

Cross-layer Floorplan Optimization For Silicon Photonic NoCs In Many-core Systems

Ayse K. Coskun³, Anjun Gu¹, Warren Jin⁴, Ajay Joshi³, Andrew B. Kahng^{1,2},

Jonathan Klamkin⁴, Yenai Ma³, John Recchio¹, Vaishnav Srinivas¹ and Tiansheng Zhang³

UCSD ¹ECE and ²CSE Departments, La Jolla, CA; ³Boston University ECE Department, Boston, MA; ⁴UCSB ECE Department, Santa Barbara, CA

Abstract—Many-core chip architectures are now feasible, but the power consumption of electrical networks-on-chip does not scale well. Silicon photonic NoCs (PNoCs) are more scalable and power efficient, but floorplan optimization is challenging. Prior work optimizes PNoC floorplans through simultaneous place and route, but does not address *cross-layer* effects that span optical and electrical boundaries, chip thermal profiles, or effects of job scheduling policies. This paper proposes a more comprehensive, cross-layer optimization of the silicon PNoC and core cluster floorplan. Our simultaneous placement (locations of router groups and core clusters) and routing (waveguide layout) considers scheduling policy, thermal tuning, and heterogeneity in chip power profiles. The core of our optimizer is a mixed-integer linear programming formulation that minimizes NoC power, including (1) laser source power due to propagation, bend and crossing losses; (2) electrical and electrical-optical-electrical conversion power; and (3) thermal tuning power. Our experiments vary numbers of cores, optical data rate per wavelength, number of waveguides and other parameters to investigate scalability and tradeoffs through a large design space. We demonstrate how the optimal floorplan changes with cross-layer awareness: metrics of interest such as optimal waveguide length or thermal tuning power change significantly (up to 4X) based on power and utilization levels of cores, chip and cluster aspect ratio, and laser source sharing mechanism. Exploration of a large solution space is achieved with reasonable runtimes, and is perfectly parallelizable. Our optimizer thus affords designers with more accurate, cross-layer chip planning decision support to accelerate adoption of PNoC-based solutions.

I. INTRODUCTION AND MOTIVATION

Over the past decade, the computing industry has regularly increased the number of cores per die [27], using parallelism to help sustain historical performance scaling. As core count continues to increase, both individual core performance and performance of the Network-on-Chip (NoC) fabric will determine the overall performance of a many-core system. It is preferable to use high-radix, low-diameter NoCs that are easier to program and provide more predictable communication. Such NoC topologies imply long global links that can be power-hungry when implemented with traditional electrical signaling circuits. Global silicon photonic link designs provide noticeably higher bandwidth density, as well as lower data-dependent energy consumption, than electrical counterparts [3], [19]. In recent years, many efforts have explored various types of high-radix, low-diameter photonic NoC (PNoC) topologies, including bus [11], [12], butterfly/Clos [9], [10], [19] topologies. A common aspect of these studies is that only one or two underlying link designs are considered as the basis of NoC architecture design decisions. Moreover, physical (layout) implementations (used for energy, performance and area evaluation) are fairly coarse. This potentially yields suboptimal PNoCs, as well as only limited comparisons versus electrical NoCs (ENoCs).

The placement and routing (P&R) solution of the various silicon photonic link components bring several concerns beyond what is commonly seen for electrical links. These concerns include high thermal sensitivity, large propagation loss in CMOS-compatible link designs, large crossing losses, and laser source inefficiencies.¹ Indeed, thorough evaluation of the PNoC design space requires true *cross-layer*² optimization that considers (i) the rich design space of silicon photonic devices, (ii) a range of network topologies, and (iii) detailed P&R solutions for PNoCs, particularly at the level of core cluster and PNoC

¹High sensitivity of optical components to manufacturing variation is an additional issue, and can be influenced by the die planning and P&R solution.

²We use the term “cross-layer” in the usual way, to connote information flow across multiple layers of the system stack. Examples: (i) considering photonic device characteristics in P&R optimizations, or (ii) optimizing the floorplan based on architecture- and application-dependent power and thermal profiles.

floorplanning. Several recent efforts have provided flavors of cross-layer design [5], [6], [28].

In this paper, we develop a cross-layer approach to floorplan optimization in many-core systems with PNoC. At the level of system organization, our optimization considers chip aspect ratio, number of cores, and clustering of the cores. With respect to network design, our optimization considers PNoC logical topology, router design and placement, number of wavelengths to be multiplexed in a waveguide, and number of waveguides. Application-dependent factors (e.g., required PNoC bandwidth and different thermal profiles of the chip) are also considered.³ Our simultaneous cluster placement and PNoC routing algorithm, based on a mixed-integer linear programming (MILP) formulation, provides a thermally-aware cross-layer global NoC (electrical and optical) optimization for a cost function of power and area. Specifically, our algorithm finds the optimal core cluster size and shape, router group placement, waveguide routing, and chip aspect ratio that together minimize total PNoC power. The main contributions of our work are as follows.

- We formulate an MILP that outputs P&R solutions for Clos PNoC with minimum power consumption, area, or some weighted combination of both. Our ILP formulation takes thermal effects of cores on photonic components into consideration and integrates all sources of power consumption, including laser power, electrical-optical-electrical (EOE) power, and thermal tuning power, required to reliably operate photonic devices.
- We develop a flow that uses a 3D extension of HotSpot [2] to precalculate thermal impacts of all cores on all potential router group locations; these are used in an optimization flow that is thermally-aware of a mix of heterogeneous power profiles.
- We propose the notion of a *power weight*, which represents the temperature impact of a core with unit power consumption on a router group, that allows us to efficiently consider a mix of high- and low-end core clusters within the router group placement optimization. This opens the door to the study of heterogeneous-core designs in the cross-layer optimization.
- We identify trends in experimental data – e.g., dominant sources of power, and impact of levers such as chip and cluster aspect ratios – that suggest future heuristic approaches to PNoC designs.

