Routing in Optical Network-on-Chip: Minimizing Contention with Guaranteed Thermal Reliability

Mengquan Li[1], Weichen Liu[2,*], Lei Yang[1], Peng Chen[1], Duo Liu[1], Nan Guan[3]

[1]College of Computer Science, Chongqing University, China
[2]School of Computer Science and Engineering, Nanyang Technological University, Singapore
[3]Department of Computing, Hong Kong Polytechnic University, Hong Kong

ABSTRACT

Communication contention and thermal susceptibility are two potential issues in optical network-on-chip (ONoC) architecture, which are both critical for ONoC designs. However, minimizing conflict and guaranteeing thermal reliability are incompatible in most cases. In this paper, we present a routing criterion in the network level. Combined with device-level thermal tuning, it can implement thermal-reliable ONoC. We further propose two routing approaches (including a mixed-integer linear programming (MILP) model and a heuristic algorithm (CAR)) to minimize communication conflict based on the guaranteed thermal reliability, and meanwhile, mitigate the energy overheads of thermal regulation in the presence of chip thermal variations. By applying the criterion, our approaches achieve excellent performance with largely reduced complexity of design space exploration. Evaluation results on synthetic communication traces and realistic benchmarks show that the MILP-based approach achieves an average of 112.73% improvement in communication performance and 4.18% reduction in energy overhead compared to state-of-the-art techniques. Our heuristic algorithm only introduces 4.40% performance difference compared to the optimal results and is more scalable to large-size ONoCs.

1 INTRODUCTION

ONoC architecture [2] provides an innovative solution for long-haul data transmission to satisfy the communication bandwidth and latency requirements with low power dissipation. State-of-the-art ONoC contains a photonic network for bulk message transmission and an electronic overlay network for the control of the photonic network. The inability of ONoCs to perform inflight buffering and processing suggests the use of *optical circuit switching*: optical messages are transmitted end-to-end in the photonic network, once an optical path is established by a control packet routed in the electronic network. This approach provides excellent communication performance for unconflicted message transmissions, in which dedicated optical circuits guarantee the full bandwidth of the channels and remain connected for the duration of the communication sessions. However, conflicted messages have to be blocked for guaranteed delivery with extra communication delays and energy overheads, which would degrade the communication performance and energy efficiency of ONoCs.

Most works [5, 14] utilize dedicated ONoC architectures and bypath links to avoid contention. They possess high efficiency while incur additional hardware overheads. Contention-aware adaptive routing [7] is an attractive solution for ONoCs, which mitigates communication conflict by utilizing its inherent flexibility. It requires no extra hardware but leverages the existing generic network architecture with lightweight software computation. Existing routing techniques generally make routing decision for every communication pair individually, which does not facilitate reducing communication contention. Large-scale ONoCs typically have multiple communications at any moment. Therefore, we focus on contention-aware adaptive routing considering multiple communications in this work.

Thermal reliability is another potential issue in ONoC designs due to the intrinsic thermal sensitivity of nanophotonic devices, which has aroused great concern in recent years [6, 9, 13, 17, 18]. On every routing path in ONoC, multiple optical switches switch signals in stages to implement data transmission. Micro-ring resonator (MR), performing as a wavelength-selective optical switch in every stage, is especially susceptible to temperature variations. The resonant wavelength of a MR red-shifts with increasing temperatures and blue-shifts with decreasing ones. The undesired mismatch between signal wavelength and the resonant wavelength of the MR will result in additional optical power loss, potentially causing performance degradation and even functional failures. For guaranteeing the thermal reliability of ONoCs, device-level thermal tuning is widely employed for switches to compensate their temperature-dependent wavelength shift. However, this technique is energy-hungry, such that the energy consumption for thermal regulation would possibly null the benefits of ONoC in energy efficiency [13, 18]. The problem will be more critical in the ONoCs who have large temperature variations and a large number of switching stages.

Separately performance improvement and thermal management are impractical for deep optimization on ONoCs. However, there are apparent contradictions between optimizing communication performance, energy efficiency and guaranteeing reliability. The local thermal tuning guarantees the thermal reliability of generic ONoCs while introduces extra energy overhead. In order to minimize the regulation energy consumed by thermal tuning, messages prefer to be routed along the path that has the lowest thermal gradient, while

*Corresponding author: Weichen Liu. Email: liu@ntu.edu.sg.

ASP-DAC 2019, Jan. 2019, Tokyo, Japan

ACM ISBN 978-1-4503-6007-4/19/01...$15.00
https://doi.org/10.1145/3287624.3287650

communication pairs prefer to select different routing paths without contention for better communication performance.

