

# Addressing Thermal and Power Delivery Bottlenecks in 3D Circuits

Sachin S. Sapatnekar  
Department of Electrical and Computer Engineering  
University of Minnesota  
Minneapolis, MN 55455  
sachin@umn.edu

**Abstract**— The enhanced packing densities facilitated by 3D integrated circuit technology also has an unwanted side-effect, in the form of increasing the amount of current per unit footprint of the chip, as compared to a 2D design. This has ramifications on two critical issues: firstly, it means that more heat is generated per unit footprint, potentially leading to thermal problems, and secondly, more current must be supplied per package pin, leading to possible power delivery bottlenecks. This paper presents an overview of the challenges and solutions in the domain of addressing these two issues in 3D integrated circuits.

## I. INTRODUCTION

One of the primary advantages of 3D chips stems from their ability to pack circuitry more densely than in 2D. However, this increased level of integration also results in side-effects, in the form of new limitations and challenges to the designer. Thermal and power delivery problems can both be traced to the fact that a  $k$ -tier 3D chip could use  $k$  times as much current as a single 2D chip of the same footprint, while using substantially similar packaging technology. The implications of this are as follows:

- First, the 3D chip generates  $k$  times the power of the 2D chip, which implies that the corresponding heat generated must be sent out to the environment. If the design technique is thermally unaware, and the package thermal characteristics for 2D and 3D circuits are similar, this implies that on-chip temperatures on 3D chips will be higher than for 2D chips. Elevated temperatures can hurt performance and reliability, in addition to introducing variabilities in the performance of the chip. Therefore, on-chip thermal management is a critical issue in 3D design.
- Second, the package must be capable of supplying  $k$  times the current through the power supply ( $V_{dd}$  and ground) pins, as compared to the 2D chip. Moreover, the power delivery problem is worsened in 3D ICs as through-silicon vias (TSVs) contribute additional resistance to the supply network. Given that reliable power grid design is a major bottleneck even for 2D designs, this implies that significant resources have to be invested in building a bulletproof power grid for the 3D chip.

It can be, therefore, argued that thermal and power delivery issues are two sides of the same coin. This paper presents an overview of these two issues and reviews available solutions to overcome these problems.

## II. THERMAL ISSUES IN 3D ICs

### A. Full-chip thermal analysis

Full-chip thermal analysis involves the application of classical heat transfer theory. The differences lie in incorporating issues that are specific to the on-chip context: for example, on-chip geometries are strongly rectilinear in nature, and involve rectangular geometric symmetries; the major sources of heat, the devices, lie in either a single layer in each 3D tier, and the points at which a user is typically interested in analyzing temperature are within the device layer(s).

On-chip thermal behavior at the macroscale is defined by the heat equation, which is a parabolic PDE [1]:

$$\rho c_p \frac{\partial T(\mathbf{r}, t)}{\partial t} = k_t \nabla^2 T(\mathbf{r}, t) + g(\mathbf{r}, t), \quad (1)$$

where  $\rho$  is the density of the material (in  $\text{kg/m}^3$ ),  $c_p$  is the heat capacity of the chip material (in  $\text{J}/(\text{kg K})$ ),  $T$  is the temperature (in K),  $\mathbf{r}$  is the spatial coordinate of the point at which the temperature is being determined,  $t$  represents time (in sec),  $k_t$  is the thermal conductivity of the material (in  $\text{W}/(\text{m K})$ ), and  $g$  is the power density per unit volume (in  $\text{W}/\text{m}^3$ ).

The thermal conductivity,  $k_t$ , in a uniform medium is isotropic, and thermal conductivity values for silicon, silicon dioxide, and metals such as aluminum and copper are fundamental material properties whose values can be determined from standard tables. In practice, in early stages of analysis and for optimization purposes, integrated circuits may be assumed to be layer-wise uniform in terms of thermal conductivity. The bulk layer has the conductivity of bulk silicon, and the conductivity of the metal layers is often computed using an averaging approach: this region consists of a mix of silicon dioxide and metal, and depending on the metal density within the region, an effective thermal conductivity may be used for macroscale analysis.

The solution to Equation (1) corresponds to the transient thermal response. In the steady state, all derivatives with respect to time go to zero, and therefore, steady-state analysis corresponds to solving the PDE:

$$\nabla^2 T(\mathbf{r}) = -\frac{g(\mathbf{r})}{k_t} \quad (2)$$

This is the well-known Poisson's equation.

The time constants of heat transfer are of the order of milliseconds, and are much longer than the subnanosecond clock periods in today's VLSI circuits. Therefore, if a circuit remains within the same power mode for an extended period of time, and its power density distribution remains relatively constant, steady-state analysis can capture the thermal behavior of the circuit accurately. Even if this is not the case, steady-state analysis can be particularly useful for early and more approximate analysis, in the same spirit that steady-state analysis is used to analyze power grid networks early in the design cycle. On the other hand, when greater levels of detail about the inputs are available, or when a circuit makes a number of changes between power modes at time intervals above the thermal time constant, transient analysis is possible and potentially useful.