In the following, Section II reviews relevant previous work, and Section III describes our floorplan optimization approach including details of the MILP formulation. Section IV gives experimental results of floorplan optimization, and we conclude in Section V.

II. PREVIOUS WORK

Floorplanning and P&R approaches for NoC designs have attracted significant research attention during the past years. The design decisions on these objects greatly impact the performance and energy efficiency of the overall many-core system. However, due to the vast design space and complexity, it is very challenging to consider all constraints during the design stage optimization of the NoC. To clearly present the large span of design constraints in our work, we make a comparison between our proposed method and representative previous work in this field, as

³As detailed in Section IV.A, some of these parameters (physical dimensions of routers, laser source sharing solution, thermal sensitivity coefficients for photonic devices, etc.) are fixed in the experiments that we report. However, it is straightforward to explore these additional axes in many-core chip optimization, e.g., using an “outer loop” around the optimization that we describe.

in Table I. In this table, we show the optimization goal of each work in the last column. With the exception of the first four papers, all papers in the table focus on PNoCs.

TABLE I: Classification of previous work and our work. OR–Optical Routing; OP–Optical Placement; ER–Electrical Routing; EP–Electrical Placement; TA–Thermally-Aware; 3D–3D Related; NoC: NoC Topology.

Work	OR	OP	ER	EP	TA	3D	NoC	Opt.
Jafari et al. [4]			✓	✓				Fault-tolerant
Yan et al. [23]			✓	✓		✓		Wire Length
Ou et al. [24]			✓	✓	✓	✓		Wire Length
Dubois et al. [22]			✓	✓		✓	✓	Fault-tolerant
2-Sided Swap [8]	✓							Signal Loss
O-Router [27]	✓							Total Power
SNAKE [7]	✓					✓	✓	Total Power
Chen et al. [13]		✓				✓	✓	Total Power
GLOW [16]	✓	✓			✓			Total Power
PROTON [14]	✓	✓				✓		Laser Power
VANDAL [5]	✓	✓				✓		Signal Loss
Our Work	✓	✓			✓	✓	✓	Total Power

In the area of ENOCs, floorplanning and P&R approaches have been widely explored, with previous work chiefly focusing on reduction of total wirelength and/or maximum on-chip temperature. Yan et al. [23] propose a hierarchical algorithm for 3D placement to achieve the above goals with fixed routing input. Dubois et al. [22] provide a 3D-NoC floorplanning method that reduces the number of vertical-link connections in 3D layout. Ou et al. [24] propose a P&R method that minimizes chip area and routing wirelength, while satisfying current-flow and current-density constraints. Jafari et al. [4] propose an algorithm allowing a simultaneous P&R search. We use a similar MILP-based formulation to simultaneously find an optimal P&R solution.

In contrast to traditional ENOCs, PNoCs have more design-space constraints stemming from use of optical devices. For example, photonic devices' P&R solution directly impacts the attenuation of the optical signal, and in turn, the laser source power consumption. Previous work [7], [8], [27] propose routing algorithms that minimize the optical losses in the PNoC given a fixed netlist. However, although the optimization in [27] provides 50% optical power reduction, it does not consider the thermal tuning power, which is a significant contributor to the total PNoC power. Chen et al. [13] investigate the impact of on-chip laser source placement and sharing on laser source power consumption; however, they do not propose a method to optimize the placement under different thermal profiles, and their work does not consider optimal routing of the PNoC. Hendry et al. [5] implement a tool flow that provides an environment to place and route the PNoC. A waveguide routing tool is provided in this flow to find the waveguide route with minimum optical signal loss. However, there are two common drawbacks of the above techniques: (1) they do not take into consideration thermal profiles, which can significantly impact thermal tuning power; and (2) the router placement is fixed before routing, which limits the design space for optimization. Li et al. [15] introduce a methodology for evaluating the thermal impact of chip designs on PNoC with VCSEL-based laser sources, but do not combine it with a comprehensive floorplan optimization. PROTON [14] and GLOW [16] both provide P&R algorithms for many-core systems. While GLOW takes thermal profile into account, it does not optimize router placement or account for thermal tuning power of the router groups, and the PNoC routing is performed with respect to a single thermal map. Typically, for many-core chips, thermal maps vary constantly depending on the system workloads and thread scheduling methods, hence it is important that the generated solution works for as many thermal maps as possible.

The key differentiating aspect of our optimizer is that it performs place and route for PNoC using a *cross-layer* approach, which simultaneously considers system design choices, network design choices, optical device choices and application-dependent factors. The optimizer solves for the best placement of on-chip optical devices *and* routing of waveguides so as to minimize total PNoC power (including EOE conversion and thermal tuning power). Our optimizer can optimize for one or more thermal profiles computed based on common workload profiles.

III. MILP-BASED FLOORPLAN OPTIMIZATION

Our floorplan optimization comprehends the many-core chip and the PNoC as follows (see Figure 1). In the chip, *cores* are grouped together to form *tiles*. All communication within a tile is local (i.e., does not go through the PNoC) and electrical. In the studies reported below, we assume a fixed bandwidth of 512 GB/s for the PNoC [13]. We also assume an 8-ary 3-stage Clos logical topology of the PNoC. The PNoC consists of *router groups*, each assigned to a set of tiles that constitute a *cluster*. All communication between tiles within a given cluster and between routers in the same router group goes through electrical links. Two router groups across clusters communicate with each other using an optical link. The connection from one router group to another is called a *net*, which we must route legally within the routing graph.⁴

A. Notation Used in the MILP

Table II gives parameters and notations that we use in formalizing our MILP. The PNoC is defined by the locations of each router group in the set C , the orientations of their corresponding clusters, and the specific waveguides used to connect the router groups according to the topology implied by the set of nets N . As shown in Figure 1, each router group is associated with a rectangular cluster of tiles around it. The cluster can be oriented vertically or horizontally, with the router group itself at the cluster's geometric center. A is the set of all available edges in the routing graph, where a_{vrq} (resp. a_{hrq}) denotes a vertical edge from vertex (r, q) to $(r + 1, q)$ (resp. a horizontal edge from vertex (r, q) to $(r, q + 1)$). N is the predefined set of nets connecting the router groups according to the logical topology of the PNoC. Each net n has a given source cluster s_n and sink cluster t_n , where $s_n, t_n \in C$.