In this paper, we address the above trade-off problem by developing novel routing techniques. Analyzing the thermal effects in ONoC, we theoretically formulate the boundary condition of the number of switching stages in routing paths for reliable data transmissions and present a routing criterion at the network level based on the current silicon photonic fabrication technology. The criterion, combined with recognized device-level thermal tuning, can implement thermal-reliable ONoC. We further propose two contention-aware adaptive routing approaches, including a MILP model and a heuristic algorithm (CAR), to minimize communication conflict based on the guaranteed thermal reliability, and meanwhile, mitigate the energy overheads of thermal regulation in the presence of chip thermal variations. By applying the criterion, our approaches can achieve excellent performance with largely reduced complexity of design space exploration. Compared to state-of-the-art techniques on synthetic communication traces and realistic benchmarks, our MILP-based approaches improves communication performance by 112.73% and reduces energy overhead by 4.18% on average. The heuristic algorithm achieves a satisfying performance comparable to the MILP model with only 4.40% difference and is more scalable to large-size ONoCs.

The rest of this paper is organized as follows: Sec. 2 analyzes the thermal effects in ONoC and presents a routing criterion to guarantee thermal reliability. In Sec. 3, we formulate the problem addressed in this paper and propose two routing approaches. The evaluation results are shown in Sec. 4 and Sec. 5 draws the conclusions.

2 THERMAL RELIABLE ONOC

An ONoC (the electronic network is omitted here for simplicity) and the logical view of an optical router are illustrated in Fig. 1(a). We provide a specific switching network design of a nonblocking 5×5 optical router [3] in Fig. 1(b), in which the wavelength-selective optical switches play a key role in implementing high-level routing policies in ONoCs. There are two different switch designs shown in Fig. 1(c): parallel switching element (PSE) and crossing switching element (CSE). When a switch is configured to be switched *off*, the optical signal from the *Input* port will be delivered to the *Through* port. Otherwise, it is resonated into the ring and delivered to the *Drop* port.

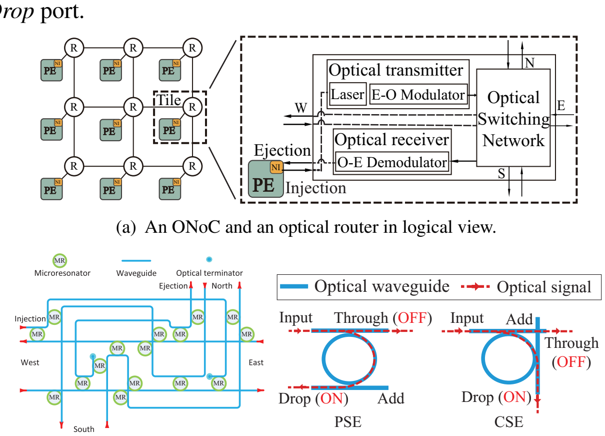

(a) An ONoC and an optical router in logical view.

(b) Router Cygnus [3]

(c) Switching elements

Figure 1: ONoC and its functional components.

The resonant wavelength of switching elements is sensitive to ambient temperature. Given the initial resonant wavelength, λ_0, at initial temperature, T_0, the formula of the resonant wavelength of a switching element, λ_{SE}, and its ambient temperature, T_{SE}, can be expressed as follow:

$$\lambda_{SE} = \lambda_0 + \rho \cdot (T_{SE} - T_0) \tag{1}$$

where ρ is the temperature-dependent resonant wavelength shift coefficient of the switching element and is about 0.06 nm/K at the 1550 nm wavelength range [18].

The undesired deviation from the peak resonant wavelength caused by temperature variation would result in additional power loss. According to the traveling wave theory [10], the optical power loss due to a switching element can be formulated as follow:

$$L_{SE} = 10log\left(\left(\frac{2\kappa^2 + \kappa_p^2}{2\kappa^2}\right)^2 \left(1 + \frac{4\left(\lambda_{in} - \lambda_0 - \rho(T_{SE} - T_0)\right)^2}{\theta^2}\right)\right) \tag{2}$$

where κ^2 is the fraction of power coupling between the waveguides and the ring. κ_p^2 is the fraction of intrinsic power loss per round-trip in the ring. λ_{in} is the wavelength of optical signals. θ is the -3dB bandwidth of the power transfer spectrum.

