

System-level optimization of Network-on-Chips for heterogeneous 3D System-on-Chips

Jan Moritz Joseph*, Dominik Ermel*, Lennart Bamberg†, Alberto García-Oritz†, Thilo Pionteck*

*Otto-von-Guericke-Universität Magdeburg, Institut für Informations- und Kommunikationstechnik, Germany

Email: {jan.joseph, dominik.ermel, thilo.pionteck}@ovgu.de

†University of Bremen Institute of Electrodynamics and Microelectronics, Germany

Email: {agarcia, bamberg}@item.uni-bremen.de

Abstract—For a system-level design of Networks-on-Chip for 3D heterogeneous System-on-Chip (SoC), the locations of components, routers and vertical links are determined from an application model and technology parameters. In conventional methods, the two inputs are accounted for separately; here, we define an integrated problem that considers both application model and technology parameters. We show that this problem does not allow for exact solution in reasonable time, as common for many design problems. Therefore, we contribute a heuristic by proposing design steps, which are based on separation of intralayer and interlayer communication. The advantage is that this new problem can be solved with well-known methods. We use 3D Vision SoC case studies to quantify the advantages and the practical usability of the proposed optimization approach. We achieve up to 18.8% reduced white space and up to 12.4% better network performance in comparison to conventional approaches.

I. INTRODUCTION

In heterogeneous 3D integrated System-on-Chips (SoCs), dies in disparate technologies (e.g. mixed-signal and digital nodes) are stacked and vertically connected. This is promising, as components with different requirements to technology can be integrated efficiently. Heterogeneity gains attention in the industry, e.g. Intel's recent architecture "Lakefield" [?].

Communication in these SoCs can be realized via Networks-on-chip (NoCs). Heterogeneity has a vast influence on the NoC design, as Ref. [2] shows at the *physical level* and Ref. [3] at the *architectural level*. At the *system level*, the locations of components (i. e. PEs), routers and vertical links, as well as the network topology are optimized. Many aspects of the system-level optimization have been solved (e.g. floorplaning [4] or NoC topology synthesis [5]), but their integrated solution has not been considered sufficiently so far. In heterogeneous 3D SoC the necessity for an integrated approach arises from the varying components' and routers' costs between the layers plus the mutual influence between the vertical interconnect planning and the router / component placement.

This paper targets such an integrated approach. Three fundamental characteristics of heterogeneity are identified for the system-level optimization (Sec. II). Next, *the system-level optimization problem is defined* (Sec. II). As for many design problems, it is impossible to calculate an exact solution in reasonable time. Thus, *we propose five design steps that efficiently generate a solution* by separation of the intralayer and the interlayer communication (Sec. IV). Each individual step is elementary to enable efficient solving; existing methods can




Fig. 1: Router model with KOZ and RD.

be used with small modifications considering the fundamental characteristics. The heuristic is applied to heterogeneous 3D SoC case studies to quantify its advantages (Sec. V).

II. THE INTEGRATION ISSUES

The following three fundamental characteristics of heterogeneous 3D integration influence the system-level optimization:

1) Routers and components have varying performance, power and area (PPA) per technology node / per layer. Some components cannot be implemented in certain nodes due to physical limitations, e.g. analog sensors in digital technologies.

2) Through-Silicon Vias (TSV) arrays are used as vertical interconnects. TSVs are commonly manufactured in the via-middle process flow and thus yield keep-out-zones (KOZs) when crossing layers. Therefore, area must be reserved.

3) It is beneficial to use redistribution (RD) that connects the TSV arrays with the routers via horizontal metal wires. This allows the routers to be vertically connected without being exactly above each other. The length of the RD is limited by the target clock frequency and the technology.

These characteristics influence the NoC: First, a partially vertically connected mesh topology is more area-efficient than a fully-connected NoC. Second, the KOZs and the RD must be modeled. If a router connects downwards, there will be a KOZ (Fig. 1a). If it connects upwards, there is none (Fig. 1b). If RD is used, the KOZ will be outside of the router (Fig. 1c).

III. THE PROBLEM FORMULATION

We define the system-level optimization of NoCs considering heterogeneous integration. Its input is divided into application and technology properties, shown in Fig. 2, upper left-hand side: The *application* is described by a core graph [6] in that the nodes are the components and the edge weights model the bandwidth between the components. The *technology* properties of the components and routers are given by PPA tables. At the example of area, the table gives $area_{c,l}$, in which




Fig. 2: Problem definitions with input and output & steps of the heuristic (green: optimal solution. blue: heuristic solution).

