

Distributed TSV Topology for 3-D Power-Supply Networks

Michael B. Healy, *Member, IEEE*, and Sung Kyu Lim, *Senior Member, IEEE*

Abstract—3-D integration has the potential to increase performance and decrease energy consumption. However, there are many unsolved issues in the design of these systems. In this work we study the design of 3-D power supply networks and demonstrate a technique specific to 3-D systems that improves IR-drop and dynamic noise over a straightforward extension of traditional design techniques. Previous work in 3-D power delivery network design has simply extended 2-D techniques by treating through-silicon vias (TSVs) as extensions of the C4 bumps. By exploiting the smaller size and much higher interconnect density possible with TSVs we demonstrate significant reduction of nearly 50% in the IR-drop and 42% in the dynamic noise of our large-scale 3-D design. Simulations also show that a 3-tier stack with the distributed TSV topology actually lowers IR-drop by 21% and dynamic noise by 32% over a non-3-D system with less power dissipation. We analyze the power distribution network of an envisioned 1000-core processor with 30 stacked dies and show scaling trends related to both increased stacking and power distribution TSVs. Finally, we examine several techniques for minimizing IR-drop and dynamic noise and their effects on our large-scale 3-D system.

Index Terms—3-D, inductive noise, power supply network, through-silicon via (TSV).

I. INTRODUCTION

3-D stacking of ICs has generated increasing interest from the VLSI community in recent years. The many potential benefits of 3-D integration include reduced power consumption from off-chip communication, reduced wirelength and delay, and lower-cost process integration. However, there are many challenges involved in the design of 3-D ICs that have not been met. Increased volumetric power density combined with increased thermal resistance between the lower layers and the heatsink imply increased operating temperatures and an associated reduction in reliability. Smaller footprints combined with larger package-level system power imply increased power delivery problems. Solutions to all of these problems are the subject of ongoing work in both academia and industry. In this work we provide a layout-level examination of the design of

3-D power delivery networks, and demonstrate that the unique environment of 3-D ICs can have a dramatic effect on IR-drop and dynamic noise in these networks.

IR-drop (sometimes referred to as ground-bounce) is the resistive voltage drop in power and ground distribution networks caused by the dynamic and leakage power of ICs. IR-drop causes many problems in modern microprocessor and application-specific integrated circuit (ASIC) designs and was one of the causes of the end of the frequency scaling era. As device scaling continues, lower and lower supply voltages are increasing total current and reducing power supply noise margins even further. These issues are causing a larger and larger percentage of available routing resources to be dedicated to power supply distribution in high-performance designs, which can add significantly to congestion problems and reduce the amount of functionality that can be packed into a unit area.

Dynamic supply noise (sometimes referred to as di/dt noise or simultaneous switching noise) is transient voltage instability in power and ground distribution networks caused by the interaction of the capacitance and inductance of those distribution networks with time-varying switching activity in ICs. Dynamic noise causes problems with timing closure and device reliability, because lower supply voltages cause transistors to switch more slowly. Decoupling capacitance (decap) is typically added to the power distribution network to mitigate the effects of dynamic noise, however, large amounts of decap can cause significant increases in leakage power. Modern designs require large amounts of decap to meet supply noise constraints. Techniques that reduce decap requirements are valuable additions to an IC designer's toolkit.

Many researchers have proposed optimization schemes for traditional IC power network design. Previous work on 3-D power delivery networks has largely assumed a straightforward extension of 2-D power delivery network design. Huang *et al.* [1] presented a physical model of 3-D power distribution networks. In their model power/ground through-silicon vias (TSVs) and power supply C4 bumps are always aligned with one another. Jain *et al.* [2] extended the work of Gu *et al.* [3] by examining the use of multi-story power delivery in 3-D ICs. In their approach there are two power domains and the ground network of one domain is the power network of the other domain. Again, the TSVs and supply bumps are always assumed to be fully aligned and they are divided among the three power distribution networks evenly. Yu *et al.* [4] demonstrate an optimization scheme for supply bump assignment and via insertion simultaneously considering both supply noise and temperature. They again assume that supply bumps are aligned with TSVs in every case. Healy and Lim [5] presented the only

Manuscript received March 01, 2011; revised August 07, 2011; accepted August 12, 2011. Date of publication October 13, 2011; date of current version July 27, 2012. This work was supported in part by the National Science Foundation under Grant CCF-0546382, by the SRC Interconnect Focus Center (IFC), and by Intel Corporation.

M. B. Healy is with TJ Watson Research Center, IBM Research, Yorktown Heights, NY 10598 USA (e-mail: mbhealy83@gmail.com).

S. K. Lim is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: limsk@ece.gatech.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TVLSI.2011.2167359

analysis of TSVs not aligned with supply bumps, however, they offer only a high-level analysis and do not validate their results using other methods.

Additionally, Sapatnekar *et al.* [6] perform decap placement optimization in 2-D and 3-D power-supply networks guided by linearized noise models. They examine stacks with up to four tiers, but do not provide an in-depth examination of the differences between 2-D and 3-D power-supply networks. Thorolfsson *et al.* [7] demonstrate the design techniques used to produce a 3-D integrated synthetic aperture radar FFT processor. They discuss the methodology used to deliver power to the upper tiers. However, the general applicability of their methodology is limited by their processing technology, which only includes three metal layers. Sun *et al.* [8] present a 3-D integrated voltage converter with a cellular architecture for reducing power delivery losses and reducing on-chip noise. Finally, Pavlidis *et al.* [9] examine the use of TSVs to bypass small local-via stacks for power-distribution in 3-D ICs and show how it can improve power-supply noise and reduce routing congestion.

The focus of this work is on improving 3-D power-distribution-network performance by placing power-distribution TSVs in a more effective manner. Previous efforts at 3-D power-grid optimization are largely orthogonal to this and may generally be combined with our proposed technique. Our goal is to explore this aspect of power delivery in 3-D ICs and how it differs from traditional designs. In support of this effort, we examine both large and small scale designs with up to 30 stacked tiers using a 130-nm process technology. The 130-nm standard cell library was chosen because it was the most advanced technology available with fully-characterized timing behavior. Designs using 30 stacked tiers are likely to be at least 10 to 15 years into the future, and are unlikely to use technology near the 130-nm node for more than a few specialized functions or tiers. However, thermal limitations will restrict power density significantly. Also, the ratio between the power density of the design, and the distribution wiring determines the IR-drop in 2-D systems, and must remain relatively constant with technology scaling to maintain the IR-drop within supply margins. It is also possible to reduce C4 pitch to achieve the same goal, however the expense of this has proven to be uneconomical. For dynamic noise, the relative amounts of on-chip decap and the package inductance are the major deciding factors, which will change significantly in 3-D systems.

Compared to prior efforts we demonstrate the benefits of re-examining the unique capabilities of TSVs relative to package-level bumps. We also perform our analysis using layout-level designs and validate our modeling results using commercial-grade sign-off IR-drop analysis software. The major contributions of this work are as follows.

- We present the first layout-level analysis of 3-D power distribution networks that is validated using commercial tools.
- We demonstrate the potential IR-drop and dynamic noise benefits of spreading power and ground distribution TSVs away from the power and ground supply bumps in designs with non-uniform power dissipation.
- We examine scaling trends in 3-D power distribution networks using this framework to demonstrate future poten-




Fig. 1. Bumps, TSVs, and wires in a 3-D P/G network.

tial for increased 3-D stacking on an envisioned 1000-core system.

- We analyze several modifications of power distribution network design unique to 3-D systems and show their effects on IR-drop and dynamic noise.

The rest of this paper is organized as follows. Section II presents an overview of 3-D and flip-chip power supply network design concepts, and introduces the novel TSV topology thoroughly examined in this work. Next, Section III presents a comparison of the area overhead of the proposed TSV topology and demonstrates power supply noise improvement using a simple example. Then, Section IV details the prototype system used for our layout-based supply noise studies. Sections V and VI discuss our analysis and validation methodologies for computing IR-drop and dynamic noise, respectively. Next, Section VII explains the results of our extensive simulations. Section VIII examines several techniques for reducing power-supply noise in our model. Finally, Section IX summarizes our conclusions and results.

II. 3-D AND FLIP-CHIP POWER NETWORKS

High performance 3-D systems will generally use flip-chip style packaging to increase off-chip interconnect density and reduce parasitics. Flip-chip power distribution systems are commonly laid out as grids. High-level metal layers are reserved for laying out a coarse-grained grid with large wires that connects a regular array of power and ground C4 bumps. A fine-grained mesh provides local distribution and connects to lower-level-metal power rings or standard-cell row distribution wiring. Most commercial products today have C4 bump pitches around 100 to 200 μm , however, researchers have demonstrated micro-bumps with pitches down to 20 μm [10], [11].