To ensure that ONoC function properly, the signal power reaches to the receiver of each routing path should be guaranteed. Otherwise, the fluctuations in received power may lead to dramatic reliability degradation and even failure [12]. We formulate the condition in (3), where P_{TX} is the input optical power of the transmitter on a path; $\sum_{i=1}^{m} L_{SE_i}$ is the total optical power loss inserted in the m active switches (i.e., the number of switching stages) in the path; L_{WG} includes the power loss due to the passive switches and the waveguides, and S_{RX} is the receiver sensitivity. Compared to the switches in the ON state, the thermal sensitivity of the passive switches and the waveguides is negligible [18].

$$P_{TX} - \sum_{i=1}^{m} L_{SE_i} - L_{WG} \geq S_{RX} \tag{3}$$

The above model indicates that there are two methods to implement reliable data transmission for ONoCs in presence of temperature fluctuations. One is to increase the input power of the transmitter to guarantee enough optical power reaching the receiver under a high loss in the path. The other way is to control the total power loss of optical switches by employing device-level thermal tuning [13], which is more energy-efficient than the former method in current silicon photonic technologies [18]. By utilizing local microheaters and photodetectors (PDs) that are readily available in a typical photonic platform, thermal tuning technique dynamically maintains the local temperature and the resonant wavelength of the switches throughout the duration of their operation. The major concern of this technique is the power consumption, which may significantly decrease the benefits of ONoCs in energy efficiency. Formally, the regulation power consumed by thermal tuning in a routing path is formulated as follow:

$$E^o_{thermal-tuning} = \sum_{i=1}^{m} \varepsilon \cdot \rho \cdot (T_{SE_i} - T_0) \cdot t_{dur} \tag{4}$$

where ε is the tuning efficiency in mW/nm; T_{SE_i} is the temperature of the i-th switch; t_{dur} is the duration time of the switches' operation. We apply thermal tuning in this work for ONoC reliability and further optimize the regulation energy overhead in Sec. 3.

However, even in the cases where every switch works stably and properly through thermal tuning, the system reliability of ONoC is also affected by the number of switching stages in routing paths. A large number of switching elements in a routing path would increase the total power loss of the path, resulting in a risk of unreliability. To ensure that the signal power received by the receiver on an optical path is not lower than its sensitivity, we deduce a boundary condition for the limit of the number of switching stages in routing paths from (3) as follow:

$$m \leq (P_{TX} - L_{WG} - S_{RX})/L_{SE}^{ideal} \tag{5}$$

where L_{SE}^{ideal} is the optical loss of switches at their desired temperature. The power loss due to the switches under temperature fluctuations is always greater than the loss of ideal switches ($\sum_{i=1}^{m} L_{SE_i} \geq m \cdot L_{SE}^{ideal}$).

To show the impact of the number of switching stages on the power loss of routing paths more intuitively, we conduct experiments to simulate the transmission processes of optical signals along the routing paths that have different number of switching stages based on the current manufacturing technology of silicon photonic devices. We employ the professional photonic component & circuit simulations [1], which are widely used by the nanophotonics community for design and verification. A recognized optical router Cygnus (see Fig. 1(b)) is employed for ONoCs in the simulations, in which at most one switching element is active for a specific transmission. We assume that all the switches in routers work at their desired temperature by applying thermal tuning.

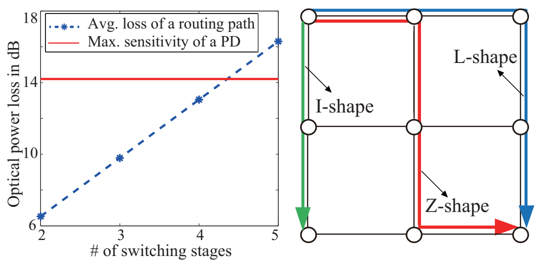

(a) Effect of # of switching stages on optical power loss. (b) Three types of routing paths.

Figure 2: Analysis on the # of switching stages in routing paths.

As shown in Fig. 2(a), the total power loss of a routing path is proportional to the number of switching stages in this path. The state-of-the-art PDs in the receivers achieve a sensitivity of -14.2 dBm [18]. Assumed the input power is 1 mW, we can observe that if the number of switching stages in a routing path is more than 4, the received power of the PD is lower than its sensitivity. Our results are consistent with those reported in [15]. This problem will be more serious with large temperature variations. Therefore, in this paper, we restrict the number of switching stages in routing paths to 4 for guaranteeing system reliability and, as a side benefit, energy-saving[1]. The less number of switching stages a route has, the less input power the transmitter requires for reliable transmission. Besides, the energy consumed by thermal tuning decreases with the decreasing number of switching elements. There are three types of shortest routing paths[2] for every communication pair: I-shape, L-shape, and Z-shape routes, as illustrated in Fig. 2(b).

Based on the theoretically formulated boundary condition of the number of switching stages in routing paths, we present a routing criterion at the network level considering the current silicon photonic fabrication technology. The criterion, combined with device-level thermal tuning, can implement thermal-reliable ONoC. In addition, it helps to exponentially reduce the search space when making routing decisions, which improves the efficiency of our approaches proposed in Sec. 3.

