

# Sparse 3-D NoCs with Inductive Coupling

Michihiro Koibuchi  
National Institute of Informatics,  
JAPAN  
koibuchi@nii.ac.jp

Lambert Leong  
University of Hawai'i at Manoa, USA  
lambert3@hawaii.edu

Tomohiro Totoki  
Keio University, JAPAN  
totoki@am.ics.keio.ac.jp

Naoya Niwa  
Keio University, JAPAN  
naoya@am.ics.keio.ac.jp

Hiroki Matsutani  
Keio University, JAPAN  
matutani@arc.ics.keio.ac.jp

Hideharu Amano  
Keio University, JAPAN  
hunga@am.ics.keio.ac.jp

Henri Casanova  
University of Hawai'i at Manoa, USA  
henric@hawaii.edu

## ABSTRACT

Wireless interconnects based on inductive coupling technology are compelling propositions for designing 3-D integrated chips. This work addresses the heat dissipation problem on such systems. Although effective cooling technologies have been proposed for systems designed based on Through Silicon Via (TSV), their application to systems that use inductive coupling is problematic because of increased wireless-communication distance. For this reason, we propose two methods for designing *sparse* 3-D chips layouts and Networks on Chip (NoCs) based on inductive coupling. The first method computes an optimized 3-D chip layout and then generates a randomized network topology for this layout. The second method uses a standard stack chip layout with a standard network topology as a starting point, and then deterministically transforms it into either a “staircase” or a “checkerboard” layout. We quantitatively compare the designs produced by these two methods in terms of network and application performance. Our main finding is that the first method produces designs that ultimately lead to higher parallel application performance, as demonstrated for nine OpenMP applications in the NAS Parallel Benchmarks.

## CCS CONCEPTS

• **Hardware** → **Temperature optimization; Network on chip; 3D integrated circuits.**

## KEYWORDS

Inductive coupling, 3-D chip layout

## ACM Reference Format:

Michihiro Koibuchi, Lambert Leong, Tomohiro Totoki, Naoya Niwa, Hiroki Matsutani, Hideharu Amano, and Henri Casanova. 2019. Sparse 3-D

NoCs with Inductive Coupling. In *The 56th Annual Design Automation Conference 2019 (DAC '19)*, June 2–6, 2019, Las Vegas, NV, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3316781.3317913>

## 1 INTRODUCTION

Although 3-D integration of chips based on Through Silicon Via (TSV) is becoming mainstream, an alternate and attractive approach is to use inductive coupling technology. Inductive coupling, with the ThruChip Interface (TCI), offers high speed (e.g., 8 Tbps [10]), low power (e.g., 0.14 pJ/b [9]) due in part to the removal of electrostatic discharge (ESD) protection devices, high integration (e.g., a 128-die NAND stack [16]), and applicability to an increasing range of computing systems including a reconfigurable processor [11]. Most recently, a high-end chip multiprocessor (CMP) stacked with 4 DRAM chips for an aggregate 2 TiB/s RAM bandwidth using inductive coupling is planned for use in the custom supercomputer that was ranked in 4th position of the November 2017 Top500 list. This system will be the first to use commercially available CMPs designed based on inductive coupling.

3-D integrated chips face a heat-dissipation problem due to high chip power density, thus motivating the development of innovative cooling technologies. Microchannel (liquid) cooling and the use of graphite sheets have been well studied for 3-D stacking using TSV. Unfortunately, their application to systems based on inductive coupling is problematic. The main reason is that these cooling technologies increase the vertical distance between pairs of communication chips, which in the case of TCI would require larger induction coils, which in turns would reduce bandwidth density. For instance for an inter-chip distance of 20  $\mu\text{m}$ , resp. 80  $\mu\text{m}$ , the necessary coil diameter is 60  $\mu\text{m}$ , resp. 240  $\mu\text{m}$ , resulting in a 80 Gbps, resp. 22 Gbps, per coil bandwidth [5]. As a result, a 4x increase in inter-chip distance leads to a 58x decrease in bandwidth density.