For 3-D systems the TSVs will fill the role of the C4 bumps for intermediate tiers. Each tier will contain its own power distribution network. Fig. 1 shows the general topology of a 3-D power distribution. The vertical resistance between adjacent tiers should be close to that of the C4 bumps to maintain reliable power and ground voltages in large-scale 3-D systems. The resistance of individual C4 bumps is on the order of 5 m Ω . Additionally, TSVs should be smaller than the C4 bumps or large amounts (25% or more) of die area will become unusable.

TSVs can be manufactured in many different sizes. Diameters near 1 μm have been shown in the literature. Power and




Fig. 2. Two TSV topologies for power distribution in a single tile of the distribution network. C4 bumps are shown in blue and P/G TSVs in red. The combined resistance of all TSVs in each topology is equal.

ground TSVs should be large to have low resistance, but signal TSVs should be small to increase interconnect density and reduce parasitic capacitance. Manufacturing multiple TSV sizes on a single die would increase cost and reduce yield. Therefore, it will likely be necessary to use a single TSV size for both power distribution and signal wiring.

In this work it is assumed that only one TSV size is available, and is optimized for signals. There are several potential combinations of TSV distribution that could be used to deliver power. Fig. 2 shows two of the basic choices we investigate thoroughly in this paper.

- **Clustered Topology:** Multiple small TSVs are clustered over the C4 pads for both power and ground distribution.
- **Distributed Topology:** Multiple small TSVs are distributed evenly throughout the die for both power and ground distribution.

For both of the topologies the combined resistance of all the TSVs is assumed to be the same. Fig. 2 depicts TSV topologies for a single tile in the power/ground network. This tile is mirrored and replicated all over the chip.

III. ANALYTICAL COMPARISON OF TSV TOPOLOGIES

The difference between the coefficients of thermal expansion (CTE) of silicon and TSV conductors causes thermal stress in the silicon die. This thermally-induced stress can affect device performance. TSV manufacturing processes may also negatively impact nearby device performance and manufacture. For these reasons, gates and transistors are generally placed outside of a keep-out region (KOR) around the TSVs. Fig. 3 shows a group of TSVs with the KOR highlighted and also defines several dimensions associated with the KOR. The figure defines K as the distance of the edge of the KOR from the TSV, T as the dimension of the TSV, and S as the TSV-to-TSV space. The area taken up by an $n \times m$ array of TSVs in the clustered topology is then

$$A_{\text{clustered}} = (2K + (n - 1) \cdot S + n \cdot T) \times (2K + (m - 1) \cdot S + m \cdot T). \quad (1)$$

Additionally, the area taken up by an $n \times m$ array of TSVs in the distributed topology is then

$$A_{\text{distributed}} = n \cdot m \cdot (2K + T)^2. \quad (2)$$




Fig. 3. Illustration of the KOR around a group of TSVs. The distance of the edge of the KOR from the TSV is defined as K , the dimension of the TSV is T , and S is the TSV-to-TSV space.

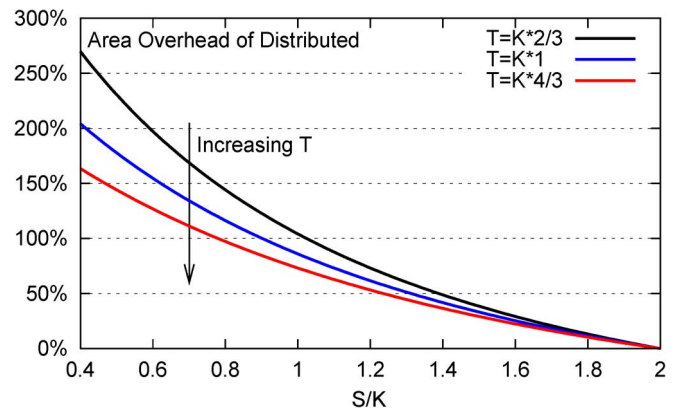


Fig. 4. Area overhead for a 5×5 array of TSVs in the distributed topology compared to the clustered topology. The ratio between S and K is varied on the independent axis. Data for several different values of T are shown.

It is obvious that the distributed topology will occupy more silicon area than the clustered topology when S is small. Fig. 4 shows the area overhead for a 5×5 array of TSVs in the distributed topology compared to the clustered topology. The ratio between S and K is varied on the independent axis. Data for several different values of T are shown. Both topologies occupy the same area, and the overhead is zero when $S \geq 2K$. The area overhead is 86% when $S = K = T$ for $n = m = 5$.

The values for S , K , and T are determined by the TSV manufacturing technology and the design rules. When $S \geq 2 \cdot K$, the distributed and clustered topologies occupy the same area. In general, S , K , and T have values that are near to one another, typically within a factor of two. It should also be noted that, for the prototype layout presented in the next section, the total area occupied by $6 \times 6 \mu\text{m}$ TSVs is less than 5% of the chip area. Additionally, the power TSVs are located under the power and ground stripes, which already constrain placement and routing. Finally, the KOR parameters mainly impact the area of the design and not the power-supply performance. The power-supply model used in this work aggregates the resistance of clustered TSVs, and so does not show any impact from the KOR parameters.

A simple 1-D example demonstrating the IR-drop of the clustered and distributed TSV topologies is shown in Fig. 5. The ground network of a two-tier system is modeled using current

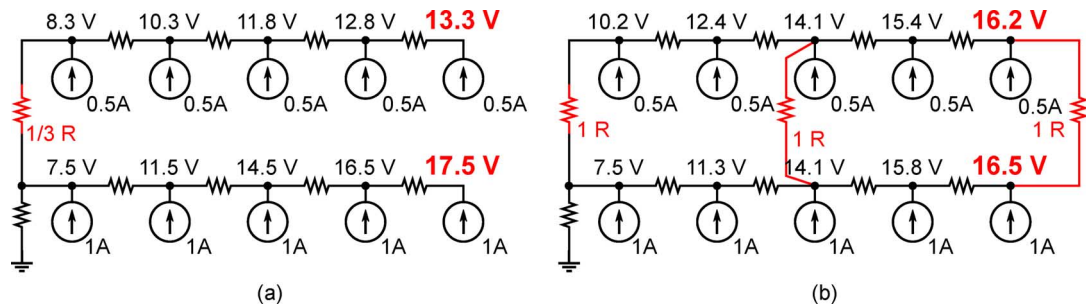


Fig. 5. Simple 1-D example that demonstrates the power-supply-noise improvement encountered when using the distributed TSV topology in systems with non-uniform per-tier power dissipation. The clustered topology shown in (a) results in a maximum IR-drop of 17.5 volts. The distributed topology with the same number of TSVs shown in (b) results in lower maximum IR-drop of 16.5 V. All resistance values are $R = 1 \Omega$, except where noted in red.

sources and resistances. The lower tier has twice the power dissipation of the upper tier. The upper tier is connected to the lower tier with three TSVs either clustered together on the left side (a) or distributed throughout the design (b). The figure shows the voltage at every node of the network for both cases. The maximum voltage is 17.5 V for the clustered case and 16.5 V for the distributed case, an improvement of 5.7%. This example demonstrates the basic reason why the distributed TSV topology generally results in better power distribution network performance. In the clustered case the difference between the maximum per-tier voltages is $17.5 - 13.3 = 4.2$ V. This represents the “slack” in the upper tier. For the distributed topology the slack in the upper tier is much lower, around 0.3 V, and the maximum drop on the lower tier is also lower than for the circuit with the clustered topology.

A final potential difference between the distributed and clustered TSV topologies is their affect on temperature. Several works, such as that by Goplen and Sapatnekar [12], demonstrate the usefulness of TSVs inserted to provide a lower resistance path for heat to escape 3DICs. The distributed and clustered topologies would result in different thermal distributions due to the high thermal conductivity of TSV materials. However, detailed examination of this difference is outside the scope of this work.

IV. PROTOTYPE LAYOUT

The prototype layout used in our simulations is based on a design targeted at demonstrating extreme memory bandwidth using 3-D interconnects. Our design is a many-core processor composed of an array of simple cores connected with a nearest-neighbor communication mesh. Each core has eight banks of dedicated SRAM directly stacked above it in two separate tiers. Each core tier contains a 10×10 array of cores. One grouping of one core tier and two SRAM tiers is defined to be one “set” of our scalable prototype layout. We envision stacking 10 sets together to form a 1000-core processor. The full 1000-core processor is shown in Fig. 6. The core architecture used in this work is very similar to that presented by Healy *et al.* [13].