3 CONTENTION-AWARE ROUTING

Utilizing local thermal tuning and restricting the number of switching stages in routes, we guarantee the thermal reliability of generic ONoC, while it also brings a challenge: reducing the energy consumed for thermal tuning. In addition, we aim at mitigating communication contention by leveraging the inherent flexibility of adaptive routing in this paper. While the flexibility would decrease when the number of switching stages is restricted, which is adverse to contention reduction.

Considering the contradiction between communication performance, energy efficiency and thermal reliability, we formulate the problem as follow: *Given the floorplan of a thermal-reliable ONoC that is implemented by using the solutions presented in Sec. 2, and a communication demand represented by a set of pairs that require data transmission, determine a route for every pair to minimize communication conflict, and meanwhile, mitigate the energy overheads of thermal regulation in presence of chip thermal variations.*

3.1 MILP-based routing approach

We first model the problem by MILP formulations. Our model includes two phases: the major objective is to minimize communication conflict; based on it, we minimize the total energy consumption with considering chip thermal variations. The notations used in MILP formulations are summarized with definition in Tab. 1.

Table 1: Variables used in MILP model

Notation	Definition
D	The communication demand.
rP_i	The set of available routes (I/L/Z-shape paths by applying the routing criterion) of pair i.
L	The set of optical links.
$reql_{i,k}$	The set of required links for constructing the routing path k of pair i.
T_r	The ambient temperature of the router r.
$SE_{i,k}$	The set of switching elements in the routing path k of pair i.
$X_{i,k}$	Binary variable, equals to 1 if communication pair i routes along the routing path k.
$N_{c\times}$	The number of communication conflicts.
E_i	The energy consumption of comm. pair i.

In the first phase, we minimize communication contention.

$$Objective\ 1: \ minimize(N_{c\times}) \tag{6}$$

The key constraints are described as follows:

Constraint for communication demand: Equation (7) guarantees that at most one routing path is selected for every pair.

$$\sum_{k \in rP_i} X_{i,k} \leq 1 \qquad \forall i \in D \tag{7}$$

Constraint for route overlap: For each optical link $l \in L$, it can be occupied by zero or one routing path. If link l is occupied by a

[1]As the development of silicon photonic technology, this restriction can be relaxed due to low-loss optical device fabric.

[2]Minimal routing considers route length and route utilization, which will facilitate to mitigate the risk of contention.

route, $C_l = 1$. Then $\sum_{k\in rP_i} X_{i,k} * reql_{i,k} \leq 1$, which guarantees no route overlap.

$$\sum_{k\in rP_i} X_{i,k} * reql_{i,k} \leq C_l \qquad \forall i \in D \tag{8}$$

Constraint for communication conflict: Equation (9) calculate the number of conflicts.

$$N_{c\times} = D - \sum_{i\in D}\sum_{k\in rP_i} X_{i,k} \tag{9}$$

In the second phase, based on the minimum number of communication conflicts obtained from the first phase, we minimize total energy overhead with consideration of the regulation energy consumed by thermal tuning, shown as follow:

$$Objective\ 2:\ minimize(sum(E_i)) \quad i \in D \tag{10}$$

In addition to the constraints (7)-(8), we add two key constraints in this phase as follows:

Constraint for minimum conflict: Equation (11) guarantees the minimum contention, based on which we can further optimize the energy overhead in presence of chip thermal variations.

$$\sum_{i\in D}\sum_{k\in rP_i} X_{i,k} = D - N_{c\times} \tag{11}$$

Constraints for thermal-induced energy consumption: In this paper, we consider a set of tasks running in a relatively stable power profile, resulting in steady temperature states of the processor cores. Thus, the ONoC is in a steady temperature state in the concerned period of time. The derived Equation (12) estimates the total energy consumption of each pair, by taking thermal variations into account.

$$\begin{aligned} E_i = & X_{i,k} * \Big(\big(E^e_{int}\cdot W^e_{ctrl}\cdot h + E^e_{cu}\cdot(h+1)\big) \\ & + \big(m\cdot E^o_{MR} + E^o_{oeeo}\cdot W^o_{payload}\big) \\ & + \sum_{r\in SE_{i,k}} \big(\varepsilon\cdot\rho\cdot(T_r - T_0)\cdot(W^o_{payload}/R_{oeeo} + d\cdot n/c)\big)\Big) \end{aligned} \tag{12}$$

where W^e_{ctrl} and $W^o_{payload}$ refer to the size of control packet in the electrical network and the size of data packet in the optical network, respectively; h denotes the number of hops from the source to the destination; R_{oeeo} is the E-O and O-E interface data rate; d is the travel distance of a data packet; n is the refractive index of silicon waveguide; c is light speed in vacuum; $(W^o_{payload}/R_{oeeo} + d\cdot n/c)$ denotes the duration time of the switches' operation; E^e_{int} is the average energy required to transfer a single bit through an electrical interconnection; E^e_{cu} denotes the average energy required by the control unit to make a decision for a single packet; E^o_{MR} is the energy consumed by an active MR to transfer a data packet; E^o_{oeeo} is the energy required by a single bit for O-E and E-O conversions.

