

A Co-Synthesis Methodology for Power Delivery and Data Interconnection Networks in 3D ICs

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

A stable voltage supply is critical for multiprocessor system-on-chips (MPSoCs) to operate at near-optimal performance levels. The problem of IR drops in a Power Delivery Network (PDN) is very severe in 3D MPSoCs with network-on-chip (NoC) fabrics where the current in the PDN increases proportionally with the number of device layers. At the same time, with the increasing core counts in today's power-hungry MPSoCs, the already hard problem of voltage island-aware Network-on-Chip (NoC) design has become even more challenging. Even though the PDN and NoC design goals are non-overlapping, both the optimizations are interdependent. Unfortunately, designers today seldom consider design of the PDN while synthesizing NoCs. In this work, for the first time, we propose a novel system-level co-synthesis methodology that minimizes 3D NoC energy while meeting performance goals; and simultaneously optimizes the 3D PDN design while satisfying IR-drop constraints. Our experimental results show that the proposed co-synthesis methodology meets IR-drop constraints while minimizing energy consumption for several real-world applications, improving upon results from traditional system-level methodologies that perform PDN design and NoC synthesis separately.

Keywords

System-level CAD, NoC synthesis, core mapping, communication energy, multi-core systems

1. Introduction

Designing a robust Power Delivery Network (PDN) is critical to the overall performance of today's multi-processor system-on-chips (MPSoCs). The PDN is required to deliver a stable power supply across the chip that is within a desired voltage range and tolerate large variations in load currents [1]. For the case of multiple voltage islands (VIs) that are used in modern designs to minimize total power dissipation, the PDN is required to supply power at different voltage levels corresponding to the VIs while keeping power loss to a minimum. Unfortunately, with increasing on-chip device density and decreasing voltage levels, the supply currents have risen, however the scaling of PDN impedance has not kept up with this trend [2]. Worsening IR-drops reduce the quality of voltage supply and negatively impact MPSoC performance. This problem is more severe in 3D MPSoCs as the current in the PDN could be as many times more as the number of device layers compared to a 2D MPSoC. Besides, the number of I/O pins on an n -layered 3D design 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 MPSoCs 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. NoCs with mesh-based (regular) topologies have been employed in recent multi-core designs, e.g. TILE64 Processor [33] and Intel's Single Chip Cloud Computer (SCC) [34]. As the power dissipated in the NoC has become a significant portion of the total on-chip power, optimizing communication power in addition to computation power is critical [8]. Several recent works have proposed techniques to synthesize power-optimized regular and custom 3D NoC topologies [9], [28]-[32]. 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 NoC configuration that meets performance goals. Performing synthesis of the PDN for the generated NoC configuration in these cases can put stringent demands on the already strained PDN, possibly making it extremely difficult to meet PDN constraints such as maximum IR-drop; or lead to over-margining for the PDN, which can be wasteful.

In our experience, the traditional approach of synthesizing a NoC fabric without considering the PDN severely constrains the PDN design space, leading to sub-optimal or even completely infeasible designs. In this work, for the first time, we propose an automated system-level methodology for the co-synthesis of PDN and mesh-based NoC fabrics in Vt -enabled 3D MPSoCs that jointly optimizes communication energy and PDN cost. Our focus on a mesh NoC fabric is motivated by their popularity in emerging MPSoC chips [33]-[35], due to their structural regularity, feature predictability, and scalability. We recognize the key insight that different instances of voltage partitioning and core-to-tile mapping can *significantly* alter the voltage distribution map seen by the PDN. Intuitively, by mapping cores with higher supply current requirements to tiles (on the 3D mesh) closer to the external input power pins, the worst-case IR-drop (max-IR-drop) can potentially be alleviated. Accordingly, our methodology considers the interdependence between a synthesized NoC configuration and its corresponding IR-drop distribution in the PDN. The novel contributions of our synthesis framework are:

- We propose a Simulated Annealing (SA) based algorithm to co-synthesis a Vt -enabled 3D mesh NoC with a PDN to generate an overall optimized MPSoC design;

- We develop a linear programming (LP) formulation to evaluate the IR-drop distribution at the grid-nodes of the 3D PDN;
- We propose a novel algorithm for 3D routing path allocation which considerably minimizes communication energy in the 3D NoC;
- Our methodology generates 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 overall energy, performance, and PDN cost based design goals.

