2012 25th International Conference on VLSI Design

A Power Delivery Network Aware Framework for Synthesis of 3D Networks-on-Chip with Multiple Voltage Islands

Nishit Kapadia, Sudeep Pasricha
Department of Electrical and Computer Engineering
Colorado State University, Fort Collins, CO, U.S.A.
nkapadia@colostate.edu, sudeep@colostate.edu

Abstract - *IR drops in a Power Delivery Network (PDN) on chip multi-processors (CMPs) can worsen the quality of voltage supply and thereby affect overall performance. This problem is more severe in 3D CMPs with network-on-chip (NoC) fabrics where the current in the PDN increases proportionally to the number of device layers. Even though the PDN and NoC design goals are non-overlapping, both the optimizations are interdependent; for instance, each new core mapping on the 3D die will change traffic patterns and have a unique distribution of IR-drops in the PDN. Unfortunately, designers today seldom consider design of PDN while synthesizing NoCs. If NoC synthesis is carried out without considering the associated PDN design cost, it can easily result in an overall sub-optimal design. In this work, for the first time, we propose a novel PDN-aware 3D NoC synthesis framework that minimizes NoC power while meeting performance goals; and optimizes the corresponding PDN for total number of Voltage Regulator Modules (VRMs), current efficiency, and grid-wire width while satisfying IR-drop constraints. Our experimental results show that the proposed methodology provides more comprehensive results compared to a traditional approach where the NoC synthesis step does not consider the PDN costs.*

1. Introduction

Designing a robust Power Delivery Network (PDN) is critical to the overall performance of today's CMPs. The PDN is required to deliver a stable power supply across the chip, which is within a desired voltage range; and tolerate large variations in load currents [1]. Multiple voltage islands (*VIs*) are generally used in modern CMPs to minimize the total power dissipation while meeting performance constraints. The PDN is required to supply power at different voltage levels corresponding to the *VIs* while keeping power loss to a minimum. With increasing device density and supply voltage levels, the supply currents have risen; however the scaling of PDN impedance has not kept up with this trend [2]. IR drops can worsen the quality of voltage supply and thereby affect the ultimate performance of the CMP. This problem is more severe in 3D CMPs as the current in the PDN could be as many times more as the number of device layers compared to a 2D CMP. Besides, the number of I/O pins on an *n*-layered 3D CMP is about *n* times smaller than its 2D counter-part, thus exacerbating the problem of a degraded voltage supply in 3D designs [3].

Another critical component at the heart of emerging 3D CMPs is the network-on-chip (NoC) architecture that enables tens to hundreds of cores to communicate with each other at the intra- and inter-layer levels. As the power dissipated in the NoC has become a significant portion of the total on-chip power, optimizing the communication power in addition to computation power is critical [6]. Several recent works have proposed techniques to synthesize regular and custom 3D NoC topologies [9][28]-[32] to optimize communication power. However, these works do not consider the design of the PDN while mapping cores and designing the NoC fabric, and typically generate a single power-optimized configuration. Performing synthesis of the PDN for the best generated configuration in these cases puts stringent demands on the already strained PDN. This can either make it extremely difficult to meet the PDN constraints such as maximum grid-width, maximum number of voltage regulator modules (VRMs), and minimum current efficiency; or lead to over-margining for the PDN, which can be wasteful. Thus, the traditional approach of synthesizing a NoC fabric without considering the PDN ends up severely constraining the PDN design space, often leading to sub-optimal or even completely infeasible designs.

In this work, for the first time, we propose an automated framework for PDN-aware synthesis of mesh-based NoC fabrics in 3D CMPs that optimizes communication power while meeting application performance constraints. We recognize the key insight that different instances of voltage partitioning and core-to-tile mapping (different configurations of the NoC synthesis process) can *significantly* alter the power/voltage distribution map seen by the PDN. Accordingly, our framework considers the interdependence between a synthesized NoC configuration and its corresponding VRM placement and power efficiency in the PDN. The novel contributions of our synthesis framework are as follows:

- We employ a novel branch and bound procedure that combines directed search and random search to produce multiple mapping solutions satisfying *VI* constraints while optimizing NoC power;
- We develop a linear programming formulation as well as a fast heuristic to synthesize a PDN comprised of a segmented power grid for cores running at multiple voltages; with a topological structure of VRMs that considers physical placement of VRMs on the 3D mesh to optimize current efficiency and VRM count;
- We generate a set of interesting design points (Pareto mappings) that allow a designer to weigh the PDN design cost against NoC design cost, and select a suitable solution that meets power, performance, and PDN design goals.

2. Related Work

Many researchers [7]-[11] have proposed custom topology synthesis techniques for NoC fabrics that improve overall performance at the cost of sacrificing the regularity of mesh-based structures. Although these custom architectures are expected to achieve better latency and area utilization, their design process is more complex and faces several challenges, such as greater crosstalk and uncertainty in link delays due to irregular interconnect structures. Thus, a conservative enough custom design may actually offset the advantages of better performance [12]; especially for medium to large sized (in terms of total number of cores) NoC architectures. The problem of NoC synthesis on regular structures with multiple supply *VIs* has been addressed in several works [6][13]-[19]. Given the promise of 3D technologies, 3D NoC synthesis in recent years has also attracted significant research efforts [9][28]-[32]. These works have proposed techniques to optimize the 3D NoC designs for power, temperature, and performance. However, none of the above approaches have considered the effects of NoC synthesis on the efficiency and overheads associated with the PDN design; in other words, these approaches are not PDN-aware.