$c \in C$ denotes a component c in the set of all components C and $l \in L$ one layer in the set of all layers. Let $C(l)$ be the subset of components in the layer l . Plus, the KOZ area K and the RD length R are given.

The result of the optimization is a network graph, in that the nodes are routers and the edges are links. Plus, rectangular bounding boxes reserve space, each for one component and one 2D router or one 3D router with a KOZ. The bounding boxes follow the mesh topology. Fig. 2 (left-hand side) shows an example with bounding boxes and a network graph.

The objective is linear with five weights for scaling:

$$C = w_1 C_{area} + w_2 C_{power} + w_3 C_{perf} + w_4 C_{peak} + w_5 C_{util}$$

The first three addends are the chip area C_{area} , the summed (component) performance C_{perf} , and the summed power consumption C_{power} . This models PPA, but does not account for the network performance. Thus, the fourth addend C_{peak} models the throughput by penalizing congested links with a higher load than available bandwidth. The fifth addend C_{util} models the latency, measured in hop distance \times bandwidth. This linear model allows for extensions, e.g. the standard methods for thermal optimization (e.g. [4], [5]) are also linear.

IV. THE HEURISTIC: FIVE DESIGN STEPS

We define a heuristic for the system-level optimization by the identification of five design steps. Efficient run-time is possible due to separation of interlayer and intralayer communication. The steps are shown in Fig. 2, right-hand side.

Step ①: The **component-to-layer assignment** minimizes the weighted sum of the maximum chip area across layers (which dominates the area of the 3D stack) and the total power: $C_1 = w_1 C_{area} + w_2 C_{power} = w_1 \max_{l \in L} \sum_{c \in C(l)} area_{c,l} + w_2 \sum_{l \in L, c \in C} power_{c,l}$. The problem is formulated as an integer LP (ILP). Since the max-function is minimized, it can be modeled with an auxiliary variable z that is constraint by the max-function's inputs ($z \geq \max(a, b) \rightarrow z \geq a, z \geq b$). The ILP is constraint such that each component is assigned to one layer only. The layer order is not optimized as this would require to consider the communication as early as in this step.

Step ②: The **component floor planning** minimizes the layer area and the interlayer communication. This step can be done individually per layer, as routers must not be placed at the same position due to RD. Per layer, $C_2 = w_1 C_{area} +$

$w_4 C_{peak} + w_5 C_{util}$ is minimized with a simulated annealing (SA). It optimizes the floorplan in a 2D mesh topology by switching the components' assignment to mesh-columns and rows. Here, the component area is not simply added as in step ①; rather, the complete layer's area is minimized for each floorplan in the SA by minimizing the width $\omega_{i,j}$ and length $\lambda_{i,j}$ for all components at columns i and rows j in the 2D mesh. This optimization is constraint by the component areas $area_{c,l}$ at column/row i, j : $\omega_{i,j} \lambda_{i,j} \geq area_{c,l}$ for all i, j . The KOZ and 3D router area cannot be accounted for, at this time. For the area minimization within the SA, a linear program (LP) and a semi-definite program (SDP) can be used. In the LP, the area must be approximated as the multiplication $\omega_{i,j} \lambda_{i,j}$ is non-linear, while the SDP is exact, as it is possible to model $\omega_{i,j} \lambda_{i,j}$. For the equations of both LP and SDP as well as a performance analysis, see Ref. [7].

Step ③: The **TSV array count** minimizes the area and maximizes the interlayer communication, as more TSVs reduce link load but require more KOZs. The exact routing, and thus the communication, is unknown; hence, it is given by the expected bandwidth b_j for the j -th TSV array. This is calculated by uniform-randomly distributing TSV arrays to the grid and use the floor plan to find the nearest-neighbored components with the Manhattan distance d_j and sum the component's upward and downward bandwidth. Thus, the objective for i TSVs is $C_3 = iK + \sum_{j=1}^i b_j d_j$. As the TSV count is limited, exhaustive search finds the global minimum.

Step ④: The **placement of 3D routers** minimizes the communication for a given TSV count. For the objective C_{util} (hop distance \times bandwidth) the routing algorithm is used. A penalty is added for links with more load than bandwidth available (C_{peak}). Routers in adjacent layers with a distance smaller than the RD-length R can be connected. The set of connections is chosen via a SA that randomly switches them.

Step ⑤: The solution is **legalized** to accommodate the area for added 3D routers and TSV arrays. The LP/SDP from the second step are used again.