2. Related Work

Techniques for optimizing PDNs in 3D ICs have been studied in several recent works [1]-[3], [20], [21]. Amelifard et al. [1] use dynamic programming to generate a multi-level tree topology of suitable voltage regulator modules to improve PDN power efficiency in multiple voltage island system-on-chip designs. 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. Khan et al. [3] analyze the impact of TSV size/spacing on IR drops and voltage droops in several 3D power delivery configurations; with SPEC benchmarks running on a 4-core architecture. 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 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 interconnection fabric and core mapping across layers.

Many researchers [7]-[11] have proposed custom 2D NoC topology synthesis techniques 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]. Due to the above mentioned drawbacks of custom topologies, recent industrial multi-core designs (e.g., Intel's Teraflop and SCC processors, and Tiler's Tile64) make use of regular topologies [33]-[35].

The problem of 2D NoC synthesis on regular structures with multiple supply V_I s has been addressed in several works [6], [13]-[19]. Ogras et al. [15] perform VI partitioning and static voltage-frequency assignment on a pre-mapped NoC to optimize communication energy while meeting task deadline constraints. By performing voltage-partitioning before core mapping, Jang et al. [19] demonstrate improvements over prior work (e.g., [15]) to achieve less power overhead by reducing total number of voltage level converters (VLC) and multiple-clock first-in first-out (MCFIFO) frequency level converters needed for inter-island communication. [6] improves upon [19] in terms of communication power by

proposing novel heuristics to generate a core-to-tile mapping and a routing scheme that more aggressively optimized inter-island communication between cores.

Given the promise of 3D technologies, 3D NoC synthesis in recent years has also attracted significant research efforts [9], [28]-[32]. A 3D-NoC synthesis algorithm called the Ripup-Reroute-and-Router-Merging is proposed by Yan et al. [28], where the ripup-reroute algorithm is used for routing flows and router-merging used to optimize topologies for communication power. Seiculescu et al. [29] utilize min-cut partitioning of cores, located within as well as across device layers; to establish core-to-router connectivity. The same authors propose a tool for synthesizing application-specific 3D NoCs in [31] that finds paths for the communication flows and performs placement of network components on to the 3D layers; while reporting savings in interconnect power consumption and delay. In [37], improvements are shown over [31] by additionally considering the impact of TSV serialization and network interface (NI) placement in application specific 3D NoCs. Zhou et al. [9] present a floorplan-aware synthesis algorithm for application-specific 3D NoCs, based on simulated allocation; generating network topologies to optimize NoC power, network latency, and chip temperature. Arjomand et al. [32] perform thermal-aware assignment of tasks onto a 3D regular mesh followed by voltage-frequency planning to minimize power dissipation and meet application constraints. However, none of the above approaches have considered the effects of 3D NoC synthesis on the efficiency and overheads associated with the 3D PDN design; in other words, these approaches are not PDN-aware.

Unlike any prior work, in this paper we present a novel co-synthesis methodology for PDN and mesh-based NoC fabrics in 3D MPSoCs. To the authors' knowledge, this is the first work which proposes a system-level co-synthesis framework that co-optimizes the 3D NoC fabric with the 3D PDN fabric to produce a more efficient overall MPSoC design.

3. PDN Design with Multiple Voltages

High circuit density in smaller footprint 3D ICs presents a unique challenge for designers of PDNs, as it requires the network to deliver significantly more current than in 2D ICs with fewer P/G bumps, while also circumventing increasingly daunting IR-drop issues. The PDN should be able to restrict the IR drops at each core-input within the set tolerance limit of the rated core voltage.

We use an MPSoC platform with a 3D mesh of tiles, where there is a one-to-one mapping of processing cores onto these tiles. In order for the PDN to handle multiple voltage domains, not unlike the 'voltage volume' concept introduced in [21], we assume the MPSoC design to be partitioned into 3D- V_I s such that cores of same voltage are vertically aligned in the 3D stack (as shown in figure 1). This enables independent and physically disjoint 3D power supply grids to connect to all cores operating at the same voltage (as shown for the blue voltage island in figure 1). The inter-layer connections in the grid can be made with help of power TSVs. Even though all V_I s are contiguous, we do not restrict the V_I shapes to rectangles (as recommended in [4]). It is

shown in [5] that sizing of pitches and widths of power grid does not change the on-chip voltage distribution, therefore, in our work, we assume fixed uniform power grids as in prior works [5][23]. Also, as we investigate the steady state effects of the PDN, time-varying network characteristics such as transient noise are not considered.

