

2

Fault-tolerant Photonic Network-on-Chip

Michael Meyer, Abderazek Ben Abdallah

*Department of Computer Science and Engineering, University of Aizu,
Aizu-Wakamatsu, Fukushima, 965-8580 Japan
e-mail: {d8161104, benab}@u-aizu.ac.jp*

Abstract

Photonic Networks-on-Chip (PNoCs) promise significant advantages over their electronic counterparts. In particular, they offer a potentially disruptive technology solution with fundamentally low power dissipation that remains independent of capacity while providing ultra-high throughput and minimal access latency. However, the major optical device in PNoC systems, microring resonators (MRs), are very sensitive to temperature fluctuation and manufacturing errors. A single MR failure may cause messages to be misdelivered or lost, which results in bandwidth loss or even complete failure of the whole system. This chapter describes a fault-tolerant PNoC architecture. The system is based on a fault-tolerant path-configuration and routing algorithm, a microring fault-resilient photonic router, and uses minimal redundancy to assure accuracy of the packet transmission even after faulty MRs are detected.

Keywords: Fault-tolerant, Micro-ring Resilient, Photonic Network-on-Chip.

2.1 Introduction

Photonic Network-on-Chip (PNoC) is becoming an attractive solution enabling ultra-high communication bandwidth in the terabits per second range, low power, and low communication latency [7, 10, 11, 9, 12]. When combined with Wavelength Division Multiplexing (WDM), multiple parallel op-

tical streams of data concurrently transfer through a single waveguide, while MRs, which can be switched as high as 40 GHz, are used to realize wavelength-selective modulators, and switches [44]. While a single-layer configuration can provide low-loss waveguides and high-performance photonic devices, it suffers from limited integration density due to waveguide crossing and limited real estate. A way to go beyond this limitation is to monolithically stack multiple photonic layers above Si as multilayered electrical interconnections realized in modern electronic circuits [8, 61]. Figure 2.2 shows a high level view of a three-dimensional PNoC (PHENIC) implemented with one electrical control layer and several photonic communication layers [39].

The main components of an PNoC include a laser source, which generates phase-coherent and equally spaced wavelengths, waveguides, which is used as a transmission medium, and modulators and photodetectors, which convert electrical digital data to and from photonic signals [32]. Figure 2.1 shows a typical on-chip optical link that uses an external laser as a light source. It is expected that the laser source could produce up to 64 wavelengths per waveguide for a Dense Wavelength Division Multiplexing (DWDM) network.

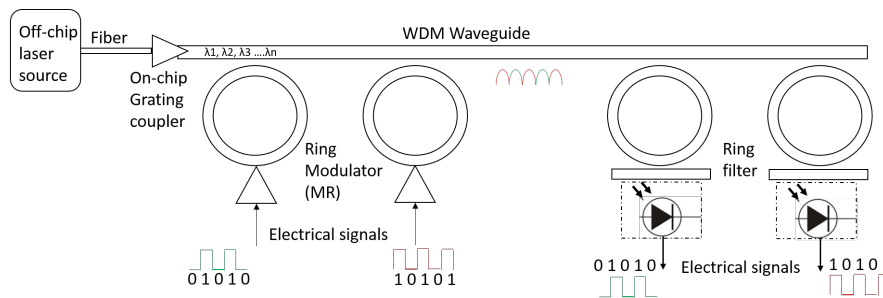


Figure 2.1: Photonic link architecture.

Fault tolerance is crucial when considering mission critical applications where the system must correctly function even when something goes wrong. One such an application is that of space travel, where repair or replacement is not a possible option, and billions of dollars would be wasted.

2.1.1 Design Challenges

The photonic domain is immune to transient faults caused by radiation [29], but is still susceptible to process variation (PV) and thermal variations (TV)




Figure 2.2: 3D-Stacked photonic network-on-chip architecture.

as well as aging. The aging typically occurs faster in active components as well as elements that have high TV [26]. In the optical domain, the faults can occur in MRs, waveguides, routers, etc. Active components, such as MRs, have higher failure rates than passive components, e.g. waveguides [26]. A single MR failure can cause messages to be misdelivered or lost, which results in bandwidth loss or even complete failure of the whole system. Together, fabrication-induced PV and TV effects present enormous performance and reliability concerns. TV causes a microring to respond to a different wavelength than intended. This can take the form of a passband shift in the MRs. When an MR heats up, it expands, changing its radius, and therefore shifting the wavelengths which it uses to the right [15]. As reported in [44], a change of as little as 1°C can shift the resonance wavelength of a microring by as much as 0.1nm. This is not permanent and will return when the temperature returns to normal. Therefore, systems' temperature must be kept at a reasonable value in order for the MRs to resonate correctly. This is challenging, especially in large complex computing system, which uses thousands of these components. Trimming technique [4] is generally used to dynamically modify the resonance frequency of a microring to overcome both thermal drift and fabrication inaccuracy. This technique can be accomplished by dynamically increasing the current in the $n+$ region or by heating the ring [22, 4, 48].

PV is the variations of critical physical dimensions, e.g. thickness of wafer, width of waveguides also affect the resonant wavelengths of MRs. This means that not all fabricated MRs can be used due to PV. As a result, network nodes that do not have all working MRs would lose some or all of wavelengths/bandwidth in communication [56]. To solve this problem, Xu et al. [58] proposed a method of flexible wavelength assignment. Because

the networks are already built with excess detectors or Modulators for each message, the node with the excess components can compensate and rematch to the components which have been affected by PV.

Over time, all silicon based ICs wear down. We refer to this phenomenon as *aging*. Some of the aging effects only apply to the active components, because of their electrical subcomponents [54], such as the MRs, while other aging affects all parts, even the waveguides.

Recent PNoCs researches (i.e. network topology, router micro-architecture design, and performance and power optimization and analysis) have resulted in several architectures capable of transmitting at a high data bandwidth and low energy dissipation [7, 10, 11, 9, 12]. In [8], we proposed an energy-efficient and high-throughput hybrid silicon-photonic network-on-chip based on a smart contention-aware path-configuration algorithm and an energy-efficient non-blocking optical switch to further exploit the low energy properties of the PNoC systems. However, little attention has been given to the aspect of fault-tolerance and reliability along the photonic interconnects.

This chapter presents a fault-tolerant PNoC architecture. The system is based on a fault-tolerant path-configuration and routing algorithm, a microring fault-resilient photonic router, and uses minimal redundancy to assure accuracy of the packet transmission even after faulty MRs are detected.

2.1.2 Fault Models

It is worth noting that the light is not sensitive to radiation or electromagnetic fields, the signals which control the optical network can be sensitive to it. The following is a list of actual possible causes that can contribute to the failure of an optical device.

2.1.2.1 PNoC Signal Strength

Typical NoCs are defined by their power consumption, delay and throughput. PNoCs also have to consider the Signal-to-Noise Ratio at the receiving end. Because they do not buffer and retransmit, the signal gets weaker based off of how many hops it jumps. This does not significantly affect the power the network consumes, but it can lead to a higher sensitivity to noise.

2.1.2.2 Electrostatic Discharge

While the waveguides are not electrically conductive, the switches and photo-detectors are. This means that they are sensitive to high currents. One thing which can ruin an IC is electrostatic discharge(ESD). This is when a current

enters in through the I/O pins of the control circuit, or it can be caused by an extremely strong magnetic field. This all results in the aforementioned extreme current, and this current causes severe damage to the silicon in the components. Possible points of damage are the dielectric, the PN junctions, and any wiring connecting to the controllers. Because of the scaling, the causing phenomena have become harder to control [24]. This can be prevented by proper packaging to the IC providing ESD protection at the pins.

2.1.2.3 Noise

This is one of the unique things that we categorize as a cause for a fault. The reason is because the noise can be caused simply by poorly matched wavelengths. It can also be caused by creating a path that is too long, or a path that crosses too many intersections. These paths tend to be caused by rerouting or non-minimalistic routing, but other factors can contribute and cause more noise. The most common factors are listed in the following subsections.