Fig. 1: (a) Example of chip floorplan to illustrate our terminology. (b) A vertex and its surrounding edges in the routing graph. (c) 3-stage Clos topology with 8 router groups per stage.

B. Formal MILP Statement

We minimize a bicriterion objective function (Equation (1)) that is a weighted combination of the PNoC area and power. In the objective, α and β are user-specified scaling factors.

$$\text{Minimize: } \alpha \cdot P_{PNoC} + \beta \cdot AREA_{PNoC} \quad (1)$$

Subject to:

$$\sum_{r \in R, q \in Q, f \in \{0,1\}} \gamma_{frq}^c = 1, \quad \forall c \in C, \gamma_{frq}^c \in \{0,1\} \quad (2)$$

$$o_{crq} = \sum_{r' \in R, q' \in Q, f \in \{0,1\}} o_{fr'q'} \gamma_{fr'q'}^c, \quad \forall c \in C \quad (3)$$

$$\sum_{c \in C} o_{crq} \leq 1, \quad \forall q \in Q, r \in R \quad (4)$$

$$2v_{rq}^n - e_{hrq-1}^n - e_{vr-1q}^n - e_{hrq}^n - e_{vrq}^n - \sum_{f \in \{0,1\}} \gamma_{frq}^{s_n} - \sum_{f \in \{0,1\}} \gamma_{frq}^{t_n} = 0, \quad (5)$$

$$\forall n \in N, r \in R, q \in Q$$

$$r_c = \sum_{r \in R, q \in Q, f \in \{0,1\}} r \cdot \gamma_{frq}^c, \quad q_c = \sum_{r \in R, q \in Q, f \in \{0,1\}} q \cdot \gamma_{frq}^c, \quad \forall c \in C \quad (6)$$

⁴Implicitly, the studies reported below consider monolithic integration [20], [29] (as opposed to TSV-based stacked-die integration) of the photonic components with serpentine routing of all waveguides together (due to the cost of the trenches on the die). We assume on-chip laser sources are placed next to the router groups on a separate layer [13] where the link begins and ends.

TABLE II: Notations.

Notation	Meaning
C, R, Q	sets of router groups, rows and columns, respectively
$G(V, A)$	routing graph that defines all possible locations and connections of the PNoC
V	set of vertices in the routing graph (= router group locations)
A	set of edges connecting pairs of vertices in the routing graph
N	set of 2-pin nets connecting the router groups
r, q	indices of rows and columns of tiles (defining possible vertex locations)
c	a router group $\in C$
n	a net $\in N$ with one source and one sink
$s_n (t_n)$	source (sink) router group of net n
$r_c (q_c)$	tile row r (column q) coordinate of router group c
γ_{frq}^c	0-1 indicator of whether router group c occupies vertex (r, q) with orientation f
$a_{vrq} (a_{hrq})$	denoting a vertical edge from vertex $(r-1, q)$ to (r, q) (horizontal edge from vertex (r, q) to $(r, q+1)$)
f_c	0-1 indicator of whether cluster of router group c is horizontal (0) or vertical (1)
$x_{cr} (y_{cq})$	0-1 indicator of whether cluster of router group c occupies row r (column q)
o_{crq}	0-1 indicator of whether router group c occupies tile (r, q)
$o_{f'r'q'}(r, q)$	precalculated two-dimensional array capturing whether a cluster placed at (r', q') with orientation f would occupy tile (r, q) ; the array has entry 1 at each location that is occupied; with other entries 0
v_{rq}^n	0-1 indicator of whether vertex (r, q) is used in the route of net n
$e_{hrq}^n (e_{vrq}^n)$	0-1 indicator that edge $a_{hrq} (a_{vrq})$ is used in the route of net n
$d_{hrq}^n (d_{vrq}^n)$	cost of using $a_{hrq} (a_{vrq})$ in the route of net n
$used_r (used_q)$	0-1 indicator of whether row r (column q) contains any router groups
H, W	height and width of the chip
H_r, W_r	height and width for each tile
H_C, W_C	height and width for each router group
$P_{PNoC}, AREA_{PNoC}$	power and area of the PNoC
$P_{laser}, P_{tuning}, P_{electrical}$	static optical laser power, thermal tuning power and electrical/EOE power of the PNoC, respectively
$P_{prop}, P_{bend}, P_{cross}$	propagation power per unit length, power per bend and power per crossing, respectively
$P_{constant}$	losses that are not affected by the MILP solution, including through loss, coupling loss, etc.
P_{loss}	sum of all optical loss terms, which then defines the laser power needed
$w_{r'q'}(r, q)$	thermal weight for vertex (r, q) due to tile (r', q')
θ_c	thermal impact at router group c due to the thermal weights and the power profiles
θ_{max}	maximum thermal impact among all router group locations
$p_{r'q'}$	power level of tile (r', q') (power profile)
p_c	power weight of cluster c (to allow for heterogeneous clusters)
$SV_{rq} (SH_{rq})$	0-1 indicator for a straight vertical (horizontal) route at vertex (r, q)
$SV_{rq}^n (SH_{rq}^n)$	0-1 indicator for a straight vertical (horizontal) route on net n at vertex (r, q)
B_{rq}	0-1 indicator for a bend at vertex (r, q)
\bar{B}_{rq}	0-1 indicator for no bend at vertex (r, q)
CR_{rq}	0-1 indicator for a cross at vertex (r, q)
n_{bend}	total number of bends
n_{cross}	total number of crossings
P_{tuning}^0	thermal tuning power per degree Kelvin
$P_{modulator}, P_{detector}, P_{SERDES}, P_{cluster}$	DSENT-calculated power of modulator, detector, SERDES and ENoC within a cluster

$$f_c = \sum_{r \in R, q \in Q, f \in \{0,1\}} f \cdot \gamma_{frq}^c, \quad \forall c \in C \quad (7)$$

Structural Constraints. A number of constraints enforce proper structure of the cluster placement and the PNoC routing. Using the 0-1 indicator variable γ_{frq}^c , Equation (2) ensures that exactly one vertex (r, q) and one orientation (horizontal or vertical) are chosen for each router group c . Equation (6) captures the vertex (r_c, q_c) in the routing graph where the router group c is placed, and Equation (7) captures the orientation f_c of the cluster of router group c .