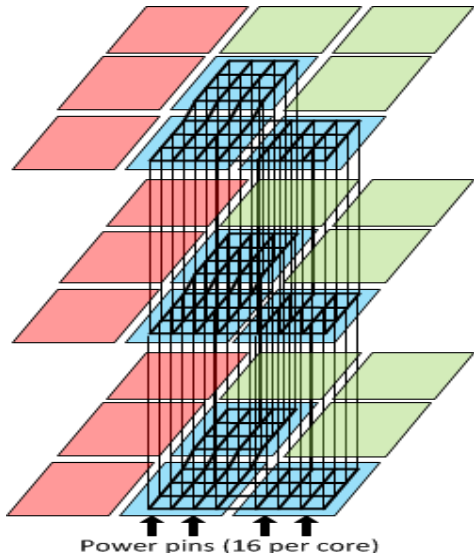


Figure 1: Example of contiguous 3D-VIs in a (3x3x3) MPSoC, cores within a single 3D-VI are vertically aligned (VIs are color-coded). The 3D power grid for just the blue VI is shown. Each VI has a separate power grid, with 16 grid-nodes per core.

4. Co-synthesis Problem Formulation

We assume 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 consists of a compute core and a NoC router;
- A core graph $G(V, E)$; with a set of T vertices $\{V_1, V_2, \dots, V_T\}$ representing homogenous cores on which application tasks have already been mapped, and the set of M edges $\{e_1, e_2, \dots, e_M\}$ that represent communication dependencies (communication volumes) between cores;
- A set of Ω supply voltage (V_{dd}) levels, which also represents the total number of VIs;
- A set of pre-assigned couplets constituting operating voltages and operating frequencies for the T cores $\{(v_1, f_1), (v_2, f_2), (v_3, f_3), \dots, (v_T, f_T)\}$, to meet compute performance requirements of tasks mapped to the cores;
- A set of pre-assigned maximum supply currents $\{i_1, i_2, \dots, i_T\}$ required by the T cores which depend on their respective task assignments;
- Ω separate regular 3D power grids (each corresponding to a VI), with multiple power input pins for every grid.

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 3D-VI contiguity constraints are satisfied (all cores are contiguously placed within individual 3D VIs), as well as the PDN IR-drop constraints are satisfied; while minimizing total energy for communication using NoC

components (routers, links, voltage level converters, mixed clock FIFOs) and the PDN cost (represented by the maximum IR-drop). Note that our proposed methodology is also applicable to 2D designs (when the vertical dimension (dim_z) of the mesh is set to 1).

5. PDN-NoC Co-synthesis Framework

We now present the detailed flow of our PDN-NoC co-synthesis framework. We use a dual-objective simulated annealing (SA) based algorithm, but instead of saving just the best solution, at each iteration, we maintain a Pareto front of two dimensional (PDN and NoC) costs that are not dominated by any solution found until then. The two optimization objectives are communication energy in the NoC (CE) and the maximum IR-drop in the entire PDN (max-IR-drop). The three possible solution perturbations are described below:

Perturbation 1: core-swap – Two cores randomly selected from the same 3D-VI are swapped. As the swap takes place within the same 3D-VI, VI-contiguity is not a concern.

Perturbation 2: 3D-VI-swap – Two 3D-VIs (of equal sizes), are randomly selected to be swapped.

Perturbation 3: tile-exchange – between two separate 3D-VIs

The VI configuration in the 3D mesh is changed by reciprocal occupation of a tile (tile-exchange) between two adjacent VIs (with at least 2 tiles in each VI being adjacent to the other VI) on every tier of the 3D MPSoC. Any tile-exchange can be valid only if both participating VIs retain their contiguity (as shown in Figure 2).

Note that entire columns of cores (across all tiers) need to be exchanged between the 3D-VIs, and if this perturbation is valid on one tier it is valid on all tiers because all 3D-VIs are vertically aligned across tiers. The pseudo-code for our framework is given below:

SA-based co-synthesis framework

inputs: as described in section 4

- 1: Generate a valid initial mapping and evaluate cost(current solution)
- 2: Set T to T_{init}
- 3: **while** ($T \leq T_{min}$) { K=0
- 4: **while** ($K < K_{max}$) {
- 5: Perturb current mapping solution in one of 3 ways to get new solution
- 6: Perform 3D_Routing() and evaluate the new 3D-NoC solution for CE
- 7: Run PDN_solver() to evaluate new mapping solution for max-IR-drop
- 8: Compute: $cost(new\ solution) = \beta * (CE) + \gamma * (max-IR-drop)$
- 9: If max-IR-drop constraint is violated, ACCEPT=0, else set value of ACCEPT according to the SA-acceptance criterion
- 10: if (ACCEPT=1), then $cost(current\ solution) = cost(new\ solution)$ and update the Pareto front
- 11: K= K+1 }
- 12: T = $\alpha * T$ }

output: final set of solutions on the Pareto front with the associated 3D NoC and 3D PDN costs

