduction. Dynamic load-balancing techniques used at runtime bring the minimal load mismatch to zero. As a result, total power can be reduced by up to 10%

in our experiments due to the improved VR efficiency and reduced I^2R power loss in the delivery network. Also, both resistive (IR) and inductive (Ldi/dt) voltage noise are reduced. Reduction of the current also results in using fewer power-dedicated pins.

The downside of *CoreUnfolding* is a 2.2% decrease in the peak performance of the processor for single-core applications, which could be compensated for by the reduced voltage noise. That is considering the on-chip VR power consumption does not cause thermal throttling. If such effects are taken into account, speedups should be expected in several benchmarks. Overall, the presented *CoreUnfolding* system allows for power savings and area savings and makes the usage of voltage stacking more feasible for future systems.

Future work may include the generalization to an n-level stack, with further savings, but that presents more challenges in the partitioning. Another possibility for modifying the scheme is to have an uneven partitioning, thus changing the $2V_{dd}/V_{mid}$ ratio to an arbitrary value. This opens the opportunity to save power by having different parts of the core at different voltages, but could also impact performance.

ACKNOWLEDGMENTS

We thank the reviewers for their valuable feedback on the article. We would also like to thank Zhou Xiazhou for his review and feedback regarding voltage regulator design.

REFERENCES

- C. Bienia, S. Kumar, J. P. Singh, and K. Li. 2008. The PARSEC benchmark suite: Characterization and architectural implications. In *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques (PACT'08)*. ACM, New York, NY, 72–81. DOI: <http://dx.doi.org/10.1145/1454115.1454128>
- E. A. Burton, G. Schrom, F. Paillet, J. Douglas, W. J. Lambert, K. Radhakrishnan, and M. J. Hill. 2014. FIVR – Fully integrated voltage regulators on 4th generation Intel®Core™SoCs. In *Proceedings of the 2014 27th Annual IEEE Applied Power Electronics Conference and Exposition (APEC'14)*. 432–439. DOI: <http://dx.doi.org/10.1109/APEC.2014.6803344>
- E. Candan, P. S. Shenoy, and R. C. N. Pilawa-Podgurski. 2014. A series-stacked power delivery architecture with isolated differential power conversion for data centers. In *Proceedings of the 2014 IEEE 36th International Telecommunications Energy Conference (INTELEC'14)*. 1–8. DOI: <http://dx.doi.org/10.1109/INTELEC.2014.6972231>
- L. Chang, R. K. Montoye, B. L. Ji, A. J. Weger, K. G. Stawiasz, and R. H. Dennard. 2010. A fully-integrated switched-capacitor 2.1 voltage converter with regulation capability and 90% efficiency at 2.3 A/mm². In *Proceedings of the 2010 IEEE Symposium on VLSI Circuits (VLSIC'10)*. IEEE, 55–56.
- N. K. Choudhary, S. V. Wadhavkar, T. A. Shah, H. Mayukh, J. Gandhi, B. H. Dwiel, S. Navada, H. H. Najafabadi, and E. Rotenberg. 2011. FabScalar: Composing synthesizable RTL designs of arbitrary cores within a canonical superscalar template. In *Proceedings of the 38th Annual International Symposium on Computer Architecture (ISCA'11)*. ACM, New York, NY, 11–22.
- X. Dong, J. Zhao, and Y. Xie. 2010. Fabrication cost analysis and cost-aware design space exploration for 3-D ICs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 29, 12 (2010), 1959–1972.
- J. Gu and C. H. Kim. 2005. 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, 192–197.
- J. L. Henning. 2006. SPEC CPU2006 benchmark descriptions. *SIGARCH Computer Architecture News* 34, 4 (Sept. 2006), 1–17.
- W. Huang, G. Schuellein, and D. Clavette. 2003. A scalable multiphase buck converter with average current share bus. In *18th Annual IEEE Applied Power Electronics Conference and Exposition, 2003 (APEC'03)*, Vol. 1. IEEE, 438–443.
- Linear Technology Inc. 2012. Dual 4A Per Channel DC/DC uModule Regulator. Retrieved from <http://www.linear.com/product/LTM4614>.
- Linear Technology Inc. 2014. Dual 8A Per Channel DC/DC uModule Regulator. Retrieved from <http://www.linear.com/product/LTM4616>.