To obtain a well-defined solution to Equation (1), a set of boundary conditions must be imposed. Typically, at the chip level, this involves building a package macromodel, and assuming that this macromodel interacts with a constant ambient temperature.

Both the finite difference method (FDM) and the finite element method (FEM) discretize the entire chip and form a system of linear equations relating the temperature distribution within the chip to the power density distribution. The major difference between the FDM and FEM is that while the FDM discretizes the differential operator, the FEM discretizes the temperature field. The primary advantage of the FDM and FEM is their capability of handling complicated material structures, particularly non-uniform interconnect distributions in a VLSI chip.

The FEM and FDM methods both lead to problem formulations that require the solution of large systems of linear equations. The matrices that describe these equations are typically sparse (more so for the FDM than the FEM, as can be seen from the individual element stamps) and positive definite. There are many different ways of solving these equations, including direct methods based on use variants of Gaussian elimination, such as LU factorization-based techniques, and iterative methods such as Gauss-Jacobi, Gauss-

Seidel, and successive overrelaxation, as well as more contemporary approaches based on the conjugate gradient method or GMRES. For FDM methods, the similarity with the power grid analysis problem invites the use of similar solution techniques, including random walk methods [2] and other methods such as multigrid approaches [3].

The finite difference method employs a standard device in heat transfer theory that builds an equivalent *thermal circuit* through the *thermal-electrical analogy*. Each node in the discretization corresponds to a node in the circuit. The steady-state equation corresponds to a network where “thermal resistors” are connected between nodes that correspond to spatially adjacent regions, and “thermal current sources” that map on to power sources. The voltages at the nodes in this thermal circuit can then be computed by solving the circuit, and these yield the temperature at those nodes.

The ground node, or reference, for the circuit corresponds to a constant temperature node, which is typically the ambient temperature. If isothermal boundary conditions are to be used, this simply implies that the node(s) connected to the ambient correspond to the ground node. On the other hand, it is possible to use a more detailed thermal model for the package and heat spreader, consisting of an interconnection of thermal resistors and thermal capacitors, or another type of compact model such as a reduced-order model, with a connection to the ambient, or ground, node.

The overall equations for the circuit may be formulated using modified nodal analysis, and we may obtain a set of equations

$$\mathbf{G}\mathbf{T} = \mathbf{P} \quad (3)$$

Here  $G$  is an  $n \times n$  matrix and  $\mathbf{T}, \mathbf{P}$  are  $n$ -vectors, where  $n$  corresponds to the number of nodes in the circuit. It is easy to verify that the  $G$  matrix is a sparse conductance matrix that is symmetric and is diagonally dominant.

For transient thermal analysis, the time-dependent left hand side term in Equation (1) is nonzero. Using a similar finite differencing strategy as above, the equation may be discretized in the space domain, and the concept of a *thermal capacitor* may be defined. Transient thermal analysis involves the solution of an RC circuit with current and voltage sources, and again, shows similarities with transient power grid analysis.

The FEM provides another avenue to solve Poisson’s equation. A succinct explanation of FEM, as applied to the on-chip case, is provided in [4]. In finite element analysis, the design space is first discretized or meshed into elements. Different element shapes can be used such as tetrahedra and hexahedra. For the on-chip problem, where all heat sources are modeled as being rectangular, a reasonable discretization [5] for the FEM divides the chip into 8-node rectangular hexahedral elements. The temperatures at the nodes of the elements constitute the unknowns that are computed during finite element analysis, and the temperature within an element is calculated using an interpolation function that approximates the solution to the heat equation within the elements.

As in the case of circuit simulation using the modified nodal formulation, stamps are created for each element and added to the global system of equations, given by:

$$\mathbf{K}_g \mathbf{T} = \mathbf{P} \quad (4)$$

where  $\mathbf{T}$  is the vector of all the nodal temperatures. This system of equations is typically sparse and can be solved efficiently. In the FEM, these stamps are called element stiffness matrices,  $K$ , and their values can be determined using techniques based on the calculus of variations. The fixed values in  $\mathbf{T}$ , corresponding to the ground nodes, can be moved to the right hand side to obtain the reduced set of equations

$$\mathbf{K}'_g \mathbf{T} = \mathbf{P}' \quad (5)$$

### B. Thermal optimization of 3D circuits

The illustration in Figure 1 shows a simple thermal model for a 3D circuit, and outlines techniques for overcoming thermal challenges in these structures. The figure shows a schematic of a 3D chip

sitting atop a heat sink: this is modeled using a distributed power source feeding a distributed resistive network, connected to a thermal resistance that models the heat sink. Although this is a coarse model, it suffices for illustrative purposes. By the thermal-electrical analogy, the voltage in this network represents the temperature on the chip. The temperature can therefore be reduced using the following schemes:

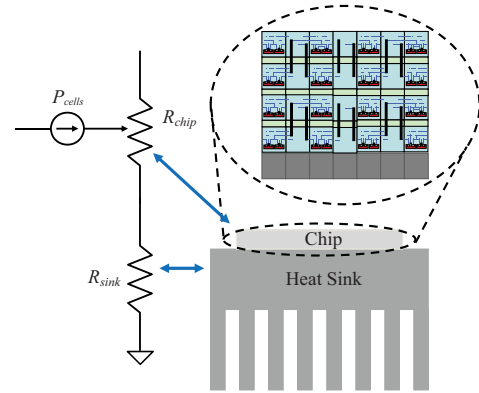


Fig. 1. A simple thermal model for a 3D chip.