Another cooling technology is submerged liquid cooling, using liquids that provide both electric insulation and high heat transfer efficiency, such as mineral oil or fluorinert. Unfortunately, even with these options, it is not feasible to operate 3-D stack layouts of high-power chips. For instance, consider a 3-D stack layout of a model of the Xeon E5-2667v4 (up to 3.6 GHz) chip, using a temperature constraint of 78°C (as specified in the processor's data sheet). Using HotSpot v6.0 [15], we find that even with submerged liquid cooling this layout cannot operate more than 2 stacked chips


Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*DAC '19*, June 2–6, 2019, Las Vegas, NV, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6725-7/19/06...\$15.00

<https://doi.org/10.1145/3316781.3317913>



**Figure 1: Connecting partially overlapping chips using inductive coupling.**

at their highest frequency. The scalability of 3-D stack layouts is thus inherently limited due to heat dissipation. Note that with TCI, the heat dissipation problem would be further exacerbated because of the need to insert glue between chips (to hold the 3-D structure), thus preventing the cooling liquid from circulating over chip surfaces.

TCI makes it possible to arrange arbitrary numbers of chips in a package after chips have been fabricated. Importantly, a wide range of potential 3-D physical chip layouts are possible. Due to high tolerance to coil misalignment [3, 13], two chips can be connected by partially overlapping them in arbitrary fashion, as shown in Figure 1. It is thus possible to design *sparse* 3-D chip layouts so as to increase heat dissipation, thereby accommodating possibly large numbers of high-frequency chips. Note that spacers (not shown in the figure) are inserted below the upper-layer (gray) chips, excluding areas in which they overlap with bottom-layer (white) chips. Glue is inserted in the overlap areas only. The number of inductors in the overlap area determines the inter-chip communication bandwidth.

Higher inter-chip bandwidth, which implies larger overlap area between connected chips, leads to a lower cooling surface area. So do lower hop counts, which imply a more tightly packed 3-D physical layout of chips. There is thus a trade-off between interconnect performance and heat dissipation. In this work we explore this trade-off, the main objective being to perform 3-D integration of chips using TCI so as to maximize parallel application performance. We make the following contributions:

- 1) We propose a method, NoCAL (NoC-after-Layout), which uses a greedy heuristic to compute a 3-D chip layout that minimizes inter-chip path lengths while respecting thermal constraints. This method then generates a randomized network topology with low average shortest path length (ASPL) while capping the length of physical on-chip links. The produced designs should achieve a good trade-off between interconnect performance and heat dissipation.
- 2) We propose another method, LANoC (Layout-after-NoC), which uses as a starting point a traditional 3-D physical stack layout with standard network topologies, such as meshes, for which a good physical layout is known, thus insuring short physical link length. This method then deterministically transforms the layout into either a “staircase” layout or a “checkerboard” layout. The produced designs should be beneficial to parallel applications in practice because a standard, regular network topology is used.
- 3) We evaluate and compare NoCAL and LANoC using 4-core Xeon-like chips as a case study. We find that both methods can produce designs that accommodate large numbers of high-frequency chips with good interconnect performance while respecting thermal

constraints. We also find that NoCAL is superior to LANoC: it produces network topologies with lower hop counts and affords higher parallel application performance, as demonstrated for nine OpenMP applications in the NAS Parallel Benchmarks.

The remainder of this paper is organized as follows. Section 2 presents background information on TCI. Section 3 describes the NoCAL and LANoC methods. Section 4 presents simulation results. Section 5 reviews related work. Section 6 concludes with a summary of findings and perspectives on future work.

## 2 INDUCTIVE COUPLING THROUGHCHIP INTERFACE

Several interconnection technologies have been developed for 3-D integrated chips, including wire-bonding, micro-bump, and through-silicon via (TSV). These technologies make it possible to integrate chips in 3-D rigid stacks. More recently, wireless capacitive- and inductive-coupling technologies have also been developed. In particular, inductive-coupling using the TruChip Interface (TCI) operates by generating a magnetic field between a pair of transmitter and receiver coils (i.e., inductors) for data transfer.