The optimal result obtained by our MILP model can be converted into a equivalent routing solution. We optimize communication performance by minimizing contention, and explore the design space to find the optimal routing solution with minimized energy consumption based on the minimum contention.

3.2 Heuristic algorithm

As the problem addressed in this paper is NP-hard, the running time of the proposed MILP model increases exponentially with the increasing sizes of network and communication demand. To effectively solve this complex issue, we further propose a heuristic contention-aware routing (CAR) algorithm.

As shown in Algorithm 1, we can firstly obtain the set of all the available routing paths for every communication pair by applying the routing criterion proposed in Sec. 2, which constitutes its communication region (denoted as *cRegion*) (Line 2–4). For a pair (src, des), there are at most $|des.x - src.x| + |des.y - src.y|$ paths in total. The route of each pair is selected from its communication region. For the pairs whose communication regions are independent of other regions, we are free to select a routing path based on the routes' temperature distribution, without the risk of contention (Line 5–7). The function `GetRouteWithMaxEnergyEfficiency()` is designed to find the route that maximizes energy efficiency among the communication region. For the pairs that there are overlaps between their regions, we should consider the risk of link conflict, potentially resulting in communication contention. To handle the potential conflict regions, we sort pairs according to the sizes of their communication regions, from small to large (Line 9). The pairs with small communication regions are prior to the pairs with large regions because it has less alternatives. It is easier for the pairs with large communication regions to find feasible routes even if some of links are occupied by other communications. In the most extreme case, no feasible route can be obtained for the pairs that there is only one available routing path and it has been occupied by others. We make routing decisions for pairs in the sorted order. For every communication pair, we select the route that achieves maximal energy efficiency on the basis of no communication contention ($cRegion \cap idleLs$), which can be realized by labeling the occupied/idle links (Line 10–15).

Algorithm 1: Contention-aware routing (CAR) algorithm.

Data: The *ONoC*, the communication demand *D*
Output: Routes for every communication pair *rPs*
1 *idleLs* = GetIdleLinks(*ONoC*);
2 **for** $cPair \in D$ **do**
3   *cRegion* = GetAvailableCommPaths(*cPair*, *idleLs*);
4   #A set of I-shape, L-shape, and Z-shape paths;
5 *IndePairs* = FindIndependPairs(*D.cRegion*);
6 **for** $cPair \in IndePairs$ **do**
7   *cPair.rP* = GetRouteWithMaxEnergyEfficiency(*cRegion*);
8 Update(*D*);
9 $D = sort(D, D.cRegion, Ascend)$;
10 **for** $cPair \in D$ **do**
11   *idleLs* = GetIdleLinks(*ONoC*);
12   **if** $src.x = des.x$ ***or*** $src.y = des.y$ **then**
13   *cPair.rP* = (src, des); #I-shape path;
14   **Continue**;
15   *cPair.rP* = GetRouteWithMaxEnergyEfficiency ($cRegion \cap idleLs$);
16 **return** (*rPs*);

Fig. 3 gives an example to illustrate our CAR algorithm. There are three communication pairs and we can obtain their communication regions. Since $cRegion_1$ is independent of the other regions, we are free to select a route for $Pair_1$ according to the temperature distributions of the routes $rP_{1,1}$ and $rP_{1,2}$. However, the communication regions of $Pair_2$ and $Pair_3$ are partly overlapped. The links $(R_6 \rightarrow R_2)$ and $(R_7 \rightarrow R_3)$ are potentially conflict when deciding routes for them. In this case, we first select a route for the pair that has smaller region ($Pair_2$ in this example). Assume the route $rP_{2,2}$ achieves maximum energy efficiency for $Pair_2$, the link $(R_7 \rightarrow R_3)$ is occupied. Thus the available communication region of $Pair_3$ is reduced to $\{rP_{3,1}, rP_{3,2}, rP_{3,3}\}$ and we can select one of routing paths based on their temperature distributions.