- *Through low power design*: By reducing the power dissipation of the chip, the thermal current injected into the network can be reduced, controlling the IR drop, and therefore, the voltage.
- *By rearranging the heat sources*: The locations of the heat sources can be altered through physical design (floorplanning and placement) to obtain improved temperatures. Coarsely speaking, this implies that high power modules should be moved away from each other.
- *By improving thermal conduits*: The temperature may also be reduced by improving the effective thermal conductivity of paths from the devices to the heat sink. An effective method for achieving this is through the insertion of *thermal vias*: thermal vias are structurally similar to electrical vias, but serve no electrical purpose. Their primary function is to conduct heat through the 3D structure and convey it to the heat sink.
- *By improving the heat sink*: An improved heat sink results in an improvement in the value of  $R_{sink}$ , which can help reduce the temperature.




Fig. 2. A placement for the benchmark ibm01 in a four-tier 3D technology [6].

As an example of thermal optimization by spreading the heat sources, it is instructive to view the result of a typical 3D thermally-aware placement [5]: a layout for the benchmark circuit, ibm01, in a four-tier 3D process, is displayed in Figure 2. The cells are positioned in ordered rows on each tier, and the layout in each individual tier looks similar to a 2D standard cell layout. The heat sink is placed at the bottom of the 3D chip, and the lighter shaded regions are hotter than the darker shaded regions. The coolest cells are those in the bottom tier, next to the heat sink, and the temperature increases as we move to higher tiers. The thermal placement method consciously

mitigates the temperature by making the upper tiers sparser, in terms of the percentage of area populated by the cells, than the lower tiers.

In this paper, we will focus our discussion primarily on the second and third items in the list; the others, namely, low-power design and heat sink design, are broad and widely studied topics, and are not described here. This discussion will focus primarily on physical design issues, specifically on floorplanning, placement, thermal via insertion, and routing, although there is significant scope for improvement at the architectural level as well.

### C. 3D floorplanning

The 3D floorplanning problem is analogous to the 2D problem, with all the constraints and opportunities that arise with the move to the third dimension. Typical cost functions include a mix of the conventional wirelength and total area costs, and the temperature and the number of intertier vias.

The approach in [7] presented one of the first approaches to 3D floorplanning, and used the TCG representation, described in Section 11.7, for each tier, and a bucket structure for the third dimension. Each bucket represents a two-dimensional region over all tiers, and stores, for each tier, the indices of the blocks that intersect that bucket. In other words, the TCG and this bucket structure can quickly determine any adjacency information. A simulated annealing engine is then utilized, with the moves corresponding to perturbations within a tier and across tiers; in each such case, the corresponding TCG(s) and buckets are updated, as necessary.

A simple thermal analysis procedure is built into this solution, using a finite difference approximation of the thermal network to build an RC thermal network. Under the assumption that heat flows purely in the  $z$  direction and there is no lateral heat conduction, the RC model obtained from a finite difference approximation has a tree structure, and Elmore-like computations (see Section 3.1) can be performed to determine the temperature. The optimization heuristically attempts to make this a self-fulfilling assumption, by discouraging lateral heat conduction, introducing a cost function parameter that discourages strong horizontal gradients. A hybrid approach performs an exact thermal analysis once every 20 iterations or so and uses the approximate approach for the other iterations.

The work in [8] expands the idea of thermally driven floorplanning by integrating thermal via insertion into the simulated annealing procedure. A thermal analysis procedure based on random walks [9] is built into the method, and an iterative formula, similar to [10], is used in a thermal via insertion step between successive simulated annealing iterations.

Another approach in [11] using a force-directed formulation to perform floorplanning. This thermally-driven approach uses a three-stage flow, starting with global optimization in 3D space, to optimization in 2.5D space with layer assignment, and ending with macro-block legalization, and shows improvements over [7] in both quality of the result and runtime.

### D. 3D placement

Several approaches to 3D placement have been proposed in the literature. The procedure in [5] presents a 3D-specific force-directed placer that incorporates thermal objectives directly into the placer. Instead of the finite difference method that is used in many floorplanners, this approach employs FEA, which discretizes the design space into regions known as elements. Another method in [12] maps an existing 2D placement to a 3D placement through transformations based on dividing the layout into  $2^k$  regions, and defining local transformations for heuristic layout refinement.