Techniques for optimizing PDNs in 3D ICs have been studied in a few recent works [1]-[3][20][21]. Amelifard et al. [1] use dynamic

1063-9667/12 $26.00 © 2012 IEEE
DOI 10.1109/VLSID.2012.81

programming to generate a multi-level tree topology of suitable Voltage Regulator Modules (VRMs) to improve the power efficiency in the PDN. Jain et al. [2] propose a multi-story power delivery technique which improves upon IR noise in the PDN by recycling current between different power supply domains. Falkenstern et al. [20] use simulated annealing to co-synthesize the floorplan and P/G network, optimizing wirelength, area, P/G routing area, and IR-drops. Chen et al. [21] propose an integrated 3D architecture of stacked-TSV, thermal and power distributed network (STDN); and use a simulated annealing floorplanner to minimize voltage drop, temperature, and other factors in STDN. None of these PDN optimization techniques considers the system level impact of the 3D NoC fabric and core mapping across the layers.

In this paper, we present novel techniques for PDN design as well as NoC synthesis; for mesh-based 3D CMPs. To the authors' knowledge, this is the first work which proposes a physically aware 3D NoC synthesis framework that also integrates PDN optimization to produce a more efficient overall CMP design.

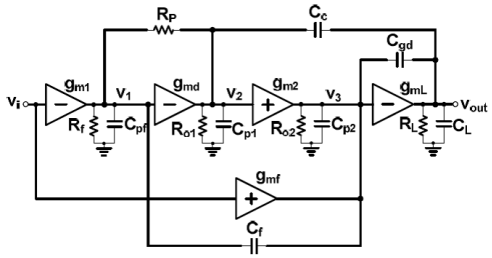

Figure 1: Schematic of an LDO-VRM [23]

3. PDN Design with Multiple Supply Voltages

High circuit density in smaller footprint 3D ICs presents a unique challenge for designers of PDNs, requiring the network to deliver significantly more current than in 2D ICs with fewer P/G bumps, while also circumventing increasingly daunting IR-drop issues. Voltage regulator modules (VRMs) are key components of any PDN, responsible for stepping down the high voltage of the power source. To cope with supply voltage variations in emerging 2D and 3D ICs, traditional off-chip voltage regulators require large decoupling (or bypass) capacitors and inductors that end up occupying excessive PCB-area. Moreover, the parasitic inductance and resistance between the regulator and the processor hinders the regulator from reacting quickly to load transients [4]. Bringing the voltage regulators on-chip (closer to the load) is one solution to the problem that would result in smaller decaps and inductors needed, as the parasitic elements fall. Additionally, an on-chip regulator can react quickly to the load transients, save on-board space, as well as reduce the number of external P/G pins needed. Low Drop-Out regulators (LDOs; Fig. 1) are particularly amenable for on-chip integration due to their small area overhead and low dropout voltage, as opposed to switching-type regulators which employ on-chip inductors that occupy valuable area [5]. The characteristics of LDO-VRMs assumed in our PDN design framework follow designs proposed in [23][24].

A PDN traditionally uses a single continuous power grid made of orthogonal interconnects (on the top wiring levels) running across the chip at the electrical potential of the external pin voltage. In systems with multiple voltage levels, VRMs can be inserted appropriately to step down from the single external voltage level to the different operating voltage levels of the cores/modules/VIs. As all voltages are stepped-down from a high external voltage level, the power conversion efficiency in this approach could be poor. Alternatively, as many parallel power grids as the number of supply voltages can be implemented [22]. But even with reasonable number of voltage levels, this approach could result in a prohibitively high PDN routing area overhead. Ultimately, the chosen design approach must cope with the IR-drop problem which is worse in 3D ICs by as much as 3.4× compared to 2D ICs [3], and becomes more severe as we move farther away (on the power grid) from the power source. This problem can either be rectified by inserting additional VRMs, which have an associated area overhead or by increasing the grid-wire width, which increases the PDN routing area on the chip.

In this work, we propose a PDN design structure with a segmented grid configuration that allows for more power efficient stepping down of the voltages derived from VRM outputs, while at the same time considering the overheads of VRM insertions and grid-wire width. To the best of our knowledge, the problem of PDN design that includes determining the locations of VRMs for 3D CMPs operating at multiple voltage levels has not been addressed before.

4. Problem Formulation

We are given the following inputs to our problem:

- A regular 3D mesh-based NoC with dimensions (dim_x, dim_y, dim_z) with the number of tiles $T = dim_x * dim_y * dim_z$ and each tile consisting of a compute core and a NoC router;
- A core graph $G(V,E)$; with a set of T vertices $\{V_1, V_2, ..., V_T\}$ representing homogenous cores on which tasks have already been mapped, and the set of M edges $\{e_1, e_2, ..., e_M\}$ that represent communication dependencies between cores;
- A set of triplets constituting operating voltages, operating frequencies and maximum supply currents for the T cores $\{(v_1, f_1, i_1), (v_2, f_2, i_2), (v_3, f_3, i_3), \ldots, (v_T, f_T, i_T)\}$;
- An external voltage supply EV to the PDN and a 3D segmented power grid;
- A set of r possible grid-wire resistance (PDN branch resistance) values: $\acute{R} = \{R_1, R_2, .., R_r\}$.