The layouts used in our experiments were designed using a 130-nm standard cell library from Global Foundries. For the physical design we used Cadence’s SOC Encounter automated place and route tool. The layouts for a single core and a single memory tile are shown in Fig. 7. We also highlight the areas in




Fig. 6. 1000-core processor that is targeted in our simulations. Our sign-off noise simulation covers the 30-tier single core stack.

the layout reserved for ground TSV connections. The distribution of connection points is irregular due to the constraints of the layout, especially the locations of the hard memory macros. For the distributed TSV topology, TSVs are located at all of the potential locations. In the clustered TSV topology, all of the TSVs are grouped into the center position, over the C4 bump. The power TSV locations are similarly distributed in an offset fashion, the main difference is that the power C4 bumps are at the corners of the core, while the ground C4 bump is in the center. Each location is capable of accepting a $6 \mu\text{m}$ diameter via-first TSV, while the locations over the C4 bumps (the center and near the corners) are capable of accepting 25 or more of these TSVs.

The single-core and single-tile layouts are both $560 \mu\text{m}$ square. The core-to-core and tile-to-tile pitch is $590 \mu\text{m}$ to accommodate the inter-core logic and communication, as well as the power distribution from the C4 bumps. The full 100-core and 100-tile layers are approximately 6 mm square. Each core tile has 21.9 pF of decoupling capacitance ($=219$ pF per tier), and each memory tile has 21.7 pF of decoupling capacitance ($=217$ pF per tier). The maximum possible amount of decoupling capacitance was placed in the white spaces left after standard cell placement and timing optimization in the core tiers, and in all of the whitespace around the memory macros




Fig. 7. Layout of a single core and single memory tile from our 1000-core processor. The possible ground distribution TSV locations are highlighted in red. The ground C4 bump in the center of the core is indicated. The power C4 bumps are near the corners of the core. The total TSV area overhead is less than 5%.




Fig. 8. Power map for one core of our processor. The maximum total power consumption per core is 65.5 mW.

in the memory tiers. The resulting distribution of decoupling capacitance is relatively uniform in the core tiers, and mostly around the edges and in the middle of the memory tiers.

The maximum total power dissipation per set (1 core tier + 2 memory tiers) is approximately 13.2 W, the 1000-core system then has a total power dissipation of 132 W. Each core dissipates 65.5 mW, which is the result of statistical power simulations from Cadence Encounter. Fig. 8 shows the power map for a single core. The power dissipation of this design is not extreme, however the high volumetric power density could be a problem for traditional heatsinks. For this case, micro-fluidic channels [14], [15] have been shown to be an effective method for cooling large-scale 3-D chip stacks.

The distribution wiring in our design is concentrated in 10 μm wide stripes that run through and around each core on the upper metal layers. Secondary wiring on Metal 1 connects to the standard cell rows. The most effective location for placing distributed power delivery TSVs is on the large distribution wiring far from the C4s, especially in areas where the IR-drop is much

lower in adjacent tiers. The distribution grid parameters and layout, as well as the 3-D powermap, have a significant impact on the effectiveness of the distributed TSV topology. This work provides both an analytical demonstration of the effectiveness of the distributed TSV topology (see Section III), and also a detailed example implementation using the design discussed in this section.

V. 3-D IR-DROP ANALYSIS

A. Methodology

Layout-level IR-drop values are computed by performing power consumption simulations, either statistical or simulation-driven, to obtain gate- and module-level power consumption values. The consumption values are then divided by the nominal supply voltage, in our case 1.5 V, to obtain gate- and module-level current consumption values. Next, parasitic extraction is performed on the layout to obtain a SPICE netlist that models the power distribution network. Our experiments were performed using Cadence's QRC transistor-level extraction tool. The current consumption values are then connected to the nodes representing the corresponding transistors belonging to the appropriate gates and modules. For traditional 2-D ICs the netlist is then simulated using a power network simulator, in our case Cadence's UltraSim. Fig. 9 shows our analysis flow for 2-D netlists. In 3-D designs, the previously described steps are performed once for each type of tier (core, memory, etc.). The tier-type SPICE models are replicated for each instance of that tier-type and then connected using a resistive TSV model.

Simulation of power distribution networks is a generally difficult problem for traditional ICs. These networks can contain tens of millions of nodes. 3-D stacking exacerbates the problem even further. Given the extreme regularity of the prototype design that is examined in this work, we mitigate some of the extreme memory and execution-time requirements of power network simulation by only simulating an area containing a single

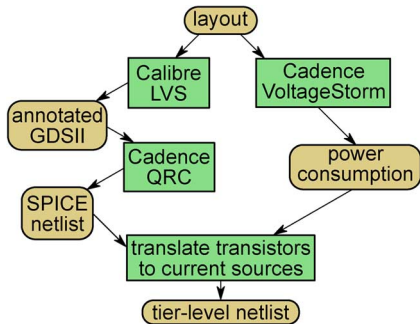


Fig. 9. Analysis flow used to obtain the tier-level netlist for IR-drop analysis. This flow is performed multiple times for each tier type, then the netlists are connected together with a TSV model for 3-D analysis.

core and the tiers directly above it as shown in Fig. 6. We stress that our design is extremely regular and so this reduction should only impact the accuracy of our analysis in a minor way.

We performed several experiments to verify the accuracy of our scaling results considering the reduced area coverage of our simulations. For example, we were able to simulate a five by five array of cores with one layer of SRAM above it. These results were then compared to a three by three array of cores and SRAM. The IR-drop results matched within 0.1%. This was repeated with successively narrower and taller stackings. All of the simulations matched within a small margin. It should be noted that the error introduced by this approach is systematic in nature, and should not affect the results of our scaling studies.

B. Validation

To validate the IR-drop analysis flow described above, we compare the results for a 2-D layout to Cadence’s VoltageStorm sign-off power noise analysis tool. The results of our analysis flow are within 4% of the values reported by VoltageStorm. We were also able to create a method for tricking VoltageStorm into performing 3-D analysis for two-tier stacks.

First we create an ICT file, a process technology description file, that contains a description of all of the metal layers in two tiers. The metal and dielectric layers are renamed so that the tier number is embedded in the name. For example, “METAL1” becomes “METAL1_1” and “METAL1_2.” Then, a techfile is created using Cadence’s TechGen based on the new ICT file. Next, we modify the LEF files provided by the foundry that describe the technology, standard cells, and macros. The DEF and instance power files for the designs of each tier are also modified in the same way. Each file is essentially duplicated so that there is one version for the first tier and one version for the second tier. The modifications basically amount to renaming the objects and metal layers in the same way that the ICT file is modified. To include detailed analysis of the macro blocks we also modify their GDSII files. We first convert the GDSII to GDT, an ascii-version of the binary GDSII data. Then we map all of the GDSII layer numbers for the metal layers into a non-overlapping number space. The modified GDT is then converted back to GDSII. The XTC extraction tool is then given a GDSII layer map file that maps the appropriate layer numbers to the correct tier’s metal layers for each macro.

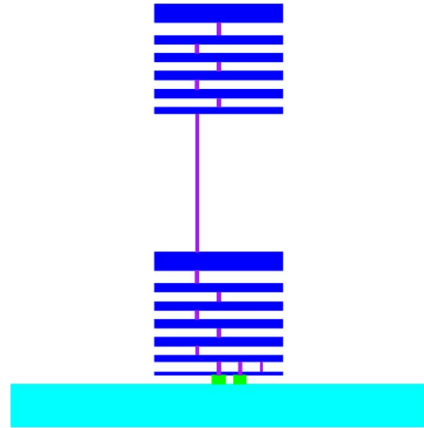


Fig. 10. Depiction of the ICT file that contains metal layers for two tiers of a 3-D stack. The ICT file is used to compile a techfile used for parasitic extraction by VoltageStorm for our 3-D IR-drop verification flow.




Fig. 11. Current waveform used for each transistor for dynamic noise analysis. A random delay is added to the start of the waveform for each gate’s transistors.