The approach in [13] observes that since 3D layouts have very limited flexibility in the third dimension (with a small number of layers and a fixed set of discrete locations), partitioning works better than a force-directed method. Accordingly, this work performs global placement using recursive bisectioning. Thermal effects are incorporated through *thermal resistance reduction nets*, which are

attractive forces that induce high power nets to remain close to the heat sink. The global placement step is followed by coarse legalization, in which a novel cell-shifting approach is proposed. This generalizes the methods in FastPlace [14] by allowing shift moves to adjust the boundaries of both sparsely and densely populated cells using a computationally simple method. Finally, detailed legalization generates a final nonoverlapping layout. The approach is shown to provide excellent tradeoffs between parameters such as the number of interlayer vias, wire length, temperature.

### E. Thermal via insertion

While silicon is a reasonably good thermal conductor, with about half the conductivity of typical metals, many of the materials used in 3D technologies are strong insulators that place severe restrictions on the amount of heat that can be removed, even under the best placement solution. The materials include epoxy bonding materials used to attach 3D tiers, or field oxide, or the insulator in an SOI technology. Therefore, the use of deliberate metal lines that serve as heat removing channels, called “thermal vias,” are an important ingredient of the total thermal solution. The thermal via insertion problem determines the optimal positions of thermal vias to provides an overall improvement in the temperature distribution. In realistic 3D technologies, the footprints of these inter-tier vias, including keep-out areas, are of the order of  $5\mu\text{m}\times 5\mu\text{m}$ .

In principle, the problem of placing thermal vias can be viewed as one of determining one of two conductivities (corresponding to the presence or absence of metal) at every candidate point where a thermal via may be placed in the chip. However, in practice, it is easy to see that such an approach could lead to an extremely large search space that is exponential in the number of possible positions; note that the set of possible positions in itself is extremely large.

To avoid unpredictable routing blockages due to thermal vias, it is necessary to enforce discipline within the design, designating a specific set of areas within the chip as potential thermal via sites. These could be chosen as specific inter-row regions in the cell-based layout, and the optimizer would determine the density with which these are filled with thermal vias.

Various published methods take different approaches to thermal via insertion. We will now describe an algorithm to post-facto thermal via insertion [10]; other procedures perform thermal via insertion during floorplanning, placement or routing will be discussed in the appropriate sections. This method takes a placed 3D circuit, and iteratively proceeds by incrementally modifying the thermal conductivities of specific FEA elements (thermal via regions).

A key observation in this work is that the insertion of thermal vias is most useful in areas with a high *thermal gradient*, rather than areas with a high temperature. Effectively, the thermal via acts as a pipe that allows the heat to be conducted from the higher temperature region to the lower temperature region; this, in turn, leads to temperature reductions in areas of high temperature. This is illustrated in Figure 3, which shows the 3D layout of the benchmark `struct`, before and after the addition of thermal vias, respectively. The hottest region is the center of the uppermost tier, and a major reason for its elevated temperature is because the tier below it is hot. Adding thermal vias to remove heat from the second tier, therefore, effectively also significantly reduces the temperature of the top tier. For this reason, the regions where the insertion of thermal vias is most effective are those that have high thermal gradients.

Therefore the method in [10] employs an iterative update formula

$$K_i^{new} = K_i^{old} \left( \frac{|g_i^{old}|}{g_{i,ideal}} \right) \quad i = x, y, z \quad (6)$$

is employed, where  $K_i^{new}$  and  $K_i^{old}$  are, respectively, the old and new thermal conductivities in each direction, before and after each iteration,  $g_i^{old}$  is the old thermal gradient, and  $g_{i,ideal}$  is a heuristically selected ideal thermal gradient.

Each iteration begins with a distribution of the thermal vias; this distribution is corrected using the above update formula, and




Fig. 3. Thermal profile of `struct` before (left) and after (right) thermal via insertion. The top four layers of the figure at right correspond to the four layers in the figure at left.

the  $K_i^{new}$  value is then translated to a thermal via density, and then a precise layout of thermal vias, using precharacterization. The iterations end when the desired temperature profile is achieved. This essential iterative idea has also been used in other methods for thermal via insertion steps that are integrated within floorplanning, placement and routing, as described in succeeding sections. This general framework has been used in several other published techniques that insert thermal vias either concurrently during another optimization, or as an independent step.

#### F. Thermally-driven routing

During routing, several objectives and constraints must be taken into consideration, including avoiding blockages due to areas occupied by thermal vias, incorporating the effect of temperature on the delays of the routed wires, and of course, traditional objectives such as wire length, timing, congestion and routing completion.

Once the cells have been placed and an initial set of locations of the thermal vias determined, the routing stage finds the optimal interconnections between the wires, and inserts further thermal vias. As in 2D routing, it is important to optimize the wire length, the delay, and the congestion. In addition, several 3D-specific issues come into play. Firstly, the delay of a wire increases with its temperature, so that more critical wires should avoid the hottest regions, as far as possible. Secondly, inter-tier vias are a valuable resource that must be optimally allocated among the nets. Thirdly, congestion management and blockage avoidance is more complex with the addition of a third dimension. For instance, a signal via or thermal via that spans two or more tiers constitutes a blockage that wires must navigate around.