- Linear Technology Inc. 2013. Ultralow Vin, 15A DC/DC uModule Regulator. <http://www.linear.com/product/LTM4611>.
- Intel Inc. 2014. Desktop 4th Generation Intel®Core™Processor Family, Desktop Intel Pentium®Processor Family, and Desktop Intel®Celeron®Processor Family - Datasheet. <http://www.intel.com/content/www/us/en/processors/core/4th-gen-core-family-desktop-vol-1-datasheet.html>.
- F. Ishihara, F. Sheikh, and B. Nikolic. 2004. Level conversion for dual-supply systems. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 12, 2 (2004), 185–195.
- J. L. Henning. 2000. SPEC CPU2000: Measuring performance in the new millennium. *IEEE Computer* 33, 7 (July 2000), 28–35.
- E. K. Ardestani and J. Renau. 2013. ESESC: A fast multicore simulator using time-based sampling. In *Proceedings of the International Symposium on High Performance Computer Architecture (HPCA'13)*.
- K. Kesarwani, C. Schaef, C. R. Sullivan, and J. T. Stauth. 2013. A multi-level ladder converter supporting vertically-stacked digital voltage domains. In *Proceedings of the 2013 28th Annual IEEE Applied Power Electronics Conference and Exposition (APEC'13)*. 429–434. DOI:<http://dx.doi.org/10.1109/APEC.2013.6520245>
- W. Kim, D. M. Brooks, and G. Y. Wei. 2011. A fully-integrated 3-level DC/DC converter for nanosecond-scale DVS with fast shunt regulation. In *Proceedings of the 2011 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC'11)*. IEEE, 268–270.
- W. Kim, D. Brooks, and G.-Y. Wei. 2012. A fully-integrated 3-level DC-DC converter for nanosecond-scale DVFS. *IEEE Journal of Solid-State Circuits* 47, 1 (2012), 206–219.
- W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks. 2008. System level analysis of fast, per-core DVFS using on-chip switching regulators. In *Proceedings of the IEEE 14th International Symposium on High Performance Computer Architecture, 2008 (HPCA'08)*. IEEE, 123–134.
- S. K. Lee, D. Brooks, and G. Y. Wei. 2012. Evaluation of voltage stacking for near-threshold multicore computing. In *Proceedings of the 2012 IEEE International Symposium on Low Power Electronics and Design (ISLPED'12)*. ACM, 373–378.
- S. K. Lee, T. Tong, X. Zhang, D. Brooks, and G.-Y. Wei. 2015. A 16-Core voltage-stacked system with an integrated switched-capacitor DC-DC converter. In *Proceedings of the IEEE Symposium on VLSI Circuits (VLSIC'15)*.
- J. Leng, Y. Zu, M. Rhu, M. Gupta, and V. J. Reddi. 2014. GPUVolt: Modeling and characterizing voltage noise in GPU architectures. In *Proceedings of the 2014 International Symposium on Low Power Electronics and Design (ISLPED'14)*. ACM, New York, NY, USA, 141–146. DOI:<http://dx.doi.org/10.1145/2627369.2627605>
- S. R. Nassif. 2008. Power grid analysis benchmarks. In *Proceedings of the Asia and South Pacific Design Automation Conference, 2008 (ASPDAC'08)*. 376–381. DOI:<http://dx.doi.org/10.1109/ASPDAC.2008.4483978>
- G. V. Piqué and H. J. Bergveld. 2012. State-of-the-art of integrated switching power converters. In *Analog Circuit Design*. Springer, 259–281.
- M. Popovich, A. V. Mezhiba, and E. G. Friedman. 2007. *Power Distribution Networks with On-Chip Decoupling Capacitors*. Springer.
- S. Rajapandian, K. L. Shepard, P. Hazucha, and T. Karnik. 2006. High-voltage power delivery through charge recycling. *IEEE Journal of Solid-State Circuits* 41, 6 (2006), 1400–1410.
- S. Rajapandian, Z. Xu, and K. L. Shepard. 2005. Implicit DC-DC downconversion through charge-recycling. *IEEE Journal of Solid-State Circuits* 40, 4 (2005), 846–852.
- Y. K. Ramadass, A. A. Fayed, and A. P. Chandrakasan. 2010. A fully-integrated switched-capacitor step-down DC-DC converter with digital capacitance modulation in 45 nm CMOS. *IEEE Journal of Solid-State Circuits* 45, 12 (2010), 2557–2565.
- C. Schaef and J. T. Stauth. 2016. Efficient voltage regulation for microprocessor cores stacked in vertical voltage domains. *IEEE Transactions on Power Electronics* 31, 2, 1795–1808. DOI:<http://dx.doi.org/10.1109/TPEL.2015.2426572>
- P. S. Shenoy, V. T. Buyukdegirmenci, A. M. Bazzi, and P. T. Krein. 2010. System level trade-offs of microprocessor supply voltage reduction. In *Proceedings of the 2010 International Conference on Energy Aware Computing (ICEAC'10)*. IEEE, 1–4.
- P. S. Shenoy, S. Zhang, R. A. Abdallah, P. T. Krein, and N. R. Shanbhag. 2011b. Overcoming the power wall: Connecting voltage domains in series. In *Proceedings of the 2011 International Conference on Energy Aware Computing (ICEAC'11)*. 1–6. DOI:<http://dx.doi.org/10.1109/ICEAC.2011.6403629>
- P. S. Shenoy, I. Fedorov, T. Neyens, and P. T. Krein. 2011a. Power delivery for series connected voltage domains in digital circuits. In *Proceedings of the 2011 International Conference on Energy Aware Computing (ICEAC'11)*. IEEE, 1–6.