2.1.2.4 Aging

Over time, all silicon based ICs wear down. Some of the aging effects only apply to the active components, because of their electrical subcomponents, while other aging affects the optical properties of the components.

Electromigration- This mainly affects the wires which control the ring resonators. It does not affect the waveguides in any way. It originally causes a delay in the wire, and can eventually lead to an open, or to a short to a nearby wire. It achieves this by thinning out the thinnest portion of the wire due to higher current density at the bottleneck [30].

Laser Degradation- After the lasers have been on for several hundred hours, they start to show signs of degradation. This shows in the form of either missing wavelengths, which can cause a channel fault, or general weakening of the original laser signal. In each of these cases, it does not become a true problem until the signal falls to a level where the worst case scenario's Signal-to-Noise ratio is too weak to receive an understandable signal [37].

Photodetector Degradation- Various studies have been done for different types of photodetectors showing that they degrade overtime, particularly from being exposed to thermal conditions or UV light. It is reasonable to assume that no matter what material photo detectors are made out of, they all seem to be vulnerable to degradation due to thermal variation, which is present in all networks [26, 54].

A lot of work has been done to combat the effects of aging. Some examples are Agarwal [2], Keane [30], and Kim [31]. These are mainly focused on

Table 2.1: Overview of Fault Causes and Effects
 (ABPV is Accelerated by Process Variation, OOHC is Optical Or Hybrid Components)

	Physical Cause	ABPV	Fault Class	Bursty	OOHC
Noise	Internal reflectance in waveguide	Y	Intermittent Logic Fault	N	Optical
	Waveguide Crossing	N	Intermittent Logic Fault	Y	Optical
	Intrachannel Noise	N	Intermittent Logic Fault	Y	Both
ESD	Build up of Static energy which is quickly released	Y	Permanent Logic Fault	N	Hybrid
	Electron Migration	Y	Intermittent->Permanent Delay or Logic Fault	N	Hybrid
	Laser Degradation	Y	Intermittent->Permanent Logic Fault	N	Hybrid
Aging	Photodetector Degradation	Y	Intermittent->Permanent Delay or Logic Fault	N	Hybrid
	Waveguide Degradation	Y	Intermittent->Permanent Logic Fault	N	Optical
	MR Degradation	Y	Intermittent->Permanent Logic Fault	N	Hybrid
	Variation in wear-out effects due to temp. differences	Y	Intermittent and Permanent delay, stuck-at and misrouting	Y/N	Hybrid
Temp. Variation	Performance variation due to temp differences	Y	Intermittent delay and misrouting	Y	Hybrid
	Variation in wear-out effects due to temp. differences	Y	Intermittent and Permanent delay, stuck-at and misrouting	Y/N	Optical

the electrical side, but the fact that these do exist show the hope for a future where optical aging can be researched and prevented. Many parameters such as the wavelengths and laser strength can possibly be modified throughout the life of a chip to counteract the aging effects in a similar manner to what Mintarno does for Electrical networks [40].

2.1.2.5 Process Variability

This can affect both the active and inactive components of the optical network. The variability accounts for material impurities, doping concentrations, and size and geometries of structures [47]. One single dimple in a particular point in the coupling region of a ring resonator can greatly affect the coupling properties and thus cause problems for the switch, or maybe just the channel. A poor geometry can also cause a certain component to be more sensitive to aging or ESD. Obviously if a variation gets bad enough, an entire link can be rendered useless. This would be considered an early permanent fault, and should be detected before a device is released. The impurities in a waveguide can cause such a block, or cause there to be a change in the reflectivity of the material, and that causes a higher amount of insertion loss, resulting in a lower signal-to-noise ratio. Other similar chains-of-events can occur from bad doping of the photodetectors. Minimizing this process variability can greatly increase the reliability of the system, even without implementing fancier and area or energy heavy redundancies. The unfortunate truth is that with recent advances in scaling, the variability continues to increase [33, 50].

2.1.2.6 Temperature Variation

For electrical components, temperature variation can cause changes in properties such as resistivity and cause more power consumption or delay, but in the optical domain, it is quite different. Ring resonators are tuned by heating up the ring, causing them to expand, which changes their passband wavelength. If the chip heats up to a point beyond the tuning, then certain channels just disappear as a whole. The increase in temperature also causes the photodetectors to degrade as mentioned in the previous section. These temperature variations also tend to speed up other forms of aging as well.

Table 2.1 summarizes the physical causes and their effects. Many of these will need to be researched further, and only time will tell exactly how reliable optical is with some other phenomena, but for now, this is a comprehensive list of all physical sources of failures within an optical network. We separated the pure optical from the hybrid components so that it can show exactly how

resilient the photons and waveguides really are, when compared with wires, but no Optical Network-on-Chip is completely free of wires.

2.2 Fault-tolerant Photonic Network-on-Chip Architecture

The Fault-tolerant Photonic Network-on-Chip (FT-PHENIC) system, shown in Fig. 2.3, is a mesh-based topology and uses minimal redundancy to assure accuracy of the packet transmission even after faulty MRs are detected. The system uses Stall-Go mechanism for flow-control, and a Matrix-arbiter as a scheduling technique [39, 3, 14, 13]. FT-PHENIC is also based on a microring fault-resilient photonic router (FTTDOR) [39] and an adaptive path-configuration and routing algorithm. As illustrated in Fig. 2.3, the proposed system consists of a Photonic Communication Network (PCN), used for data communication, and an Electronic Control Network (ECN), used for path configuration and routing. Each PE (Processing Element) is connected to a local electrical router and also connected to the corresponding gateway (modulator/detector) in the PCN [8]. Messages generated by the PEs are separated into control signals and payload signals. Control signals are routed in the ECN and used for path configuration and routing. The payloads are converted to optical data and transmitted on the PCN.

2.2.1 Microring Fault-resilient Photonic Router

The block diagram of the Microring Fault-resilient Photonic Router (FTTDOR) is shown in Fig. 2.4. It consists of a non-blocking fault tolerant photonic switch (Fig. 2.4 (a)) and a light weight control router (Fig. 2.4 (b)). Redundant MRs are carefully placed at special locations on the switch to assure fault tolerance even if one of the MRs on the backup path has a fault. The backup route for the NEWS (North-East-West-South) directions is to actually use the waveguide connected to the core ports as a master backup; therefore, the redundant MRs are all chosen at the locations which connect the NSEW ports to the core.

For a majority of faults, the design of the switch allows for an alternate, slightly less power efficient route. In fact, the backup route is less power-efficient because the packets travel across more waveguide distance, go through more active MRs, and cross more waveguides. However, the switch still maintains all of its functionality. Because backup routes are only intended for use in the switches in which faults have occurred, the extra loss will have minimal effect on the message' signal strength across the whole network.

