

A Backup Resource Customization and Allocation Method for Wavelength-Routed Optical Networks-on-Chip Topologies

Zhidan Zheng, You-Jen Chang, Liaoyuan Cheng, Tsun-Ming Tseng, Ulf Schlichtmann
{zhidan.zheng,youjen.chang,liaoyuan.cheng,tsun-ming.tseng,ulf.schlichtmann}@tum.de
Technical University of Munich, Munich, Germany

ABSTRACT

Wavelength-routed networks-on-chip (WRONoCs) are known for providing high-speed and low-power communication. Despite those advantages, the key components, microring resonators (MRRs), are prone to process and thermal variations, which cause signals to fail to reach their intended destinations. Thus, several WRONoC fault-tolerant methods propose to prepare a constant number of backups, which often leads to inefficient resource allocation, i.e. insufficient backups for the signals that are prone to errors, while more than enough backups for the signals that are barely affected, resulting in much power waste. In this work, we propose a dynamical backup resource allocation method for reliability maximization and power minimization in WRONoCs. Precisely, our method starts with accurately modeling the WRONoC faults, which considers the deviation of an MRR's default behavior as a Gaussian Distribution. Since signal paths consist of different numbers of MRRs, and the signals have different probabilities of deviating from their designated paths, our method customizes the number of backup paths for every signal and automatically allocates the minimum resources to optimize the reliability.

KEYWORDS

Wavelength-routed optical networks-on-chip, fault-tolerant design, backup resource allocation, fault model.

ACM Reference Format:

Zhidan Zheng, You-Jen Chang, Liaoyuan Cheng, Tsun-Ming Tseng, Ulf Schlichtmann. 2025. A Backup Resource Customization and Allocation Method for Wavelength-Routed Optical Networks-on-Chip Topologies. In *30th Asia and South Pacific Design Automation Conference (ASPAC '25)*, January 20–23, 2025, Tokyo, Japan. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3658617.3697783>

1 INTRODUCTION

To meet the extraordinary interconnect requirements of many-core chips, optical networks-on-chips (ONoCs) are considered as a next-generation solution for several critical advantages [1–7]: high bandwidth thanks to the wavelength-division multiplexing technology, low transmission latency in waveguides, high modulation speed (10 – 40 Gb/s), and low power consumption with microring resonators (MRRs). Therefore, ONoCs attract increasing research

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 the author(s) 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.
ASPAC '25, January 20–23, 2025, Tokyo, Japan

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0635-6/25/01
<https://doi.org/10.1145/3658617.3697783>

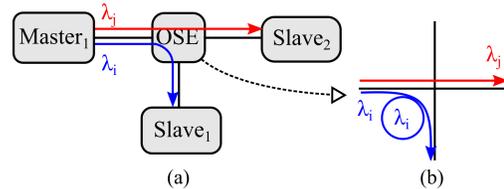


Figure 1: (a) A WRONoC topology, where a master sends two signals modulated on wavelengths λ_i and λ_j to two slaves, Slave₁ and Slave₂, respectively. (b) An implementation of the OSE in (a). The signal on λ_i is on-resonance to the MRR, and the signal on λ_j is off-resonance to the MRR in the OSE.

interest and industrial investment [1–5].

Among all categories of ONoCs, *wavelength-routed* ONoCs (WRONoCs) are well-known for supporting collision-free and reconfiguration-free communication. On WRONoCs, a master talks to a slave at any time without data collision using a pre-defined signal path, i.e. the routing of the signal path and the configuration of the MRRs along the signal path are reserved and fixed during the design phase. In other words, no time and energy for arbitration are required in WRONoCs, which makes WRONoCs a promising option for high-speed on-chip communication [3, 8–12].

In WRONoCs, a signal path that supports a communication from a master to a slave is established by connecting the network components by waveguides and configuring the wavelength of MRRs [12, 13]. Figure 1(a) shows that a master and two slaves are connected by waveguides to an optical switching element (OSE), which is formed by an MRR and two orthogonal waveguides, as shown in Figure 1(b). The MRR of this OSE is configured to resonate to the wavelength λ_i . Both signals from the master travel along the horizontal waveguide until the signal on λ_i is demultiplexed by the MRR of the OSE. Specifically, when the signal on λ_i approaches this MRR, it is coupled to the MRR and leaves the MRR via the vertical waveguide. On the other hand, the signal on λ_j ignores the MRR and keeps its original propagation direction along the horizontal waveguide to reach its planned designation, Slave₂.

Despite the advantages of WRONoCs, MRRs, the key components of WRONoCs, are highly sensitive to process and thermal variations [14–17], which raises the reliability concern in WRONoCs. For example, a change of 1 °C in temperature can shift the resonant wavelength of an MRR by 0.1 nm [15]. If the resonant wavelength of an MRR shifts to a different wavelength than was intended, the signals that should be on-resonance to the MRR cannot be coupled to the MRR and thus fail to reach their designated destinations. As a result, communications relying on those signal paths are lost, which severely degrades the reliability of WRONoCs. As shown in Figure 2(a), the resonant wavelength of the MRR shifts from λ_i to $\lambda_i + \Delta\lambda$. Due to the shift, the signal on λ_i fails to be coupled to the

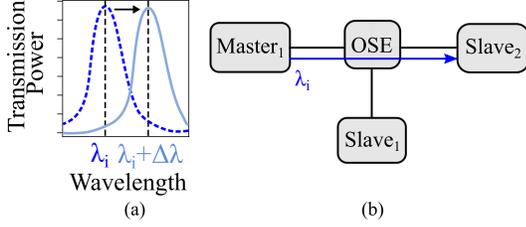


Figure 2: (a) The shift of a transmission spectrum for an MRR, which is resonant to $\lambda_i + \Delta\lambda$ rather than its planned wavelength λ_i . (b) The signal on λ_i cannot be coupled to the MRR and does not reach its designated destination, Slave₁.