Equation (3) captures which tiles on a chip are occupied by which cluster. A given o_{crq} indicates whether tile (r, q) is occupied by the

cluster of router group c . $o_{f'r'q'}(r, q)$ is a precalculated two-dimensional array that indicates whether tile (r, q) would be occupied by a cluster of a router group placed at (r', q') with orientation f . The array has an entry of one at each location that is occupied, and zero everywhere else. Equation (4) enforces the constraint that no tile on the chip can belong to more than one cluster. This ensures legal placement of clusters. If a tile is not in the footprint of any placed cluster (implying whitespace in the floorplan, e.g., for components other than the cores that communicate through the PNoC), then for that tile we will have $\sum_{c \in C} o_{crq} = 0$. Equation (5) [4] imposes flow conservation, i.e., a well-formed path of routing graph edges for each net n from its source s_n to its sink t_n . The 0-1 indicator variable v_{rq}^n captures the use of vertex (r, q) in the routing of net n ; $d_{h/vrq}^n$ is a 0-1 indicator of whether edge $a_{h/vrq}$ is used in the routing of net n .

Equations for Area and Power. The area component of our objective function is determined by the following constraints. Equation (8) uses γ_{frq}^c to identify a binary indicator for the row (x_{cr}) and column (y_{cq}) the router group c is in. There is only one γ_{frq}^c that can be non-zero, and there is only one value in an array of all rows and an array of columns that is non-zero for each router group. Equations (9) and (10) indicate which rows and columns have router groups assigned to them. Router group locations cause extra area to be taken up in the chip, so by counting the number of rows and columns that are occupied we can obtain a figure of merit for how much area is required for photonic components.

$$x_{cr} = \sum_{q \in Q, f \in \{0,1\}} \gamma_{frq}^c, \quad y_{cq} = \sum_{r \in R, f \in \{0,1\}} \gamma_{frq}^c \quad \forall c \in C \quad (8)$$

$$used_r = \begin{cases} 1 & \text{if } \sum_{c \in C} x_{cr} \geq 1, \forall r \in R \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

$$used_q = \begin{cases} 1 & \text{if } \sum_{c \in C} y_{cq} \geq 1, \forall q \in Q \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

$$\Delta H = H_C \cdot \sum_{r \in R} used_r \quad (11)$$

$$\Delta W = W_C \cdot \sum_{q \in Q} used_q \quad (12)$$

$$AREA_{PNoC} = (H + \Delta H) \cdot (W + \Delta W) - H \cdot W \quad (13)$$

The power component of the objective is determined by the following constraints. We convert P_{loss} (dbm) to P_{laser} (mW) using a piecewise-linear approximation and the laser's wall plug efficiency (WPE) to obtain the electrical input power required to operate the laser. P_{tuning} is the thermal tuning power needed to keep the ring groups at a similar temperature. $P_{electrical}$ is the power required for EOE conversion. $P_{modulator}$, $P_{detector}$, P_{SERDES} and $P_{cluster}$ values are obtained from code extracted from DSENT [1].

$$P_{PNoC} = P_{laser} + P_{tuning} + P_{electrical} \quad (14)$$

$$P_{loss} = P_{prop} \sum_{n \in N} \sum_{a_{h/vrq} \in A} d_{h/vrq}^n e_{h/vrq}^n + P_{cross} \cdot n_{cross} + P_{bend} \cdot n_{bend} + P_{constant} \quad (15)$$

$$P_{electrical} = P_{modulator} + P_{detector} + P_{SERDES} + P_{cluster} \quad (16)$$

Thermal tuning power is proportional to the difference between the thermal impact of a given router group (θ_c) and the maximum thermal impact (θ_{max}) over all router groups. Equation (17) calculates the thermal impact of each router group using the power profile of the system, with each tile's power level contributing a thermal weight $w_{r'q'}(r, q)$ to the router group at (r, q) . Given that θ_c is a product of two binary variables, we must linearize it using the following technique (repeated in how we handle bends and crossings below).

$$\theta_c = \sum_{r \in R, q \in Q, f \in \{0,1\}, r' \in R, q' \in Q} \gamma_{frq}^c \cdot w_{r'q'}(r, q) \cdot p_{r'q'}, \quad \forall c \in C \quad (17)$$

$$p_{r'q'} = \sum_{c \in C} o_{cr'q'} \cdot p_c, \quad \forall r' \in R, q' \in Q, p_c \text{ is fixed} \quad (18)$$

$$P_{\text{routing}} = P_{\text{routing}}^0 \sum_{c \in C} (\theta_{\text{max}} - \theta_c) \quad (19)$$

Accounting for Optical Bends and Crossings. We include the number of bends and crossings that exist in any given routing solution. The 0-1 indicator variable SV_{rq}^n (respectively, SH_{rq}^n) captures the existence of a straight vertical (respectively, horizontal) route through vertex (r, q) for net n . We derive SV_{rq}^n and SH_{rq}^n from $e_{h/vrq}^n$:

$$SV_{rq}^n \leq e_{vr-1q}^n; SV_{rq}^n \leq e_{vrq}^n; SV_{rq}^n \geq e_{vr-1q}^n + e_{vrq}^n - 1, \forall n \in N, r \in R, q \in Q \quad (20)$$

$$SH_{rq}^n \leq e_{hrq-1}^n; SH_{rq}^n \leq e_{hrq}^n; SH_{rq}^n \geq e_{hrq-1}^n + e_{hrq}^n - 1, \forall n \in N, r \in R, q \in Q \quad (21)$$