An initial mapping solution is arbitrarily generated, which satisfies the 3D-VI contiguity constraints. The cost of this initial solution is computed with PDN_solver() and 3D_Routing() functions which are described in sub-sections 5.1 and 5.2 respectively. The initial solution now becomes the current solution in the SA process (**step1**). To initiate the SA process, the SA-temperature parameter (T) is set to T_{init}

(step2). At each iteration, one of the 3 perturbations is randomly chosen to perturb the current solution. To generate enough mapping solutions for every 3D- VI configuration, we choose perturbations 1, 2 and 3 with probabilities 10/13, 1/13 and 2/13 respectively **(step5)**. To evaluate the new (perturbed) mapping for communication energy as well as max-IR-drop, our 3D_Routing() and PDN_solver() are performed **(steps 6, 7)**. Then, the combined cost **(step8)** is computed (cost(new_solution)). All new solutions violating the PDN max-IR-drop constraint are rejected outright. Whereas, solutions which do satisfy the max-IR-drop constraint with better (lesser) cost than the current solution are accepted, and the solutions with worse costs are either accepted or rejected according to the following SA-acceptance criterion **(step9)**:

$$if (r < exp([cost(current\ solution) - cost(new\ solution)]/T)) \\ ACCEPT = 1 \\ where, r is a random number between 0 and 1$$

The new solutions which do not have both PDN cost and NoC cost greater than any solution on the current Pareto front (non-dominated solutions) are inserted into the Pareto front, and the Pareto solutions which get dominated as a result, are discarded from the front. Thus, the Pareto front of solutions (with non-dominated costs), is checked for an update at every iteration of SA **(step10)**. At the end of every K_{max} iterations, the SA-temperature is scaled by the scaling parameter α . Finally, when the temperature parameter becomes smaller than the pre-determined T_{min} , the SA process terminates, and a set of Pareto design points are produced, each optimized for the PDN and NoC design objectives by varying degrees.




Figure 2: Example of a tile-exchange perturbation (a) an invalid perturbation where the contiguity of VI s is not retained (b) a valid perturbation

5.1. PDN-solver Formulation

As discussed previously, there are Ω regular 3D power grids in our PDN, one for each 3D- VI , and the individual VI s (on the 2D plane) are not necessarily rectangular in shape. We assume equal number of grid-nodes (in an $n \times n$ 2D grid) supplying to each core on the MPSoC (as shown in figure 1). Therefore, the total number of grid-nodes in the PDN are $T * n^2$, and total number of external inputs are $(T * n^2) / dim_z$. Note that the grid-nodes of adjacent power grids are unconnected. Finally, we assume that the power grids are located above the top tier, so that all the vertical PDN branch currents flow in the downward direction. We use a linear programming (LP) formulation to solve for the grid-node

voltages and currents flowing in the branch resistances of the PDN with multiple 3D grids. Our final metric of interest is the percentage max-IR-drop in the entire PDN, which we obtain from all the grid-node voltages.

We are given a set of T tile co-ordinates $T_{i,j,k}$, for $0 \leq i \leq dim_x - 1$, $0 \leq j \leq dim_y - 1$, $0 \leq k \leq dim_z - 1$ on a 3D mesh with dimensions $\{dim_x, dim_y, dim_z\}$. Also, 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. Values of horizontal and vertical branch resistances (R_h, R_v) are defined for a uniform grid, where R_v is typically higher due to the added TSV resistance. The design variables considered in our problem formulation are as follows:

- Horizontal PDN branch currents leaving the grid-nodes in $d \{pos_x, neg_x, pos_y, neg_y\}$ direction: $I_{i,j,k,l,m-d}$, where co-ordinates $\{i, j, k\}$ represent the 3D location of the corresponding tile on the 3D mesh and $\{l, m\}$ represent the location of the grid-node on the tile;
- Vertical PDN branch currents entering or leaving the grid-nodes: $I_{i,j,k,l,m}$ or $I_{i,j,k+1,l,m}$
- PDN grid-node voltages: $V_{i,j,k,l,m}$

The key constraints of the LP solver are as follows:-

Constraint 1: Grid Nodal Equation by Kirchoff's Current Law (KCL):

$$I_{i,j,k,l,m} - I_{i,j,k,l,m,pos_x} - I_{i,j,k,l,m,neg_x} - I_{i,j,k,l,m,pos_y} - I_{i,j,k,l,m,neg_y} - I_{i,j,k+1,l,m} = CI_{i,j,k} / n^2$$

where n^2 is the number of grid-nodes supplying to each core

Constraint 2: Grid-node voltages in the top-tier (0^{th} tier) is equal to the rated core supply voltages:

$$V_{i,j,0,l,m} = C_{i,j,k}$$

Constraint 3: Voltage drops across all vertical branch resistances are defined by the value of R_v used:

$$I_{i,j,k+1,l,m} = \frac{V_{i,j,k,l,m} - V_{i,j,k+1,l,m}}{R_v}$$

Similar equations exist for horizontal branch resistances using the value R_h .