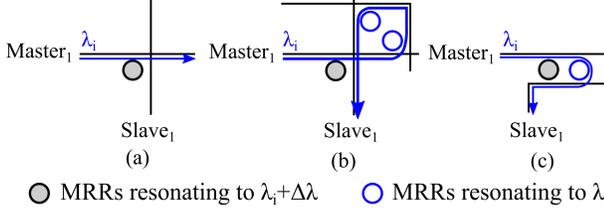


Figure 3: (a) The signal on λ_i fails to be coupled to the MRR due to the shift of the MRR’s resonant wavelength. (b) Two backup MRRs and a waveguide are inserted so that Master₁ can talk to Slave₁ again. (c) The backup MRR is placed right next to the original MRR without forming extra crossings.

MRR and deviates from its planned path, as shown in Figure 2(b), which causes the loss of the communication from Master₁ to Slave₁.

To date, several fault-tolerant WRONoC design methods have been proposed to prepare backup signal paths [18–20]. Specifically, they add extra waveguides and MRRs to a WRONoC topology and construct one backup path for every master-slave communication pair. In this case, if a fault occurs on a signal path, as shown in Figure 3(a), the master can use a backup path to communicate with the slave, as shown in Figure 3(b), so that the system can function correctly. In particular, the latest fault-tolerant design, LightR [19], proposes an efficient way of inserting MRRs by taking advantage of the topology built with parallel switching elements (PSEs), where MRRs are placed between two parallel waveguides. As shown in Figure 3(c), LightR places a backup MRR adjacent to each original MRR in a topology, which is easy to implement and avoids the power penalties caused by the extra waveguide crossings that appear in other fault-tolerant designs.

Despite the effective ways of establishing backups, the backup resources in current fault-tolerant designs are not efficiently allocated and utilized. On one hand, there are insufficient backup paths for the communications relying on the error-prone signal paths. Precisely, all communications in current designs [18–20] always have a constant number of backup paths. Thus, for the communications supported by the error-prone signal paths, more backup paths should be considered to maximize their reliability. On the other hand, some communications use the signal paths that are already reliable, and thus their backup paths are usually redundant. This redundancy introduces a waste of resources and an unnecessary increase in power overhead.

Besides, current methods have overlooked many potential faults. Typically, they assume that only one single fault will appear, regardless of the network scale, and no backup MRRs can cause

faults [18, 20]. That assumption is rather unrealistic because multiple MRRs can cause faults in reality, especially for large-scale networks [15, 16]. Moreover, backup MRRs, similar to the regular MRRs, can also lead to the deviation of signal paths. Their fault models have ignored those potential faults, and thus can hardly predict the reliability of current methods accurately.

The main contributions of this paper are summarized as follows:

- We propose an accurate fault model that considers the shift of an MRR’s transmission spectrum as a Gaussian Distribution, i.e. which models if a signal path can correctly behave or fail. Given a WRONoC topology, our fault model outputs the probability of each signal path failing to reach its designated destination, which reflects the reliability of the communication relying on that path.
- We propose the first method that can dynamically allocate backup resources for fault-tolerant WRONoC topologies. To optimize reliability and avoid redundant backups, we customize the number of backup signal paths for the communications according to the outputs of our fault model. To further improve the energy efficiency, our method takes advantage of PSEs and establishes the topologies correspondingly with a path search algorithm to look for the signal paths with the minimum increase in resource usage and insertion loss.

We compare our method to four state-of-the-art fault-tolerant design methods: RobustONoC [18], LightR [19], Actin-STAR [20], and Zygo-STAR [20]. The experimental results demonstrate our superiority in improving reliability. For example, for a large-scale network, our method increases the worst-case probability of a communication not being lost by 4% – 12% compared to the three latest methods: LightR, Actin-STAR, and Zygo-STAR.

2 BACKGROUND

2.1 ONoC Faults

In ONoCs, MRRs are very susceptible to process and temperature variations [15–17]. When the transmission spectrum of an MRR shifts to a different wavelength than was intended, it can cause faults in signal paths. There are typically two types of faults: *on-resonance* and *off-resonance* faults. Specifically, if a signal that should be on-resonance to an MRR cannot be coupled to this MRR, as shown in Figure 2, it suffers an on-resonance fault at this MRR. On the other hand, if a signal that should be off-resonance to an MRR is erroneously coupled with this MRR, it suffers an off-resonance fault. Both faults can result in transmission errors, i.e., a signal cannot reach its designated destination, which concerns the reliability of ONoCs.

2.2 Typical WRONoC Topologies

In WRONoCs, a topology specifies the logic connections among OSEs and the resonant wavelengths of the MRRs and the signals [3, 21]. Figure 4(a) and (b) show two topologies built with two typical designs of OSEs: a crossing switching element (CSE) and a parallel switching element (PSE), respectively. In both topologies, each master-slave communication has only one signal path. For example, master m_1 talks to slave s_4 in a 4×4 Light topology using the signal path represented by the green arrow shown in Figure 4(b). Without backups, the communications in those topologies can be lost when

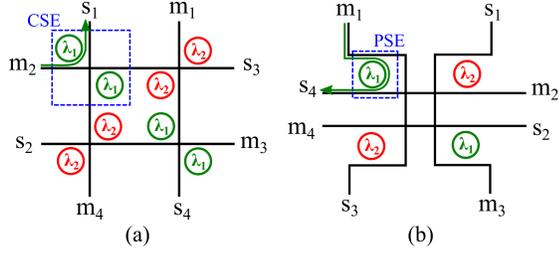


Figure 4: (a) A 4×4 GWOR topology built with four CSEs. Each CSE consists of a pair of orthogonal waveguides and two MRRs resonating to the same wavelength. Signals on that wavelength will experience a 90° change in their propagation directions when they are coupled to the MRRs, such as the signal represented by the green arrow. (b) A 4×4 Light topology built with four PSEs. Each PSE consists of a pair of parallel waveguides and an MRR, which can change the propagation directions of on-resonance signals by 180° .

the MRRs along any signal path cannot work normally and the signals fail to reach the designated slaves. Therefore, backup paths are important to improve the reliability of WRONoCs.