A possible concern with TCI is the impact of coil misalignment. Experiments with commercial products show that coil misalignment up to 40% of the coil diameter is tolerable in normal operating conditions [13]. Furthermore, increasing power improves coil misalignment tolerance. Increasing energy consumption by only 4.1%, when compared to perfectly aligned coils, leads to as much as 50% coil misalignment tolerance [3]. As a result, instead of fabricating rigid monolithic chip stacks in which chips are perfectly overlapped (e.g., using TSV), TCI makes it possible to arrange chips, once they have been fabricated, in arbitrary 3-D physical layouts in which chips only partially overlap (and in which the unavoidable misalignment of chips can be tolerated). This dramatic increase in 3-D layout design space is the main motivation for this work. One difficulty for constructing 3-D layouts with TCI is that one must avoid “crosstalk”, i.e., interference between wireless vertical links that are physically located directly above each other. Crosstalk can be avoided altogether by doubling the footprint of the induction area. In this work, instead, we construct 3-D layouts with the constraint that no two vertical links can be placed directly above each other.

Another attractive feature of TCI is its power consumption. The energy consumption of a TCI transceiver (which includes serializer, deserializer, differential conversion circuits, and a coil) is proportional to  $IV/t$ , where  $I$  is the current,  $V$  is the voltage, and  $t$  is the data rate per channel [4]. When the device size is scaled down by a factor  $\alpha$ ,  $IV$  is reduced by a factor  $1/\alpha^2$  and  $t$  is increased by a factor  $\alpha$ ; thus, the energy consumption is decreased by a factor  $1/\alpha^3$ . A 1pJ/bit inductive-coupling interface implemented with a Fujitsu 65nm CMOS process is considered in [4]. In this work we assume that chip thickness is  $15\mu\text{m}$  and that the vertical inter-chip gap is  $20\mu\text{m}$ . Using a  $225\mu\text{m}$  coil, the energy-per-bit is then only slightly higher than 1pJ/bit, which is in line with the assumptions in [4].

While TSV-based 3-D designs have become mainstream, the above features make TCI-based designs compelling competitors for 3-D integration of chips.

### 3 LAYOUT GENERATION METHODS

#### 3.1 NoC-After-Layout (NoCAL)

This approach constructs a 3D chip layout that optimizes the inter-chip network topology under thermal constraints (Section 3.1.1), and then generates a network topology for this layout (Section 3.1.2).

**3.1.1 Layout Generation.** The goal is to solve the following constrained multi-objective optimization problem. Given a chip with known physical dimensions, operating frequency and power profile, given  $n$  such chips, and given a required overlap fraction  $\omega$  between a pair of connected chips to achieve some minimum bandwidth, the objective is to arrange the chips in a 3-D physical layout in which the path lengths in the inter-chip topology (in which chips are vertices and TCI links are edges) are minimized while remaining below a given temperature threshold.

The typical approach for computing good solutions for such optimization problems is to rely on meta-heuristics that evaluate large numbers of candidate solutions. In our problem, the temperature constraint is not defined analytically but is evaluated empirically: the maximum temperature of a candidate layout is computed via a HotSpot [15] simulation, which implements numerical models derived from thermodynamics principles. Unfortunately, HotSpot invocations are compute intensive due to a large sparse LU factorization operation, implemented with the (sequential) SuperLU library. We have explored several options to alleviate this computational bottleneck. Although a multi-threaded SuperLU version exists, using it in HotSpot would require a substantial software re-engineering effort beyond the scope of this work (due to the need to change the sparse matrix storage formats and memory layouts). We have attempted to reduce the resolution of the numerical methods in HotSpot so as to reduce computational complexity (while retaining good output accuracy), but performance improvements were marginal. Ultimately, we simply invoke HotSpot concurrently for different candidate layouts. However, this approach has poor parallel efficiency on a multi-core machine due to competition for the lowest-level cache. This limits the number of concurrent invocations that can be used effectively on a single machine. We conclude that, for now, we must construct layouts using relatively few numbers of HotSpot invocations, which precludes the use of sophisticated meta-heuristics.