Constraint 4: Maximum IR-drop constraint is set to 4.5%, i.e. grid-node voltages can be no smaller than 95.5% of the corresponding rated core voltage:

$$V_{i,j,k,l,m} \geq 0.955 * C_{i,j,k}$$

To avoid inter-connecting grid-nodes of adjacent 3D- VI s, we compute a set of binary valued constants for a given core-to-tile mapping to define the adjacent grid-nodes in separate 3D-grids. The above constraints are fed to an LP solver [25] to extract exact values of all grid-node voltages in the PDN, from which the percentage max-IR-drop for a given core-to-tile mapping can easily be calculated.

5.2. Energy Aware 3D Routing (3D_Routing())

The core-to-tile mapped mesh consists of multiple VI s that run on different voltage levels as well as different frequencies. Therefore, for inter-island communication, voltage level converters (VLCs) and frequency level converter resources (MCFIFOs) are required in the corresponding routers. Whenever a low voltage core transmits to a higher voltage core, a VLC is needed on the

outgoing port of the source router. Also, for any inter-island link, an MCFIFO is needed for the higher frequency core as the connecting link works at the lower frequency. These frequency and voltage conversion components incur an overhead in terms of power dissipation and delay. Also, in a general pipelined routing architecture, the link delay on the slowest link on the path becomes the lower bound on path latency. Thus, the main objective of our routing path allocation is to find a path for each communication flow such that the number of inter-island links and the bottleneck volume (highest communication volume allocated on any link along the path) are reduced.

The order in which the communication flows are routed, is determined in the following way in our framework. The communication flows with longer minimal paths (MDs) have more choices for routing and thus have a larger scope for optimization. Also, flows with smaller volumes have a smaller overall transfer latency footprint for NoC communication. Therefore, communication flows are sorted in the increasing order of their path lengths, in decreasing order of their communication volumes for the same path length; and considered for routing in that order in this third step of our framework.

For each communication flow, we consider all candidate minimal paths. Note that, the number of all possible minimal paths between two cores on a 3D mesh, which are N hops apart ($N = x + y + z$; where x , y and z are the number of x -hops, y -hops and z -hops on the 3D path) is given by $\{^N C_x\} * \{^{(N-x)} C_z\}$. Here, $\{^N C_x\}$ represents the number of possible ways the x -path can be constructed; and $\{^{(N-x)} C_z\}$ represents the number of possible ways the z -path can be constructed, for a given x -path. Out of these candidate minimal paths, we choose a routing path based on the following optimization objectives (in that order):

- 1) Minimize energy overhead of MCFIFOs and VLCs; and
- 2) Minimize latency bottleneck on the communication path

To meet the above objectives, we first choose paths that need the minimum total number of inter-island links. Then, out of the chosen paths (with the same number of inter-island links), we choose the one which has the lowest bottleneck volume. When a path is chosen, the current communication flow (volume) is allocated to its constituent links. We also perform a post-processing design time cyclic dependency analysis based on the approach from [22], to ensure freedom from cyclic dependencies which can lead to deadlock at runtime. After the completion of this step, a mapped and routed 3D NoC-based MPSoC is obtained for an application.

6. Experiments

6.1. Experimental Setup

We use the ARM Cortex-A9 multi-core processors [26] as the baseline MPSoC compute cores in our experiments, which support four operating voltage levels ($\Omega=4$): 0.8V, 0.9V, 1.0V and 1.1V; and corresponding operating frequencies of 1055MHz, 1310MHz, 1550MHz and 1775Mz. The maximum current requirements for the processing cores range from 0.8A to 1.6A, based on the level of compute intensity of the tasks assigned to the respective cores. We use a 64-core 3D-mesh for the MPSoC with dimensions $4 \times 4 \times 4$.