Using the above method we were able to match the 3-D IR-drop results from VoltageStorm within 4%. Fig. 10 shows a depiction from Cadence’s ViewICT tool of the modified ICT file containing metal layers in two dies. Note that the second die does not have a substrate layer. This is a limitation of the tool, due to the fact that it was not designed with 3-D designs in mind. However, for power/ground network analysis the substrate can largely be ignored. For these experiments we created a face-to-back style 3-D design, however, this technique is general enough to apply to face-to-face 3-D designs as well.

VI. 3-D DYNAMIC NOISE ANALYSIS

Layout-level dynamic noise values are computed using the power consumption values and parasitic extracted networks obtained for IR-drop analysis with added decoupling capacitors. We model the power grid using an RC network for each tier, which are connected together using RLC TSV and C4 models. We create triangular current demand waveforms [16] (see Fig. 11) for each transistor such that the average power consumption matches the value obtained for IR-drop analysis. The triangular waveforms for each gate are delayed by a random amount such that the majority of them start near the beginning of the cycle. The random delays are distributed in a Gaussian fashion about zero and then the absolute value of the delay is used for the real current waveform. Our dynamic noise numbers are obtained by performing transient simulation of a repeating pattern of current demand with a cycle time of 3 ns. The peak of the voltage swing is recorded as the noise value after the swings have stabilized.

Ultrasim’s power network simulation engine does not handle large-scale transient simulation well, so we used a custom SPICE simulator based on Modified Nodal Analysis [17],

TABLE I
EFFECTIVE INDUCTANCE VALUES IN pH FOR POWER DISTRIBUTION TSVs.
THE TSV DIMENSIONS ARE IN μm

TSV Dimensions	Clustered	Distributed
$3 \times 3 \times 10 \mu\text{m}$	0.829pH	0.014pH
$6 \times 6 \times 20 \mu\text{m}$	1.600pH	0.027pH
$10 \times 10 \times 33 \mu\text{m}$	2.500pH	0.041pH
$15 \times 15 \times 50 \mu\text{m}$	3.600pH	0.058pH

which returns results within 2% of HSPICE. For our simulations we use a step size of 1 ps. TSV inductance may also be an important contributor to dynamic noise. We modeled several sizes and arrangements of TSVs using Synopsys' inductance extractor Raphael [18]. Similar simulations using Ansys Q3-D finite element analysis software examining the RLC parasitics of signal TSVs in the MIT Lincoln Labs' process were performed by Savidis and Friedman [19]. However, our power-grid simulations cover a larger array of TSVs, and so need to be simplified to reduce simulation time. We calculated an effective inductance for each TSV based on the complex set of self and mutual inductances resulting from the following set of simulations.

We created a large array of conductors representing TSVs in the various dimensions to simultaneously represent both power and ground networks. The inductance matrix containing both self and mutual inductance was then calculated using the tool. These calculated values were then used in a SPICE netlist in a voltage divider configuration to determine the effective inductance of the TSV array. The effective inductance value of the array was then divided among the number of TSV conductors in the simulated array to determine the per-TSV effective inductance value. No mutual inductance exists between the TSVs and the on-chip wiring because the TSVs lie on an axis orthogonal to the horizontal die routing. Our method of calculating TSV inductance is somewhat conservative because it ignores the reduction in effective inductance caused by mutual inductances that could exist between neighboring signal and power distribution TSVs.

The results of these TSV inductance simulations are shown in Table I. The distributed TSV topology results in effective inductance values about two orders of magnitude smaller than the clustered TSV topology. The default TSV size used in most of our simulations is $6 \times 6 \times 20 \mu\text{m}$. Additionally, we performed the same validation experiments described in the previous section to examine the accuracy of our limited-area simulations. The simulations matched within 5% for dynamic noise.

VII. EXPERIMENTAL RESULTS

For our baseline analysis we assume copper via-first TSVs with $6 \mu\text{m}$ square diameter, $20 \mu\text{m}$ depth, $35 \text{ m}\Omega$ resistance, and 1.6 pH inductance. We chose to use a square aspect ratio TSV merely for convenience of calculation and simulation. Square $6 \mu\text{m}$ diameter TSVs have the same cross-sectional area as cylindrical $6.7 \mu\text{m}$ diameter TSVs. For simplicity, we present only the results for the ground distribution network. Simulations show that the power distribution network has the same trends, only the location of the maximum IR-drop peak is shifted. In real designs the difference between the actual supply and ground voltages are what determine the performance of

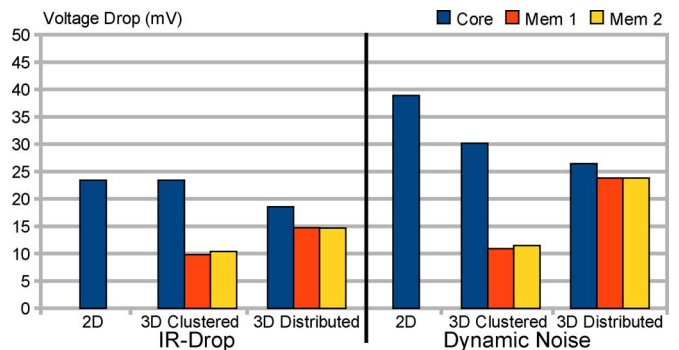


Fig. 12. Per-tier IR-drop and dynamic noise results for a 2-D design with one layer of cores only, a 3-D design using the clustered TSV topology, and a 3-D design with the distributed TSV topology. Both 3-D designs consist of three stacked tiers, one core and two memories (one set of the scalable prototype).

the gates. Given that we only simulate a single core and the tiers above it, we utilize a lumped package model for the C4 bumps. The C4 resistance and inductance in our simulations is $5 \text{ m}\Omega$ and 200 pH , respectively. Each of the memory tiers in our simulations consume about $0.7 \times$ the power value of the core tiers, so the term “low-power tier” is somewhat relative.

A. Power Supply Noise Comparison: Clustered vs Distributed

Fig. 12 shows IR-drop and dynamic noise results comparing a 2-D design with just cores to 3-D designs using one set of our scalable prototype with both the clustered and distributed TSV topologies. The 3-D design with the clustered TSV topology results in the same amount of IR-drop as the 2-D design, but lower dynamic noise than the 2-D design. The dynamic noise improvement is caused by the increase in on-chip decap present in the memory tiers. The 3-D design with the distributed TSV topology results in the lowest IR-drop and dynamic noise of all three cases shown for the reasons discussed in Section III. The distributed TSV topology improves IR-drop by 21% and dynamic noise by 32% over the 2-D system, even though the 3-D system consumes more total power.

Fig. 13 shows the effect on IR-drop of stacking more sets of the scalable prototype together. The distributed TSV topology provides a much lower IR-drop value as the number of sets stacked together becomes large. The distributed topology also allows up to six more tiers to be stacked together before crossing the 10% noise margin of 150 mV compared to the clustered topology. The basic reason for this improvement in IR-drop is that the distributed TSV topology allows the tiers with the most IR-drop to accept current through the networks with lower IR-drop. The distributed topology effectively utilizes the “IR-drop slack” of the low-power tiers to lower the maximum system-level IR-drop.

The clustered and distributed topologies result in very similar IR-drop values for systems with fewer numbers of sets stacked, as shown in Fig. 13. Fig. 14 shows the actual percentage improvement of the distributed topology over the clustered topology IR-drop for a few TSV site resistance values. These resistances are the resistance of each possible TSV location, called a TSV site. Fig. 17 contains a representation of this arrangement. For the distributed topology, there are 25 such resistances spread throughout the core layout. For the clustered

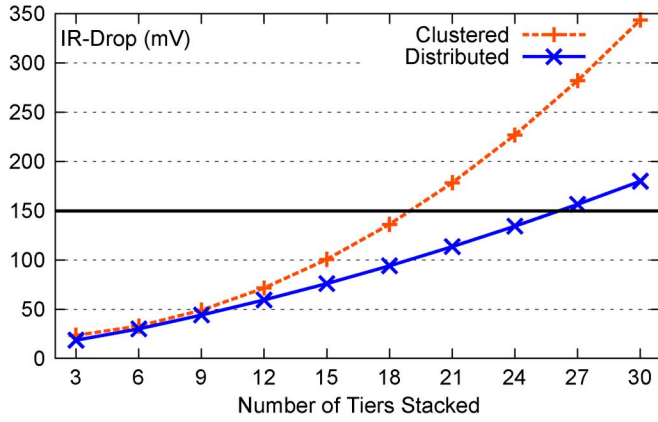


Fig. 13. Change in IR-drop as more sets of the scalable prototype layout are added. The line at 150 mV represents a 10% noise margin.