We generate layouts using a simple greedy heuristic that first constructs the 3-chip layout shown in Figure 1(a). The heuristic then incrementally adds chips to the layout. It randomly picks an existing chip in the current layout to which a new chip could be connected. It then generates up to  $r$  randomly generated candidate positions for this new chip so that (i) the candidate chip and the existing chip are in adjacent layers (i.e., one is in the layer directly above that of the other); (ii) the percentage of overlap between the existing chip and the candidate chip is at least  $\omega$ ; (iii) the candidate chip does not collide with other existing chips in the layout; and (iv) no TCI crosstalk occurs. If no candidate position has been generated after some configurable number of attempts, the heuristic fails. Otherwise, each candidate chip is putatively added to the current layout to generate a candidate new layout. HotSpot is then invoked to reject candidate layouts that would exceed the temperature threshold. If all candidate layouts are rejected the heuristics fails. Otherwise, it sorts the feasible candidate layouts lexicographically by diameter,

ASPL, and number of edges of their inter-chip network topologies. The heuristic picks the best candidate and repeats until it generates a layout with  $n$  chips. The larger the number of candidates,  $r$ , the better the results produced by the heuristic, but the longer the layout generation time. In all our results, we use  $r = 34$ , which, for example, makes it possible to generate a layout with 15 chips in just under 30 minutes using 7 3.5GHz cores on a machine with 50GiB of RAM and a 25 MiB L3 cache. Almost all the execution time is spent performing (concurrent) HotSpot invocations.

**3.1.2 NoC Generation.** Given a produced 3-D chip layout, we generate a randomized network topology with the goal of minimizing topology diameter and average shortest path length (ASPL), while enforcing a maximum physical link length  $L$ . For this purpose, we extend the randomized  $L$ -restricted grid-graph generation method in [12]. A baseline 3-D grid graph is a graph  $G = (V, E)$  such that  $V$  is set of nodes (each node  $u$  is described by three spatial  $u_x, u_y, u_z$  coordinates) and  $E$  is a set of edges between these nodes. Let  $l(u, v)$  denote the Manhattan distance between two nodes  $u$  and  $v$ , i.e.,  $l(u, v) = |u_x - v_x| + |u_y - v_y| + |u_z - v_z|$ . Given a constant  $L$ , a 3-D grid graph  $G = (V, E)$  is  $L$ -restricted if  $l(u, v) \leq L$  for all edges  $(u, v) \in E$ .

We first generate an  $L$ -restricted 3-D grid graph  $G$  in which the coordinates of each vertex correspond to the physical position of each chip tile in the 3-D layout. We impose the  $L$ -restriction on all links, considering links with at most one TCI hop. For the sake of simplicity, to model the TCI communication conversion overhead we consider that the “length” of a TCI hop is equivalent to the chip tile length. We then randomly perturb  $G$  by applying “2-toggle” operations, and then applying “2-opt” operations. The 2-toggle operation randomly picks two edges,  $(u_1, u_2)$  and  $(v_1, v_2)$ . If  $l(u_1, v_1) > L$  and  $l(u_2, v_2) > L$ , then it replaces these two edges with edges  $(u_1, v_1)$  and  $(u_2, v_2)$ . The 2-opt operation is similar, but only replaces the edges if the topology’s ASPL is decreased. In all our experiments, each operation is applied 10,000 times. Note that the above procedure can produce duplicate edges, which are then removed. The above approach, based on the 2-toggle/2-opt operations, is general for generating graphs with low diameter and ASPL given an edge length constraint, and its near-optimality is demonstrated in [12].

#### 3.2 Layout-After-NoC (LANoC)

This approach starts with a rigid stack with a 3-D or 4-D network topology and “physically shifts” chips to improve heat dissipation.

**3.2.1 3-D Network Topology (LANoC-3D).** Consider a rigid stack layout with a standard 3-D network topology, such as a 3-D mesh. To obtain a layout in which chips partially overlap in a “staircase” pattern, we shift each chip diagonally in the horizontal plane. This is depicted in Figure 2. More formally, assuming that each chip is an  $l \times l$  square and that chips need to overlap by a fraction  $\omega$  of their area, we shift each chip by  $\lambda = l(1 - \sqrt{\omega})$  in each horizontal dimension. After this operation, paths with one TCI hops require on-chip links of length at most  $2\lambda$  (so as to reach the overlap area). If the maximum link length is  $L$ , then  $\omega$  cannot be smaller than  $(1 - \frac{L}{2l})^2$ . If chips do not overlap by more than 25%, it is possible to construct this layout using only two layers, as shown in Figure 2(a), thus avoiding all TCI crosstalk. If chips overlap by more




Figure 2: LANAoC-3D construction.