To account properly for all bends in the routing solution, we define a 0-1 indicator B_{rq} to capture the existence of a bend at vertex (r, q) , and \hat{B}_{rq} as a binary indicator for a vertex used with no bends. SH_{rq} , SV_{rq} , and v_{rq} respectively indicate straight vertical routes, straight horizontal routes, and vertex used at each (r, q) coordinate, for the superposition of all routed nets $n \in N$. Finally, we add the number of bends across all (r, q) to obtain the total number of bends in the routing solution.

$$\hat{B}_{rq} \leq SH_{rq} + SV_{rq}; \hat{B}_{rq} \geq SH_{rq}; \hat{B}_{rq} \geq SV_{rq}; \quad (22)$$

$$B_{rq} + \hat{B}_{rq} - v_{rq} + \sum_{c \in C, f \in 0,1} \gamma_{fq}^c = 0, \forall r \in R, q \in Q$$

$$n_{\text{bend}} = \sum_{q \in Q, r \in R} B_{rq} \quad (23)$$

We also include all straight-straight crossings in our power loss equation, using some of the same variables. The 0-1 indicator variable CR_{rq} captures the existence of a crossing at vertex (r, q) , enabling us to obtain the total number of crossings across all (r, q) .

$$CR_{rq} \geq SH_{rq} + SV_{rq} - 1; CR_{rq} \leq SH_{rq}; CR_{rq} \leq SV_{rq} \quad (24)$$

$$\forall r \in R, q \in Q$$

$$n_{\text{cross}} = \sum_{q \in Q, r \in R} CR_{rq} \quad (25)$$

C. MILP Instance Complexity and Scalability

Using the notation and from the formulation given above, the complexity of an instance of our MILP is as follows.

- The number of variables: $8NRQ + 3CRQ + 4RQ + C + CR^2Q^2$.
- The number of constraints: $3CR^2Q^2 + NRQ + 14RQ + 5C + 1$.

For a typical instance that we study in the experiments reported below, $C = 8$, $R = 8$, $Q = 8$ and $N = 7$, implying 38152 variables and 99689 constraints. Both the number of variables and the number constraints have terms that scale (i) linearly with the number of router groups, and (ii) quadratically with the number of tiles (RQ) . If we assume that the number of cores per tile is fixed, then these parameters respectively translate to (i) the number of cores, and (ii) the size of the chip. For instances of this complexity, runtimes of ILOG CPLEX v12.5.1 [31] range from 10 seconds to several minutes on a 2.8 GHz Xeon server.

D. Optimization Flow Details and HotSpot Approximation

We conclude this section by pointing out several key details of our optimization flow and setup.

(1) Our floorplan optimizer takes as input a *.param file* with the following contents: (i) CoreParams (N_{cores} , W_{core} , H_{core} , Core Power); (ii) AspectRatio (AR_{min} , AR_{max}); and (iii) OpticalParams (loss mechanisms, waveguide dimensions and spacing, ring dimensions and spacing, and photodetector sensitivity). The *.param file* is used during setup of the MILP.

(2) Quite importantly, while we use HotSpot-3D [2] (embedded in HotSpot v6.0) as our thermal simulator, it is not practical to run HotSpot inside a high-dimensional optimization of floorplan, placement, routing and other solution attributes. Even more, the MILP approach is fundamentally incompatible with running a thermal simulator “in the loop”, as might be contemplated with annealing or other iterative optimization frameworks. We work around this issue by precharacterizing a *core impact matrix* that captures the steady-state

temperature impact of each running core on each possible router group location. The core impact matrix contains the thermal impact in K/W due to a $1W$ core at each core location. We assume a linear superposition of core impacts due to all cores to calculate a final temperature at each vertex. We compare the temperature profile based on superposition with the data from HotSpot (with all the cores active simultaneously) and confirm less than 3% error. (Briefly: (i) a script generates potential HotSpot-compatible floorplan files (.flp) based on the core dimensions, number of cores, and router group dimensions; (ii) a dummy router group is placed at every valid vertex for each .flp file; and (iii) a *core impact matrix* is generated that consists of $(N_{\text{rows}} - 2) \times (N_{\text{columns}} - 2)$ values. Figure 2 illustrates the concept of core impact matrix generation. For the simple floorplan shown in part (a), two example core impact calculations are shown in parts (b) and (c) (for the two router group locations (1,3) and (2,2), respectively). The core impacts for each router group location (r, q) are summarized in the array of core impacts (d).) (3) We extract code from the current DSENT [1] distribution to calculate the EOE power (modulator, detector, SERDES) and the electrical power for the NoC within the clusters. We would like to note that for the three link bandwidths considered in our analysis (see Table III), we leverage DSENT’s capability to perform datapath power optimization by balancing insertion loss and extinction ratio with modulator/receiver and laser power.




Fig. 2: Core impact matrix generation: (a) illustrative floorplan with 16 tiles (64 cores) and nine potential router group positions; (b) sample core impact calculation for router group (1,3); (c) sample core impact calculation for router group (2,2); (d) a 1x9 core impact array generated for the floorplan.