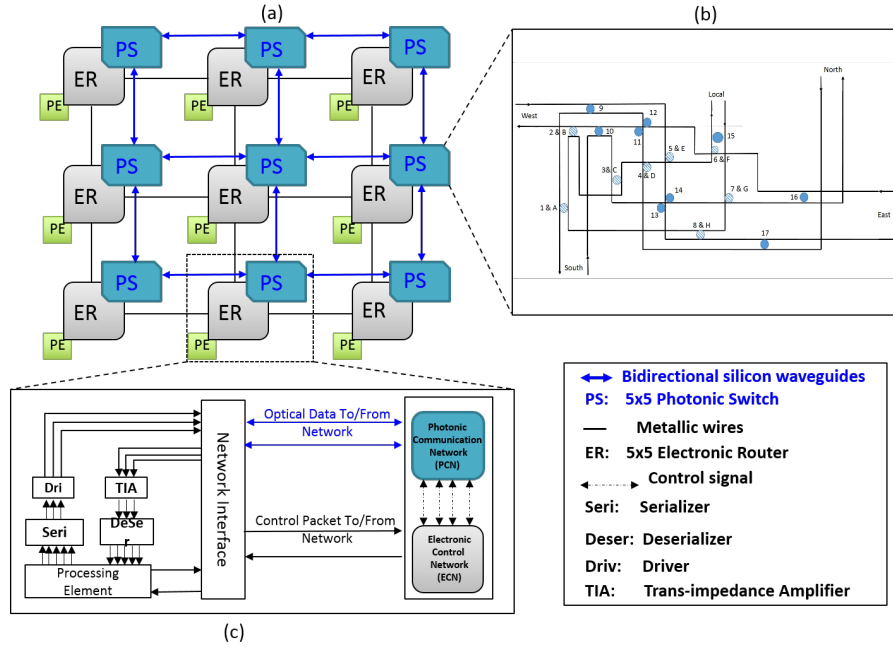
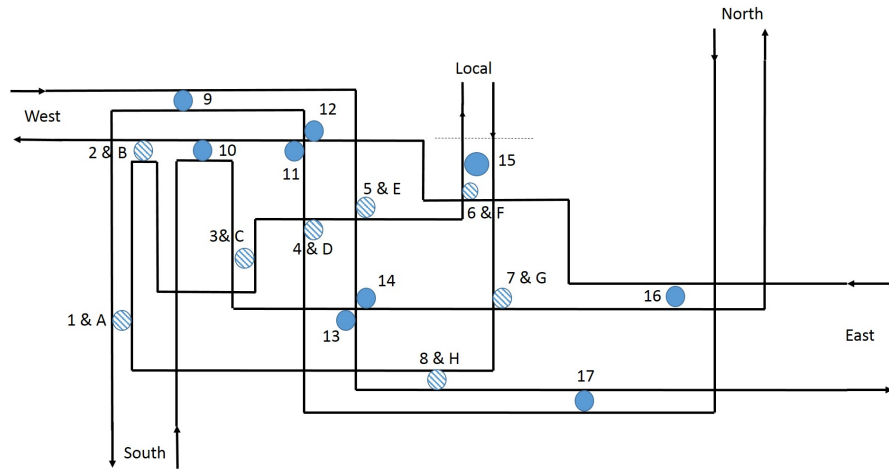


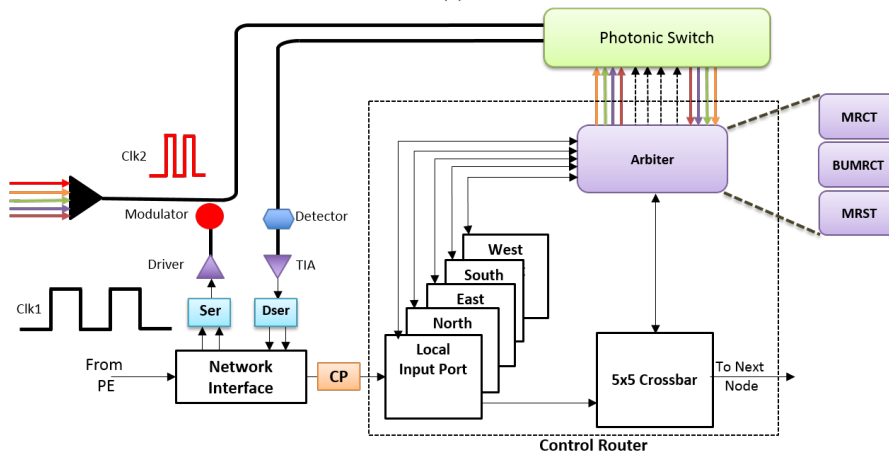
Figure 2.3: FT-PHENIC system architecture. (a) 3x3 mesh-based system, (b) 5x5 non-blocking photonic switch, (c) Unified tile including PE, NI and control modules.

The FTDDOR was designed to require no MRs from East-West and North-South traffic. Since this kind of traffic accounts for a majority of the traffic of the PCN [39], such design will save on power and continue to function in the case of any MR fails. Assuming that a single location of redundant MRs does not fail all together, the switch is able to maintain all functionality at slowed speeds.

Figure 2.5 shows a reconfiguration example of how MR 9 can be backed up by MRs 5, 15 and 1. Additionally, the MRs which connect parallel waveguides are replaced with racetracks [41]. This allows for a wider pass-band of light frequencies, makes them less sensitive to physical faults, such as reduced sensitivity to thermally-caused passband shifting. Racetracks also have a larger Mean Time Between Failures (MTBF) [41].



(a)



(b)

Figure 2.4: Microring fault-resilient photonic router (FTTDOR): (a) Non-blocking fault tolerant photonic switch, (b) Light-weight control router.

The original form of FTTDOR switch is a five-port non-blocking switch, meaning that it allows for routing from any available port to any other available port. Once a fault is detected, the switch recovers, but there is a chance that it may turn into a blocking switch; however, it should be able to maintain all functionality as long as none of the redundant MRs fail. Because the redundant MRs lie dormant, they do not require much power other than the boost in signal strength required to compensate for the signal loss, caused by passing by an inactive MR, which is minimal. As all rerouting in the switch occurs on the core waveguide, traffic certainly increases on this one waveguide as too many faults occur, which is why it should be treated as a node failure after a threshold of failed MRs is reached.

In addition to tolerating faults, FTTDOR is able to handle the *ACK* signals and the resulting regeneration process of the *Tear-down* signal at each hop. To accomplish this goal, a hybrid switching policy is used: *Spacial-switching* for the data signals by manipulating the state of the broadband switching elements and a *Wavelength-selective switching* for the *Tear-down* signals by using detectors and modulators. Moreover, since the *Tear-down* signals should be checked and regenerated at each hop, it is crucial that their manipulation be automatic and not interfere with data signals, nor cause a blockage inside the switch. When the *Tear-down* is generated at the source NI (Network Interface), it is first sent to the control router. Then, the *Photonic Switch Controller* releases the corresponding MRs and generate another *Tear-down* which is sent to the output-port modulator in the PCN where it continues its path in a hop-by-hop basis until it reaches its destination. At the destination node, the *Tear-down* is detected in the input-port and sent to the *Photonic Switch Controller* in the corresponding electronic router. In this fashion, we can omit the overhead of an additional gateway which becomes significant when we increase the number of cores. Table 2.2 shows the MRs

Table 2.2: Microring configuration for normal data transmission.

output/Input	Core	North	East	South	West
Core	-	4	6	3	5
North	7	-	16	None	14
East	8	17	-	13	None
South	1	None	12	-	9
West	2	11	None	10	-

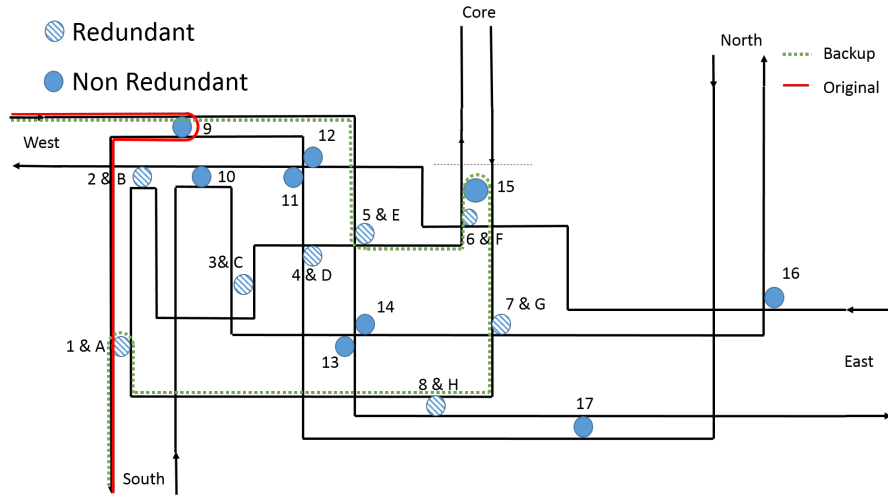


Figure 2.5: Example of how a non redundant MR’s functionality can be mimicked by redundant ones.