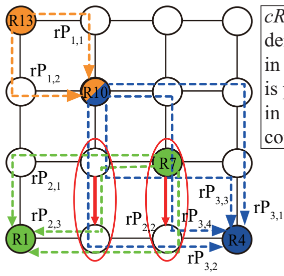

cRegion₁ (paths in orange) is independent of other regions; *cRegion₂* (paths in green) and *cRegion₃* (paths in blue) is partly overlapped, and the red links in the red circles constitute a potential conflict region.

	comm. pair	comm. Region
$cRegion_1$	$R_{13} \rightarrow R_{10}$	$\{rP_{1,1}, rP_{1,2}\}$
$cRegion_2$	$R_7 \rightarrow R_1$	$\{rP_{2,1}, rP_{2,2}, rP_{2,3}\}$
$cRegion_3$	$R_{10} \rightarrow R_4$	$\{rP_{3,1}, rP_{3,2}, rP_{3,3}, rP_{3,4}\}$

Figure 3: An example of the CAR algorithm.

Complexity analysis: The complexity of Algorithm 1 is $\mathcal{O}(M \times N)$, where M is the number of communication pairs and N^2 is the ONoC size. The complexity of looping through the $|des.x - srx.x| + |des.y - src.y|$ alternative paths to find a route with maximal energy efficiency for a pair is $\mathcal{O}(N)$ in the worst case, and there are M communication pairs in total.

Our approaches can be integrated into the centralized thermal-aware task scheduler in the operating system (OS). The scheduler gathers chip temperature information when making decisions for task assignment and scheduling. The communication requests during the period of task scheduling are known by the scheduler in advance. Therefore, the scheduler can calculate the routes for every pairs and give the routing information to the source nodes of the pairs during the period of task mapping. The setup process of routing path is conducted in the electronic control network. The electronic routers can transfer a control packet that contains route information to establish the desired routing path. We can achieve deadlock-free electronic control network with additional efforts, such as using virtual channel flow control. Both of our approaches are suitable to be applied for multiple communication requests, while we do not restrict the heuristic CAR algorithm to be applied for the cases with single communication request.

4 PERFORMANCE EVALUATION

We consider 2D mesh-based ONoCs with size range from 8×8 to 15×15 as target platforms. A simulator in MATLAB is built to evaluate our approaches on the target platforms. Considering the DVFS capability of modern processors, we model computation power under different voltage/frequency levels by McPAT v1.0 [8] for out-of-order Alpha 21346 cores in 22 nm technology. The generated processor power traces are used to estimate chip thermal distribution with HotSpot v5.02 [4]. We perform the approaches based on the generated chip thermal variations. Gurobi optimization with CVX v2.1 is employed as the MILP solver in this paper, which is designed to be the fastest, most powerful solver available for MILP problems. We base our experiments on a large number of synthetic communication traces and a set of realistic benchmarks that includes well-known StreamIt benchmarks [11] (*audiobeam*, *tde_pp*, *fmradio*, *filterbank*, and *beamformer*) as well as DSP-stone benchmarks [16] (*IIR filter* and *4_stagelattice*). The communication traces of the benchmarks are generated by using the task mapping algorithm proposed in [16]. Experiments are run on a workstation with Intel(R) Core(TM) i5-4590 at 3.30 GHz and 8 GB memory.

Evaluation results on average communication energy consumption, communication latency, network throughput, and link utilization are compared with two state-of-the-art techniques, including a thermal-sensitive routing algorithm proposed in [17] and an efficient contention-aware routing approach (DyXY) [7]. Work [17] focuses on the effect of thermal variations on the energy efficiency of ONoCs and the proposed routing mechanism is able to find the routing solution that has the minimum total energy consumption.

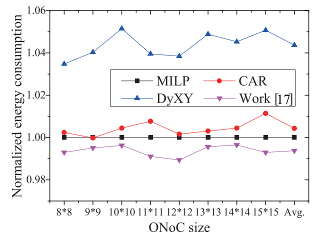

(a) Avg. energy consumption.

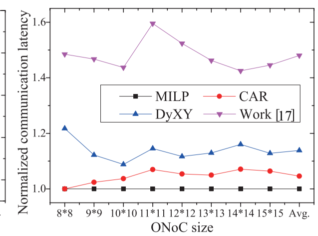

(b) Avg. comm. latency.

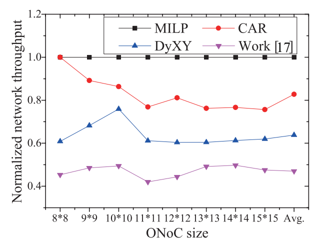

(c) Avg. network throughput.

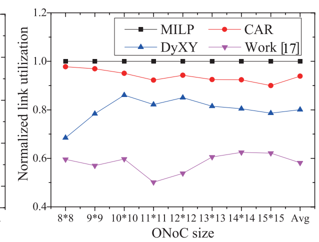

(d) Avg. link utilization.

Figure 4: Comparison of evaluation results obtained by different approaches based on synthetic communication traces.