- Infineon Technologies. 2013. High Performance DrMos TDA21220. <http://www.infineon.com/cms/en/product/power/dc-dc-converter/dc-dc-integrated-power-stage/drmos-integrated-power-stage/TDA21220/productType.html?productType=db3a3044243b532e0124de3165386adc>.
- Z. Toprak-Deniz, M. Sperling, J. Bulzacchelli, G. Still, R. Kruse, S. Kim, D. Boerstler, T. Gloekler, R. Robertazzi, K. Stawiasz, T. Diemoz, G. English, D. Hui, P. Muench, and J. Friedrich. 2014. 5.2 distributed system of digitally controlled microregulators enabling per-core DVFS for the POWER8™ microprocessor. In *Proceedings of the 2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC'14)*. 98–99. DOI : <http://dx.doi.org/10.1109/ISSCC.2014.6757354>
- V. Sprague. 2013. *Solid Tantalum Surface Mount Chip Capacitors*. Vishay Sprague, Santa Clara, CA.
- J. Wei. 2004. *High Frequency High-Efficiency Voltage Regulators for Future Microprocessors*. Ph.D. Dissertation. Virginia Polytechnic Institute and State University.
- S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta. 1995. The SPLASH-2 programs: Characterization and methodological considerations. In *Proceedings of the 22nd Annual International Symposium on Computer Architecture (ISCA'95)*. ACM, New York, NY, 24–36. DOI : <http://dx.doi.org/10.1145/223982.223990>
- Y. Zhan and S. S. Sapatnekar. 2008. Automated module assignment in stacked-Vdd designs for high-efficiency power delivery. *ACM Journal on Emerging Technologies in Computing Systems (JETC)* 4, 4 (2008), 18.
- W. Zhao and Y. Cao. 2006. New generation of predictive technology model for sub-45 nm early design exploration. *IEEE Transactions on Electron Devices* 53, 11 (2006), 2816–2823.
- X. Zhou, P.-L. Wong, P. Xu, F. C. Lee, and A. Q. Huang. 2000. Investigation of candidate VRM topologies for future microprocessors. *IEEE Transactions on Power Electronics* 15, 6 (2000), 1172–1182.

Received April 2015; revised September 2015; accepted October 2015