Table 2.3: Microring backup configuration for data transmission.

output/Input	Core	North	East	South	West
Core	15	D	F	C	E
North	G	-	6,15,7	None	5,15,7
East	H	4,15,8	-	3,15,8	None
South	A	None	6,15,1	-	5,15,1
West	B	4,15,2	None	3,15,2	-

configuration for data transmission, where 16 MRs are used in a non-blocking fashion. Table 2.3 shows the backup paths for each transmission.

We use the first six wavelengths in the optical spectrum starting from 1550 nm, with a wavelength spacing equal to 0.8 nm to maintain a low cross-talk as reported in [46]. For the acknowledgment signals, we use the first five wavelengths in the optical spectrum starting from 1550 nm: four wavelengths for the *Tear-down* signal where each one is dedicated for each port except the local one. In addition, a single wavelength is used for the *ACK*. The remaining available wavelengths are used for data transmission. The five wavelengths used to control the *ACK* and *Tear-down* signals are notably con-

stant regardless of the network size, in contrast with the fully optical where the number of wavelength used for control and arbitration grows with the network size. Thus, cutting these wavelengths from the available spectrum to be used for control, would not degrade the system bandwidth. These five wavelengths will be negligible especially when DWDM is used providing up to 128 wavelengths per waveguide [16]. The wavelength assignment for each port is shown in Table 2.4.

Should the *Tear-down* signals enter the switch, they need to be redirected to the corresponding electronic router. Since these signals are coming from different ports, and are modulated with different wavelengths, detectors capable of switching all of the four wavelengths are placed in front of the input-ports to intercept the signals. The converted optical signal will be redirected to the electronic router to be processed. According to the included information, the corresponding MRs will be released. For the *ACK*, when the PSCP reaches the destination, 1-bit optical signal is modulated starting from the output port (i.e., opposite direction) and travels back to the source. With

Table 2.4: Wavelength assignment for acknowledgment signal (Mod: Modulator, and Det: Photo-detector).

	Core	North	East	South	West
Input	Mod_{λ_0}	Det_{λ_3}	Det_{λ_2}	Det_{λ_1}	Det_{λ_4}
Output	Det_{λ_0}	Mod_{λ_1}	Mod_{λ_4}	Mod_{λ_3}	Mod_{λ_2}

this smart hybrid switching mechanism, we take advantage of the low-power consumption of the optical link by using optical pulses modulated with the adequate wavelength instead of propagating the acknowledgment signals in the ECN. Second, we take advantage of the WDM properties by separating the acknowledgment packets and the data signals and let them coexist in the same medium without interfering with each other. This contrasts with the electronic domain where these acknowledgment packets travel for a several hops consequently blocking (preventing) the waiting cores from sending their PSCP packets. Finally, we are able to tolerate faults due to the arrangement of the MRs, and allowing for redundancy at critical locations.

As a primary comparison, we performed a study on the routers, and the loss that they would each have on average, and in their worst case. The results can be seen in Table 2.5. As expected, the Crux [59] performs the best, as its only design goal was to minimize loss and noise, sacrificing a lot of function-

Table 2.5: Various switches and their estimated losses. AL: Average Loss, WL: Worst Loss

Router	Cros.	MRs	Termi.	AL(dB)	WL(dB)	WL(faulty)(dB)
Crossbar	25	25	10	1.12	1.60	∞
Crux	9	12	2	.657	1.11	∞
PHENIC	27	18	0	1.315	1.615	∞
FT-PHENIC	19	16+9	0	.965	1.115	2.215

Table 2.6: Insertion loss parameters for 22nm process.

Parameter	Value
Through Ring Loss	.5dB [59]
Pass By Ring Loss	.005dB [19]
Bending Loss	.005dB [19]
Crossing Loss	.12dB [59]
Terminator	.01 dB [19]

ality. Values for the calculation were taken from various authors, and can be seen in Table 2.6.

2.2.2 Light-weight Electronic Control Router

Figure 2.4 (b) illustrates the control router architecture, which is based upon OASIS-NoC router [6, 14, 13, 5]. As shown in the above figure, the arbiter receives the detected *Tear-down* from the above switch (colored arrows). According to the information encoded in this signal, the corresponding MRs are released and a new *Tear-down* is generated for the next hop until it reaches its final destination and all MRs involved in this communication are released. The figure shows also the connection between the network interface (NI) and the local port, where a configuration packet (CP) is sent from the NI to the local port. The CP could be a setup packet or a path blocked packet. The NI is connected also to the data switch (i.e., PCN). When the source node receives the ACK, the payload is processed by a serializer bank (if needed), a high speed driver, and a modulator to convert the electrical signal to an optical one. At the source node, the optical data leaves the data switch and go through a detection step, a high speed Trans-Impedance-Amplification step, and a deserialization step. At the end the NI's receiver, receives the payload data with its original clock speed.

2.2.3 Fault-tolerant Path-configuration and Routing

The key feature of the Fault-tolerant Photonic Path-configuration algorithm (FTPP) is that it can handle faulty MRs within the photonic switches. When a fault occurs, the algorithm checks for the secondary MRs on the list, and checks their status. The backup MR table can be very simple in the cases of a redundant MR failing, where it is simply replaced by its redundancy, or it can be slightly more complicated, as seen in Figure 2.5.

The FTTP algorithm must meet certain requirements to work with the FT-PHENIC system. It should be also able to remove the dependency between the ECN and PCN which causes a significant latency overhead in conventional hybrid-PNoC systems. In addition, the latency caused by the path blocking, which requires several cycles for the path dropping and the new path setup packet generation is considerably decreased. Another key feature of the configuration algorithm is the efficiency of the ECN resources' utilization. By moving the acknowledgment signals to the upper layer, we can reduce the buffer depth to only 2 slots, since half of the network traffic is eliminated. This reduction is a key factor to design a light-weight router, highly optimized for latency and energy.

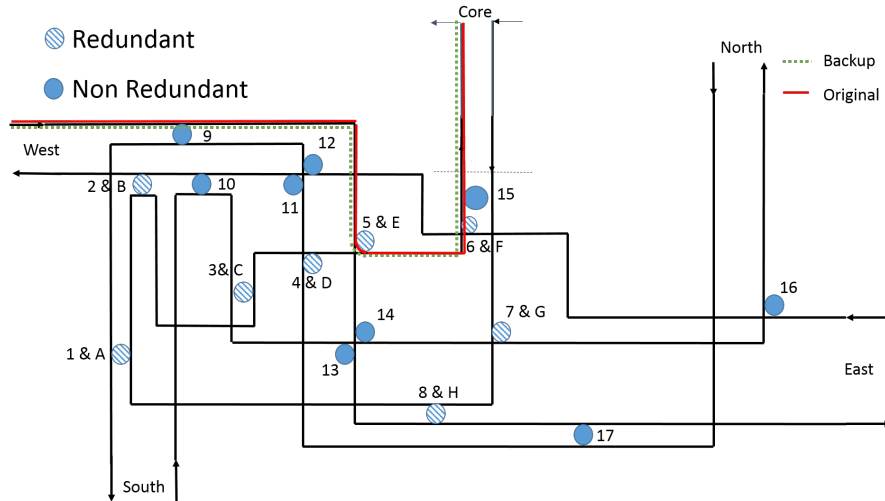


Figure 2.6: Example of how a redundant MR's functionality can be mimicked by its redundancy.