Given the above inputs, our goal is to obtain a core to die mapping and synthesize a regular 3D mesh NoC for a specific application, such that all application performance requirements as well as PDN IR-drop constraints are satisfied; while minimizing the total communication power in the NoC components (routers, links, voltage level converters or VLCs, mixed clock FIFOS or MCFIFOs), the external current (EI) drawn by the PDN, and the number of VRMs.

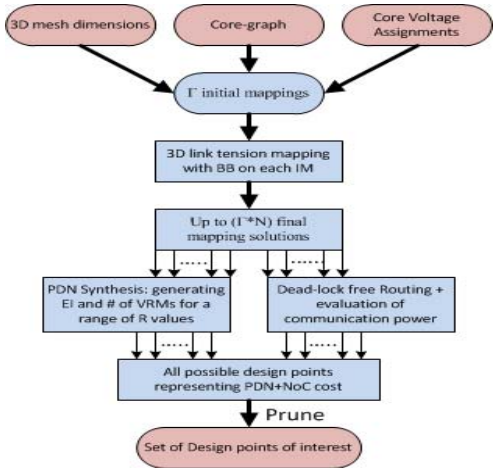

Figure 2: Flow of the PDN-aware NoC Synthesis Framework

5. PDN-aware NoC Synthesis Framework

The flow of our PDN-aware NoC synthesis framework is shown in Fig. 2. As a first step, Γ initial mappings are generated from the communication dependencies defined by the core-graph and the voltage assignments, where Γ is the total number of allowed voltage levels. A link tension based Branch & Bound (BB) procedure is then run on each Initial Mapping (IM) to generate multiple final mapping

candidates. On each of these candidates, deadlock free YXZ routing is employed to compute the total power of all the NoC components: routers, links, as well as MCFIFOs and VLCs (which are needed when crossing *VIs*). Next, the PDN synthesis step generates a PDN design and evaluates corresponding PDN costs for a candidate, given the set $\acute{R}$ of grid-wire resistance values. The communication power represents the cost of the NoC whereas, the PDN cost includes the total number of VRMs needed, grid-wire width, and the external current (*EI*) drawn. This cost is computed for all the final mapping candidates to generate a set of final solution points. The points which have both the PDN costs and the NoC cost greater than some other point are pruned to finally produce a set of final design points, each optimized for the PDN and NoC design objectives by varying degrees. In the following subsections, we describe the steps in detail.

5.1 Initial Mapping (IM)

In this first step, we generate an initial core-to-tile mapping by traversing an Inter-Island Communication Graph $IICG(V_{isl},E_{isl})$, where the vertices constitute the entire islands and the edges represent the aggregate communication bandwidth between the respective islands. A breadth-first search (BFS) starting with each of the Γ islands as the root node would produce Γ distinct sequences of islands, each of length Γ. The order of islands in each sequence is based on decreasing communication bandwidths with the island selected as the root node. Subsequently, the cores are mapped onto the tiles of the NoC in order of the island $sequence_j$ to generate IM_j.

We follow a pre-defined sequence of tile co-ordinates to generate the initial mapping as illustrated in Fig. 3 for a 64 core NoC. Notice in the figure that the mapping starts from the top-right corner of the topmost layer (layer 1) and ends at the top-left corner of the same layer. Such an ordering grows in x, y and z directions in a symmetrical way, thereby keeping the Manhattan distances between the currently placed core and recently placed cores short, at the same time, guaranteeing *VI* integrity (i.e., every core in a *VI* has at least one neighbor of the same voltage level as itself).

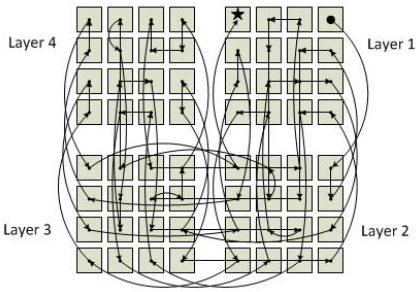

Figure 3: The order of placement of cores to generate an initial mapping (IM) for a 64-core 3D-mesh NoC

5.2 3D Link-Tension Mapping with BB

We propose a branch and bound (BB) technique that combines random search with directed search to generate multiple mapping candidates. Link-tension is the product of the communication bandwidth and post-mapping Manhattan distance for any edge in the core graph *G(V,E)*. The directed swaps are geared to reduce the highest tensions in the NoC in order to reduce communication power. The 'best swap' swaps the core under most tension (on the NoC link-tension map) in the direction of most tension. We combine the directed swaps with random swaps for effective exploration of the solution search space. Any swap which does not disintegrate *VIs* is considered valid; where we consider swaps between adjacent cores in horizontal and vertical planes, as well as horizontal-diagonal swaps.

Let *n* be the maximum branching degree and *N*, an upper bound on the total number of candidates which is a multiple of *n*; $\boldsymbol{\alpha}$ be a positive fraction which governs the weight of the random component in BB, and *C* be the number of current candidates. The pseudo code for the BB procedure is given below, which is run on each IM.