2.3 Related Works

RobustONoC [18] and LightR [19] are two fault-tolerant design methods that prepare backup paths in GWOR and Light topologies, respectively. They have different ways of inserting backup MRRs: RobustONoC adds extra waveguides and MRRs in a GWOR to form CSEs, as shown in Figure 5(a), while LightR makes full use of the PSEs by placing a backup MRR adjacent to an original MRR in Light topology, as shown in Figure 5(b). Thus, the results of RobustONoC usually have many extra MRRs and waveguide crossings. The 4×4 GWOR topology generated by RobustONoC shown in Figure 5(a) increases the numbers of MRRs and crossings by 67% and 80%, respectively, compared to the original GWOR topology shown in Figure 4(a). That introduces a large increase in insertion loss. In contrast, LightR avoids extra crossings and high MRR usage, which reduces the insertion loss compared to RobustONoC.

Actin-STAR and Zygo-STAR are two fault-tolerant WRONoC topologies [20]. They have new topological structures and avoid forming the waveguide crossings outside the CSEs, as shown in Figure 5(c) and (d). Therefore, compared to RobustONoC, those two topologies reduce the number of waveguide crossings, which decreases insertion loss. Moreover, they optimize their wavelength usage compared to other fault-tolerant design methods. However, they still have higher MRR usage than LightR.

The performance of current methods is challenged by inaccurate fault models and inefficient backup resource allocation. Specifically, they always prepare a constant number of backup paths for all communications, which is inefficient in matching the different reliability requirements. Moreover, current methods assume a fixed number of faults in a topology and randomly decide the malfunctioning MRRs for their fault models, which can hardly reflect the reliability accurately and comprehensively.

3 METHODOLOGY

In this paper, we propose a fault model to identify the reliability of a WRONoC topology and a method to customize the number

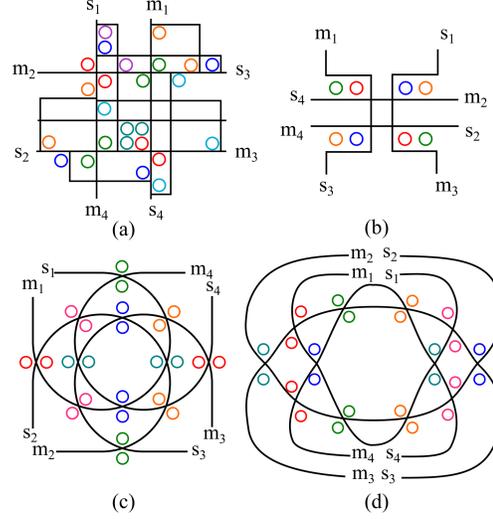


Figure 5: (a) A 4×4 fault-tolerant GWOR topology generated by RobustONoC [18]. (b) A 4×4 fault-tolerant Light topology generated by LightR [19]. (c) A 4×4 Actin-STAR topology [20]. (d) A 4×4 Zygo-STAR topology [20].

of backup paths and dynamically allocate backup resources according to the different reliability requirements. Specifically, we construct backup paths and take the PSE-based and scalable Light topology [13] as the starting point for reliability maximization and resource usage minimization.

3.1 Fault Model

To model the faults, we first introduce the transmission spectrum of an MRR. Specifically, an MRR's *on-resonance transmission*^a, denoted as $T_{on,res}(\lambda)$, can be calculated as follows [6]:

$$T_{on,res}(\lambda) = \frac{(1 - r^2)^2 a}{1 - 2ar^2 \cos(n_{eff}(\lambda)L) + (r^2 a)^2} \quad (1)$$

where a is the single-pass amplitude transmission, r is the self-coupling efficiency, and $n_{eff}(\lambda)L$ is the single-pass phase shift with L the round trip and $n_{eff}(\lambda)$ the wavelength-dependent effective index^b. For example, Figure 6(a) shows the transmission spectrum of an MRR, where the wavelengths corresponding to the peaks represent the resonant wavelengths, and the wavelengths at the troughs are the non-resonant wavelengths. When the process or thermal variation happens, the transmission spectrum of an MRR shifts, which causes the wavelengths corresponding to the peaks and troughs to differ from their initially intended wavelengths, denoted as *target* wavelengths.

Since it has been proved that wavelength shifts under process variation can be described by a Gaussian distribution [23], we model the shift as $S \sim \mathcal{N}(\mu = 0, \sigma^2)$, where μ represents the mean value and σ denotes the standard deviation of the wavelength shift. We set σ as the minimum positive difference between the target resonant wavelength and the wavelengths at the midpoints of the transmission spectrum^c, given by $\frac{T_{on,res}(\lambda_p) - T_{on,res}(\lambda_t)}{2}$, where $T_{on,res}(\lambda_p)$

^aThe power transmission happens when a signal on λ is coupled to an MRR.