Figure 3: LANAoC-4D construction.

than 25%, but no more than 50%, this layout can be constructed by placing each chip on a different layer, i.e., using  $n$  layers, as shown in Figure 2(b), still avoiding all TCI crosstalk.

**3.2.2 4-D Network Topology (LANAoC-4D).** The LANAoC approach can also be applied to 4-D network topologies. Assume that the 4-D network topology consists of several 3-D network topologies arranged in rigid stack layouts, connected so that they form a linear topology in which each vertex is a 3-D network topology, as seen on the left-hand side of Figure 3. Like in the previous section, we shift chips so that they partially overlap (see Figure 3(a)), provided that chips do not overlap by more than 25%. The resulting layout is similar to the “flatten and spread” scheme used in a two-layer checkerboard layout [6]. This layout avoids all TCI crosstalk because it only has two layers. Another option is to construct this layout using more layers, so as to increase heat dissipation. More precisely, considering  $m$   $n$ -chip 3-D network topologies, one can build this layout using  $n + m - 1$  layers (see Figure 3(b)).

## 4 SIMULATION EVALUATION

In this section we compare LANAoC and NoCAL in terms of network characteristics and of parallel (OpenMP) application performance. Note that we do not include a rigid 3-D stack in our evaluation because, even with only two chips, this design cannot operate given the thermal constraints assumed in our experimental scenarios.

### 4.1 Target Chip Multi-Processor

We use standard simulation tools to evaluate our designs quantitatively: McPAT v1.3 [7] to compute power consumption; HotSpot v6.0 [15] to compute maximum temperature (based on the power trace generated by McPAT simulations); and gem5 [1] to simulate application executions so as to estimate execution times. For simplicity and generality, we consider a baseline 16-tile chip based on the Xeon processor description file provided by McPAT, the details

Table 1: Specification of the baseline 2-D CMP.

Processor family	x86-64
Number of cores	4
L1 I/D cache size	32/128 KiB (line:64B)
L1 cache latency	1 cycle
L2 cache bank size	12 MiB (assoc:8)
L2 cache latency	6 cycles
Memory size	4 GiB
Memory latency	160 cycles
Area	169 mm <sup>2</sup>
Maximum Power (case 1)	47.2 Watts @ 2.0 GHz
Maximum Power (case 2)	56.8 Watts @ 3.6 GHz
Router pipeline	[RC][VSA][ST/LT]
Buffer size	5-flit per VC
Protocol	MOESI directory
# of VCs	3 (one VC for each message class)
Control / data packet size	1 flit / 5 flits

of which are given in Table 1. We assume the worst case design in which each module has maximum heat dissipation. The bottom part of the table shows NoC specifications.

Note that we consider two cases for the maximum chip frequency, 2.0GHz and 3.6GHz. We obtain area and maximum (static and dynamic) power distribution of a chip using McPAT for 22nm technology with physical gate lengths configured for high-performance applications. For the 2.0GHz, resp. 3.6GHz chip, McPAT computes a 47.2 Watt, resp. 56.8 Watt, power consumption, as shown in Table 1.

Note that McPAT does not account for the power consumption of TCI. However, this power consumption is less than 0.3 W for a 256 Gbps (128 bit×2.0 GHz) vertical link and so we neglect its impact (both in terms of power consumption and heat dissipation). HotSpot simulation parameters are listed in Table 2. Finally, although transient performance behaviors of parallel applications are of interest, in this work we assume maximum steady-state use of hardware resources.

Table 2: HotSpot v6.0 simulation parameters.