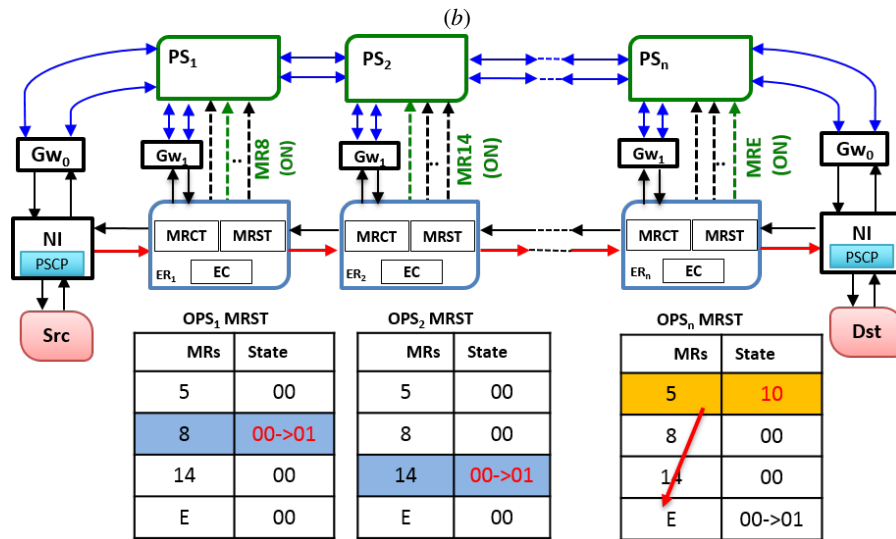
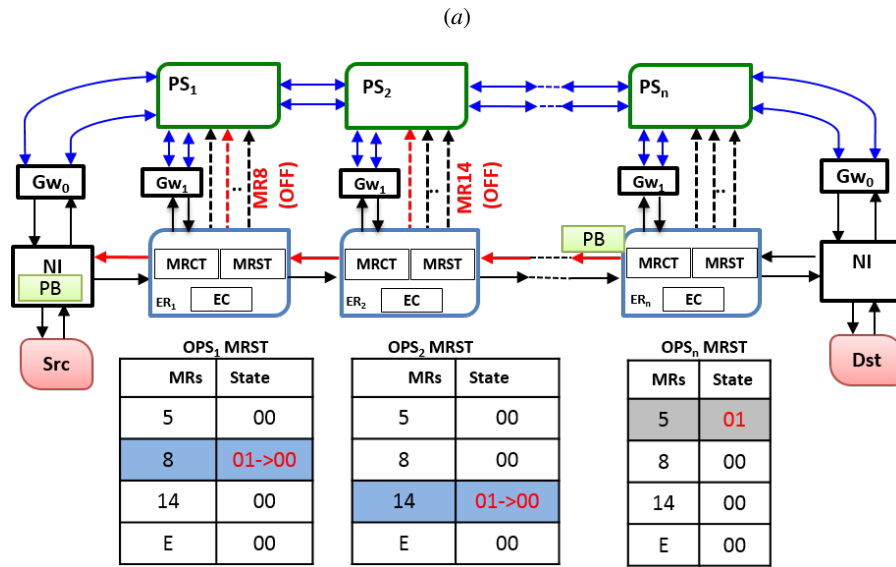
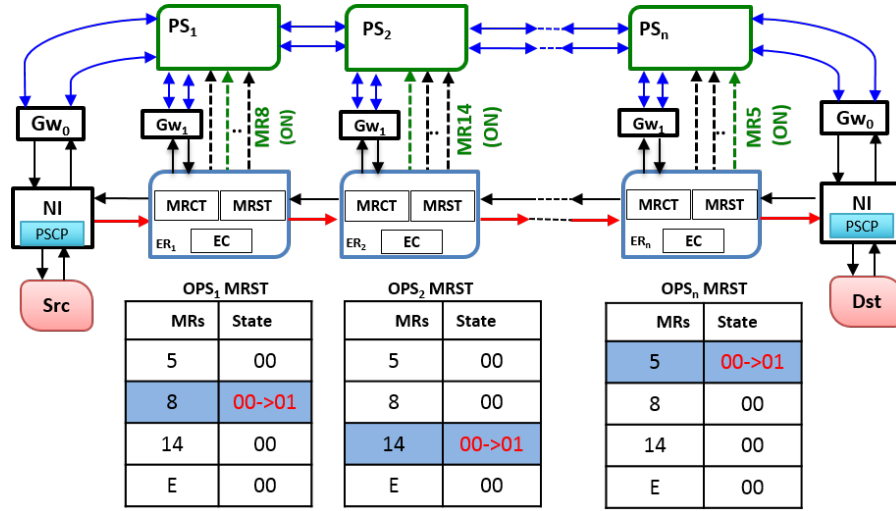


Figure 2.7: Microring fault-resilient path-configuration: a) Path-setup, (b) Path-blocked, (c) Faulty MR with Recovery. GW_0 : Gateway for data, GW_1 : Gateway for acknowledgment signals, PS: photonic Switch, MRCT: Micro

Figure 2.7 (a) shows an example of a successful path-setup process where all the necessary resources between a given source-destination pair are reserved. The corresponding pseudo code is given in Algorithm 1. Before optical data transmission, the source node issues a *Path-setup-Control-Packet* (PSCP) which is routed in the ECN and includes information about the destination and source addresses. In addition to the source and destination addresses, other information is included. For example, one-bit is used for the Packet-type field. This field can be "00" for a PSCP and "01" when this configuration packet is a Path-blocked. Other information to ensure Quality-of-Service and fault-tolerance, such as Message-ID, Fault-status, Error-Detection-Code, can also be included. For each electrical router, the output-port is calculated according to Dimension-Order routing [6]. Every time the PSCP progresses to the next router, the optical waveguides between the previous and current routers are reserved. Depending on the output port of the electrical router, the corresponding photonic router is configured by switching ON/OFF one or more MRs using the MRs configuration table shown in Table 2.2. In the example shown in Fig. 2.7 (a), the packet is entering the local input-port attached to the Network Interface (NI) and requesting the east output-port. According to Table 2.2, MR 8 is required and its availability is checked in the (Micro Ring State Table) MRST. In this table, the MR's state is "00" (free and not faulty). Therefore, the switch controller reserves the MR and changes its states from "00" (free and not faulty) to "01" (not free and not faulty). After this successful reservation (hop based), the PSCP continues its path to the next hop and the same procedure is repeated until all necessary MRs are reserved for the complete path. This process is illustrated in *lines* 1–10 of Algorithm 1. In a case where the requested MRs at a given optical switch along the path are not available, blocking occurs. This can be seen in Fig. 2.7 (b) where MR 5, which is necessary for the ejection to the local output-port from the west input-port, is used by another communication. In this case, the PSCP is converted into a *Path_blocked* packet (PB). The PB, then, travels back to the source node and releases the already reserved resources. The release is done by re-updating the corresponding entries in the MRST to "00" and by sending an electrical "OFF" signal to the corresponding MRs in the PCN. This process is illustrated in *lines* 11–15 of Algorithm 1.

If a fault is encountered along the way, denoted by a state of "10", seen in figure 2.7 (c), then the switch attempts to use its backup route within the switch to maintain the intended port-to-port communication. This allows for recovery without requiring the whole system to change the route of a packet, and can save on costly retransmission and multiple attempts at setting up

the path. Assuming that the backup path is being used for a recovery path, then the algorithm proceeds with sending the standard path blocked packet. When the *PSCP* arrives successfully at the destination node, the NI modulates one-bit acknowledgment (ACK) signal to travel back to the source via the PCN. This can be seen in lines 16 – 20 of Algorithm 1. Upon the arrival of this *ACK* signal, the source node modulates the payload through the data modulators and sends it to the destination node via the PCN. Lines 21 – 25 of Algorithm 1 depicts this data/payload transfer phase. The last process of the proposed path-configuration algorithm is the *Tear – down* step as shown in lines 26 – 31 of Algorithm 1. When the entire payload is transmitted, it is necessary to release the reserved optical resources. This is handled by the source node which sends a *Tear – down* packet to the destination after pre-determined number of cycles depending on the source-destination addresses, transmission bandwidth and message size.