Fig. 14. IR-drop improvement of the distributed TSV topology over the clustered TSV topology as the number of tiers increases. Data for several values of TSV site resistances are shown.

topology, there are 25 clustered at the center TSV location over the C4 pad. The resistances can represent multiple TSVs at each location in parallel. The results in Fig. 14 show that TSV site resistance can have a significant impact on the relative IR-drop of the two topologies. However, for large numbers of sets stacked together, the distributed topology always eventually provides lower IR-drop than the clustered topology. Fig. 15 shows the improvement of the distributed TSV topology over the clustered TSV topology for both IR-drop and dynamic noise. In general, the IR-drop and dynamic noise improvement show roughly the same trend. The IR-drop improvement is slightly higher in most cases.

Section VII-C presents results relating to a wide range of TSV parasitic resistances. Dynamic noise simulations are much more time-consuming than IR-drop simulations. Additionally, inductance is not a simple scalable quantity like resistance. For example, adding more TSVs to a TSV site to reduce the parasitic resistance will not reduce the parasitic inductance in a linear fashion. For these reasons this subsection presents dynamic noise results for a small set of TSV sizes using both types of TSV topology. The TSV sizes and their inductance values are listed in Table I. Fig. 16 shows the dynamic noise improvement of the distributed TSV topology over the clustered TSV

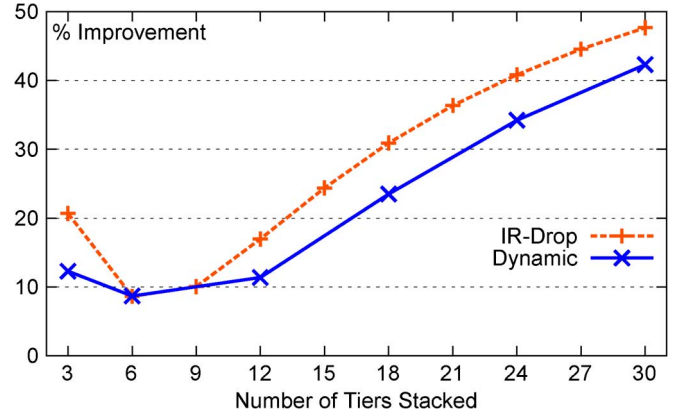


Fig. 15. Improvement of the distributed TSV topology over the clustered TSV topology as the number of tiers stacked together increases. Both dynamic noise and IR-drop improvement are shown.

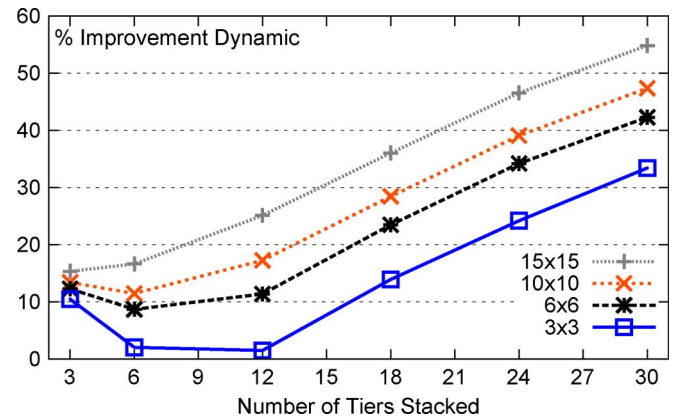


Fig. 16. Dynamic noise improvement of the distributed TSV topology over the clustered TSV topology as the number of tiers increases. Data for several TSV sizes (and associated parasitics) are shown.

topology for the four TSV sizes (and associated parasitics) examined. The dynamic noise shows very similar trends as the IR-drop (see Fig. 14).

Fig. 17 shows a 3-D representation of the IR-drop over the surface of the core farthest from the C4 supply pads for a system with two sets of our scalable prototype layout stacked together. The figure shows that the clustered TSV topology produces an IR-drop map that has a much larger spread between the maximum and minimum values. The large dip in the center of the clustered TSV topology mesh indicates the position of the TSV connected most directly to the ground network C4 bump for this core. The TSVs in the distributed topology help to pull down the IR-drop of the power distribution grid nodes that are farther from the C4 bump. The overall shape of the mesh demonstrates both that the TSVs are effective for lowering the maximum IR-drop and that more TSVs should be even more effective for that purpose.

To underscore the variation between the various tiers in the two TSV topologies, Fig. 18 shows the maximum IR-drop values in each core tier of a system with ten sets of our scalable prototype stacked together. The TSV site resistance is set to the baseline case, 35 mΩ. The difference between the maximum and minimum IR-drop in the system with the clustered TSV

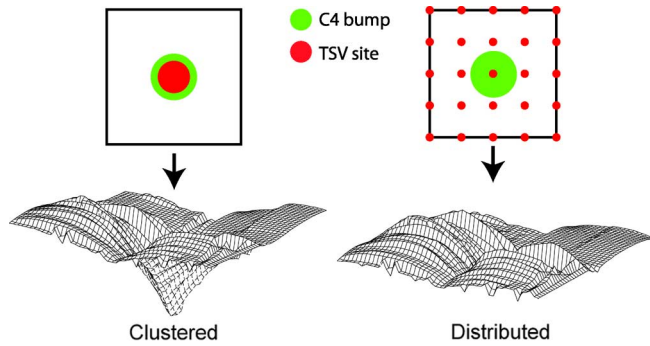


Fig. 17. IR-drop meshes for a single core in the highest core tier of two sets of our prototype layout stacked together. The left graph shows the results for the clustered TSV topology and the right graph shows the results for the distributed TSV topology. Both meshes are plotted using the same scale.

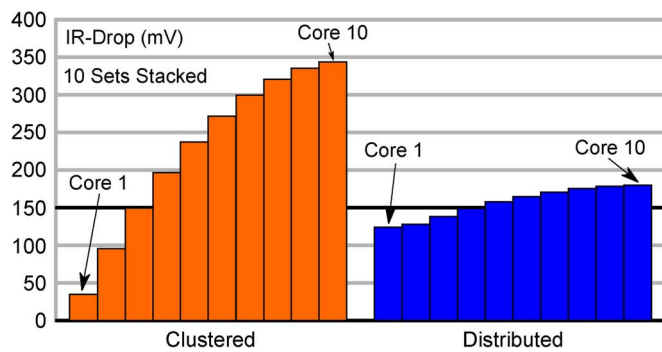


Fig. 18. Maximum per-core-tier IR-drop for ten sets of our prototype layout stacked together. The spread in values of the clustered TSV topology is much larger than for the distributed TSV topology.

topology is more than 300 mV, while the difference is less than 60 mV in the system with the distributed TSV topology. For the system with the clustered topology the transistors on the lower tiers would switch significantly faster than the transistors on the upper tiers.

B. Impact of Power Discrepancy Among Tiers

Next, we examine the effects of the power dissipation ratio between the memory tiers and the core tiers. As a reminder, the distributed TSV topology gains its IR-drop benefit from using the IR-drop slack of the low-power tiers to provide power to the high-power tiers. This implies that the total slack available (controlled by the power dissipation ratio) should effect the improvement of the distributed topology over the clustered topology. Fig. 19 shows the effect of setting the ratio at 0.5 and 1.4, as well as the default 0.7. An interesting feature of the graph is that as the number of tiers stacked increases the improvement of the distributed topology over the clustered topology becomes nearly identical for all cases. This indicates that the TSV resistance is more of a factor than the ratio of power dissipation between the low- and high-power tiers for these large-scale cases. For the case when power ratio is set to 0.5 there is extra slack available, so the distributed topology shows increased improvement. For the case when power ratio is set to 1.4 the core tiers are providing slack to the memory tiers.

Fig. 20 shows the effect of varying the power dissipation ratio on the dynamic noise improvement of the distributed topology




Fig. 19. IR-drop improvement of the distributed TSV topology over the clustered TSV topology as the number of tiers increases. The power dissipation ratio between the memory tiers and the core tiers is varied. The default ratio is 0.7 and the TSV site resistance for all cases shown is 35 m Ω .