3D Link-tension Mapping with BB

input: core graph G(V, E) and an IM solution

1: **while** ((*C*<*N*) && (at least one non-leaf node exists in C));**do**∀non-leaf nodes on the current BB level {
2: Compute *B*: $B = (n+1) - \lceil (n-1) * (C+1)/N \rceil$
3: Compute *R* and *D*: $R = \lfloor \alpha * B/2 \rfloor$; $D = B - R$
4: Find out the D best swaps (directed search) and check their validity
5: If one or more valid swaps found, proceed to step 8
6: Find the best valid swap while considering all cores
7: If no swap is valid, mark this candidate (node) as a leaf; else execute swap (branch out child) and delete current node; then, goto next iteration
8: Compute the *R* random valid swaps
9: Execute computed random and directed swaps, branching out a new child for each swap and delete current node; then, goto next iteration }
10: ∀ non-leaf nodes, perform only the best swaps until equilibrium is attained on each candidate

output :Up to N final mapping candidates

At any level of a *B*-way search tree (*B* is variable representing current degree of branching) of intermediate mapping candidates, *D* best swaps and *R* random swaps are considered for each node. With *N* as an upper bound on the total number of final candidates, the branching degree proportionally decreases with increasing number of intermediate candidates. An intermediate candidate node becomes a leaf node (signifying a final mapping candidate) when no more directed swaps are possible on it and is never again considered for further swaps. When the existing number of candidates reach the upper bound of *N*, only the best swaps are made on all the non-leaf solutions (*B* is reduced to 1) until they converge to equilibrium (a state where valid swaps which reduce total NoC tension are no longer available). Alternatively, if no non-leaf solutions remain, BB terminates as no random swaps are allowed on leaf-nodes. Finally, a set of up to *N* final mapping candidates are obtained from a single initial mapping.

5.3 PDN Synthesis

PDN synthesis is performed on each of the mapping candidates produced by the BB procedure. We propose a segmented power distribution grid for a 3D-mesh CMP with *VIs*. We address the PDN design problem for the global grid, where each grid-node supplies to a core in the 3D mesh and VRMs are integrated to scale down voltage to cores in *VIs*. Besides overheads of chip area and power dissipation of the VRM components, proper placement of these VRMs on a 3D-mesh is critical for better supply efficiency as well as for minimizing the grid-wire width needed to satisfy the IR-drop constraints. The performance of any core is highly dependent on the quality of voltage supply; besides, we do not evaluate IR-drops in the power grid at the sub-core level (one and only one core is supplied to by a grid-node); thus, a tolerance of just 1% in the voltage supply level is assumed. In this work, as we investigate the steady state effects of the PDN, time-varying network characteristics such as transient noise are not considered. We also assume that all PDN braches have uniform resistance *R*, which is the norm for PDNs. Fig. 4 shows an example of a 3D-mesh CMP with *VIs* and corresponding PDN with VRMs. Note that the orange branches run at the external voltage supply *EV*. The PDN should be able to restrict the IR drops at each core within the set tolerance limit of the rated core voltage.

We propose a novel topological structure of VRM placements for better current efficiency; where the stepped-down voltages from outputs of VRMs are used as either of the following: *(i)* as voltage supplies to cores of the same voltage, such that the IR-drop constraints are satisfied (e.g., from Core {001} to Core {002} in Fig. 4); or *(ii)* as inputs to the VRMs for cores of lower voltages (to

further step-down the voltage level), such that the minimum drop-out voltage requirements of VRM are satisfied (e.g., from Core {011} to Core {012} in Fig. 4).

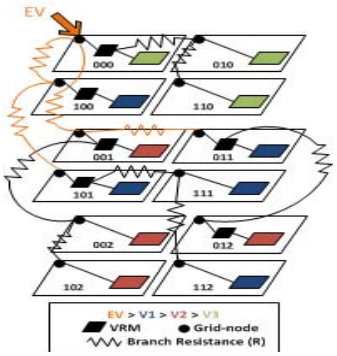

Figure 4: An Illustration of a Segmented PDN Structure

5.3.1 Linear Programming Formulation

We formulate the PDN synthesis problem as an exact Linear Programming problem with the following goal:

$$\textbf{\textit{Minimize}}: [\alpha.\sum B_{i,j,k} + \psi.(EI)]$$

where $\Sigma B_{i,j,k}$ is the total number of VRMs used and EI is the external current drawn from the power supply. We are given a set of T tile co-ordinates $T_{i,j,k}$, for $0 \le i \le dim_x\text{-}1$, $0 \le j \le dim_y\text{-}1$, $0 \le k \le dim_z\text{-}1$; on a 3D mesh with dimensions $\{dim_x, dim_y, dim_z\}$. The external voltage source (EV) is located at $T_{0,0,0}$ (upper left corner of the topmost layer in the 3D mesh). For a given core to tile mapping solution, $C_{i,j,k}$ and $CI_{i,j,k}$ are the operating voltage levels and the maximum current requirements of the cores at the respective co-ordinates. The design variables considered in our problem are as follows:

- VRM placements are represented with binary variables:
  - $B_{i,j,k} = 1$, if VRM is present at co-ordinates {i,j,k}
  - $B_{i,j,k} = 0$, if VRM is absent at co-ordinates {i,j,k}
- Branch currents emanating from the grid-nodes in d {x, y or z} direction: $I_{i,j,k\text{-}d}$ (Fig. 5(a))
- Branch currents emanating from VRMs in d {x, y or z} direction: $IV_{i,j,k\text{-}d}$ (Fig. 5(b))
- Grid-node voltages: $V_{i,j,k}$
- PDN branch resistance R, can take one of the r values from set $\acute{R} = \{R_1, R_2, .., R_r\}$

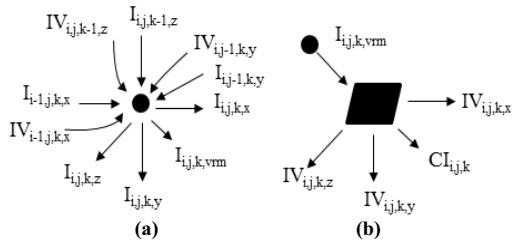

Figure 5: (a) Input/Output currents through a grid-node {i,j,k} (b) current across a VRM at co-ordinates {i,j,k}

Due to lack of space, we now present only the key program constraints.

Constraint 1: As voltage can only be stepped-down, voltage is derived from at least one neighboring grid-node with voltage no smaller than the current grid-node {i,j,k}. Thus,

$$[V_{i,j,k} \le V_{i-1,j,k}] \| [V_{i,j,k} \le V_{i,j-1,k}] \| [V_{i,j,k} \le V_{i,j,k-1}]$$
$$\forall (\{i,j,k\} - \{0,0,0\})$$

$$x_{i,j,k} = 1; \ if \ V_{i,j,k} \le V_{i-1,j,k}$$
$$= 0; \ otherwise$$

LP-Representation:

$$(V_{i-1,j,k} - V_{i,j,k}) - x_{i,j,k}.MAXVALUE < 0$$
$$(-V_{i-1,j,k} + V_{i,j,k}) - x'_{i,j,k}.MAXVALUE \le 0$$

where MAXVALUE is a large positive value, and $x'_{i,j,k}$ is the inverse of $x_{i,j,k}$: $x'_{i,j,k} + x_{i,j,k} = 1$. Constraints in y and z directions are defined similarly. Thus,

$$x_{i,j,k} + y_{i,j,k} + z_{i,j,k} \ge 1.$$

Constraint 2: VRMs need to be placed wherever the core voltage is less than the corresponding grid-node voltage

$$B_{i,j,k} = 1; \ if C_{i,j,k} < V_{i,j,k}$$
$$= 0; \ otherwise$$

$B'_{i,j,k}$ is the respective inverse: $B'_{i,j,k} + B_{i,j,k} = 1$

LP-Representation:

$$[V_{i,j,k} - C_{i,j,k}] - B_{i,j,k}.MAXVALUE \le 0$$
$$[C_{i,j,k} - V_{i,j,k}] - B'_{i,j,k}.MAXVALUE < 0$$

Constraint 3: The VRM minimum drop-out voltage constraint is defined as:

$$V_{i,j,k} - C_{i,j,k} \ge 0.2; \text{ if } B_{i,j,k} = 1$$

LP-Representation:

$$V_{i,j,k} - C_{i,j,k} + B'_{i,j,k}.MAXVALUE \ge 0.2$$

Constraint 4: The core voltages must be no greater than the respective grid-node voltages (with a tolerance of 1% of the rated core voltage):

$$(0.99) * C_{i,j,k} \le V_{i,j,k}$$

Constraint 5: Current across a VRM is shown in Fig. 5(b). Power efficiency of a VRM is defined by the following equation [1]:

$$\eta_p = \frac{I_{out}}{I_{in}} * \frac{V_{out}}{V_{in}}$$

As current efficiency is basically the power efficiency at constant voltage values; to linearize, we consider current efficiency as:

$$I_{out} = \eta_c * I_{in} = IV_{i,j,k,x} + IV_{i,j,k,y} + IV_{i,j,k,z} + CI_{i,j,k}$$

as shown in Fig. 5(b). A minimum drop-out voltage of 0.2V is assumed for the VRMs based on [23][24].

$$I_{out} = 0.98 * I_{in}; \quad \text{if } 0.4 > V_{out} - V_{in} \ge 0.2$$
$$I_{out} = 0.95 * I_{in}; \quad \text{if } 0.6 > V_{out} - V_{in} \ge 0.4$$
$$I_{out} = 0.90 * I_{in}; \quad \text{if } V_{out} - V_{in} \ge 0.6$$

Also, the input current of the VRM is defined as:

$$I_{i,j,k,vrm} = CI_{i,j,k}; \quad \text{if } B_{i,j,k} = 0$$
$$I_{i,j,k,vrm} = I_{in}; \quad \text{if } B_{i,j,k} = 0$$

The LP representations for the constraint are omitted for brevity.