The order of the five design steps is motivated as follows: For fast optimization, the step ① splits the design in layers before floor-planning each layer in the step ②. Deciding the step ③ as late as possible is beneficial for network design as then floor plans can be accounted for. This avoids overalloca-

TABLE I: Case Study PPA table.

Nodes:	Area [mm ²]		Perf. [relative]		Power [relative]	
	28 nm	45 nm	28 nm	45 nm	28 nm	45 nm
CPU (RISC-V)	35.8	62.2	1	1.34	1	1.34
ADC [9]	n.a.	53	n.a.	1	n.a.	1
SIMD (nu+) [10]	71	125	1	1.34	1	1.34
2D Router	1.3	2.25	1	1.34	1	1.34
3D Router	1.8	3.15	1	1.34	1	1.34

KOZs: 2 mm². Maximum length of RD: 5 mm.

tion of network resources. The steps ③ and ④ are separated to allow for comparison against standard approaches. Both the steps are interdependent and their separation is far from trivial; we opt for this to balance performance and accuracy. If desired, heuristic can be iterated for further improvement.

V. RESULTS

The heuristic is implemented in MATLAB. The LPs are solved with CPLEX 12.8.0 and the SDPs with Mosek 8.1. The source code is available at <https://github.com/jmjos/A-3D-NoC-DSE>. For the evaluation we use three case studies, based on a 3D Vision SoC (VSoC) [8], a typical application for heterogeneity. The PPA for 45nm mixed-signal and the 28nm digital nodes are given in Tab. I. ADCs must be implemented in a mixed-signal node. For the sake of generality, we only consider the relative differences of the power (measured in static and dynamic power) and the performance (measured in timing) generated from synthesis. The routing algorithm uses the shortest path. We assume face-to-back bonding in all cases.

Case study I: The **tiny 3D SoC** has 5 CPUs and two 28 nm digital layers. The application graph has bidirectional links with 1 Mb/s bandwidth between subsequent nodes.

Case study II: The **small 3D VSoC** has one 28 nm digital layer one 45 nm mixed-signal layer. It implements 9 analog-digital converters (ADCs) and 9 CPUs. It captures a 720p60 video stream, conducts AD-conversion and processes a convolution filter. In a conventional design, the ADCs are located in a 3×3 mesh NoC in the mixed-signal layer and the CPUs in a 3×3 mesh NoC in the digital layer. As RD is not used conventionally, the mesh sized in both layers are identical and routers are located at the same positions.

Case study III: The **large 3D VSoC** implements 9 ADCs, 18 CPUs and 3 SIMD cores. The chip has one mixed-signal layer in 45 nm node and two digital layers in 28 nm node. The VSoC runs Viola-Jones, Shi and Tomasi and KLT algorithms for face recognition and tracking from a 720p60 video. In a conventional design, the ADCs are located in a 3×3 mesh NoC in the mixed-signal layer, the CPUs and SIMD cores in a 3×3 mesh NoC in the first digital layer and the remaining CPUs are in a 4×3 NoC below.

A. Comparison to an optimal solution

The **tiny SoC** provides a small example that can be solved optimally using an analytical mixed-integer linear model (MILP) to compare against the heuristic.

We execute the heuristic using both the LP and the SDP to optimize area. The LP allows to compare against the MILP, while the SDP removes the linearization error. The results are shown in Tab. II. Considering *runtime*, the MILP is naturally

TABLE II: Optimal MILP vs. Heuristic with LP and SDP.

Method: tiny 3D SoC	Optimal MILP	LP	Heuristic		
			Δ	SDP	Δ
Optimization Runtime [s]	599	27.67	-95%	152.4	-75%
Area [mm ²] – upper layer	158.0	158.0	0%	117.0	-36%
Area [mm ²] – lower layer	77.4	77.4	0%	77.0	-0.5%
Bandwidth×Distance[mm ² Mb/s]	18.51	31.08	+41%	26.53	+15%

Parameters: Step 2: initial temp. 20, 120 iters, 0.97 cooling; Step 4: initial temp. 100, 50 iters, 0.97 cooling; Weights ω_i set to 1.

TABLE III: Execution time of heuristic for random inputs.