^bIt is proposed in [22] that n_{eff} is equal to $2.57 - 0.85 * (\lambda[\mu\text{m}] - 1.55)$.

^cBy setting the distance between these two wavelengths as σ , the probability of the shift being within this range in a Gaussian distribution is 68.2% [24], thereby ensuring a high production yield.

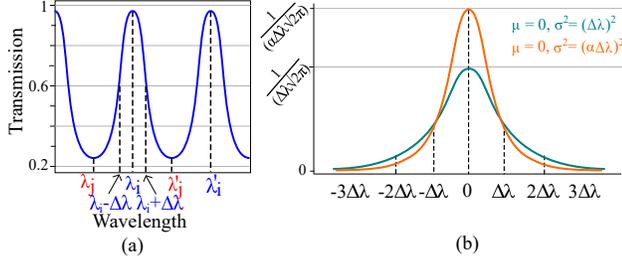


Figure 6: (a) The transmission spectrum of an MRR, which resonates to λ_i and λ_i' and does not resonate to λ_j and λ_j' . (b) Two models $S \sim \mathcal{N}(0, (\Delta\lambda)^2)$ and $S \sim \mathcal{N}(0, (\alpha\Delta\lambda)^2)$ for the deviation from λ_i and λ_j , respectively.

and $T_{on,res}(\lambda_t)$ are the maximum and minimum transmission values, respectively. For an MRR with a target resonant wavelength λ_i , we introduce a set Λ_{mid} for the wavelengths corresponding to the middle points and set $\sigma = \min\{|\lambda_{mid} - \lambda_i|\}$, where $\forall \lambda_{mid} \in \Lambda_{mid}$. As shown in Figure 6(a), the wavelengths of the two middle points closest to the target resonant wavelength λ_i are $\lambda_i + \Delta\lambda$ and $\lambda_i - \Delta\lambda$. In this case, we set σ to $(\lambda_i + \Delta\lambda) - \lambda_i = \Delta\lambda$. The model for λ_i $S \sim \mathcal{N}(0, (\Delta\lambda)^2)$ is represented by the green curve in Figure 6(b).

For an MRR with a target non-resonant wavelength λ_j , we determine σ similarly. For instance, the wavelength of the middle point closest to λ_j is $\lambda_i - \Delta\lambda$, as shown in Figure 6(a). Thus, we set σ as $\lambda_i - \Delta\lambda - \lambda_j$. The resulting shift model, $S \sim \mathcal{N}(0, (\alpha\Delta\lambda)^2)$, where $\alpha = \frac{\lambda_i - \lambda_j}{\Delta\lambda} - 1$, is represented by the orange curve in Figure 6(b).

Based on the models, we calculate the probability that a signal suffers an on-resonance or an off-resonance fault at an MRR. For an MRR with a shift model $S \sim \mathcal{N}(0, (\Delta\lambda)^2)$, if the resonant wavelength of the MRR deviates from the λ_i and most signal power cannot be coupled to the MRR, the signals on λ_i suffer on-resonance faults at this MRR. To determine the value, we introduce a threshold transmission denoted as T_{th} and a set Λ_{th} of the wavelengths corresponding to T_{th} . For a target resonant wavelength λ_i , if the resonant wavelength of the MRR shifts to the wavelength greater than the closest largest value, denoted as $\lambda_{th,i}^r$, or smaller than the closest smallest value, denoted as $\lambda_{th,i}^l$, we consider that the signal on λ_i suffers an on-resonance fault at this MRR. Based on the shift model $S \sim \mathcal{N}(0, (\Delta\lambda)^2)$, when the shift of resonant wavelength is out of the range $(-d_1, d_2)$, where $d_1 = \lambda_i - \lambda_{th,i}^l$ and $d_2 = \lambda_{th,i}^r - \lambda_i$, we consider that the on-resonance fault happens and calculate the probability $P_{on} = P(x \leq -d_1) + P(x \geq d_2)$ using the Cumulative Distribution Function (CDF) of the Gaussian distribution:

$$P(x \leq -d_1) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{-\frac{d_1}{\Delta\lambda}} e^{-\frac{1}{2}t^2} dt \quad (2)$$

$$P(x \geq d_2) = 1 - P(x \leq d_2) = 1 - \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{d_2}{\Delta\lambda}} e^{-\frac{1}{2}t^2} dt$$

Similarly, we apply the same way to calculate the probability of a signal suffering an off-resonance fault at an MRR, denoted as P_{off} .

We set the parameters of our fault model as follows. Assuming that an MRR with a radius $30 \mu\text{m}$ has a target resonant wavelength 1503.99 nm and a target non-resonant wavelength 1505.63 nm , we can calculate their transmission values using the Eq. 1^d and have

^dWe set the single-pass amplitude transmission $a = 1$ and the self-coupling efficiency $r = \sqrt{(1 - 0.85^2)} = 0.53$ according to [25].

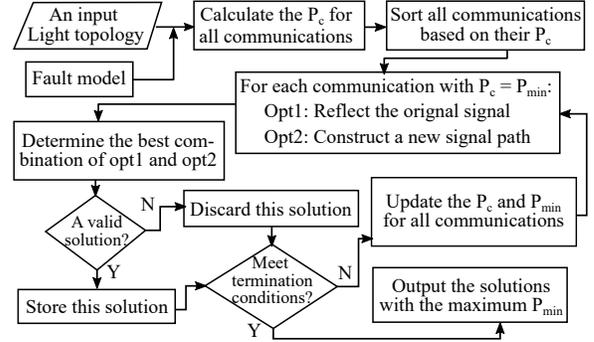


Figure 7: The flow chart of our method.

our shifts models for 1503.99 nm and 1505.63 nm : $S \sim \mathcal{N}(0, 0.58^2)$ and $S \sim \mathcal{N}(0, 0.36^2)$, respectively. By setting $T_{th} = 0.35^e$, we calculate the probabilities of causing an on-resonance fault and an off-resonance fault using Eq. 2. Finally, we have the $P_{on} = 4.2\%$ and $P_{off} = 0.5\%$. In the rest parts of this paper, we use these two values to check the reliability of signal paths.