Heatsink	12×12×3 cm, 400 W/mK, 0.3024 m <sup>2</sup>
Heat spreader	6×6×0.1 cm, 400 W/mK
TIM / Glue	20 μm, 0.25 W/mK
Outside temp.	25°C
Heat transfer efficiency	14W/(m <sup>2</sup> K)
Temperature threshold	80 °C

### 4.2 Overlap Area

As discussed in [4], a 675μm × 675μm area is needed for a single 64 Gbps vertical TCI link. Since each of the 16 on-chip routers is equipped with a vertical link, given the dimension of our chip a 25% overlap between two TCI-connected chips allows for just under 3 Tbps inter-chip bandwidth.

Figure 4 plots the maximum feasible overlap fraction between two chips vs. number of chips for each layout when each chip operates at its highest frequency, while remaining under the temperature threshold. A larger overlap is preferable, as it enables higher bandwidth between the chips. In all these results we only consider layouts with at most 25% overlap between TCI-connected chips, and for the




Figure 4: Maximum overlap area between two chips vs. number of chips, when chips operate at their highest frequency: 2.0GHz (left) and 3.6GHz (right).

LANoC approach we only consider layouts that use 2 layers (i.e., as in Figures 2(a) and 3(a)).

For the LANoC-4D layout (the checkerboard), the overlap fraction decreases as the number of chips increases, due to the temperature constraint. By contrast, the LANoC-3D (the staircase) and NoCAL layouts can maintain 25% overlap across the board. We conclude that these two sparse 3-D layouts can accommodate large numbers of chips with high-bandwidth vertical TCI links while remaining below the temperature threshold.

### 4.3 Network Statistics

Table 3 shows for each layout the network diameter, average shortest path length (ASPL), and bisection bandwidth for an 8-chip system designed using the NoCAL and LANoC methods. The NoCAL layout using 8 chips is illustrated in Figure 5. In these results the LANoC method uses mesh topologies as starting points, of degree 6 for LANoC-3D and of degree 7 for LANoC-4D. The mesh used by LANoC-4D has degree 7 (instead of 8) because the system consists of only 2 interconnected 4-chip LANoC-3D systems. Results are shown for NoCAL topologies with degree 6 and 7, so as to allow fair comparison to the LANoC method. As each chip consists of 16 tiles, each of which has a router, these results are for a 128-node NoC.

We set the maximum link length to 5 tiles. These results show that in NoCAL, since we remove duplicate edges, the number of links is smaller than in LANoC topologies. However, NoCAL provides higher bisection bandwidth, lower network diameter and ASPL.

Table 3: Network metrics for an 8-chip system.

	Diam.	ASPL	Bi. BW	# of Links
NoCAL(deg=6)	9	4.26	22	275
NoCAL(deg=7)	8	3.91	26	318
LANoC-3D (deg=6)	13	5.17	16	304
LANoC-4D (deg=7)	10	4.28	32	352

### 4.4 Application Performance

We perform full-system simulation of an 8-chip system, using gem5 to simulate the execution of 9 parallel programs from the OpenMP implementation of the NAS Parallel Benchmarks. These programs were compiled with GCC 4.4.7 and executed (in simulation) with a Linux kernel 2.6.22.9 using 32 threads, i.e., using all processor cores.




Figure 5: The resulting NoCAL layout using 8 chips.




Figure 6: Benchmark execution time, relative to LANoC-3D with degree 6 (2.0GHz processors).




Figure 7: Benchmark execution time, relative to LANoC-3D with degree 6 (3.6GHz processors).

Figures 6 and 7 show application execution times for 8-chip systems built using the NoCAL and the LANoC methods, for chips operating at 2.0GHz and 3.6GHz frequencies. For each application, results are shown for NoCAL topologies with degree 6 and 7, so as to allow fair comparison to the LANoC-3D (degree 6) and LANoC-4D (degree 7) topologies. As expected, for a given design method, a higher degree leads to lower application execution time. The main observation in all these results is that NoCAL leads to lower application execution times than LANoC.

Recall that the results in Section 4.3 show that NoCAL produces topologies with fewer edges but better diameter, ASPL and bisection than LANoC. The results in this section show that these improvements in topology characteristics, in spite of a lower total number of edges, directly translate to improvements in parallel application time execution. Furthermore, although the network topology in LANoC-generated designs has a structure that can in principle

benefit parallel applications, our results show that this benefit is offset by the lower hop counts of the randomized network topology in NoCAL-generated designs.