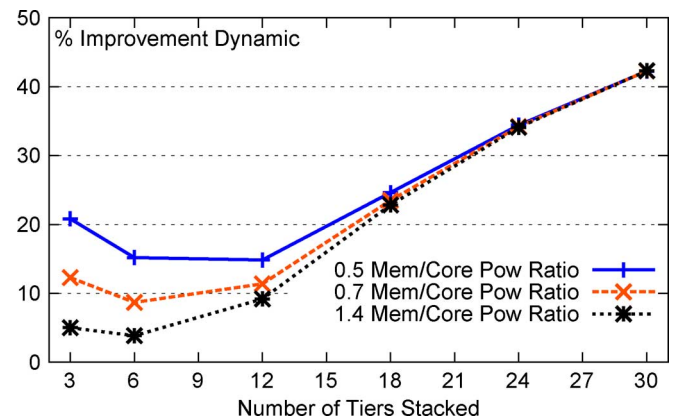


Fig. 20. Dynamic noise improvement of the distributed TSV topology over the clustered TSV topology as the number of tiers increases. The power dissipation ratio between the memory tiers and the core tiers is varied. The default ratio is 0.7.

over the clustered topology. The dynamic noise improvement exhibits similar trends to the IR-drop improvement shown in Fig. 19. The improvement of the distributed topology over the clustered topology becomes nearly identical for all power dissipation ratios as the number of tiers stacked increases. Again, this indicates that the TSV parasitics play a more important role in determining the improvement than the power dissipation ratio.

C. Impact of TSV Site Resistance on IR-Drop

Now we examine the effect of TSV site resistance on the IR-drop of the two topologies. The various values plotted could be created from longer or shorter TSV (and silicon) depth, different materials (tungsten, copper, etc.), varying interposer materials and contact resistances between stacked TSVs, or small arrays of TSVs in close proximity.

Fig. 21 shows the per-tier maximum IR-drop trends resulting from scaling TSV site resistance in the case where one "set" of our scalable prototype layout is stacked together (one core tier and two memory tiers). The graph shows several trends. First, because only the TSV site resistance is being scaled, the package resistance remains constant, and the core tier in the clustered TSV topology maintains the same IR-drop irrespective of TSV

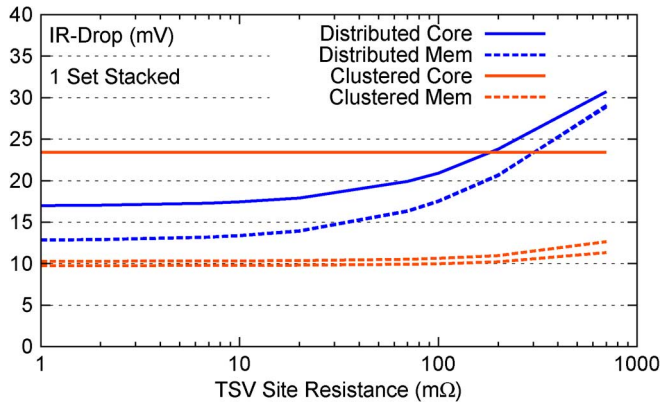


Fig. 21. Effect of TSV site resistance on the maximum IR-drop of one set of our prototype layout. Note that TSV site resistance is on a log scale. The solid lines represent the core tiers and the dashed lines represent the memory tiers.




Fig. 22. Effect of TSV site resistance on the maximum IR-drop of two sets of our prototype layout. Note that TSV site resistance is on a log scale. The solid lines represent the core tiers and the dashed lines represent the memory tiers.

site resistance. Second, the IR-drop scaling between the tiers of the distributed TSV topology shows a much stronger correlation, i.e. the maximum IR-drop of the various tiers in the system have values that are much closer together than for the clustered TSV topology. This indicates that the power networks of the neighboring tiers are tied together more strongly, and thus are able to support one another. Finally, while the distributed TSV topology has nearly 30% better IR-drop than the clustered topology for low TSV site resistance, the crossover point between the two styles occurs at around 200 mΩ TSV site resistance. For higher resistances the distributed TSV topology begins to suffer from much higher IR-drops.

Increasing the number of sets stacked to two, we repeat the site resistance scaling simulations in Fig. 22. To reduce visual clutter we only show the results for the two core tiers as well as the top (furthest from the supply bumps) two memory tiers. In this graph the TSV site resistance begins to affect the IR-drop scaling behavior of the clustered TSV topology. Again, the more correlated nature of the IR-drop between the various tiers in the distributed TSV topology is evident. Also of note is that the crossover point when the distributed topology produces higher




Fig. 23. Effect of TSV site resistance on the maximum IR-drop of ten sets of our prototype layout. Note that TSV site resistance is on a log scale. The IR-drop of only cores 1, 5, and 10 are shown.

IR-drop than the clustered topology has shifted to the left compared to Fig. 21. This effect can also be seen in the context of Figs. 13 and 14.

Further increasing the number of sets to the maximum examined, ten sets stacked together, we again repeat the resistance scaling simulations in Fig. 23. To maintain readability we only show the results for three core tiers, cores 1, 5, and 10. The same trends as in the previous graphs remain evident, though there are several interesting observations to be made. It is interesting to note the change in scale on the dependent axis (IR-drop) between Figs. 21 and Fig. 23. The maximum IR-drop plotted increases from 40 to 800 mV. Also, the variation between the tiers in the clustered TSV topology become even more extreme in the case with 30 tiers stacked together. The difference between the maximum and minimum is more than 700 mV when the TSV site resistance is 700 mΩ. These results are very unrealistic, but show trends that result from pushing our model to its limits.

D. Possible Electro-Migration Issues

Electro-migration has become an increasingly important consideration in deep sub-micron IC design. Figs. 24 and 25 show the current density in the TSVs for the clustered and distributed TSV topologies, respectively. The baseline TSV dimensions are assumed: 6 μm square diameter and 20 μm length. The maximum current density in the TSVs configured in the clustered topology is much lower than for the TSVs in the distributed topology. This is the result of the increased resistance between the C4s and the TSVs that are distributed further away. Current follows the path of least resistance, thus the TSVs with the least resistance, the ones directly over the C4 or nearby, will have higher current. In a realistic implementation the higher current density of the TSVs in the distributed case would be mitigated by increasing the TSV count in the TSV sites over the C4 bumps. In this work we have used the same number of TSVs for the clustered and distributed cases for fairness of comparison. The clustered topology does not exhibit spikes in TSV current density because all the TSVs have virtually identical resistance to the C4, so the current density in each TSV of the cluster is uniform. It should also be noted that for this particular case the




Fig. 24. TSV current density for the clustered topology. The current density numbers are sorted in increasing order from left to right.




Fig. 25. TSV current density for the distributed topology. The current density numbers are sorted in increasing order from left to right.

current density is far below the limits of most commonly-used TSV conductors.

VIII. TECHNIQUES FOR DECREASING NOISE

A. Decreasing C4 Bump Pitch

The Euclidean distance between neighboring C4 bumps in our default layout is $400 \mu\text{m}$. Given the low power dissipation of a single core this is sufficient for low-tier systems, however, for our 1000-core system the IR-drop is still above the 10% noise margin, even using the distributed TSV topology. In this and the following subsections we examine several methods to reduce the IR-drop and dynamic noise for our 1000-core system to meet the requirements. Fig. 26 shows the IR-drop when the C4 bump pitch is reduced to allow 6 and 10 bumps per core. As the figure shows, it is possible to reduce the IR-drop for our 1000-core system below the 10% noise margin, 150 mV, by adding 6 or more C4 bumps per core, which translates to a C4 bump pitch below $200 \mu\text{m}$. Also of note, lower IR-drop is achieved by using the distributed TSV topology than by halving the C4 bump pitch. Fig. 27 shows the same comparison for dynamic noise. The trends are very similar to the trends for IR-drop.

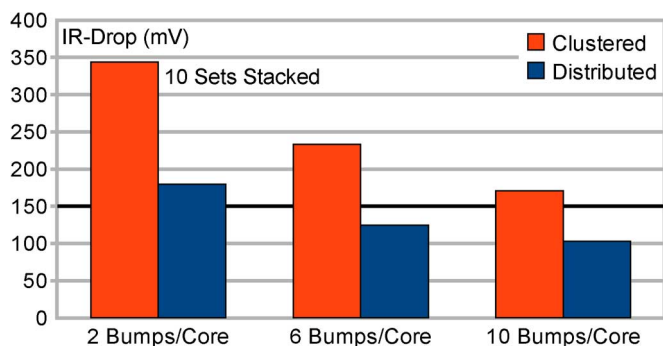


Fig. 26. Maximum IR-drop for ten sets of our prototype layout stacked together with increasing numbers of C4 bumps per core added.