Our experiments were conducted using real-world application task-graphs (*Sparse_96*, *Robot_88*, *Fppp_334*) taken from the Standard Task Graph (STG) repository [27], where each vertex (task) corresponds to processing time (in units of hundreds of ms) and the edge weights represent inter-task communication volumes (with values ranging between 20 Gb and 220 Gb). In order to convert task-graphs into core-graphs, we iteratively shrank the inter-task edges with the lowest sum of processing-times of the two tasks it connects, until the number of tasks (G) equaled the number of tiles or cores (T) on the 3D-mesh (assuming $G \geq T$). In this way, the benchmarks *Robot_88*, *Sparse_96* and *Fppp_334* are converted to a core-graph of 64 cores.

The power values of routers and links (32-bit wide) for different voltages, frequencies and router complexities at full load are obtained from ORION 2.0 [38]. The active times of NoC-components are used for computing their energy values,

$$E_{component} = Power_{component} * AT_{component}$$

where active time (AT) of a NoC-component is the time period for which the component needs to be active (at full load) in a single traversal of the application (task-graph). As the VLCs and MCFIFOs required to interact between VIs incur a power overhead that is proportional to their voltage supply, we consider the power overhead of these components, with the actual power values based on reported overheads from existing literature [36].

In our implementation of the SA-based algorithm, we simulate for 2000 iterations (with $T_{init}=100$, $T_{min}=14$ and $K_{max}=100$) and the SA-temperature scaling factor used is $\alpha=0.9$. The normalizing coefficient values of β and γ are set in a manner to make both terms in the SA cost function approximately equal. Our regular 3D-PDN power grids are modeled based on the guidelines provided in [24]. With 16 cores on each tier, a total of 256 input power pins are used with $n^2=16$ grid-nodes for each core. Thus, the total number of grid-nodes in the entire PDN is equal to 1024. The values of $R_h=40m\Omega$ and $R_v=83m\Omega$ are used for the horizontal and vertical branch resistances, based on [3][24].

6.2 Results

To the authors' knowledge, this is the first work which proposes a co-synthesis framework that co-optimizes the 3D NoC fabric with the PDN fabric to produce a more efficient overall 3D MPSoC design. Therefore, to evaluate the efficacy of the proposed 3D PDN-NoC co-synthesis methodology, we compare the results obtained from our co-synthesis framework with our implementation of a similar SA-based (PDN-unaware) synthesis technique which optimizes exclusively for NoC communication energy. For a fair comparison of results, we simulate the SA for the same number of iterations in both the cases. Our `PDN_solver()` and `3D_Routing()` are utilized to evaluate the max-IR-drop and the communication energy in both the implementations. We note here that, from our experiments, we have found our `3D_Routing()` algorithm to produce more than 20% savings (on average) in total communication energy compared to a 3D implementation of the genetic algorithm based routing scheme used in [13].

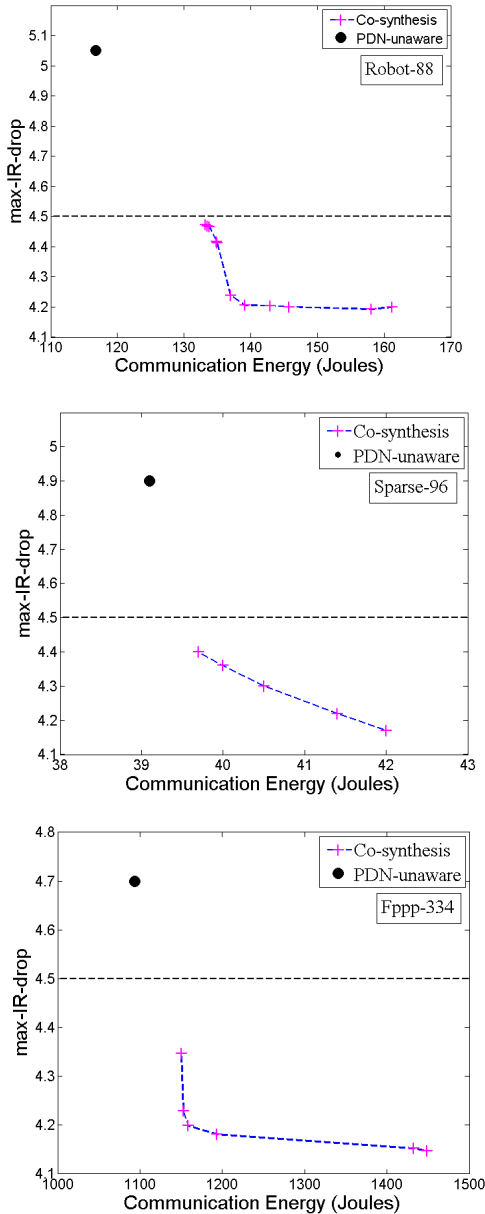


Figure 3: Results of the PDN-NoC co-synthesis framework for the 3 application benchmarks – Robot_88, Sparse_96, Fppp_334.