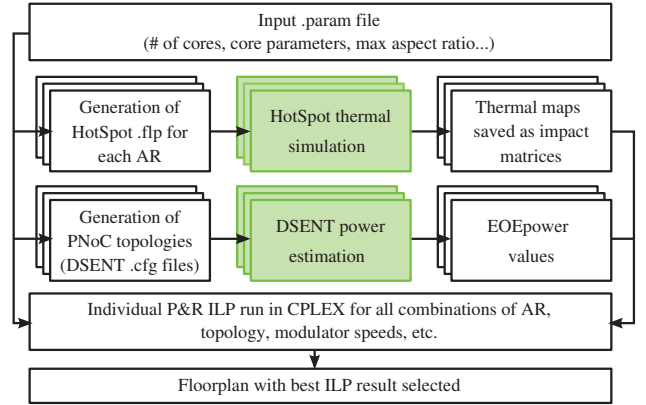


Fig. 3: The floorplan optimization flow.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

A. Simulation Infrastructure

To test our optimization model, we use technology parameters corresponding to a 22 nm SOI CMOS process, and cores whose architecture is similar to the IA-32 in Intel Single-Chip Cloud Computer (SCC) [30], for our many-core systems. Each core has a 16 KB I/D L1 cache along with a 256 KB private L2 cache. After scaling the IA-32 core to 22 nm, each core (including caches) is reduced to a square shape with a side of 1.129 mm. We use HotSpot’s default configuration file settings, but scale the heat spreader and heat sink lengths to be 2X and 4X the longest chip side length, respectively, for each floorplan. We also modify the configuration file in DSENT to match our experiments as follows: 22 nm technology; 1 GHz operating frequency; 2, 4, 8 Gbps link data rates for all the test cases; 3-stage 8-ary Clos or 3-stage 16-ary

Clos topology according to different test cases; 64, 128, or 256 cores according to different test cases.

B. Design of Experiments

From above, our floorplan optimizer finds the optimal packing of clusters and routing of waveguides, based on given design inputs. To validate our optimization approach over a large experiment space, we use the set of configurations shown in Table III. We consider WPE values of 5% and 15%, considering WPEs of current and future on-chip laser sources.

TABLE III: Experimental configurations studied.

#cores	Clos size	(chip AR, cluster AR)	optical datarate (Gbps)	#waveguides
64	8-ary (1 core/tile)	(1:1,1:2), (1:4,1:2)	8	8,16,32,64,128
			4	16,32,64,128
			2	32,64,128
128	8-ary (2 core/tile)	(1:2,1:1), (1:2,1:4)	8	8,16,32,64,128
			4	16,32,64,128
			2	32,64,128
256	8-ary (4 core/tile)	(1:1,1:2), (1:1,1:8)	8	8,16,32,64,128
			4	16,32,64,128
			2	32,64,128
	16-ary (1 core/tile)	(1:1,1:1), (1:1,1:4), (1:4,1:1), (1:4,1:4)	8	32,64,128
			4	64,128
			2	128

Workloads are intrinsically different from each other, which results in potential variations in their power profiles. Especially in a many-core system, it is common to have multi-program workloads and thus imbalanced power profiles [17], [18]. Furthermore, the emergence of heterogeneous systems exacerbates the imbalance within the power profiles. Thus, optimizing for known imbalances in power profiles may work as a viable goal for many real-life systems. Our experiments consider the power profiles in Figure 4(a)-(f).⁵ We include these power profiles in our design of experiments to demonstrate that the optimal floorplan is sensitive to the power profile, and that designers can potentially determine the floorplan based on a power profile of a use case or combination of use cases (average, weighted-average, or worst-case). We assume the optical loss coefficients listed in Table IV.




Fig. 4: Six power profiles studied. Darker tiles indicate higher power cores.

TABLE IV: Losses in PNoCs. [19]

Loss Mechanism	Loss Contribution
Splitter Through Loss	0.2 dB per split
Waveguide Propagation Loss	2 dB per cm
Waveguide Crossing Loss	0.05 dB per crossing
Ring Drop Loss	1.5 dB per wavelength per ring
Ring Insertion Loss	0.1 dB per wavelength per ring
Ring Through Loss	0.01 dB per wavelength per ring
Photodetector Loss	0.1 dB per photodetector
Merge Loss	5 dB per merge

C. Results and Discussion

We now present our experimental results. In all cases that we consider, the logical topology is a chain from router group $c = 0$ to router group $c = |C|$, with $|N| = |C| - 1$ nets. Figure 5 shows how the accumulated thermal weight profile and the optimal floorplan vary with change in N_{cores} for a given NoC topology, optical data rate and number of waveguides. We see that although waveguide lengths increase with N_{cores} , the thermal tuning power (which depends on the thermal weight profile as highlighted in Figure 5(a)) tends to flatten out in larger chips due to more symmetry.

⁵Our power profiles capture possibilities arising from both floorplanning of high- vs. low-power cores and thermally-aware task allocation. Profile (f) is a “Pringle’s chip” that mimics systematic variation (slow cores on die edge requiring boosted supply voltage) seen in large, full reticle field ICs starting around the 65 nm node [25], [26].

Figure 6 shows how the thermal weight profile and floorplan vary with the aspect ratio (AR) of the chip. In general, a skewed chip AR leads to a larger periphery, creating more asymmetry in the thermal weight profile as shown in Figure 6(a), but at the same time allowing for a shorter waveguide length.

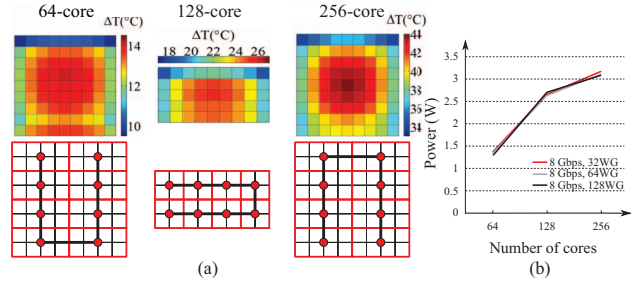


Fig. 5: Accumulated thermal weight profile and optimal floorplan vs. N_{cores} .

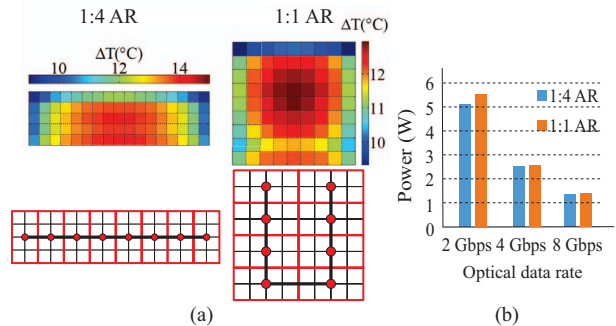


Fig. 6: Accumulated thermal weight profile and optimal floorplan vs. AR.