For a signal path, denoted as s , consisting of N_s^{on} on-resonance MRRs and N_s^{off} off-resonance MRRs, we calculate the possibility of reaching its designated destination as $P_s = (1 - P_{on})^{N_s^{on}} * (1 - P_{off})^{N_s^{off}}$. If any of the N_s^{on} MRRs causes an on-resonance fault or any of the N_s^{off} MRRs causes an off-resonance fault, s can fail to reach its designated slave.

Given a topology, we create a set \mathbb{C} to store all master-slave communication pairs. If $c \in \mathbb{C}$ relies on only one signal path s , its probability of not being lost, denoted as P_c , is equal to P_s . On the other hand, if c relies on more than one signal path, its P_c is related to the P_s of every signal path s that supports c . We introduce a set \mathbb{S}_c to store the signal paths of c . A communication c is lost only when all of its signal paths fail to reach the planned slave. Therefore, we calculate its P_c as: $1 - \prod_{s \in \mathbb{S}_c} (1 - P_s)$. Among all communications, we denote the minimum P_c as P_{min} , which reflects the worst-case reliability of this topology.

3.2 Backup Resource Customization and Allocation

Figure 7 shows the flow of our method to customize backup resources for the communications in an input Light topology. By applying our fault model and calculating the P_c for every $c \in \mathbb{C}$, we identify the P_{min} and sort the communications in ascending order. For the communication with $P_c = P_{min}$, we have two approaches to increase their P_c values: the first is to reflect its original signal in the topology, and the second is to construct a new signal path. We will introduce the details of both approaches later. There are three possible combinations of them: only using the first, or the second, or using both of them. For each combination, we check the resulting new P'_c and the resource usages. Among the combinations resulting in the new P'_c larger than the previous P_c , we choose the one that leads to the least increase in MRR usage and the largest improvement in P_c . For the chosen combination, we determine its resulting P'_{min} . If it is larger than or equal to the previous P_{min}

^eThe 0.35 transmission is a crucial threshold for power compensation. A transmission below 0.35 can lead to significant power compensation, resulting in excessive crosstalk noise and a low signal-to-noise ratio (SNR).

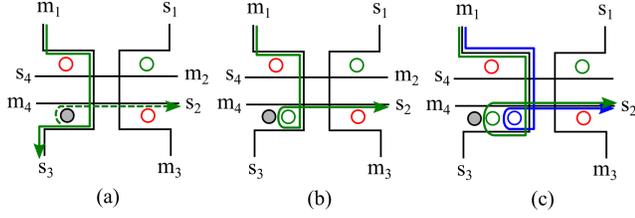


Figure 8: (a) The signal that intends to support the communication between m_1 to s_2 deviates from its planned path represented by the green dashed line and suffers an on-resonance fault. (b) A backup MRR, which is configured to resonate to the wavelength of the signal path from m_1 to s_2 , is placed right next to the original on-resonance MRR of this signal path. (c) A new signal path represented by the blue line for the communication from m_1 to s_2 .

minus a user-defined tolerance value ϵ , we mark this solution valid and store it; otherwise, we discard the solution and repeat the process for the other communications with P_{min} until we meet the termination conditions: either the latest resulting P_{min} is greater than or equal to 99.9%, the largest possible value, or we have not found a valid solution for X consecutive times. X is a user-defined coefficient to control the computational effort.

The first approach is inserting MRRs to reflect the deviating signal to its original path. Specifically, we apply the way shown in Figure 3(c) to place an MRR adjacent to the on-resonance MRR along a signal path and configure the backup MRR to resonate to the wavelength of this signal path. Figure 8(a) and (b) show an example. Note that the backup MRRs must be placed after the position where the signal meets its original on-resonance MRR; otherwise, they will cause additional faults.

The second approach is constructing a new signal path as a backup for a communication, which means introducing a new wavelength and a set of MRRs configured to resonate to this wavelength. For example, Figure 8(c) shows a signal path represented by the blue line for the communication from m_1 to s_2 . To establish a new signal path with minimized resource usage and insertion loss, we propose a two-step path-establishment method:

Step1: we model the interconnection of a Light topology as a graph and apply the breadth-first search (BFS) algorithm to find all possible signal paths for a communication. For simplicity, we consider multiple sequentially connected PSEs as a *PSE array*, as highlighted by the red dashed square shown in Figure 9(a), and treat it as a single component with two inputs and two outputs. Then, we model a PSE array with a small graph with four vertices and four edges, as shown in Figure 9(b). To generate the graph modeling the interconnection of an $N \times N$ Light, we identify all $N(\frac{N}{2} - 1)$ PSE arrays, index them as $1, \dots, N(\frac{N}{2} - 1)$, and model each of them as the small graph. Moreover, we introduce $2N$ vertices to indicate the masters and the slaves and add directed edges to connect the masters and the slaves to the inputs and outputs of the PSE arrays aligned with the logic scheme of the Light topology. For example, Figure 9(c) shows a graph for a 4×4 Light topology. Based on the graph, we search for all possible signal paths for a communication using BFS and set them as candidates.