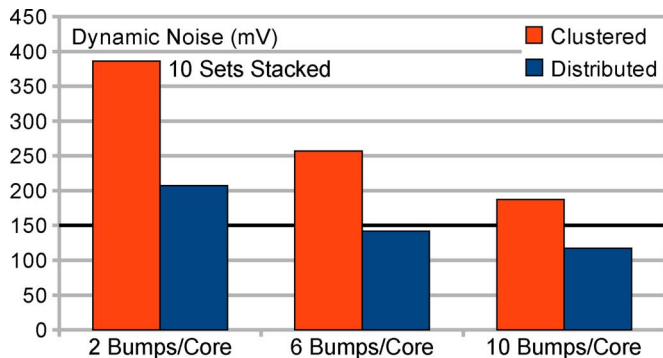


Fig. 27. Maximum dynamic noise for ten sets of our prototype layout stacked together with increasing numbers of C4 bumps per core added.

B. Adding Decap Tiers

This subsection provides results related to the addition of layers containing only decoupling capacitance. We created the layout for this tier using the same 130-nm process used in the rest of our analysis, however, we expect that real systems incorporating decap tiers will use processes that enhance the capacitance per unit area while reducing production cost. Our decap tier contains a total of nearly 0.18 nF per core, for a total of 18 nF per tier, at a density of $0.57 \text{ fF}/\mu\text{m}^2$.

Fig. 28 shows the improvement generated by adding either a single decap tier *per system* or one decap tier *per set* over the same system without decap tiers. In this case the decap tier is always added to the set on the side that is farthest from the power supply bumps (i.e., the ordering would be bump, core, memory, memory, decap, core, memory, etc.). The figure shows that the distributed TSV topology always benefits more from extra decap than the clustered topology. It also shows that beyond 4 sets stacked (12 tiers) the clustered topology results in *worse* power supply performance when adding one decap tier *per set*. The extra decoupling capacitance of the decap tiers cannot overcome the additional TSV parasitics that are introduced by adding decap tiers to the system. Also, adding one decap tier per set for systems with eight (24 tiers) or more sets stacked results in lower improvement than adding a single tier per system for the distributed TSV topology.

Next, Fig. 29 shows the impact on dynamic noise of adding increasing numbers of decap tiers for the 10 sets stacked system. Each decap tier is distributed throughout the stack nearest one set of the scalable design. The tiers are added to the sets nearest

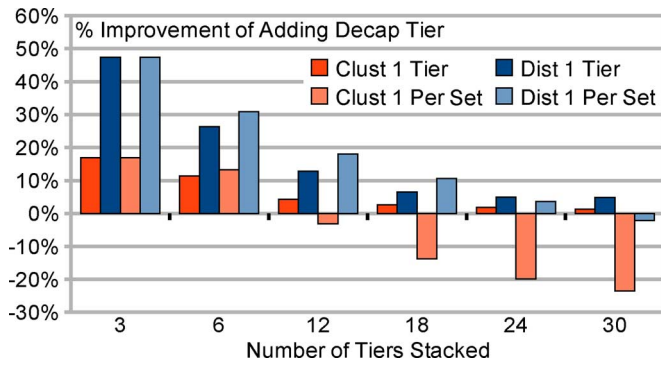


Fig. 28. Improvement in dynamic noise created by adding either a single decap tier per system (1 Tier) or one decap tier per set (1 Per Set) to our scalable processor.

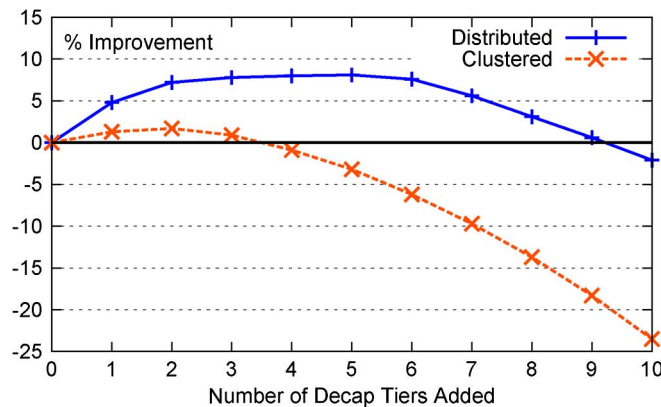


Fig. 29. Improvement in dynamic noise over the baseline case when adding increasing numbers of decap tiers for the 10 sets stacked system.

the heatsink first, and then further down towards the sets near the supply bumps. Like, the earlier discussion, the figure shows that small numbers of decap tiers can decrease dynamic noise, but adding too many decap tiers increases overall TSV parasitics and results in increased dynamic noise.

C. Pass-Through TSVs

Finally, we examine TSVs that pass-through the lower tiers without connecting to their power grids. These TSVs are meant to supply power only to the higher tiers in the stack. This trades off some additional lateral IR-drop in the lower tiers for lower maximum IR-drop in the system as a whole. A physical depiction of this design approach is shown in Fig. 30 for the two TSV topologies. Fig. 31 shows the impact on maximum IR-drop of increasing stacking. Fig. 32 shows the maximum dynamic noise with increased stacking. The simulations are for a case with 6 bumps/core. The results show that this technique is beneficial for large stacks with the clustered TSV topology, and reduces IR-drop for the 10 sets stacked case by nearly 18% and dynamic noise by nearly 17%. This style of implementing pass-through TSVs does not improve noise for systems with the distributed TSV topology.

There are many possible connection topologies for pass-through TSVs in combination with the distributed TSV topology. After some searching we found a set of connections that result in lower supply noise for the distributed TSV case.




Fig. 30. Side view of a 3-D stack with pass-through power distribution TSVs. The TSVs connected to the C4 bump on the right do not connect to the distribution wiring on the lower two tiers.




Fig. 31. Maximum IR-drop with pass-through power distribution TSVs. The results are for a case with 6 bumps/core.




Fig. 32. Maximum dynamic noise with pass-through power distribution TSVs. The results are for a case with six bumps/core.

Fig. 33 depicts this topology. In the distributed topology, the memory layers provide noise slack and decap to the core layers. Passing through some of the lower-level core layers therefore allows more of the memory layers to lower the noise level of the highest core layer, which has the maximum noise in the system. There remain several parameters, such as number of TSVs that pass through and the depth of the stack that they pass through, that are related to this pass-through TSV connection topology and that effect the final noise performance of the system. The affect of these parameters is relatively small, but we demonstrate the use of this design style to show that pass-through TSVs can still be beneficial for the distributed TSV topology. Fig. 34 shows the dynamic noise results of using various number of TSVs passing through various numbers of core layers below the uppermost core layer. The figure shows that passing through more core layers always results in decreasing dynamic noise. Additionally, there is an optimal number of pass-through TSVs that should be used. If too few

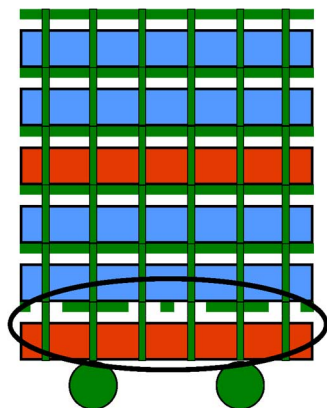


Fig. 33. Side view of a 3-D stack with an alternative connection topology of pass-through power distribution TSVs for the distributed TSV topology. The TSVs not connected to the C4 bumps do not connect to the distribution wiring on the lower core tier (orange).




Fig. 34. Maximum dynamic noise with alternative pass-through power distribution TSVs for the distributed TSV topology. The number of core layers that are passed through varies, as well as the number of non-C4 TSVs that pass-through. The results are for the 10-sets stacked case with six bumps per core.

are used the benefits are very small, and if too many are used the lower core tiers begin to exhibit higher maximum noise.

IX. CONCLUSION

In this work we have explored 3-D power delivery network design and shown that both IR-drop and dynamic noise can be improved in these systems by exploiting the particular attributes of power supply TSVs that are unique compared to those of C4 supply bumps. Previous works have assumed a straightforward extension of traditional power supply network design in which the TSVs are treated as an extension of the C4 bumps. We advocate a design style in which power network TSVs are distributed with small pitch (relative to the package bumps) throughout the entire surface of the layout to increase the level of coupling between the power distribution networks of the various tiers in the 3-D stack. This allows the utilization of decoupling capacitance and IR-drop and dynamic noise slack in the lower-power tiers to reduce maximum system-level IR-drop and dynamic noise.