The source's NI sends the electronic *Tear – down* packet (TD) to the first electronic router ER_1 . The Electronic Controller (EC) in this router indexes the MRCT with input-output ports information and determines the MRs that need to be released. As we can see in this figure, the state of MR 8, previously reserved in the path-setup process, is reset to *Free* (state="00") and electrical "OFF" signals are sent to the MR.

After the MRs are deactivated, a new optical Tear-down signal is generated according to the used wavelength. It is sent through the PCN to the next hop where it is converted back to electrical and redirected to the EC in the corresponding electronic router to be processed. After this process, the MRs are released and a new optical Tear-down signal is generated. This process is repeated until the *Tear-down* reaches the destination and all optical resources are released. It is important to mention that the path-setup and path-blocked processes of the proposed algorithm are very similar to the conventional ones [7, 10, 1, 25, 18, 52]. The main difference is that the MRST in our proposal contains only two states: *Free* and *Active*. The MRs are set "ON" as soon as the *PSCP* succeeds to reserve them. In the conventional mechanisms, three states are necessary: *Free*, *Reserved*, and *Active*. When the *PSCP* finds the requested MRs *Free*, it updates their states in the MRST to *Reserved* without turning them "ON". When the complete path-setup process is completed, the *ACK* signal travels back to the source node and sets the corresponding MRs "ON" by updating their states in the MRST to *Active*. With the proposed algorithm, some portions of the reserved path might be set "ON" and then "OFF" due to the unavailability of the resources. However, it enables the fast *ACK* transmission in the PCN.

Algorithm 1: Fault-tolerant path-configuration algorithm.

```

// Path Setup Control Packet for communication  $i$ , PSCP $_i$ 
// Path Blocked Packet for communication  $i$ , PB $_i$ 
Input:  $S_i, D_i$ 
// From ACK detector
Input:  $Det_{ACK_s}$ 
// To ACK modulator
Output:  $Mod_{ACK_s}$ 
// From Teardown detector
Input:  $TeardMod_i$ 
// To Teardown modulator
Output:  $TeardMod_i$ 
// To Microring resonator
Output:  $MRS_{j=0..n}$ 
// Buffer writing and routing computation stages
1 initialization;
2 while ( $Path\text{-}Setup\text{-}Control\text{-}Packet (PSCP) \neq 0$ ) do
3   DestAdd  $\leftarrow$  PSCP $_i$ ;
4   PortIn  $\leftarrow$  PSCP $_i$ ;
5   if ( $resource\ are\ available$ ) then                                /* check MRs state */
6     if MR is not faulty then
7       Grant $_i$   $\leftarrow$  Arbiter;
8     else if Backup MR is Not Faulty then                          /* check backup MRs state */
9       GrantBackup $_i$   $\leftarrow$  Arbiter;
10    else                                                            /* no possible path */
11      Blocked $_i$   $\leftarrow$  Arbiter;
12      FaultyNodeList  $\leftarrow$  Node;
13    end
14  else                                                            /* generate path blocked */
15    Blocked $_i$   $\leftarrow$  Arbiter;
16  end
17 end
// Path blocked
18 initialization;
19 while ( $PB \neq 0$ ) do                                            /* Path blocked arrives */
20   if ( $MRS_i$  state is reserved) then                             /* release reserved MRs */
21     release  $\leftarrow$  MRS $_i$ ;
22 end
// Generate ACK
23 initialization;
24 while ( $NI\ receiver \leftarrow PSCP_i$ ) do                          /* PSCP arrives to Dest */
25   if ( $PSCP$  arrives to NI) then                                  /* generate ACK to Src */
26     ACK $_i$   $\leftarrow$  To modulator ACK ( $\lambda_0$ );
27 end
// Receives ACK
28 initialization;
29 while ( $NI\ receiver \leftarrow ACK_i(\lambda_0)$ ) do                  /* ACK arrives to Src  $\lambda_0$  */
30   if ( $ACK$  arrives to the NIsender) then                       /* modulate the data */
31     Data $_i$   $\leftarrow$  To Data's Modulator;
32 end
// Identify and Generate Teardown $_i$ 
33 initialization;
34 while ( $From\ detector\ signal = Teardown_i$  with  $\lambda_i$ ) do
35   findInport  $\leftarrow$   $\lambda_i$ ;                                       /* find In-port according to the wavelength */
36   free  $\leftarrow$  MRS $_i$ ;                                           /* Free involved MRs */
37   Teardwon $_i$   $\leftarrow$  To modulator  $\lambda_i$ ; /* generate new Tear-down according to  $\lambda_i$  */
38 end

```

In conventional path-configuration algorithms, the ACK and Tear-down packets are transmitted in the ECN and have to go through all the buffering, routing computation, and arbitration stages. With the proposed algorithm, they are carried via the PCN. As a consequence, the ETE latency can be significantly reduced in addition to the dynamic energy saving that can be achieved. Additionally, conventional path-configuration algorithms do not check for faulty MRs. This will allow the system to tolerate more MR failures, and take advantage of the fault tolerant switch.

2.3 Evaluation

We evaluate the FT-PHENIC system using a modified version of PhoenixSim which is developed in the OMNeT++ simulation environment [19]. The simulator incorporates detailed physical models of basic photonic building blocks such as waveguides, modulators, photodetectors, and switches. Electronic energy performance is based on the ORION simulator [27]. We evaluate the bandwidth performance and energy consumption for 16, 64 and 256 cores systems.

We compare the performance of the FT-PHENIC systems with the baseline PHENIC [8], and the system using the algorithm proposed by Xiang et al. [56]. Xiang’s network was chosen over other typical systems [53, 45, 55, 17], because it uses some form of fault tolerance, and most of their results would mimic the baseline PHENIC. For the fault related data, we disabled a

Table 2.7: Configuration parameters.

Network Configuration	Value
Process technology	32nm
Number of tiles	256,64,16
Chip area (equally divided amongst tiles)	400 mm^2
Core frequency	2.5GHz
Electronic Control frequency	1GHz
Power Model	Orion 2.0
Buffer Depth	2
Message size	2 kilobytes
Simulation time	10ms (25 10^8 cycles)

certain number of MRs at random, and recorded the data. To get better results, we would run each system at each fault rate 10 times, and then averaged each

Table 2.8: Photonic communication network energy parameters.

Network Configuration	Value
Datarate (per wavelength)	2.5GB/s
MRs dynamic energy	375fJ/bit
MRs static energy	400 μ W
Modulators dynamic energy	25fJ/bit
Modulators static energy	30 μ W
Photodetector energy	50fJ/bit
MRs static thermal tuning	1 μ W/ring

test's total energy, average bandwidth, and average latency. Currently, the MR is disabled for the whole test, and thus models either a permanent or intermittent fault. Dealing with passband shift or temporary overheating of an MR is outside of the scope of this paper, beyond redundancy as a solution. The fault rates were chosen to span from 0 to 30% due to the fact that at this point, all of the tested networks were in deadlock.

2.3.1 Complexity Evaluation

The complexity evaluation considers the number of used rings and the resulting static thermal tuning. The number of used MRs is given by equation 2.1, where $Mod/Detc_{(ring)}$ is the number of rings required to modulate/detect the payload signal. $Switch_{(ring)}$ is the number of ring required for the photonic switch to route the optical data. Finally, the $ACKs_{(ring)}$ is the number required to handle the acknowledgment signal.

$$Total_{(ring)} = Mod/Detc_{(ring)} + Switch_{(ring)} + ACKs_{(ring)} \quad (2.1)$$

Tables 2.9 and 2.10 show the comparison results for 64 and 256 cores system,

Table 2.9: MR requirement comparison results for 64 cores systems.