A router typically grids tile layout into rectangular tiles, each with a horizontal and vertical capacity that determines the number of wires that can traverse the tile, and an inter-tier via capacity that determines the number of free vias available in that tile. For a single net, the degrees of freedom that are available are in choosing the locations of the inter-tier vias, and selecting the precise routes within each tier. The locations of inter-tier vias will depend on the resource contention for vias within each grid. Moreover, critical wires should avoid the high-temperature tiles, as far as possible.

The work in [15] presents a thermally conscious router, using a multilevel routing paradigm with integrated intertier via planning and incorporating thermal considerations. An initial routing solution is constructed by building a 3D minimum spanning tree (MST) for each multipin net, and using maze routing to avoid obstacles.

Another approach to 3D routing, presented in [16], combines the problem of 3D routing with heat removal by inserting thermal vias in the  $z$  direction, and introduces the concept of *thermal wires*. Like a thermal via, a thermal wire is a dummy object: it has no electrical function, but is used to spread heat in the lateral direction.

The global routing scheme goes through two phases. In Phase I, an initial routing solution is constructed. A 3D MST is built for each multipin net, and based on the corresponding two-pin decomposition, the routing congestion is statistically estimated over each lateral routing edge. A recursive bipartitioning scheme is then used to assign

intertier vias. Signal intertier via assignment is then performed across the cut in each recursive bipartition. At each level of the hierarchy, the problem of signal intertier via assignment is formulated as a min-cost network flow. Once the intertier vias are fixed, the problem reduces to a 2D routing problem in each tier, and maze routing is used to route the design.

Next, in Phase II, a linear programming approach is used to assign thermal vias and thermal wires. A thermal analysis is performed, and fast sensitivity analysis using the adjoint network method. The benefit of adding thermal vias, for relatively small perturbations in the via density, is given by a product of the sensitivity and the via density, a linear function. The objective function is a sum of via densities and is also linear. Additional constraints are included in the formulation to permit overflows, and a sequence of linear programs is solved to arrive at the solution.

### III. POWER DELIVERY IN 3D ICs

Despite the recent surge in 3D IC research, there has been little work from the circuit design and automation community on power delivery issues for 3D ICs. On-chip power supply noise has worsened in modern systems because scaling of the Power Supply Network (PSN) impedance has not kept up with the increase in device density and operating current due to the limited wire resources and constant RC per wire length, and this situation is worsened in 3D ICs. The increased IR and Ldi/dt supply noise in 3D chips may cause a larger variation in operating speed leading to more timing violations. The supply noise overshoot due to inductive parasitics may aggravate reliability issues such as oxide breakdown, Hot Carrier Injection (HCI) and Negative Bias Temperature Instability (NBTI) (which are also negatively affected by elevated temperatures). Consequently, on-chip power delivery will be a critical challenge for 3D ICs.

#### A. The basics of power delivery

According to scaling roadmaps, future high performance ICs will need multiple, sub-1V supply voltages, with total currents exceeding  $100 A/cm^2$  even for 2D chips [17]. Conventional power delivery methods for high performance ICs employ a DC-DC converter, known as a voltage regulator module (VRM). The VRM is typically mounted on the motherboard, with external interconnects providing the power to the chip. The package parasitics, contributed by the I/O pads and bonding wires are modeled as an inductance and resistance in series. The intrachip power delivery network is an R(L)C network, including decoupling capacitors (decaps) that are intended to damp out transient noise.

The chip acts as a distributed noise source drawing current in different locations and at different frequencies, causing imperfections in the delivered supply. The supply that reaches the chip is affected by IR and Ldi/dt drop across the package constituting the supply noise: the package impedance has largely remained unaffected by technology scaling. Scaling does, however, result in some unwanted effects on-chip, namely, in increased currents and faster transients from one technology node to the next. The former aggravate the IR drop, while the latter worsen the Ldi/dt drop. Over and above these effects is the issue of global resonant noise, in which the supply impedance gets excited to produce large droops on supply at or near the resonant frequency. With these increased levels of noise and reduced noise margins, as Vdd levels scale down, reliable power delivery has become a major challenge.

The presence of severe power delivery bottlenecks necessitates a look at entirely novel power delivery schemes for 3D chips. In this section, we introduce two possible approaches for this purpose.

#### B. On-chip voltage regulation

One way of dealing with the power delivery problem in 3D ICs (and also in conventional 2D ICs) is to bring the DC-DC converter

module closer to the processor. A boosted external voltage and its local down-conversion, ensures that the current through external package,  $I_{ext}$ , is small, and relaxes the scaling requirement on external package impedance. Moreover, this point of load regulation isolates the load from global resonant noise from external package and decap.