Fig. 4 illustrates the result comparison based on synthetic communication traces. We conduct 100 groups of experiments for each size of ONoC. In every group of experiment, the sources, the destinations and the volume of communications among processors in a communication demand are randomly generated. The voltage/frequency levels of processors are randomly assigned, consequently resulting in a randomly-generated chip thermal distribution. Our approaches significantly improve communication performance in virtue of the minimized contention. The MILP-based approach achieves an average of 32.44% and 12.15% reduction in communication latency than work [17] and the DyXY, respectively. Our heuristic only increases the latency by 4.40% compared to the MILP model. Regarding to the network throughput that denotes the rate of optical packet transmitted through ONoC (measured in packets per second), the MILP-based approach achieves an average of 112.73% and 56.78% improvements than work [17] and the DyXY, respectively. Compared to the two techniques, our heuristic also averagely improves the network throughput by 76.04% and 29.74%. We denote the average ratio of busy links to the total links in an ONoC during a period of time as the average link utilization. Our MILP-based approach increases average link utilization by 71.89% and 24.87% compared to work [17] and the DyXY, respectively. The heuristic only has 6.52% utilization difference compared to the MILP model. In addition, compared to the work [17] that achieves excellent energy efficiency with the minimum power consumption, our MILP-based approach and the CAR algorithm only consume an average of 6.49 pJ/pck and 11.1 pJ/pck more energy, resulting in 0.62% and 1.06% extra energy overhead, respectively. Compared to the DyXY, our MILP model reduces 4.18% energy overhead on average.

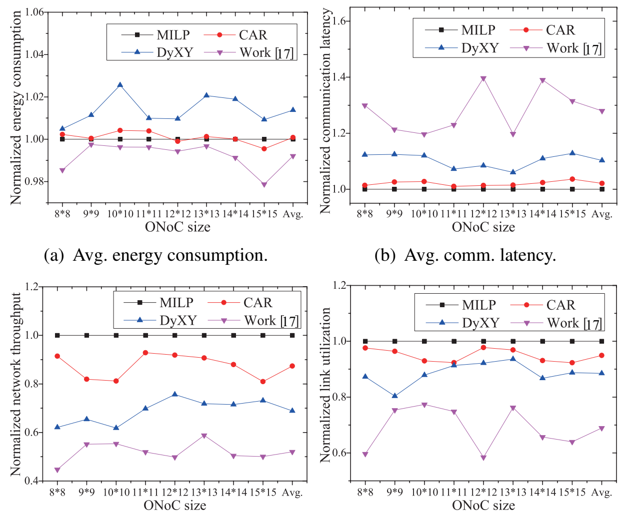

(a) Avg. energy consumption. (b) Avg. comm. latency.

(c) Avg. network throughput. (d) Avg. link utilization.

Figure 5: Comparison of evaluation results obtained by different approaches based on the realistic benchmarks.

Fig. 5 illustrates the evaluation results based on the realistic applications. Compared to work [17] and the DyXY, our MILP-based approach achieves an average of 21.87% and 9.32% reduction in communication latency, an average of 92.17% and 45.13% improvement in network throughput, and an average of 45.05% and 12.81% improvement in link utilization, respectively. The proposed heuristic only has an average of 7.33% performance difference compared to the MILP model. Regarding to the communication energy consumption, the MILP-based approach and the CAR algorithm consume an average of 8.03 pJ/pck and 8.55 pJ/pck more energy than work [17], resulting in only 0.80% and 0.88% extra energy overhead, respectively. The experimental results based on the realistic benchmarks are consistent with those based on the randomly-generated communication traces.

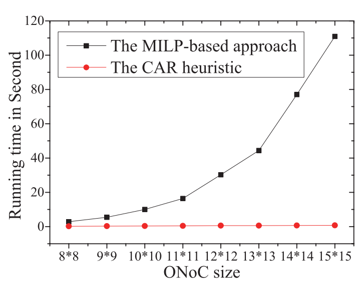

Figure 6: Comparison of the proposed approaches on algorithm running time.

Finally, we illustrate the running time of the proposed approaches in Fig. 6. Our MILP-based approach, as a formal method, can perform systematic exploration in the design space and provide a theoretically guaranteed optimal solution. Thanks to the exponentially reduced search space by applying the routing criterion proposed in Sec. 2, the time cost of our MILP model is acceptable. Meanwhile, our proposed polynomial-time algorithms are highly efficient with negligible running time (< 1s) compared with the MILP method, especially in large-scale ONoCs.

5 CONCLUSION

In this paper, we propose a routing criterion and two routing approaches that are designed for ONoC to guarantee reliability and improve communication performance and energy efficiency. Conflict-induced communication performance degradation and the energy overheads resulted from thermal tuning are minimized through our approaches based on guaranteed thermal reliability. Evaluations based on synthetic communication traces and realistic benchmarks show the effectiveness of our approaches, compared to state-of-the-art techniques.