## 5 RELATED WORKS

**NoC Topologies with Low Hop Counts** – Previous work has exploited the small-world effect for on-chip network topologies. By introducing long-range links, a small-world topology that adds wired shortcut links to  $k$ -ary 2-mesh can be built in order to reduce path hop counts [14]. Another approach that has been proposed to reduce hop counts in NoC topologies is randomization [8, 12]. We use this approach in the second step of our NoCAL method.

**Thermal Management** – Many techniques have been proposed to manage heat dissipation at chip design time and at runtime. In terms of thermal-aware chip design, 3-D floorplan algorithms have been proposed [15], as well as algorithms that compute microchannel layouts. Dynamic Thermal Management (DTM) at runtime has also been investigated [2], by which thermal constraints are met by lowering performance based on various dynamic triggers. Clock/power gating and Dynamic Voltage and Frequency Scaling (DVFS) can be used to implement DTM strategies at the hardware level. The results in this work assume that all chips operate at their maximum frequencies, i.e., the worst-case design for temperature. Our work is orthogonal to DTM techniques but widens the overall design space for thermal-aware systems, making the co-design of sparse layouts and DTM techniques a compelling, and combinatorial, challenge.

## 6 CONCLUSIONS

Inductive coupling technology, such as TCI, provides an attractive alternative to TSV for 3-D integration of chips. Although traditional stack layouts cannot scale due to heat dissipation concerns, TCI enables a wide range of 3-D chip layout options. In this work we have explored the layout design space, focusing on the trade-off between interconnect performance and heat dissipation. We have proposed two methods, NoCAL and LANoC, to generate chip layouts and network topologies for these layouts. NoCAL generates a chip layout with low inter-chip path lengths and then generates a randomized network topology for this layout. LANoC, instead, perturbs standard layouts that use standard network topologies. We have found that both methods can generate designs that support large number of high-frequency chips. As expected, NoCAL produces network topologies with lower hop counts than LANoC. More importantly, as demonstrated experimentally for nine OpenMP applications in the NAS Parallel Benchmarks, NoCAL-generated designs lead to higher parallel application performance than LANoC-generated designs. We conclude that attempting to provide applications with standard network topologies (such as 3-D or 4-D meshes) is not necessarily the right approach. Instead, optimizing the layout for a good trade-off between inter-chip interconnect and heat dissipation, and then generating the network topology, leads to better results in practice.

Our results show that NoCAL is effective even though it employs a simple greedy heuristic. For larger numbers of chips, however, the objective function has likely many local minima. It may thus be necessary to use meta-heuristics which, as explained in Section 3.1.1, may not be practical due to the HotSpot computational bottleneck. A future work direction is thus to accelerate HotSpot simulations

and/or to use learning algorithms to quickly approximate HotSpot simulation results. In this work we have considered the case in which all chips operate at the same frequency. Another future work direction is to allow for heterogeneous chip frequencies (e.g., so that chips at the edge of the layout can operate at higher frequencies than chips toward the center of the layout). Allowing different frequencies would lead to designs with higher interconnect performance and/or aggregate compute power, which could benefit applications that can operate on heterogeneous platforms. Finally, we have assumed that chips operate at the same frequency throughout application execution. Another, more ambitious, future work direction is to remove this assumption and, as mentioned in Section 5, explore the co-design of sparse layouts and DTM techniques.

## ACKNOWLEDGMENTS

This work was supported by JSPS KAKENHI Grant Number 18H03215.