Traditionally, the efficiency of monolithic DC-DC converters has been limited by the small physical inductors allowed on-chip. Typical off-chip DC-DC conversion requires high-Q inductors of the order of  $1 - 100\mu H$  [18], which are difficult to implement on-chip due to their area requirements. With growing power delivery problems, the focus has been on building compact inductors through technologies like thin film inductors [19], or on more efficient, but costly, DC-DC converters through multiphase/interleaving topologies. Clearly, there is an onus to incorporate these on-chip, which calls for a different process altogether. The possibility to stack different wafers with heterogeneous technologies, as offered by three dimensional wafer-level stacking in 3D ICs thus is the natural solution for realizing on-chip switching converters.

### C. Multistory power delivery

A promising technique for achieving high efficiency on chip DC-DC conversion and supply noise reduction is the multistory power delivery (MSPD) scheme [20], [21]. It has been demonstrated in [22] that the idea becomes particularly attractive for 3D IC structures involving stacked processors and memories.

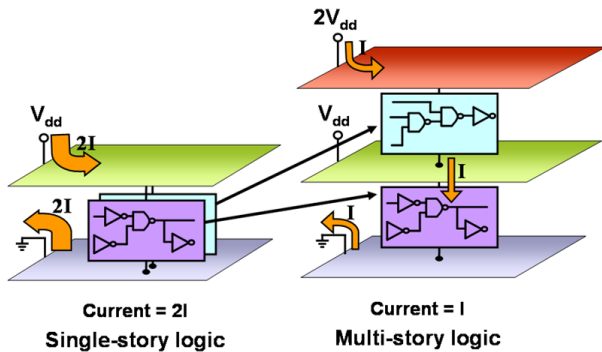


Fig. 4. Conventional and multistory power delivery schemes.

Figure 4 demonstrates the basic concept of MSPD. A schematic of a conventional supply network is shown in Figure 4(a), where all circuits draw current from a single power source. Figure 4(b) shows the multistory supply network, with subcircuits operating between two supply stories. The concept of a “story” is merely an abstraction to illustrate the nature of the power delivery scheme, as opposed to the 3D IC architecture, where circuits are physically stacked in tiers. In this scheme, current consumed in the “ $2V_{dd}-V_{dd}$  story” is subsequently recycled in the “ $V_{dd}-Gnd$  story.” Due to this internal recycling, half as much current is drawn compared to the conventional scheme, with almost the same total power consumption. A reduced current is beneficial since it cuts down the supply noise. This results in minimal noise on that rail.

The main issue with this technique is requirement of separate body islands. This may be difficult in typical bulk processes. However, if we consider 3D ICs, the tiers are inherently separated electrically, which makes MSPD particularly attractive.

The MSPD idea described in section C can be automated by solving an optimization problem [23], which formulates the two-story problem as one of module assignment between the two stories. An important consideration in the design of an MSPD circuit is to locally maintain the current balance between logic blocks operating in different Vdd domains because otherwise, the imbalance current will flow through voltage regulators and be wasted. Another important

issue that has to be considered is the design stage at which the circuit should be partitioned into different Vdd domains. Note that a level shifter is required at the output of a logic block if it is used to drive another logic block operating in a different Vdd domain. Level-shifters occupy silicon area and cause extra delays in the circuit.

The module assignment problem is addressed at the floorplanning level where the number of modules is usually not very large, and their area is largely ignored. It is assumed that  $K$  voltage regulators are distributed across the chip: these regulators are well-designed and output stable voltage levels at  $V_{dd}$ .

Each regulator is represented by the point it taps into the  $V_{dd}$  grid. As shown in Figure 5(a), the chip is divided into  $K$  regions accordingly such that there is one regulator in each region and the  $i^{th}$  region contains all the points on chip that primarily draw (sink) currents from (to) the  $i^{th}$  regulator. The division of the chip into these nonoverlapping regions can be achieved by meshing the die area using a fine grid, and determining which grid cell each block belongs to, e.g., each cell can be said to belong to the region controlled by the nearest voltage regulator.




Fig. 5. Graph construction (a) partitioning the chip into disjoint regions each of which is controlled by a voltage regulator, and (b) constructing a graph where node  $V_i$  corresponds to module  $M_i$ .

Once the chip is partitioned into disjoint regions, it is assumed that any “imbalanced” current, i.e., current that is not recycled to the next story in a particular region, goes through the regulator in the same region and is wasted. If a module is located at the boundary between multiple regions, it will be decomposed into several submodules with one submodule in each region it overlaps and with the constraint that all submodules must be assigned to the same Vdd domain.

Let us focus on a particular region corresponding to a particular voltage regulator. Assume the modules located in this region are  $M_1, M_2, \dots, M_n$ , where the current flowing through module  $M_i$  as a function of time  $t$  is given by  $I_i(t)$ . Because voltage regulators can only respond to the low to mid frequency components of the imbalance currents while the high frequency components are usually handled by on-chip decaps, we preprocess the input current traces obtained through cycle-accurate power simulations to smooth out the high frequency components in the current signals. Therefore,  $I_i(t)$  should be understood as containing only the low to mid frequency components of the current flowing through module  $M_i$ .