We explore the results generated by our SA-based 3D PDN-NoC co-synthesis approach for a large 64 (4×4×4) core MPSoC. Figure 3 shows the 2D solution space, with each candidate solution characterized by its communication energy and max-IR-drop. With the PDN-unaware 3D NoC synthesis approach, where the optimization objective is NoC communication energy exclusively, the single solution point represents the PDN cost (max-IR-drop) corresponding to the optimal NoC energy cost. On the other hand, our co-synthesis approach produces a 2D Pareto front, enabling the designer to choose a design point for an appropriate trade-off between NoC energy and PDN costs. The PDN unaware approach, which is representative of system-level NoC synthesis approaches proposed in literature to date, does not optimize for PDN cost. As a result, the solution generated optimizes for NoC energy cost, but cannot meet the 4.5% max-IR-drop

constraint for all three of the applications that we considered, rendering the final solutions infeasible. In such cases, more effort and resources are required in a subsequent PDN design phase, e.g., by changing grid wire sizing or increasing supply voltages and currents, to satisfy core current requirements.

Our 3D NoC-PDN co-synthesis framework proposed in this work, in contrast, produces feasible core-to-die mapping solutions and communication routes which together co-optimize the max-IR-drop explicitly, along with communication energy. Improving upon the worsening IR-drops in 3D MPSoCs without incurring additional PDN resources, as made possible by our co-synthesis framework, could improve MPSoC performance quite significantly, and reduce back-end designer effort required in creating modern 3D MPSoC designs with very stringent PDN constraints.

7. Conclusion

In contrast to the traditional MPSoC design approach where PDN design is done in a separate phase, after a core mapping and routing solution which is optimized exclusively for NoC costs (e.g. communication energy) is obtained, this work proposes an automated system-level framework for the co-synthesis of PDN and NoC fabrics in 3D MPSoCs. Our framework enables the designer to trade-off the PDN design costs against the NoC design costs, which can lead to a more efficient overall solution. By co-synthesizing PDN and NoC fabrics, our co-synthesis framework investigates a wider solution search space compared to a PDN-unaware synthesis approach which exclusively optimizes for NoC costs. In doing so, it generates more desirable solutions that can ease design effort and improve time-to-market of emerging 3D MPSoC designs.

Acknowledgement

This research is sponsored in part by grants from NSF (CCF-1252500) and SRC.

References

- [1] B. Amelifard, M. Pedram, “Optimal design of the power-delivery network for multiple voltage-island system-on-chips,” *IEEE TCAD* 28(6), pp. 888-900, May 2009.
- [2] P. Jain, T. Kim, J. Keane, C. Kim, “A multi-story power delivery technique for 3D integrated circuits,” *Intl. Symp.on ISLPED* pp. 57-62, Aug. 2008.
- [3] N. Khan, S. Alam, S. Hassoun, “System-level comparison of power delivery design for 2D and 3D ICs,” *IEEE Intl. Conf. 3D-IC*, Sept. 2009.
- [4] W. Lee, D. Marculescu, Y. Chang, “Post-floorplanning power/ground ring synthesis for multiple-supply-voltage designs,” *ISPD*, 2009.
- [5] Q. Zhou, J. Shi, B. Liu, Y. Cai, “Floorplanning considering IR drop in multiple supply voltages island designs,” *TVLSI* 19(4), Apr. 2011.
- [6] N. Kapadia, S. Pasricha, “VISION: A framework for voltage island aware synthesis of interconnection networks-on-chip,” *Proc. GLSVLSI*, pp. 31-36, May 2011.