Step2: for every candidate, we check its resource usage and insertion loss. Firstly, we check if the candidate signal path can reuse

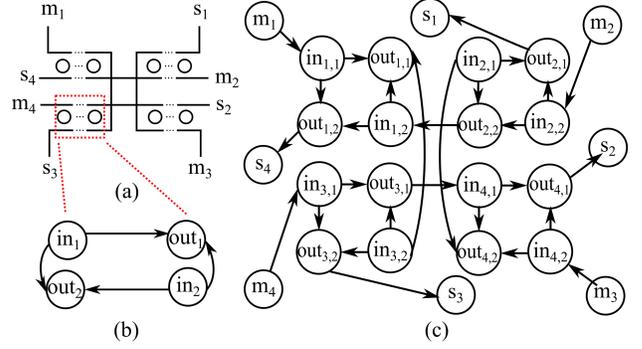


Figure 9: (a) A 4×4 Light topology with four PSE arrays. (b) A small graph for a PSE array, where each vertex represents an input or an output, and each directed edge represents the path from an input to an output. The signal paths from in_1 to out_2 and from in_2 to out_1 rely on MRRs. (c) The graph representing the interconnection of a 4×4 Light, where $in_{i,j}/out_{i,j}$ means the input/output j for the i -th PSE array.

the wavelengths and MRRs that have been used by the existing signal paths. To this end, we propose a vertex coloring model, where all existing signal paths and the candidate signal path are represented as vertices and wavelengths as a set of colors. An undirected edge is added between two vertices representing the overlapping signal paths that should not use the same wavelengths to avoid data collision. To solve the problem, we apply the Recursive Largest First (RLF) algorithm [26]. When we finish assigning the wavelengths to the signal paths, we configure the on-resonance MRR along the paths to resonate to the corresponding signals and calculate the resource usage. Note that the signal paths on the same wavelengths can share the same MRRs if they use the paths from $in_{i,1}$ to $out_{i,2}$ and from $in_{i,2}$ to $out_{i,1}$ of the i -th PSE array. For the insertion loss of a signal path, we apply the same way in [13] to calculate it. Among the candidate signal paths, we select the one that minimizes resource usage. If there are multiple candidates, we choose the one with minimum insertion loss.

4 RESULTS

We implemented our method in C++ and conducted all experiments in this paper on a computer with an Intel 8-core 2.6 GHz CPU. To evaluate the performance of our method, we compare it to four state-of-the-art fault-tolerant WRONoC methods: RobustONoC [18], LightR [19], Actin-STAR [20], and Zygo-STAR [20], in terms of three aspects: the worst-case reliability, energy efficiency, and scalability.

4.1 Discussion: The Worst-Case Reliability

We test all methods for five networks: a 4×4 , 6×6 , 8×8 , 12×12 , and 16×16 network^f. In an $N \times N$ network, there are $N(N - 1)$ communications, i.e. every master communicates to all other slaves. For the topologies synthesized by all methods, we apply our fault model and identify the worst-case reliability among all communications, i.e. P_{min} . For all networks, our method prepares backup paths for the communications with the setting $\epsilon = 1\%$ and

^fNote that for the comparison to RobustONoC, we only consider the 4×4 network since the MRRs in RobustONoC in larger sizes of networks have different settings than ours and the other state-of-the-art methods.

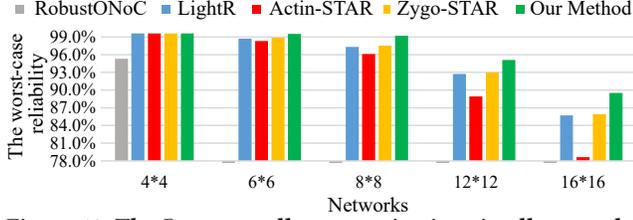


Figure 10: The P_{min} over all communications in all networks.

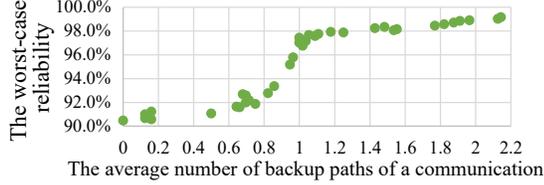


Figure 11: The average number of backup paths for each communication and the resulting P_{min} in our 8×8 network.

$X = 1000^{\text{§}}$. Among all valid solutions, we choose the one with the largest P_{min} and compare the results to that of the other methods.

Figure 10 shows the comparison results. In all cases, our method improves the reliability compared to the state-of-the-art methods. The improvement becomes more significant when the network scales up. For example, for a 16×16 network, our method increases the P_{min} by 12.2%, 4.2%, and 4.0% compared to Actin-STAR, LightR, and Zygo-STAR, respectively. The improvement of reliability in our fault-tolerant topologies is mainly driven by effectively customizing the number of backups for each communication. As shown in Figure 11, for the 8×8 network, the worst-case reliability is maximized when each communication has, on average, 2.14 backup paths. In contrast, state-of-the-art methods always prepare one backup path for every communication. The constant number of backup paths can hardly ensure high reliability, especially for large-scale networks.

4.2 Discussion: Energy Efficiency

Constructing different numbers of backup paths results in not only different reliability but also various power overhead. To evaluate the energy efficiency of all methods, for each $c \in \mathbb{C}$ of a network, we identify the worst-case insertion loss value over all its signal paths, denoted as il_c^{max} , and calculate the power of this signal path as $10^{(il_c^{max} + E_{pd})/10}$ [27], where E_{pd} is a constant value and denotes the photodetector sensitivity. We apply the loss and power parameters proposed from [20] and [9]. By summing up the signal power for each $c \in \mathbb{C}$, we obtain the total signal power. For the comparison in this subsection, we focus on the large-scale networks: 12×12 , and 16×16 , and show the comparison results in Figure 12(a). In particular, we present our all valid solutions with the P_{min} that is higher than the P_{min} of all other methods.