To support our claims we designed a 1000-core 3-D processor across 30 stacked tiers at the layout level. Our 3-D IR-drop analysis method was verified against commercial-grade sign-off IR-drop analysis software from a major EDA vendor at both the 2-D and two-tier 3-D level. Detailed and extensive simulations

of the stacking scaling and TSV resistance scaling demonstrate that the distributed TSV topology generally provides much lower IR-drop and dynamic noise. In our baseline system with 30 stacked tiers the distributed topology provides nearly 50% lower IR-drop and 42% lower dynamic noise than the clustered topology. For low-tier systems the savings are still significant. In fact, the distributed TSV topology lowers IR-drop for a 3-tier system compared to a non-3-D system by 21%, and dynamic noise by 32%, even though the total power consumption is higher in the 3-tier system. We also examine several more complex techniques to reduce power supply noise and their effects on both the clustered and distributed TSV topologies.

The optimal location of TSVs distributed to improve power-supply performance is completely dependent on the topology of the power-distribution network under consideration. It is also constrained by the floorplanning and placement requirements of the design. Designers wishing to use this technique in practical applications must remember that the benefit occurs when the distribution wiring between high- and low-power tiers is shared with finer than C4 pitch. As long as the lateral distribution drop is larger than the drop through the TSVs, then moving even a few TSVs away from the C4 bumps may be beneficial. Additionally, it should be noted that designs with large hard macros or array structures are not well suited to placing distributed TSVs with small enough pitch. The most ideal designs for the distributed topology contain seas of standard cells on all shared dies with regular small-pitch patterns of distribution wiring.

It should also be noted that distributing TSVs occurs between neighboring C4s connected to the same power net. For example, in a design with a regularly alternating pattern between power and ground, a potentially beneficial place to locate distributed TSVs for the power net is very close to the ground C4. The largest macro sizes need not be below the C4 pitch itself to benefit from distributed TSVs; the macro sizes need only be below the power-to-power or ground-to-ground pitch.

REFERENCES

- [1] G. Huang, M. Bakir, A. Naeemi, H. Chen, and J. Meindl, "Power delivery for 3D chip stacks: Physical modeling and design implication," in *Proc. IEEE Conf. Elect. Perform. Electron. Packag.*, 2007, pp. 205–208.
- [2] P. Jain, T.-H. Kim, J. Keane, and C. H. Kim, "A multi-story power delivery technique for 3D integrated circuits," in *Proc. Int. Symp. Low Power Electron. Design*, 2008, pp. 57–62.
- [3] J. Gu and C. H. Kim, "Multi-story power delivery for supply noise reduction and low voltage operation," in *Proc. Int. Symp. Low Power Electron. Design*, 2005, pp. 192–197.
- [4] H. Yu, J. Ho, and L. He, "Allocating power ground vias in 3D ICs for simultaneous power and thermal integrity," *ACM Trans. Design Autom. Electron. Syst.*, vol. 14, no. 3, pp. 1–31, 2009.
- [5] M. B. Healy and S. K. Lim, "Power delivery system architecture for many-tier 3D systems," in *Proc. IEEE Electron. Components Technol. Conf.*, 2010, pp. 1682–1688.
- [6] S. Sapatnekar, P. Zhou, and K. Sridharan, "Power grid optimization in 3D circuits using MIM and CMOS decoupling capacitors," *IEEE Design Test Comput.*, vol. 26, no. 5, pp. 15–25, 2009.
- [7] T. Thorolfsson, K. Gonsalves, and P. Franzon, "Design automation for a 3DIC FFT processor for synthetic aperture radar: A case study," in *Proc. IEEE/ACM Design Autom. Conf.*, 2009, pp. 51–56.
- [8] J. Sun, J.-Q. Lu, D. Giuliano, T. P. Chow, and R. J. Gutmann, "3D power delivery for microprocessors and high-performance ASICs," in *Proc. IEEE Appl. Power Electron. Conf.*, 2007, pp. 127–133.
- [9] V. F. Pavlidis and G. De Micheli, "Power distribution paths in 3-D ICs," in *Proc. Great Lakes Symp. VLSI*, 2009, pp. 263–268.

- [10] M. Bakir, H. Reed, P. Kohl, K. Martin, and J. Meindl, "Sea of leads ultra high-density compliant wafer-level packaging technology," in *Proc. IEEE Electron. Components Technol. Conf.*, 2002, pp. 1087–1094.
- [11] A. Aggarwal, P. Markondeya Raj, M. Sacks, A. Tay, and R. Tumala, "Ultra fine-pitch wafer level packaging with reworkable composite nano-interconnects," in *Proc. IEEE Electron. Packag. Technol. Conf.*, 2004, pp. 132–137.
- [12] B. Goplen and S. S. Sapatnekar, "Placement of thermal vias in 3-D ICs using various thermal objectives," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 25, no. 4, pp. 692–709, 2006.
- [13] M. B. Healy, K. Athikulwongse, R. Goel, M. M. Hossain, D. H. Kim, Y.-J. Lee, D. L. Lewis, T.-W. Lin, C. Liu, M. Jung, B. Ouellette, M. Pathak, H. Sane, G. Shen, D. H. Woo, X. Zhao, G. H. Loh, H.-H. S. Lee, and S. K. Lim, "Design and analysis of 3D-MAPS: A many-core 3D processor with stacked memory," in *Proc. IEEE Custom Integr. Circuits Conf.*, 2010, pp. 1–4.
- [14] D. Sekar, C. King, B. Dang, T. Spencer, H. Thacker, P. Joseph, M. Bakir, and J. Meindl, "A 3D-IC technology with integrated microchannel cooling," in *Proc. IEEE Int. Interconnect Technol. Conf.*, 2008, pp. 13–15.
- [15] M. B. Healy and S. K. Lim, "A study of stacking limit and scaling in 3D ICs: An interconnect perspective," in *Proc. IEEE Electron. Components Technol. Conf.*, 2009, pp. 1213–1220.
- [16] Q. K. Zhou, *Power Distribution Network Design for VLSI*. New York: Wiley-IEEE, 2004.
- [17] C.-W. Ho, A. E. Ruehli, and P. A. Brennan, "The modified nodal approach to network analysis," *IEEE Trans. Circuits Syst.*, vol. 22, no. 6, pp. 504–509, Jun. 1975.
- [18] Synopsys, San Jose, CA, "Raphael," 2009. [Online]. Available: <http://www.synopsys.com>
- [19] I. Savidis and E. G. Friedman, "Electrical modeling and characterization of 3-D vias," in *Proc. IEEE Int. Symp. Circuits Syst.*, 2008, pp. 784–787.



Michael B. Healy (M'11) received the B.S., M.S., and Ph.D. degrees from the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, in 2004, 2006, and 2010, respectively.

From 2006 to 2007, he interned with Intel and worked in the Penryn architecture team on their first 45 nm processor. After earning the Ph.D. degree, he moved to a post-doctoral research position with IBM. His overall technical interest is on 3-D integration technology and the rewards and challenges it brings to the semiconductor industry. His research has

spanned a broad range of work and includes placement for configurable architectures, thermal/performance trade-offs in 2-D and 3-D microarchitectural floorplanning, power-supply-noise-aware microarchitectural floorplanning, multi-granularity thermal-aware floorplanning for multi-core architectures, and power and thermal issues in large-scale 3-D systems. His current interest is in enabling multiple types of detailed analysis at the early stages of system design.



Sung Kyu Lim (S'94–M'00–SM'05) received the B.S., M.S., and Ph.D. degrees from the Computer Science Department, University of California, Los Angeles (UCLA), in 1994, 1997, and 2000, respectively.

From 2000 to 2001, he was a Post-Doctoral Scholar with UCLA, and a Senior Engineer at Aplus Design Technologies, Inc. He joined the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, in 2001, where he is currently an Associate Professor. His research focus

is on the physical design automation for 3-D ICs, 3-D system-in-packages, microarchitectural physical planning, and field-programmable analog arrays. He is the author of *Practical Problems in VLSI Physical Design Automation* (Springer, 2008).

Dr. Lim was a recipient of the National Science Foundation Faculty Early Career Development (CAREER) Award in 2006. He was on the Advisory Board of the ACM Special Interest Group on Design Automation (SIGDA) during 2003–2008 and received the ACM SIGDA Distinguished Service Award in 2008. He was an Associate Editor of the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS during 2007–2009. His work is nominated for the Best Paper Award at ISPD'06, ICCAD'09, CICC'10, and DAC'11. He is a member of the Design International Technology Working Group of the International Technology Roadmap for Semiconductors (ITRS). He has been leading the Cross-Center Theme on 3-D Integration for the Focus Center Research Program (FCRP), Semiconductor Research Corporation (SRC), since 2010.