5.3.2 PDN Synthesis Heuristic

We also propose a more scalable and near-optimal solution to the PDN synthesis problem which basically does a breadth first search (BFS) starting from the farthest node (tile) from the external power source (root node located at the lower right corner of the bottom-most layer) while assigning grid-node voltages and branch currents. The heuristic computes the placement of VRMs at grid co-ordinates while satisfying all the constraints discussed in the LP-formulation, including IR-drop constraints. In the BFS procedure, each level in the breath-first tree is termed as a *front*, thus, the node {0, 0, 0} becomes the final front. Also, any node in the current front derives the input current from its 'upstream neighbor(s)'. The heuristic attempts to minimize the number of VRM insertions at every front and at the same time, chooses as many incoming currents as possible at each node for a better IR-drop distribution. Let $CD_{i,j,k}$ be the current demand at the node {i,j,k} which is the sum of outgoing currents at the grid node. The basic flow of our heuristic is as follows:

PDN Synthesis Heuristic

input: $CI_{i,j,k}$ and $C_{i,j,k}$ values for the 3D mesh

1: Put root node in the front and assign: $V_{root}=C_{root}$ and $CD_{root}=CI_{root}$
2: **while** front is non-empty, **do{**
3: Assign upstream neighbors and processing priorities to all the nodes in the front by calling *Priority_assign()*
4: Sort all cores in the front in order of their processing priorities; for each core in the sorted-list, do{
5: Distribute $CD_{i,j,k}$ over all (one or more) of the incoming branches
6: Place VRMs at upstream or current node by calling *VRM_insert()*
7: Assign grid-node voltages at upstream nodes based on Ohms law
8: Advance the front, i.e. current front is deleted and the set of all upstream neighbors of the old front becomes the new front }

output : EI, $B_{i,j,k}$, $V_{i,j,k}$ and branch currents

The *VRM_insert()* function inserts a VRM at the upstream/current node when the upstream neighbor has a lower/higher voltage than the current node. In our PDN heuristic, we assume that the IR drop constraint will not be violated when at least one of any two consecutive nodes on the current path contains a VRM; therefore, an IR drop constraint violation is possible only when the voltages of upstream and current nodes are similar. In such a situation, an IR drop constraint violation (voltage required at the upstream node is higher than the 1% tolerance range) is rectified by the *VRM_insert()* function by inserting a VRM at the upstream node.

Given $B_{i,j,k}$ (VRM presence bit), $V_{i,j,k}$ and $CD_{i,j,k}$, the grid-node voltages and the current demands for all the nodes in the current front, the *Priority_assign()* function computes a set of upstream neighbors and processing priorities for each core in the current front. As VRM insertion is not needed (in absence of IR-drop violation) for upstream neighbor(s) of similar voltage as the current grid-node; whenever one or more upstream nodes with similar voltage are available, ($C_{upstream}=V_{i,j,k}$) they are used exclusively to supply the current node [Rule 1]. Also, all available incoming currents from upstream neighbors (with similar voltage levels) are utilized to reduce the overall effective resistance of the PDN. If no upstream nodes of similar voltage are available, and one or more prospective upstream neighbors have higher core voltages than the current grid-node, ($C_{upstream}>V_{i,j,k}$) the one with the lowest voltage amongst them is chosen to minimize current loss in the corresponding VRM [Rule 2]. Finally, if upstream nodes of only lower voltages are available, ($C_{upstream}<V_{i,j,k}$) the one with the highest voltage amongst them is chosen to minimize current loss in the corresponding VRM [Rule 3].

After a non-zero set of upstream neighbors are assigned to each node in the front, the relative order of processing of these nodes is determined by *Priority_assign()*. Any node in the current front which has already been assigned a VRM has a rigid voltage requirement because its grid-node voltage is assigned to supply to down-stream nodes. Therefore, nodes in the front which have VRMs inserted are given highest processing priority of 0 to be able to use the unassigned upstream neighbors. For the rest of the nodes; nodes assigned upstream neighbors through [Rule 3], [Rule 2] and [Rule 1] are assigned the processing priorities of 1, 2 and 3 respectively. The nodes in the front with processing priorities of 1 and 2 derive current from a single upstream neighbor and thus are given precedence over the ones with processing priority of 3 in the order of processing.

5.4 3D Routing

We employ YXZ routing in our 3D NoC fabric, which is not only deadlock-free but also uses minimal routes and has a low area footprint for implementation, thus enabling power efficient routing. The YXZ routing scheme is used on each of the mapping candidates produced by the BB procedure to compute NoC power. During the routing process, link-insertions are performed as needed to support application bandwidths, and router sizes are simultaneously updated. Also, VLCs and MCFIFOs are inserted for inter-*VI* links appropriately. After routing is done for the entire mesh, total power dissipation in the NoC is computed taking into consideration the number of links inserted, link loads, router sizes, number of VLCs and MCFIFOs used, and the corresponding voltage/frequency values.

5.5 Solution Pruning

Once a set of solutions is output by the synthesis flow, a penultimate pruning is performed to remove solution points that are not relevant. For instance, if two solutions have the same number of VRMs and communication power, but different external current (EI) values, then the solution with the higher EI value is pruned.

6. Experiments

We used the ARM11MPCore multi-core processors [26] as the base compute cores in our experiments, which support three operating voltage levels. Our experiments were conducted on applications based on pseudo-random core graphs derived using TGFF [27] with edge weights annotated with bandwidths representing inter-core communication requirements. We conservatively assume that the square of the voltage scales linearly with the frequency, as in previous works [6]. The three core voltages and their corresponding frequencies and maximum supply current values we employed are: (0.8V, 195MHz, 0.52A), (1.0V, 304MHz, 0.50A) and (1.2V, 437MHz, 0.48A). The maximum supply current values are derived from the rated maximum compute core power values. Also, the value of external voltage source of the PDN is assumed to be 1.5V. The architecture and power values of routers and links as well as MCFIFOs and VLCs operating at different voltages and frequencies are taken from [6]. The branch resistance values used in the PDN design are: 43, 53 and 63 mΩ (based on [3]). In the BB procedure for 3D mapping, values of n=9 and N=800 are used so that the maximum values that variables *B*, *D* and *R* can take are 9, 5 and 4 respectively, with α=1. The PDN linear programming formulation is solved using the open-source tool lp_solve 5.5.2 [25].