	FT-PHENIC	PHENIC	Xiang
Mod/Detc	64	64	64
Switch	1152	1152	1600
ACKs	640	640	-
Redundant MRs	384	-	-
Total	2240	1856	1664
Sta. Power(mW)	44	37	33

respectively. We can see that the optimized networks have the lowest number of rings. In fact, this kind of network is even more sensitive to MR faults as each MR is critical for the functionality of the node. In addition, with minimal number of rings, the resulting insertion loss is lower than the fault tolerant design. For the proposed FT-PHENIC system, it has an additional rings used for acknowledgment signal, compared to the other networks, as well as for fault-tolerance. This increase can reach 33% when compared to the optimized crossbar and PHENIC systems. We also observe the same behavior when evaluating the required static thermal tuning, which is required to maintain the functionality of the ring, under 20K temperature with $1\mu\text{W}$ for each ring.

Table 2.10: MRs requirement comparison results for 256-core systems.

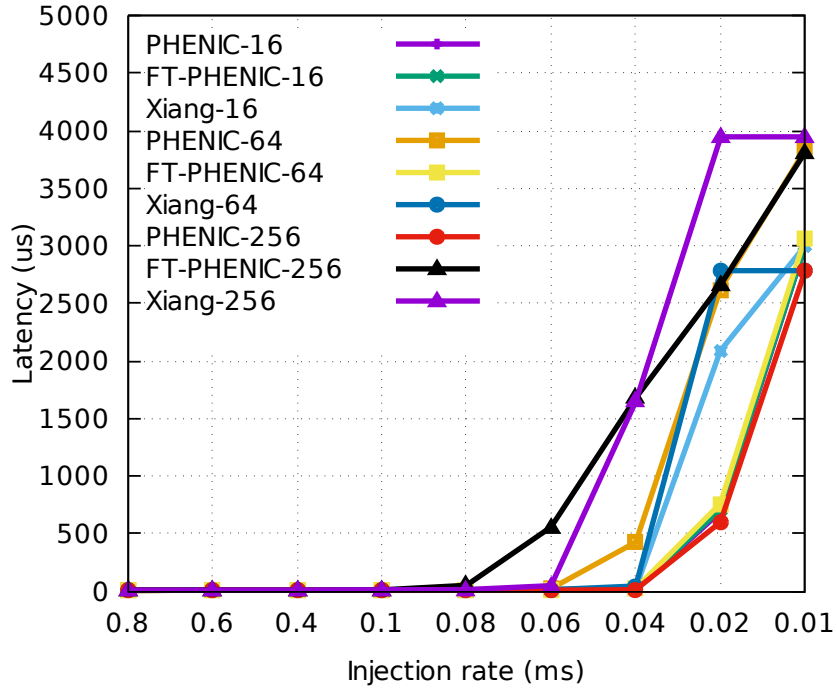
	FT-PHENIC	PHENIC	Xiang
Mod/Detc	256	256	256
Switch	4608	4608	6400
ACKs	2560	2560	-
Redundant MRs	1536	-	-
Total	8960	7424	6656
Sta. Power(mW)	179	149	133

2.3.2 Latency and Bandwidth Evaluation


Figures 2.8 (a) and (b) show the overall average latency and the average latency near the saturation region, respectively. We can see that for zero-load latency, all networks behave in the same way. Near saturation, PHENIC shows more flexibility and scalability in 256 cores when compared to the other networks. For the 64 cores configuration, the crossbar-based system slightly outperforms both PHENIC systems in terms of latency. This can be explained by the use of Optical-to-Electronic conversion of the *Teardown* which affects the overall latency of small networks.

The latency is heavily affected by the failure rate of MRs, and as the systems fail more, the latency increases until the whole system fails. This has a lot to do with failed path setup. Figure 2.9 shows the results of the latency test when adding in varying amounts of MR failures. The FT-PHENIC demonstrates its ability to withstand MR failures over all other systems.

For the achieved bandwidth, Fig. 2.10 shows that the bandwidth is increased by about 51% when compared to Xiang's system, for both 64 and 256 cores configurations. When compared to the crossbar, torus and PHENIC sys-



(a)



(b)

Figure 2.8: Latency comparison results under random uniform traffic: (a) Overall Latency, (b) Latency near-saturation.




Figure 2.9: Latency results of each system as faults are introduced.

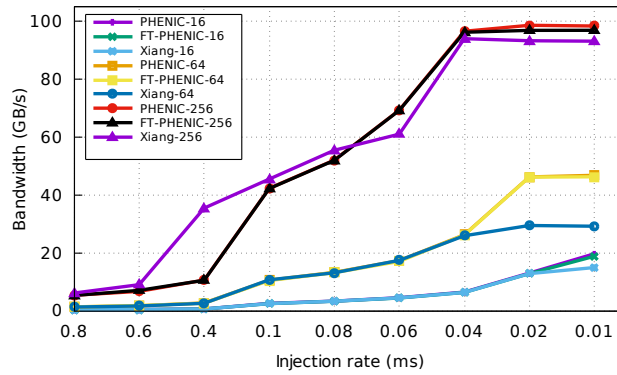


Figure 2.10: Bandwidth comparison results under random uniform traffic.

tems, we see that the four systems behave similarly. While the torus system has the capability of setting the path with less hop count, the FT-PHENIC system can achieve the same performance without the need for an extra network access which is required for the torus. This behavior is observed for 16, 64 and 256 core systems.

The latency increase caused by failed MRs will in turn cause the bandwidth to decrease. The effects of the failures on the bandwidth can be seen




Figure 2.11: Bandwidth comparison results as faults are introduced.

in figure 2.11. As with the latency, only FT-PHENIC and Xiang show any tolerance to faults, with FT-PHENIC outperforming Xiang.

2.3.3 Energy Evaluation

Figure 2.12 shows the total energy and the energy efficiency comparison results for 16, 64 and 256 cores systems. For the 256 cores configuration, the proposed system outperforms all other networks. This is illustrated by an improvement in terms of energy efficiency reaching 26% when compared the crossbar-based (non blocking). When compared to the torus-based architecture, FT-PHENIC improves the energy efficiency by upwards of 70%. The torus-based architecture offers high bandwidth thanks to the connection between edges leading to short communications. On the other hand, it comes at high energy cost. This can be explained by the fact that the additional input-ports, required for the edge connections established in the torus-based system, incur increased area and consequently an energy overhead.

Figure 2.13 shows the total energy and energy efficiency of the systems when 4% of their MRs have failed. Some systems were not able to complete simulation, and so their energy is marked as 0J, and an efficiency of 0pJ/bit, just so the functioning ones remain visible. The extra energy comes from the extra run time. It is important to notice how much the scale has changed for the energy efficiency between the fault-free and 4% fault results.

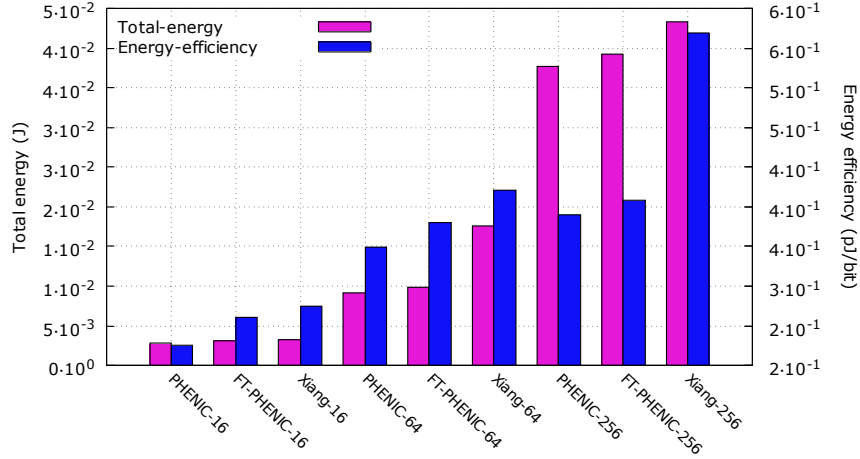


Figure 2.12: Total energy and energy efficiency comparison results under random uniform traffic near-saturation.