## REFERENCES

- [1] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoib, N. Vaish, M. D. Hill, and D. A. Wood. 2011. The gem5 Simulator. *ACM SIGARCH Computer Architecture News* 39, 2 (May 2011), 1–7.
- [2] D. M. Brooks and M. Martonosi. 2001. Dynamic Thermal Management for High-Performance Microprocessors. In *International Symposium on High-Performance Computer Architecture (HPCA)*. 171–182.
- [3] S. Gopal, S. Das, D. Heo, and P. P. Pande. 2017. Energy and Area Efficient Near Field Inductive Coupling: A Case Study on 3D NoC. In *International Symposium on Networks-on-Chip (NOCS)*. 5:1–5:8.
- [4] T. Kagami, H. Matsutani, M. Koibuchi, Y. Take, T. Kuroda, and H. Amano. 2016. Efficient 3-D Bus Architectures for Inductive-Coupling ThruChip Interfaces. *IEEE Trans. VLSI Syst.* 24, 2 (2016), 493–506.
- [5] T. Kuroda. 2016. Near-Field Coupling Integration Technology. In *International Symposium on Networks-on-Chip (NOCS)*, Keynote, [http://www.arc.ics.keio.ac.jp/nocs16/files/Kuroda\\_NOCS2016.pdf](http://www.arc.ics.keio.ac.jp/nocs16/files/Kuroda_NOCS2016.pdf).
- [6] W. Stephen Lacy, Jose L. Cruz-Rivera, and D. Scott Wills. 1998. The Offset Cube: A Three-Dimensional Multicomputer Network Topology Using Through-Wafer Optics. *IEEE Trans. on Parallel Distrib. Syst.* 9, 0 (Sept. 1998), 893–908.
- [7] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi. 2009. McPAT: An Integrated Power, Area, and Timing Modeling Framework for Multicore and Manycore Architectures. In the *IEEE/ACM International Symposium on Microarchitecture*. 469–480.
- [8] H. Matsutani, M. Koibuchi, I. Fujiwara, T. Kagami, Y. Take, T. Kuroda, P. Bogdan, R. Marculescu, and H. Amano. 2014. Low-latency wireless 3D NoCs via randomized shortcut chips. In *Design, Automation Test in Europe Conference Exhibition (DATE)*. 1–6.
- [9] N. Miura, H. Ishikuro, K. Niitsu, T. Sakurai, and T. Kuroda. 2008. A 0.14 pJ/b Inductive-Coupling Transceiver With Digitally-Controlled Precise Pulse Shaping. *IEEE Journal of Solid-State Circuits* 43, 1 (Jan 2008), 285–291.
- [10] N. Miura, K. Kasuga, M. Saito, and T. Kuroda. 2010. An 8Tb/s 1pJ/b 0.8mm<sup>2</sup>/Tb/s QDR inductive-coupling interface between 65nm CMOS GPU and 0.1μm DRAM. In *IEEE International Solid-State Circuits Conference (ISSCC)*. 436–437.
- [11] N. Miura, Y. Koizumi, Y. Take, H. Matsutani, T. Kuroda, H. Amano, R. Sakamoto, M. Namiki, K. Usami, M. Kondo, and H. Nakamura. 2013. A Scalable 3D Heterogeneous Multicore with an Inductive ThruChip Interface. *IEEE Micro* 33, 6 (2013), 6–15.
- [12] K. Nakano, D. Takafuji, S. Fujita, H. Matsutani, I. Fujiwara, and M. Koibuchi. 2016. Randomly Optimized Grid Graph for Low-Latency Interconnection Networks. In *International Conference on Parallel Processing (ICPP)*. 340–349.
- [13] K. Niitsu, Y. Kohama, Y. Sugimori, K. Kasuga, K. Osada, N. Irie, H. Ishikuro, and T. Kuroda. 2010. Modeling and Experimental Verification of Misalignment Tolerance in Inductive-Coupling Inter-Chip Link for Low-Power 3-D System Integration. *IEEE Trans. VLSI Syst.* 18, 8 (2010), 1238–1243.
- [14] U. Y. Ogras and R. Marculescu. 2006. It’s a Small World After All: NoC Performance Optimization via Long-Range Link Insertion. *IEEE Transactions on Very Large Scale Integration Systems* 14 (2006), 693–706.
- [15] R. Zhang and M. R. Stan and K. Skadron. 2015. HotSpot 6.0: Validation, Acceleration and Extension. University of Virginia, Tech. Report CS-2015-04.
- [16] M. Saito, N. Miura, and T. Kuroda. 2010. A 2Gb/s 1.8pJ/b/chip inductive-coupling through-chip bus for 128-Die NAND-Flash memory stacking. In *IEEE International Solid-State Circuits Conference (ISSCC)*. 440–441.