6.1 Results

Our first set of experiments focus on the PDN synthesis problem and compare the fidelity of solutions obtained from the LP formulation and the heuristic. The PDN synthesis results (*EI* and # of VRMs needed for a range of *R* values) obtained using our heuristic are found to be within 15% of the corresponding LP-results for small benchmarks with less than 10 cores. Table 1 summarizes the results for 3D CMPs with 4 (2×2), 8 (2×2×2), and 9 (3×3) cores. The key advantage of using the heuristic is that it generates a solution in a matter of a few seconds, whereas the LP formulation takes in the order of several hours. For problem sizes with greater than 10 cores, the LP solver did not return a solution even after being left to run for an entire day. Thus the heuristic provides a more scalable and thus practical solution to the PDN synthesis problem.

Table 1: Comparison between results obtained from LP and the PDN Heuristic Implementations

#of cores	3D core config.	R (mΩ)	LP		Heuristic	
			# of VRMs	EI (Amp)	# of VRMs	EI (Amp)
4	2×2 (1 layer)	43	2	0.203	2	0.208
		63	2	0.203	2	0.208
8	2×2×2 (2 layers)	43	3	0.409	4	0.421
		63	4	0.411	4	0.422
9	3×3 (1 layer)	43	6	0.464	6	0.480
		63	6	0.469	7	0.483

Next we explore the results generated by our complete PDN-aware 3D NoC synthesis approach for a large 64 (4×4×4) core CMP. Fig. 6 shows the 3D solution space, with each candidate solution characterized by its communication power, number of VRMs, and external current (EI) drawn. The results are shown for the grid wire resistance values of 43 mΩ and 63 mΩ (results for the 53 mΩ case

not shown). The most important insight from these results is that the solution with the minimum communication power (highlighted by red dots in each of the figures) does not necessarily have optimum PDN cost. For instance in Fig. 6(b), the solution (shown as red '*') with the lowest communication power (NoC cost) of 2374 mW, not only requires the highest number of VRMs (31) but also has a very high *EI* of 3.972 A. Traditional 3D NoC synthesis approaches output the lowest communication power solution. By the time the PDN is designed in the back end for this solution, it is too late to trade-off PDN complexity with communication power. For the lowest communication power solutions shown in Fig. 6 (a)(b), it would be quite a significant challenge to meet the area and power design-constraints during PDN design and designers may need to resort to over-margining which is wasteful and increases system cost.

In contrast, our proposed PDN aware NoC synthesis framework can produce a set of interesting design points that can enable trade-offs between NoC power dissipation, VRM count, and external current drawn. For instance, possible solutions which optimize all three cost metrics are the pentagrams highlighted in blue. Additionally, if either the minimal number of VRMs or the minimum EI is required as the final design solution, other design points (a square or a diamond in black, respectively) could be chosen. Ultimately, our framework allows designers to explore trade-offs in the PDN and NoC design space early on at the system level, and is invaluable to achieving better quality designs.

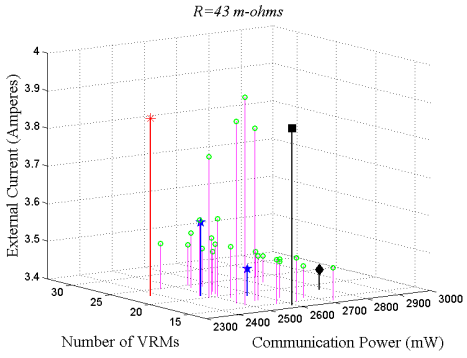

(a)

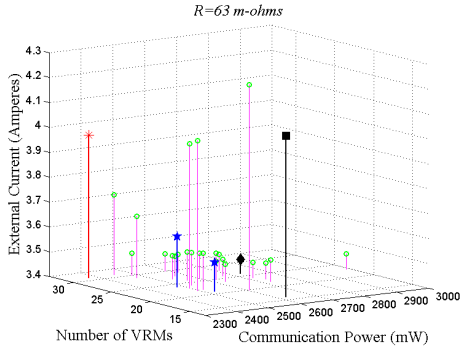

(b)

Figure 6: Results of the PDN-aware NoC synthesis framework for (a) R=43 mΩ (b) R=63 mΩ

7. Conclusion

In contrast to the traditional CMP design approach where PDN design is done on a mapping instance which is optimized exclusively for NoC costs (e.g. communication power) this work advocates an automated framework for PDN-aware synthesis of NoCs in 3D CMPs. Our framework enables the designer to weigh the PDN design costs against the NoC design costs and thereby obtain a more efficient overall solution. The experimental results show that the solution space uncovered by our framework can allow system level trade-off analysis of PDN and NoC design decisions which has never been attempted before. By accepting solutions with less than optimal NoC power dissipation characteristics, our framework reveals that it is possible to significantly reduce PDN design cost.