Figure 7 shows the thermal weight profiles and floorplans for the different power profiles described in Figure 4. We note that the PNoC power varies by nearly 1.7X across the different power profiles. We also note that the optimal floorplans vary when we change WPE from 5% to 15%. A poorer laser source efficiency tends to favor the U-shaped floorplan. In comparison to a baseline vertical U-shaped floorplan, the floorplan in Figure 7(e) saves up to 15% power under the heterogeneous power profile in Figure 4(e).




Fig. 7: Accumulated thermal weight profile on the first row, and optimal floorplan with WPE of 5% and 15% on the second and third row respectively for power profiles (a) - (f) in Figure 4.

Figure 8 addresses the impact of optical data rate on thermal weight profile, floorplan, and PNoC power consumption for a given NoC topology, N_{cores} and number of waveguides. We note that the power increases as optical data rate decreases, owing to the larger number of wavelengths needed to maintain the PNoC bandwidth.

Figure 9 shows how the optimal floorplan varies when we assign a power weight (p_c) to each cluster. This allows us to evaluate a mix of heterogeneous clusters. Part (a) of the figure shows the optimal floorplan for the case where the first half of router groups in the logical topology have low-power clusters (indicated by a blue router group), and the second half (red) have high-power clusters (indicated by a

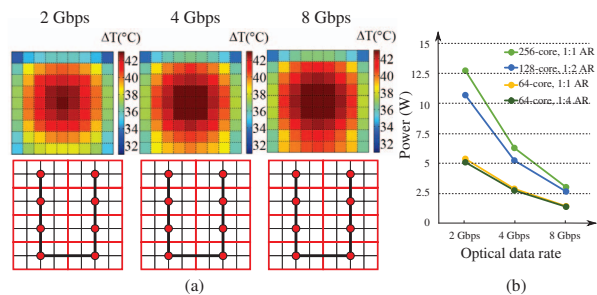


Fig. 8: Accumulated thermal weight profile and optimal floorplan vs. optical data rate.

red router group). Part (b) shows the case where the low-power and high-power clusters alternate. The optimal floorplan attempts to pack as many high-power clusters as possible into regions of relatively similar thermal weight. Such heterogeneous clusters could also allow us to study best options for “dark silicon” during thermal throttling. In the future, we plan to evaluate how such heterogeneous clusters allow for an independent cross-layer knob that is distinct from job scheduling.




Fig. 9: Optimal floorplan for different cluster power weights; blue (red) router-group indicates a low- (high-) power cluster respectively.

From our experiments, we arrive at the following general conclusions.

- Both thermal tuning power and laser power are important sources of power in the PNoC. Sensitivity to thermal weight profiles is especially important for cases with better WPE.
- Larger chips present an economy of scale for the PNoC power due to the more symmetric thermal weight profiles of larger chips.
- Skewed chip aspect ratios provide larger periphery and create asymmetry in the thermal weight profiles.
- The maximum achievable optical data rate is always preferred.
- It is important to consider different power profiles during the design time, since heterogeneous power profiles expose inherent weaknesses to certain router group locations. Being thermally aware of runtime management issues during floorplan optimization provides a key cross-layer advantage to such an optimization. Weighting the power profiles based on duty cycle and benchmarking metrics could provide a way to choose an optimal floorplan that is aware of the heterogeneous runtime power profiles.
- Allowing for power weights associated with clusters provides an additional knob to investigate the best mix and locations for high- and low-performance clusters, and the impact of dark silicon considerations on the optimal floorplan.

V. CONCLUSIONS

This paper has proposed a cross-layer, thermally-aware optimizer for floorplanning of PNoCs. Our simultaneous placement of router groups and core clusters, along with routing of waveguides, comprehends scheduling policy, thermal tuning of photonic devices, and heterogeneity in application-dependent chip power profiles. Our optimizer uses an MILP formulation that minimizes PNoC power by explicitly considering laser source power (aware of propagation and other losses), EOE conversion power, and thermal tuning power. We also introduce the concept of a power weight, associated with each core cluster, which allows optimal placement of heterogeneous clusters, accounting for designs with heterogeneous cores. We achieve very reasonable running times for optimization of large many-core chips (on the order of a few minutes), which gives users the capability to make quick design-time decisions. We verify scalability and accuracy of the optimizer over a large design space. The results show that the optimal PNoC power is sensitive to thermal weight profiles and power profiles (1.7X variation),

optical data rate (3-4X variation), number of cores (with chip edge for laser power and inversely for tuning power) and chip aspect ratios (up to 10%). Also, compared to thermally-agnostic solutions, our optimizer saves up to 15% PNoC power. We expect to integrate the floorplan optimizer with more extensive EDA flows in the future.

ACKNOWLEDGMENT

This work has been partially funded by the NSF grants CCF-1149549 and CNS-1149703. Work at UCSD has been supported by NSF, Samsung and the IMPACT+ Center.