If we associate a 0/1 integer variable  $x_i$  with module  $M_i$ ,

$$x_i = \begin{cases} 0 & \text{if } M_i \text{ operates between the } 2V_{dd} \text{ and } V_{dd} \text{ rails} \\ 1 & \text{if } M_i \text{ operates between the } V_{dd} \text{ and } GND \text{ rails} \end{cases} \quad (7)$$

then the total current flowing through the voltage regulator at time  $t$  will be approximated by

$$I_R(t) = \left| \sum_{i=1}^n I_i(t) \cdot (1 - x_i) - \sum_{i=1}^n I_i(t) \cdot x_i \right| = \left| \sum_{i=1}^n I_i(t) \cdot (1 - 2x_i) \right| \quad (8)$$

This problem can be shown to map on to one of graph partitioning to maximize the cut between the partitions, where the weights of the edges are given by

$$w(V_i, V_j) = \left[ \sum_{k=1}^K \frac{S_{ik} S_{jk}}{S_i S_j} \right] \overline{I_i(t) I_j(t)} \quad (9)$$

Here,  $S_i$  represents the area of the  $i^{th}$  module, and the overlap area between the  $i^{th}$  modules and the  $k^{th}$  region is denoted by  $S_{ik}$ .

The intuition behind Equation (9) is that for any pair of modules, only the portions that are located in the same region count toward the calculation of the correlation between them. If modules  $M_i$  and  $M_j$  are completely separated into two disjoint regions, the weight  $w(V_i, V_j)$  will be zero, and therefore, the corresponding edge can be removed from the graph. An example of graph construction is shown in Figure 5(b). A Fiducia-Mattheyses-like approach is used to speedily find the optimal partition that maximizes the cut.

Experimental results demonstrate that the method is effective in building partitions for multistory power grids in both 2D and 3D chip under an SOI-based process, where blocks from multiple stories may coexist on the same tier. It is shown that the partitioning-based method is successful in recycling a large amount of power through the system, and the quality of results of the partitioning-based method compare favorably with an annealing approach.

#### D. Optimization of 3D power grids

Several techniques are available to increase the reliability of power grids and control power grid noise, such as wire widening, grid topology optimization, and decoupling capacitor (decap) insertion. Of all these techniques, the decaps are arguably the most powerful method for reducing transient noise. Decaps serve as local current reservoirs, and can be used to satisfy sudden surges in current demand by the functional blocks/cells, while keeping supply voltage levels relatively stable. Active/passive damping methods for resonant noise using decaps have also been proposed [24], [25].

Conventional technologies for implementing decaps are based on SiO<sub>2</sub>-based structures that are widely used in robust power delivery network design. 3D power grid optimization has been studied in, for example, [26], [27]. Unlike the 2D case, new considerations come into play while optimizing a 3D power grid using CMOS decaps, specifically related to congestion and leakage issues.

The approach in [28] presents an approach for decap allocation in 3D power grids, using both conventional CMOS decaps and metal-insulator-metal (MIM) decaps. MIM capacitors are fabricated between metal layers, and have high capacitance density and low leakage current density. However, they cannot be used unconditionally to replace CMOS decaps, since their use incurs a cost: they present routing blockages to nets that attempt to cross them. In [28], the decap budgeting problem, using both CMOS and MIM decaps, is formulated as a Linear Programming (LP) problem, and an efficient congestion-aware algorithm is proposed to optimize the power supply noise. An iterative flow is used to solve the decap allocation problem, based on a sequence of linear programs formulation. Experimental results demonstrate that the use of CMOS decaps alone is insufficient to overcome the violations; the use of MIM decaps results in high levels of congestion; and the optimal mix of the two meets both congestion and noise constraints, with low leakage.

#### IV. ACKNOWLEDGMENTS

The author would like to acknowledge the work of Brent Goplen, Chris Kim, Yong Zhan, Tianpei Zhang, and Pingqiang Zhou, which has contributed to the content of this article.