References

[1] B. Amelifard, et al., "Optimal Design of the Power-Delivery Network for Multiple Voltage-Island System-on-Chips," IEEE TCAD 28(6), May 2009.
[2] P. Jain et al. "A multi-story power delivery technique for 3D integrated circuits," Intl. Symp.on ISLPED pg. 57, Aug. 2008.
[3] N. Khan et al. "System-level comparison of power delivery design for 2D and 3D ICs," IEEE Intl. Conf. 3D-IC, Sept. 2009.
[4] W. Kim et al. "Enabling OnChip Switching Regulators for Multi-Core Processors using Current Staggering," Proc. ASGI, 2007.
[5] Z. Zenget al. "Tradeoff analysis and optimization of power delivery networks with on-chip voltage regulation," Proc. DAC pp. 831, June 2010.
[6] N. Kapadia, et al., "VISION: A Framework for Voltage Island Aware Synthesis of Interconnection Networks-on-Chip," Proc. GLSVLSI 2011.
[7] S. Murali et al., "Designing Application-Specific Networks on Chips with Floorplan Information," Proc. ICCAD. 2006.
[8] C. Seiculescu et al., "NoC Topology Synthesis for Supporting Shutdown of Voltage Islands in SoCs," Proc. DAC, July, 2009.
[9] P.Zhou et al., "Application-Specific 3D Network-on-Chip Design Using Simulated Allocation," Proc. ASPDAC. Jan., 2010.
[10] K. Srinivasan, K. Chatha, "A low complexity heuristic for design of custom network-on-chip architectures," Proc. DATE, 2006.
[11] K. Srinivasan et al., "Linear-Programming-Based Techniques for Synthesis of Network-on_Chip Architectures," IEEE TVLSI, 14(4) Apr 2006.
[12] U. Ogras, et al., "It's a Small World After All": NoC Performance Optimization Via Long-Range Link Insertion," IEEE TVLSI 14(7), Jul 2006.
[13] L.Leung, C. Tsui; "Energy-Aware Synthesis of Networks-on-Chip Implemented with Voltage Islands," Proc. DAC, 2007.
[14] P. Ghosh, et al. "Efficient Mapping and Voltage Islanding Technique for Energy Minimization in NoC under Design Constraints," Proc. SAC, 2010.
[15] U. Ogras et al, "Voltage-Frequency Island Partitioning for GALS-based Networks-on-Chip", Proc. DAC, pp: 110–115, 2007.
[16] J. Hu, R. Marculescu; "Communication and task scheduling of application-specific networks-on-chip," IEE CDT 152(5), Sep. 2005.
[17] C Chou et al., "Energy and Performance-Aware Incremental Mapping for Networks on Chip with Multiple Voltage Levels," IEEE Transactions on CAD (TCAD), 27(10): 1866–1879, Oct. 2008.
[18] J. Hu, R. Marculescu; "Energy and Performance-Aware Mapping for Regular NoC Architectures," IEEE TCAD 24(4), Apr. 2005.
[19] W. Jang, D. Ding, D. Pan, "A Voltage-Frequency Island Aware Energy Optimization Framework for Networks-on-Chip," Proc. ICCAD, pp: 264 – 269, Nov. 2008.
[20] P.Falkenstern et al., "Three-dimensional integrated circuits (3D IC) Floorplan and Power/Ground Network Co-synthesis," Proc. ASPDAC, 2010.
[21] H. Chen et al., "A New Architecture for Power Network in 3D IC," Proc. DATE, March 2011.
[22] M. Popovich et al., "On-Chip Power Distribution Grids with Multiple Supply Voltages for High-Performance Integrated Circuits," IEEE Transactions on VLSI Systems vol. 16(7), pp. 908-921, July 2008.
[23] Y. Lee, K. Chen; "A 65nm sub-1V multi-stage low-dropout (LDO) regulator design for SoC systems," IEEE Intl. MWSCAS pg. 584; Aug. 2010.
[24] J. Guo, K. Leung; "A 6-µW Chip-Area-Efficient Output-CapacitorlessLDO in 90-nm CMOS Technology," IEEE Journal of Solid State Circuits 45(9), pg. 1896, Sept. 2009.
[25] http://lpsolve.sourceforge.net/5.5/
[26] http://www.arm.com/products/processors/selector.php
[27] http://ziyang.eecs.umich.edu/~dickrp/tgff/
[28] X. Jiang, T. Watanabe, "An efficient 3D NoC synthesis by using genetic algorithms," Proc. TENCON Nov. 2010.
[29] C. Seiculescu et al., "SunFloor 3D: A tool for Networks On Chip topology synthesis for 3D systems on chips," DATE 2009.
[30] K. Siozios et al., "A High-Level Mapping Algorithm Targeting 3D NoC Architectures with Multiple Vdd," ISVLSI 2010.
[31] C. Seiculescu et al., "Comparative Analysis of NoCs for Two-Dimensional Versus Three-Dimensional SoCs Supporting Multiple Voltage and Frequency Islands," IEEE TCAS, May 2010.
[32] M. Arjomand et al., "Voltage-Frequency Planning for Thermal-Aware, Low-Power Design of Regular 3-D NoCs," VLSID 2010.