REFERENCES

- [1] C. Sun, C. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L. Peh and V. Stojanovic, “DSNT - A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-on-Chip Modeling”, *Proc. NOCS*, 2012, pp. 201-210.
- [2] J. Meng, K. Kawakami and A. Coskun, “Optimizing Energy Efficiency of 3-D Multicore Systems with Stacked Dram under Power and Thermal Constraints”, *Proc. DAC*, 2012, pp. 648-655.
- [3] A. Shacham, K. Bergman and L. P. Carloni, “On the Design of a Photonic Network-on-Chip”, *Proc. NOCS*, 2007, pp. 53-64.
- [4] R. Jafari, F. Dabiri and M. Sarrafzadeh, “An Efficient Placement and Routing Technique for Fault-tolerant Distributed Embedded Computing”, *Proc. RTCSA*, 2005, pp. 1-9.
- [5] G. Hendry, J. Chan, L. P. Carloni and K. Bergman, “VANDAL: A Tool for the Design Specification of Nanophotonic Networks”, *Proc. DATE*, 2011, pp. 1-6.
- [6] C. Batten, A. Joshi, V. Stojanovic and K. Asanovic, “Designing Chip-Level Nanophotonic Interconnection Networks”, *IEEE JETCAS* 2(2) (2012), pp. 137-153.
- [7] L. Ramini, P. Grani, S. Bartolini and D. Bertozzi, “Contrasting Wavelength-Routed Optical NoC Topologies for Power-Efficient 3D-stacked Multicore Processors using Physical-Layer Analysis”, *Proc. DATE*, 2013, pp. 1589-1594.
- [8] C. Condrat, P. Kalla and S. Blair, “Channel Routing for Integrated Optics”, *Proc. SLIP*, 2013, pp. 1-8.
- [9] H. Gu, J. Xu and W. Zhang, “A Low-power Fat Tree-based Optical Network-On-Chip for Multiprocessor System-on-chip”, *Proc. DATE*, 2009, pp. 3-8.
- [10] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang and A. Choudhary, “Firefly: Illuminating Future Network-on-chip with Nanophotonics”, *Proc. ISCA*, 2009, pp. 429-440.
- [11] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apse, M. A. Watkins and D. H. Albonese, “Leveraging Optical Technology in Future Bus-based Chip Multiprocessors”, *Proc. MICRO*, 2006, pp. 492-503.
- [12] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil and J. H. Ahn, “Corona: System Implications of Emerging Nanophotonic Technology”, *Proc. ISCA*, 2008, pp. 153-164.
- [13] C. Chen, T. Zhang, P. Contu, J. Klamkin, A.K. Coskun and A. Joshi, “Sharing and Placement of on-chip Laser Sources in Silicon-Photonic NoCs”, *Proc. NOCS*, 2014, pp. 88-95.
- [14] A. Boos, L. Ramini, U. Schlichtmann and D. Bertozzi, “PROTON: An Automatic Place-and-route Tool for Optical Networks-on-chip”, *Proc. ICCAD*, 2013, pp. 138-145.
- [15] H. Li, A. Fourmigue, S. Le Beux, X. Letartre, I. O’Connor, and G. Nicolescu, “Thermal Aware Design Method for VCSEL-based On-chip Optical Interconnect”, *Proc. DATE*, 2015, pp. 1120-1125.
- [16] D. Ding, B. Yu and D. Z. Pan, “GLOW: A Global Router for Low-power Thermal-reliable Interconnect Synthesis using Photonic Wavelength Multiplexing”, *Proc. DAC*, 2012, pp. 621-626.
- [17] S. Lu, R. Tessier and W. Bursleson, “Reinforcement Learning for Thermal-aware Many-core Task Allocation”, *Proc. GLSVLSI*, 2015, pp. 379-384.
- [18] A. K. Coskun, T. S. Rosing, K. Whisnant and K. Gross, “Static and Dynamic Temperature-Aware Scheduling for Multiprocessor SoCs”, *IEEE Trans. on VLSI Systems* 16(9) (2008), pp. 1127-1140.
- [19] A. Joshi, C. Batten, Y. Kwon, S. Beamer, I. Shamim, K. Asanovic and V. Stojanovic, “Silicon-photonic Clos Networks for Global on-chip Communication”, *Proc. NOCS*, 2009, pp. 124-133.
- [20] J. S. Orcutt, B. Moss, C. Sun et al., “Open Foundry Platform for High-performance Electronic-photonic Integration”, *Optics Express*, 2012, pp. 12222-12232.
- [21] T. Zhang, J. Abellan, A. Joshi and A. Coskun, “Thermal Management of Manycore Systems with Silicon-photonic Networks”, *Proc. DATE*, 2014, pp. 1-6.
- [22] F. Dubois, A. Sheibanyrad, F. Petrot and M. Bahmani, “Elevator-First: A Deadlock-Free Distributed Routing Algorithm for Vertically Partially Connected 3D-NoCs”, *IEEE Trans. on Computers* 62(3) (2011), pp. 609-615.
- [23] H. Yan, Q. Zhou, X. Hong and Z. Li, “Efficient Hierarchical Algorithm for Mixed Mode Placement in Three Dimensional Integrated Circuit Chip Designs”, *Tsinghua Science and Technology* 14(2)(2009), pp. 161-169.
- [24] H. Ou, H. Chien and Y. Chang, “Simultaneous Analog Placement and Routing with Current Flow and Current Density Considerations”, *Proc. DAC*, 2013, pp. 1-6.
- [25] B. Hargreaves, H. Hult and S. Reda, “Within-die Process Variations: How Accurately Can They Be Statistically Modeled?”, *Proc. ASPDAC*, 2008, pp. 524-530.
- [26] A. Lesea, Xilinx, Inc., *personal communication*, 2008.
- [27] D. Ding, Y. Zhang, H. Huang, R. T. Chen and D. Pan, “O-Router: An Optical Routing Framework for Low Power On-Chip Silicon Nano-Photonic Integration”, *Proc. DAC*, 2009, pp. 264-269.
- [28] Z. Li, A. Qouneh, M. Joshi, W. Zhang, X. Fu and T. Li, “Aurora: A Cross-Layer Solution for Thermally Resilient Photonic Network-on-Chip”, *IEEE Trans. on VLSI Systems* 23(1) (2015), pp. 170-183.
- [29] M. Georgas, B. Moss, C. Sun et al., “A Monolithically-integrated Optical Transmitter and Receiver in a Zero-change 45nm SOI Process”, *Proc. Symposium on VLSI Circuits Digest of Technical Papers*, 2014, pp. 1-2.
- [30] J. Howard, S. Digne, S. Vangal et al., “A 48-core IA-32 Processor in 45 nm CMOS Using On-die Message-passing and DVFS for Performance and Power Scaling”, *IEEE JSSC* 46(1) (2011), pp. 173-183.
- [31] IBM ILOG CPLEX, www.ilog.com/products/cplex/