ACKNOWLEDGMENT

This work is supported by NTU NAP M4082282 and SUG M4082087, HP-NTU Digital Manufacturing Corporate Lab, Singapore and NSFC 61772094, China.

REFERENCES

[1] [n. d.]. Lumerical Inc. https://www.lumerical.com/.
[2] Keren Bergman, Luca P Carloni, Aleksandr Biberman, Johnnie Chan, and Gilbert Hendry. 2014. *Photonic network-on-chip design*. Springer.
[3] Huaxi Gu, Kwai Hung Mo, Jiang Xu, and Wei Zhang. 2009. A low-power low-cost optical router for optical networks-on-chip in multiprocessor systems-on-chip. In *IEEE Computer Society Annual Symposium on VLSI*. 19–24.
[4] Wei Huang, Shougata Ghosh, Sivakumar Velusamy, Karthik Sankaranarayanan, Kevin Skadron, and Mircea R Stan. 2006. HotSpot: A compact thermal modeling methodology for early-stage VLSI design. *IEEE TVLSI* 14, 5 (2006), 501–513.
[5] Sébastien Le Beux, Jelena Trajkovic, Ian O'Connor, Gabriela Nicolescu, Guy Bois, and Pierre Paulin. 2011. Optical ring network-on-chip (ORNoC): Architecture and design methodology. In *Pro. of DATE*.
[6] Hui Li, Alain Fourmigue, Sébastien Le Beux, Xavier Letartre, Ian O'Connor, and Gabriela Nicolescu. 2015. Thermal aware design method for vcsel-based on-chip optical interconnect. In *Pro. of DATE*. 1120–1125.
[7] Ming Li, Qing-An Zeng, and Wen-Ben Jone. 2006. DyXY: a proximity congestion-aware deadlock-free dynamic routing method for network on chip. In *Pro. of DAC*. 849–852.
[8] Sheng Li, Jung Ho Ahn, Richard D Strong, Jay B Brockman, Dean M Tullsen, and Norman P Jouppi. 2009. McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures. In *IEEE/ACM International Symposium on Microarchitecture*. 469–480.
[9] Zhongqi Li, Amer Qouneh, Madhura Joshi, Wangyuan Zhang, Xin Fu, and Tao Li. 2015. Aurora: A cross-layer solution for thermally resilient photonic network-on-chip. *IEEE TVLSI* 23, 1 (2015), 170–183.
[10] Weichen Liu, Peng Wang, Mengquan Li, Yiyuan Xie, and Nan Guan. 2017. Quantitative Modeling of Thermo-Optic Effects in Optical Networks-on-Chip. In *Pro. of GLSVLSI*. ACM, 263–268.
[11] Weichen Liu, Juan Yi, Mengquan Li, Peng Chen, and Lei Yang. 2018. Energy-Efficient Application Mapping and Scheduling for Lifetime Guaranteed MPSoCs. *IEEE TCAD* (2018).
[12] Moustafa Mohamed, Zheng Li, Xi Chen, Li Shang, and Alan R Mickelson. 2014. Reliability-aware design flow for silicon photonics on-chip interconnect. *IEEE TVLSI* 22, 8 (2014), 1763–1776.
[13] Kishore Padmaraju and Keren Bergman. 2014. Resolving the thermal challenges for silicon microring resonator devices. *Nanophotonics* 3, 4-5 (2014), 269–281.
[14] Sebastian Werner, Javier Navaridas, and Mikel Luján. 2015. Amon: An advanced mesh-like optical noc. In *Pro. of HOTI*. 52–59.
[15] Yiyuan Xie, Mahdi Nikdast, Jiang Xu, Wei Zhang, Qi Li, Xiaowen Wu, Yaoyao Ye, Xuan Wang, and Weichen Liu. 2010. Crosstalk noise and bit error rate analysis for optical network-on-chip. In *Pro. of DAC*. 657–660.
[16] Lei Yang, Weichen Liu, Peng Chen, Nan Guan, and Mengquan Li. 2017. Task mapping on smart noc: Contention matters, not the distance. In *Pro. of DAC*. 1–6.
[17] Kang Yao, Yaoyao Ye, Sudeep Pasricha, and Jiang Xu. 2017. Thermal-sensitive design and power optimization for a 3D torus-based optical NoC. In *Pro. of ICCAD*. 827–834.
[18] Yaoyao Ye, Jiang Xu, Xiaowen Wu, Wei Zhang, Xuan Wang, Mahdi Nikdast, Zhehui Wang, and Weichen Liu. 2013. System-level modeling and analysis of thermal effects in optical networks-on-chip. *IEEE TVLSI* 21, 2 (2013), 292–305.