Generally, Actin-STAR has the highest total signal power and the least P_{min} among the four methods. LightR and Zygo-STAR have a similar performance in reliability, but LightR decreases signal power versus Zygo-STAR. When our method prepares a similar number of backups as LightR and results in almost the same power consumption, the worst-case reliability of our method surpasses that of LightR. With more backups, our method achieves the highest P_{min} among the state-of-the-art methods and requires slightly

[§]This setting can effectively avoid accepting worse solutions and excessive program runtime.

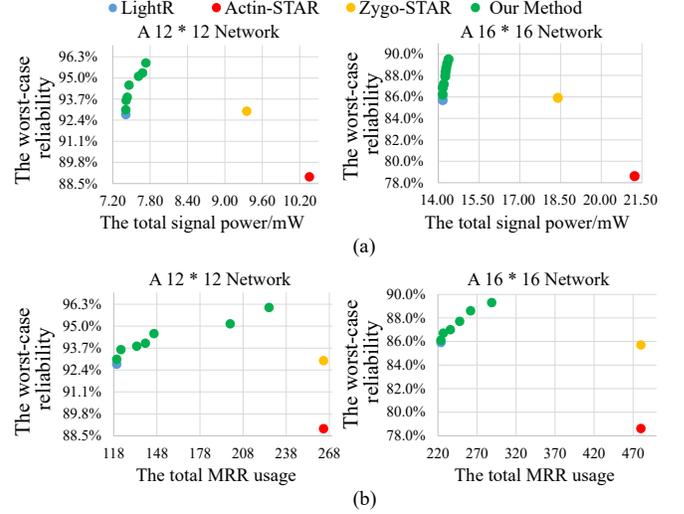


Figure 12: (a) The total signal power and (b) the total MRR usage corresponding the worst-case reliability in large-scale networks.

more signal power (< 1 mW) than LightR. Since we minimize the insertion loss during backup allocation, our method can achieve a good balance between reliability and energy efficiency.

4.3 Discussion: Scalability

We count the total MRR usage in each fault-tolerant topology and evaluate the scalability of the methods for the large-scale networks. Figure 12(b) shows the comparison results. Similar to the performance in signal power, Actin-STAR and Zygo-STAR have the highest MRR usage among the four methods, while LightR minimizes the number of MRRs. When our method uses the same number of MRRs as LightR, our P_{min} is still higher than the result of LightR. Even with the largest P_{min} , our method can decrease the MRR usage by 20% – 40% compared to Actin-STAR and Zygo-STAR. Taking advantage of PSEs and the scalable topological structure, our method can generate efficient fault-tolerant topologies at any scale.

5 CONCLUSION

In this work, we propose a design method to synthesize fault-tolerant WRONoC topologies. Specifically, we propose a fault model, which can accurately model the faults caused by the shift of an MRR's spectrum and reflect the reliability of a network. Instead of preparing a constant number of backup paths, our method considers the different reliability requirements, customizes the number of backup paths for the communications, and automatically allocates the minimum backup resources. The comparison to the state-of-the-art fault-tolerant methods shows that our method can effectively improve reliability. The communications in our synthesized networks have the least probability of being lost compared to the current methods. More importantly, since our method constructs backups, minimizing the increase in resource usage and insertion loss, our method outperforms the state-of-the-art methods, Actin-STAR and Zygo-STAR, significantly reducing total signal power and MRR usage.

ACKNOWLEDGMENTS

This work is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project Number 496766278.