#### REFERENCES

- [1] M. N. Özışık, *Heat Transfer: A Basic Approach*. New York, NY: McGraw-Hill, 1985.
- [2] H. Qian, S. R. Nassif, and S. S. Sapatnekar, "Random walks in a supply network," in *Proceedings of the ACM/IEEE Design Automation Conference*, 2003, pp. 93–98.
- [3] P. Li, L. T. Pileggi, M. Asheghi, and R. Chandra, "IC thermal simulation and modeling via efficient multigrid-based approaches," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 25, no. 9, pp. 1763–1776, Sep. 2006.
- [4] B. Goplen, "Advanced placement techniques for future VLSI circuits," Ph.D. dissertation, University of Minnesota, Minneapolis, MN, 2006.
- [5] B. Goplen and S. S. Sapatnekar, "Efficient thermal placement of standard cells in 3D ICs using a force directed approach," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, 2003, pp. 86–89.
- [6] C. Ababei, Y. Feng, B. Goplen, H. Mogal, T. Zhang, K. Bazargan, and S. Sapatnekar, "Placement and routing in 3D integrated circuits," *IEEE Design & Test*, vol. 22, no. 6, pp. 520–531, November–December 2005.
- [7] J. Cong, J. Wei, and Y. Zhang, "A thermal-driven floorplanning algorithm for 3D ICs," in *Proceedings of the International Symposium on Physical Design*, 2004, pp. 306–313.
- [8] E. Wong and S. K. Lim, "3D floorplanning with thermal vias," in *Proceedings of the Design, Automation & Test in Europe*, 2006, pp. 878–883.
- [9] H. Qian and S. S. Sapatnekar, "A hybrid linear equation solver and its application in quadratic placement," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, 2005, pp. 905–909.
- [10] B. Goplen and S. S. Sapatnekar, "Thermal via placement in 3D ICs," in *Proceedings of the International Symposium on Physical Design*, 2005, pp. 167–174.
- [11] P. Zhou, Y. Ma, Z. Li, R. P. Dick, L. Shang, H. Zhou, X. Hong, and Q. Zhou, "3D-STAF: Scalable temperature and leakage aware floorplanning for three-dimensional integrated circuits," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, 2007, pp. 590–597.
- [12] J. Cong, G. Luo, J. Wei, and Y. Zhang, "Thermal-aware 3D IC placement via transformation," in *Proceedings of the Asia-South Pacific Design Automation Conference*, 2007, pp. 780–785.
- [13] B. Goplen and S. S. Sapatnekar, "Placement of 3D ICs with thermal and interlayer via considerations," in *Proceedings of the ACM/IEEE Design Automation Conference*, 2007, pp. 626–631.
- [14] N. Viswanathan and C. C.-N. Chu, "Fastplace: Efficient analytical placement using cell shifting, iterative local refinement and a hybrid net model," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 24, no. 5, pp. 722–733, May 2005.
- [15] J. Cong and Y. Zhang, "Thermal-driven multilevel routing for 3-D ICs," in *Proceedings of the Asia-South Pacific Design Automation Conference*, 2005, pp. 121–126.
- [16] T. Zhang, Y. Zhan, and S. S. Sapatnekar, "Temperature-aware routing in 3D ICs," in *Proceedings of the Asia-South Pacific Design Automation Conference*, 2006, pp. 309–314.
- [17] Semiconductor Industry Association, "International Technology Roadmap for Semiconductors (ITRS)," 2007, <http://public.itrs.net/>.
- [18] P. Hazucha, G. Schrom, J. Hahn, B. A. Bloechel, P. Hack, G. E. Dermer, S. Narendra, D. Gardner, T. Karnik, V. De, and S. Borkar, "A 233-MHz 80%–87% Efficient Four-Phase DC-DC Converter Utilizing Air-Core Inductors on Package," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 4, pp. 838–845, Apr. 2005.
- [19] K. H. Kim, J. Kim, H. J. Kim, S. H. Han, and H. J. Kim, "A megahertz switching DC/DC converter using FeBN thin film inductor," *IEEE Transactions on Magnetics*, vol. 38, no. 5, pp. 3162–3164, Sep. 2002.
- [20] J. Gu and C. H. Kim, "Multi-story power delivery for supply noise reduction and low voltage operation," in *International Symposium on Low Power Electronics and Design*, 2005, pp. 192–197.
- [21] S. Rajapandian, K. Shepard, P. Hazucha, and T. Karnik, "High-tension power delivery: Operating 0.18  $\mu\text{m}$  CMOS digital logic at 5.4V," in *Proceedings of the IEEE International Solid-State Circuits Conference*, 2005, pp. 298–599.
- [22] P. Jain, T. Kim, J. Keane, and C. H. Kim, "A multi-story power delivery technique for 3D integrated circuits," in *Proceedings of the ACM International Symposium on Low Power Electronics and Design*, 2008, pp. 57–62.
- [23] Y. Zhan and S. S. Sapatnekar, "High efficiency Green function-based thermal simulation algorithms," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 26, no. 9, pp. 1661–1675, September 2007.
- [24] J. Gu, H. Eom, and C. H. Kim, "A switched decoupling capacitor circuit for on-chip supply resonance damping," in *IEEE Symposium on VLSI Circuits*, 2007, pp. 126–127.
- [25] J. Xu, P. Hazucha, M. Huang, P. Aseron, F. Paillet, G. Schrom, J. Tschanz, and C. Zhao, "On-die supply-resonance suppression using band-limited active damping," in *Proceedings of the IEEE International Solid-State Circuits Conference*, 2007, pp. 286–603.
- [26] E. Wong, J. Minz, and S. K. Lim, "Decoupling capacitor planning and sizing for noise and leakage reduction," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, 2006, pp. 395–400.
- [27] G. Huang, M. Bakir, A. Naeemi, H. Chen, and J. D. Meindl, "Power delivery for 3D chip stacks: Physical modeling and design implication," in *Proceedings of Electrical Performance of Electronic Packaging*, 2007, pp. 205–208.
- [28] P. Zhou and S. S. Sapatnekar, "Congestion-aware power grid optimization for 3D circuits using MIM and CMOS decoupling capacitors," in *Proceedings of the Asia-South Pacific Design Automation Conference*, 2009.