HEURISTIC'S PART execution times	5 cmp. 2 layers	40 cmp. 4 layers	80 cmp. 4 layers	1000 cmp. 3 layers
① Comp. to layer assignment	0.4 s	0.4 s	0.4 s	0.4 s
② Layer floor planning	146 s	375 s	802 s	0.3 h
③ Number of TSV arrays	0.1 s	0.2 s	0.2 s	5 min
④ Placement of TSV arrays	3.8 s	298 s	1300 s	30 h
⑤ Legalization	0.2 s	0.5 s	0.5 s	0.5 s
COMPLETE	152 s	675 s	2104 s	31 h

Parameters: Step 2: initial temp. 20, 120 iters, 0.97 cooling; Step 4: initial temp. 100, 50 iters, 0.97 cooling; Weights ω_i set to 1. Intel i7-7740X, Ubuntu 16.04 LTS.

very slow; it finds a solution within ~9.5 min. The heuristic is much faster. Its efficient runtime with a linear model allows to use an SDP to remove the linearization error for area. The heuristic with SDP is still 75% faster than the MILP. Considering *area*, the MILP and heuristic with LP yield the same results as the LP used for area optimization in the SA in step ② uses a subset of the MILP's constraints. Thus, area is identical for the same floorplan, as given in this simple case. The SDP improves area by up to 36%. Considering *network performance*, the MILP has the best result as it does not separate interlayer and intralayer communication. The SDP outperforms the LP as a side-effect of better area. Regarding *power and (component) performance*, the results of MILP and heuristic are equivalent, since both use the same LP.

B. Heuristic run-time performance

Tab. III gives the exemplary runtime performance of the heuristic with SDP. A large realistic input set with 80 components and 4 layers finishes in 35 min. The largest set suffers from the brute-force approach in step ③, but a method with higher performance can be found easily. The implementation of all steps is reasonably fast to optimize even very large inputs sets, so the runtime of the heuristic is adequate.

C. Advantages of optimization with redistribution

RD allows connecting routers vertically that are not located exactly above each other. This enables more flexible vertical interconnections and thus increases the interconnection

TABLE IV: Effect of RD on communication.

Small 3D VSoC MAXIMUM LENGTH OF RD	USED RD		BANDWIDTH×DISTANCE IN [mm Mb/s]
	LENGTH [mm]		
0 mm (0%) (BASELINE)	0	47.78	
90 mm (25%)	39.47	44.57	-6.72%
180 mm (50%)	95.54	42.84	-10.34%
360 mm (100%)	354.36	41.85	-12.41%
Large 3D VSoC			
0 mm (0%) (BASELINE)	0	46.49	
90 mm (50%)	78.42	43.04	-7.42%
180 mm (100%)	131.39	40.48	-12.93%

Parameters: Step 2: initial temp. 30, 200 iters, 0.98 cooling; Step 4: initial temp. 1000, 50 iters, 0.97 cooling; Weights ω_i set to 1, fixed number of TSVs.

TABLE V: Conventional vs. optimized design.

Small 3D VSoC	BASELINE	INTEGRATED	DIFFERENCE
BANDWIDTH \times DISTANCE	2390 mm Mb/s	2830 m Mb/s	+18.41%
MAXIMUM LINK LOAD	120 Mb/s	190 Mb/s	+58.33%
WHITESPACE	30.48 mm ²	25.77 mm ²	-15.44%
Large 3D VSoC	BASELINE	INTEGRATED	DIFFERENCE
BANDWIDTH \times DISTANCE	2118 mm Mb/s	2599 mm Mb/s	+22.68%
MAXIMUM LINK LOAD	115.8 Mb/s	149.1 Mb/s	+28.77%
WHITESPACE	42.92 mm ²	34.86 mm ²	-18.79%

Parameters: Step 2 and 4: initial temp. 30, 500 iters, 0.97 cooling

TABLE VI: Optimized network loads.

Bandwidth \times distance [mm Mb/s]	BASELINE	INTEGRATED	Δ
SMALL VSoC	10.46	10.46	$\pm 0.00\%$
LARGE VSoC	39.74	37.99	-4.40%

Parameters: Step 2: initial temp. 30, 200 iters, 0.97 cooling, step 4: initial temp. 1000, 50 iters, 0.97 cooling; Weights ω_i set to 1, fixed number of TSVs.

efficiency. For our experiments, the length of the RD is calculated from the vendor models and is reduced gradually to demonstrate the positive effect of RD on the application communication's hop distance. The RD decreases the hop distance by up to 12.93%. For this improvement, the integrated approach is essential because both the floorplan and the horizontal NoC topology are required to find optimized vertical links for the application. Summarizing, RD improves the communication by $\sim 13\%$ over a baseline without RD.