REFERENCES

- [1] Sudip Shekhar, Wim Bogaerts, Lukas Chrostowski, John E. Bowers, Michael Hochberg, Richard Soref, and Bhavin J. Shastri. Roadmapping the next generation of silicon photonics. *Nature Communications*, 15(1):751, Jan 2024.
- [2] Sebastian Werner, Javier Navaridas, and Mikel Luján. A Survey on Optical Network-on-Chip Architectures. *ACM Computing Surveys*, 50(6), December 2017.
- [3] Tsun-Ming Tseng, Alexandre Truppel, Mengchu Li, Mahdi Nikdast, and Ulf Schlichtmann. Wavelength-Routed Optical NoCs: Design and EDA — State of the Art and Future Directions: Invited Paper. In *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 1–6, 2019.
- [4] Scott Van Winkle, Avinash Karanth Kodi, Razvan Bunescu, and Ahmed Louri. Extending the Power-Efficiency and Performance of Photonic Interconnects for Heterogeneous Multicores with Machine Learning. In *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 480–491, 2018.
- [5] Near Margalit, Chao Xiang, Steven M. Bowers, Alexis Bjorlin, Robert Blum, and John E. Bowers. Perspective on the future of silicon photonics and electronics. *Applied Physics Letters*, 118(22):220501, 06 2021.
- [6] W. Bogaerts, P. De Heyn, T. Van Vaerenbergh, K. De Vos, S. Kumar Selvaraja, T. Claes, P. Dumon, P. Bienstman, D. Van Thourhout, and R. Baets. Silicon microring resonators. *Laser & Photonics Reviews*, 6(1):47–73, 2012.
- [7] Aditya Narayan, Yvain Thonnart, Pascal Vivet, César Fuguet Tortolero, and Ayse K. Coskun. WAVES: Wavelength Selection for Power-Efficient 2.5D-Integrated Photonic NoCs. In *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 516–521, 2019.
- [8] Zhidan Zheng, Mengchu Li, Tsun-Ming Tseng, and Ulf Schlichtmann. ToPro: A Topology Projector and Waveguide Router for Wavelength-Routed Optical Networks-on-Chip. In *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, pages 1–9, 2021.
- [9] Marta Ortín-Obón, Mahdi Tala, Luca Ramini, Víctor Viñals-Yuferra, and Davide Bertozzi. Contrasting Laser Power Requirements of Wavelength-Routed Optical NoC Topologies Subject to the Floorplanning, Placement, and Routing Constraints of a 3-D-Stacked System. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, PP:1–14, 03 2017.
- [10] Mengchu Li, Tsun-Ming Tseng, Davide Bertozzi, Mahdi Tala, and Ulf Schlichtmann. CustomTopo: A Topology Generation Method for Application-Specific Wavelength-Routed Optical NoCs. In *2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, page 1–8, New York, NY, USA, 2018.
- [11] Zhidan Zheng, Liaoyuan Cheng, Kanta Arisawa, Qingyu Li, Alexandre Truppel, Shigeru Yamashita, Tsun-Ming Tseng, and Ulf Schlichtmann. Multi-Resonance Mesh-Based Wavelength-Routed Optical Networks-on-Chip. In *Proceedings of the 61st Annual Design Automation Conference (DAC)*, pages 1–6, 2024.
- [12] Xianfang Tan, Mei Yang, Lei Zhang, Yingtao Jiang, and Jianyi Yang. On a Scalable, Non-Blocking Optical Router for Photonic Networks-on-Chip Designs. In *2011 Symposium on Photonics and Optoelectronics (SOPO)*, pages 1–4, 2011.
- [13] Zhidan Zheng, Mengchu Li, Tsun-Ming Tseng, and Ulf Schlichtmann. Light: A Scalable and Efficient Wavelength-Routed Optical Networks-On-Chip Topology. In *2021 Asia and South Pacific Design Automation Conference (ASP-DAC)*, page 568–573, 2021.
- [14] Fengxian Jiao, Sheqin Dong, Bei Yu, Bing Li, and Ulf Schlichtmann. Thermal-Aware Placement and Routing for 3D Optical Networks-on-Chips. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–4, 2018.
- [15] Christopher J. Nitta, Matthew K. Farrens, and Venkatesh Akella. Resilient Microring Resonator Based Photonic Networks. In *Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO-44, page 95–104, New York, NY, USA, 2011. Association for Computing Machinery.
- [16] Michael Conrad Meyer, Akram Ben Ahmed, Yuichi Okuyama, and Abderazek Ben Abdallah. FTTDOR: Microring Fault-resilient Optical Router for Reliable Optical Network-on-Chip Systems. In *2015 IEEE 9th International Symposium on Embedded Multicore/Many-core Systems-on-Chip*, pages 227–234, 2015.
- [17] Michael Conrad Meyer, Akram Ben Ahmed, Yuki Tanaka, and Abderazek Ben Abdallah. On the Design of a Fault-Tolerant Photonic Network-on-Chip. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*, pages 821–826, 2015.
- [18] Yu-Kai Chuang, Yong Zhong, Yi-Hao Cheng, Bo-Yi Yu, Shao-Yun Fang, Bing Li, and Ulf Schlichtmann. RobustONoC: Fault-Tolerant Optical Networks-on-Chip with Path Backup and Signal Reflection. In *2021 22nd International Symposium on Quality Electronic Design (ISQED)*, pages 67–72, 2021.
- [19] Zhidan Zheng, Mengchu Li, Tsun-Ming Tseng, and Ulf Schlichtmann. LightR: A Fault-Tolerant Wavelength-Routed Optical Networks-on-Chip Topology. *Applied Sciences*, 13(15), 2023.
- [20] Yan-Lin Chen, Wei-Che Tseng, Wei-Yao Kao, and Yao-Wen Chang. A General Wavelength-Routed Optical Networks-on-Chip Model with Applications to Provably Good Customized and Fault-Tolerant Topology Designs. In *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, pages 1–7, 2023.
- [21] Zhidan Zheng, Mengchu Li, Tsun-Ming Tseng, and Ulf Schlichtmann. XRing: A Crosstalk-Aware Synthesis Method for Wavelength-Routed Optical Ring Routers. In *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1–6, 2023.
- [22] Lukas Chrostowski and Michael Hochberg. *Silicon Photonics Design: From Devices to Systems*. Cambridge University Press, 2015.
- [23] Sai Vineel Reddy Chittamuru, Ishan G. Thakkar, and Sudeep Pasricha. Libra: Thermal and process variation aware reliability management in photonic networks-on-chip. *IEEE Transactions on Multi-Scale Computing Systems*, 4(4):758–772, 2018.
- [24] Trent McConaghy, Kristopher Breen, Jeffrey Dyck, and Amit Gupta. *Variation-Aware Design of Custom Integrated Circuits: A Hands-on Field Guide*. 01 2013.
- [25] Asif Mirza, Febin Sunny, Sudeep Pasricha, and Mahdi Nikdast. Silicon Photonic Microring Resonators: Design Optimization Under Fabrication Non-Uniformity. In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 484–489, 2020.
- [26] Frank Thomson Leighton. A graph coloring algorithm for large scheduling problems. *J. Res. Natl. Bur. Stand.* (1977), 84(6):489–506, November 1979.
- [27] Anja Von Beuningen, Luca Ramini, Davide Bertozzi, and Ulf Schlichtmann. PROTON+: A Placement and Routing Tool for 3D Optical Networks-on-Chip with a Single Optical Layer. *ACM Journal on Emerging Technologies in Computing Systems*, 12(4), 12 2015.