- [7] S. Murali et al., "Designing application-specific networks on chips with floorplan information," Proc. ICCAD, pp. 355-362, Nov. 2006.
- [8] C. Seiculescu, S. Murali, L. Benini, G. De Micheli, "NoC topology synthesis for supporting shutdown of voltage islands in SoCs," Proc. DAC, pp. 822-825, July 2009.
- [9] P. Zhou, P. Yuh, S. Sapatnekar, "Application-specific 3D network-on-chip design using simulated allocation," Proc. ASPDAC, pp. 517-522, Jan. 2010.
- [10] K. Srinivasan, K. Chatha, "A low complexity heuristic for design of custom network-on-chip architectures," Proc. DATE, pp. 130-135, March 2006.
- [11] K. Srinivasan, K. Chatha, G. Konjevod, "Linear-programming-based techniques for synthesis of network-on-chip architectures," IEEE TVLSI 14(4), pp. 407-420, April 2006.
- [12] U. Ogras, R. Marculescu, "It's a small world after all": NoC performance optimization via long-range link insertion," IEEE TVLSI 14(7), pp. 693-706, July 2006.
- [13] L. Leung, C. Tsui, "Energy-aware synthesis of networks-on-chip implemented with voltage islands," Proc. DAC, pp. 128-131, June 2007.
- [14] P. Ghosh, A. Sen, "Efficient mapping and voltage islanding technique for energy minimization in NoC under design constraints," Proc. SAC, pp. 535-541, March 2010.
- [15] U. Ogras, R. Marculescu, P. Choudhary, D. Marculescu, "Voltage-frequency island partitioning for GALS-based networks-on-chip", Proc. DAC, pp. 110-115, June 2007.
- [16] J. Hu, R. Marculescu, "Communication and task scheduling of application-specific networks-on-chip," IEE CDT 152(5), pp. 643-651, Sep. 2005.
- [17] C. Chou, U. Ogras, R. Marculescu, "Energy and performance-aware incremental mapping for networks on chip with multiple voltage levels," IEEE TCAD 27(10), pp. 1866-1879, Oct. 2008.
- [18] J. Hu, R. Marculescu, "Energy and performance-aware mapping for regular NoC architectures," IEEE TCAD 24(4), pp. 551-562, 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, Y. Xie, Y. Chang, Y. Wang, "Three-dimensional integrated circuits (3D IC) floorplan and power/ground network co-synthesis," Proc. ASPDAC, pp. 169-174, Jan. 2010.
- [21] H. Chen, H. Lin, Z. Wang, T. Hwang, "A new architecture for power network in 3D IC," Proc. DATE, pp. 401-406, March 2011.
- [22] D. Starobinski et al., "Application of network calculus to general topologies using turn-prohibition," Trans. on Networking, 11, 3, (June 2003), 411-421.
- [23] S. Kose, E. Friedman, "Fast Algorithms for power grid analysis based on effective resistance," ISCAS, 2010.
- [24] N. Khan, S. Alam, S. Hassoun, "Power delivery design for 3-D ICs using different through-silicon via (TSV) technologies," TVLSI, Apr. 2011.
- [25] lp_solve 5.5, <http://lpsolve.sourceforge.net/5.5/>
- [26] ARM processor selector, <http://www.arm.com/products/processors/selector.php>
- [27] Standard Task Graph Set, <http://www.kasahara.elec.waseda.ac.jp/schedule/>
- [28] S. Yan, B. Lin, "Design of application-specific 3D networks-on-chip architectures," in Proc. ICCD, pp. 142-149, Oct. 2008.
- [29] C. Seiculescu, S. Murali, L. Benini, G. De Micheli, "SunFloor 3D: A tool for networks on chip topology synthesis for 3D systems on chips," Proc. DATE, pp. 9-14, April 2009.
- [30] K. Siozios, I. Anagnostopoulos, D. Soudris, "A high-level mapping algorithm targeting 3D NoC architectures with multiple Vdd," Proc. ISVLSI, pp. 444-445, July 2010.
- [31] C. Seiculescu, S. Murali, L. Benini, G. De Micheli, "Comparative analysis of NoCs for two-dimensional versus three-dimensional SoCs supporting multiple voltage and frequency islands," IEEE TCAS 57(5), pp.364-368, May 2010.
- [32] M. Arjomand, H. Sarbazi-Azad, "Voltage-frequency planning for thermal-aware, low-power design of regular 3-D NoCs," Proc. VLSID, pp. 57-62, Jan. 2010.
- [33] Tiler Corporation, TILE64™ Processor. Product Brief. 2007.
- [34] Intel SCC, <http://techresearch.intel.com/>, 2010.
- [35] Intel TeraFlops Processor, <http://techresearch.intel.com/2007>.
- [36] T. Chelcea, S. Nowick, "A low-latency for mixed-clock systems," Proc. CSW, pp. 119-126, Apr. 2002.
- [37] S. Pasricha, "A framework for TSV serialization-aware synthesis of application specific 3D networks-on-chip", Proc. VLSID, Jan. 2012.
- [38] A. Kahng, B. Li, L.-S. Peh, K. Samadi, "ORION 2.0: A fast and accurate NoC power and area model for early-stage design space exploration," Proc. DATE, pp. 423-428, Apr. 2009.