D. Advantages of the integrated approach

1) *Whitespace reductions*: The conventional design of the case studies is packaging-inefficient because of the size difference between ACDs, CPUs and SIMD cores. This yields whitespace. It is reduced by our integrated approach due to a higher degree of freedom in the placement of the components and the routers. The area and the network performance of conventional and optimized designs are given in Tab. V. We achieve 15.44% and 18.79% reductions in whitespace for the small and the large 3D VSoC, respectively. The algorithm does not consider interlayer communication and thus the communication is worse by 18-58%. The results show the typical limitations of the conventional approach without RD and grid size variability. However, these features reduces the whitespace but, naturally, have a negative effect on the network performance. Whether this is an acceptable compromise depends on the design targets.

2) *Communication improvements*: To show the advantages of the integrated approach, we compare it against a solution without application information. Therefore, we use the cases' traffic vs. uniform random traffic, which does not provide application-specific traffic patterns. Tab. VI gives the results. We do not see an improvement for the small VSoC, because the application's traffic is almost uniformly distributed. The communication is improved by 4.4% for the large VSoC.

E. Comparison against related work

1) *Component floor planning*: We compare our floor plan against [6], as this work also accounts for different core sizes. We show different multimedia benchmarks in Tab. VII. For the first two benchmarks, we improve the area at costs of the communication. However, for the H263enc Mp3 dec

TABLE VII: Comparison vs. Ref. [6].

Benchmark	Area [mm ²]			BW \times Dist. [mm ² Mb/s]		
	[6]	SA	Δ	[6]	SA	Δ
H256decMp3dec	11301	8244	-27.1%	19858	21280	+7.16%
mp3decMp3dec	8568	8516	-0.61%	17546	17572	+0.15%
H263encMp3dec	12535	10474	-16.4%	255324	250187	-2%

Parameters: initial temp. 30, 15000 iters, 0.98 cooling; average over 30 reruns.




Fig. 3: SA vs. PSO [11] for third step.

benchmark we see a general improvement: We achieve up to 16.4% better area and 2% better communication using the integrated approach.

2) *Placement of 3D routers*: Using SA within our framework is justified by comparison against partial swarm optimization (PSO) [11] to determine the optimal number of TSV arrays. The difference in bandwidth \times distance is shown in Fig. 3. In average, SA is 3.125% and 2.563% better than PSO for Video Object Plane Detection (VOPD) and Double VOPD benchmarks, respectively. For 2 TSV arrays, SA is even 15% better than PSO for the VOPD benchmark. These positive results make SA a reasonable method.

VI. SUMMARY

The system-level optimization improves the floor plan and the network topology of NoCs for heterogeneous 3D SoCs. Therefore, the proposed integrated approach considers properties of both the application and the technology. We propose a heuristic with five design steps for an efficient optimization that splits interlayer and intralayer communication. Whitespace is reduced by up to 18.8%. Communication is better by up to 4.4% through early consideration of traffic patterns.

ACKNOWLEDGMENTS

This work is funded by DFG projects PI 447/8 and GA 763/7.

REFERENCES

- [1] "Intel Previews New Hybrid CPU Architecture with Foveros 3D Packaging," <https://newsroom.intel.com/video-archive/video-intel-previews-new-hybrid-cpu-architecture-with-foveros-3d-packaging/>, accessed: 2019-05-17.
- [2] Bamberg, L. et al., "Coding-aware Link Energy Estimation for 2D and 3D Networks-on-Chip with Virtual Channels," *PATMOS*, 2018.
- [3] Joseph, J.M. et al., "Area and power savings via asymmetric organization of buffers in 3D-NoCs for heterogeneous 3D-SoCs," *MICPRO*, 2017.
- [4] Cong, J. et al., "A thermal-driven floorplanning algorithm for 3D ICs," in *ICCAD*, 2004.
- [5] Seiculescu, C. et al., "SunFloor 3D: A Tool for Networks on Chip Topology Synthesis for 3-D Systems on Chips," *IEEE TCAD*, 2010.
- [6] Srinivasan, K. et al., "Linear-programming-based techniques for synthesis of network-on-chip architectures," *IEEE TVLSI*, vol. 14, 2006.
- [7] Joseph, J.M. et al., "Area optimization with non-linear models in core mapping for system-on-chips," in *MOCAS*, 2019.
- [8] Á. Zarándy, *Focal-plane sensor-processor chips*. Springer, 2011.
- [9] Lyu, T. et al., "A 12-Bit High-Speed Column-Parallel Two-Step Single-Slope ADC for CMOS Image Sensors," *Sensors*, 2014.
- [10] Flich, J. et al., "Exploring manycore architectures for next-generation HPC systems through the MANGO approach," *MICPRO*, 2018.
- [11] Manna, K. et al., "Integrated Through-Silicon Via Placement and Application Mapping for 3D Mesh-Based NoC Design," *ACM TECS*, 2016.