From these results, we can see that FT-PHENIC outperforms systems with either non-blocking or blocking switches. In addition, it provides heightened energy efficiency, far greater than the torus-based which can offer the same bandwidth as the proposed system. We conclude that the obtained improvement by FT-PHENIC is the result of the association of three main factors together: (1) the non-blocking switch supporting optical acknowledgment signals, (2) the light-weight router with reduced buffer size, (3) and the path setup algorithm to adopt hybrid switching inside the photonic switch.

2.4 Related Literature

There are three main types of optical fault tolerance that we were able to find. The first one is various methods of adaptive routing. The second one is techniques involving redundancy, which is commonly implemented in the network interface by using WDM as a redundancy technique. The third one involves buffering, checking, and proceeding like a standard electronic NoC.

Adaptive routing [36, 57, 49] is the most common method for fault-tolerance in mesh based architectures because of the large amount of possible minimal paths. It does require some extra logic in the routing decision, but

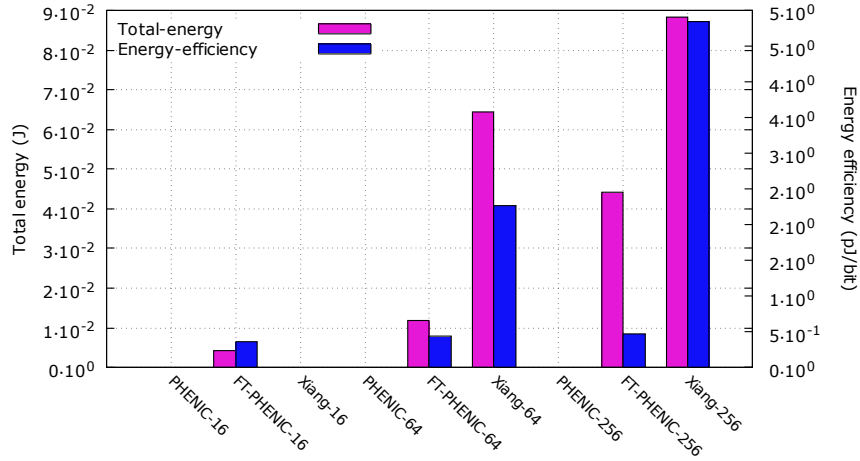


Figure 2.13: Total energy and energy efficiency comparison results under random uniform traffic with 4% of MRs acting faulty.

this is minimal compared to an extra interconnect at each location. For it to truly support multiple faults, it must also support non-minimal routing in order to avoid a non-reserved deadlock situation. It should also be noted that implementing fault tolerance on a deadlock free algorithm can negate that feature. This is not troublesome to optical networks as deadlock is a non-issue due to the fact that end-to-end is reserved before the transmission can start, and is only an issue during path setup.

Ramesh et al. proposed a method [49] of determining and using backup routes. The algorithm determines the least cost path. This path will be used unless there is a fault detected, in which case the backup path is used. Ramesh proposed using a set of probe packets. When the destination receives one of the probe packets, it then sends a PACK signal for each probe packet. If a packet is dropped due to faults, then a NACK signal is sent. This is a solution of off-chip optical networks though.

Loh breaks his algorithm [36] into a similar fashion to Ramesh. It has a Default Routing algorithm and a backup routing method. His two methods are called Logical Route and Adaptive Route. The Logical Route in his paper is a few sets of dimension order routing. The adaptive algorithm determines which of the deterministic routings to use. This method simply checks for

faults along the way, and if it can be detected, then it tries to switch to the other form of dimension order routing. This is an attempt to shift from X to Y when a problem is found in the X direction. This results in a routing algorithm which is minimal and adaptive, deadlock-free, and livelock-free.

Fault Regions [57] is a form of adaptive routing where each node keeps track of the permanent faults of its neighbors. This then allows for the path making decision to be educated with respect to faults up to a certain distance away. It can then guarantee that no old permanent faults are going to cause problems with the transmission. One such an algorithm is proposed by Xingyun [57]. He proposed a quite interesting optical network. It comes in the form of a torus which only allows data in two directions. This allows for some unique fault tolerance ideas. While they may not be minimalistic routing it will switch directions, go under the chip and come back from the top and reroute to avoid a bad crossing. This could possibly cause large amounts of insertion loss from routing around the network's length multiple times. This loss would translate to high power cost, and not yield any true benefits to converting to optical. This is still only monitoring its own outputs though.

Look Ahead Routing [56] is another type of adaptive routing which is most interesting to implement in a nanophotonic setting. This is where a node has knowledge of its neighbors' faulty links, and possibly its neighbors' neighbors' links. With this data at hand, the routing can protect a path and guarantee its success. The only issue would be implementing one of the detection algorithms mentioned at the beginning of this section. Although it hasn't been implemented in a photonic chip yet, there is no obvious reason preventing it from being translated over. Xiang's method [56] uses a Minus-first routing algorithm as a basis. The author does not detail how to detect a faulty link, but once a faulty link is discovered, it runs a Minus-first algorithm, checking each step along the way. This method attempts to find all paths from the source to the destination from the problematic node, and then determines which one requires the least amount of time. This switch shows that only the links are optical, and the switches themselves are electrical. This also allows for the implementation of buffers, which allow for a few more fault tolerance options which can be detailed in Radetzki's paper [47].

Modular redundancy uses WDM (Wavelength Division Multiplexing) as a fault tolerance tool [38, 51, 60]. The general idea is that if a certain wavelength is causing problems, either through noise or a manufacturing defect, and this problem can be detected, then certain wavelengths can be disabled and enabled. This is highly effective for modulator and photo-detector based faults. These focus on permanent and intermittent faults, because a transient

fault would occur far too late for a wavelength to be switched. *Noise* has been




Figure 2.14: Example of photonic switches. From left to right: PHENIC's original [8], Crossbar, and Crux [59].

a large source of faults within optical networks. Currently, there are many different forms of optical switches which are used in networks. The main goal of these switch designs is to reduce the area, when compared to the crossbar switch. We will only focus on the non-blocking switches because of their performance benefits. Three examples of optical switches can be seen in figure 2.14. The first is an example of a typical optimized switch, which reduces crossings and MRs. The second is the five-port crossbar switch, which uses the maximum number of MRs, crossings, and terminators, but is a simplistic non-blocking design. The last, Crux by Ye et al. [59], is a switch which is optimized for XY-deterministic routing. This allows it to drop some extra MRs, but it no longer maintains the functionality to travel from the Y-direction to X-direction, such as North to East. This does greatly reduce the noise, when compared to other switches which can perform all network routing operations. Many other switches and networks were proposed to improve the SNR [41, 43, 28, 21, 20]. The reason this noise is so heavily researched is explained by Nikdast et al. [42].

Additionally, various authors have looked into the affect of *thermal variance*, and how to combat it [35, 34]. There are various ways to combat it, but the most common way is to cool down the ring to normal temperatures, which can be done by keeping it inactive, or by thermal tuning [35]. Trimming [4] was also one solution, which was mentioned in the introduction, and appears to be a promising answer to the problem. To the best of our knowledge, none of the existing solutions proposed so far take advantage of switch structure to provide fault tolerance. The focus of all other research has been on the routing algorithms or different locations to provide modular redundancy, or noise reduction.

2.5 Chapter Summary and Discussion

This chapter presents a fault-tolerant Photonic Network-on-Chip architecture, which uses minimal redundancy to assure accuracy of the packet transmission even after faulty microrings (MRs) are detected. The system is based on a fault-tolerant path-configuration and routing algorithm, and a microring fault-resilient photonic router. Simulation results show that FT-PHENIC enjoys about 50% increase in bandwidth and about 60% decrease in energy related to the typical crossbar unit, versus other reported architectures. Additionally, the FT-PHENIC tolerates MR faults quite well up until around when 20% of the MRs have failed. These encouraging results highlight the potential of using photonic on chip and the FT-PHENIC hybrid PNoC architecture to meet the design and performance challenges of future generations of many-core systems.

One key thing holding back the reliability of optical switches is the reliability of the basic MR unit. This reliability is based off of the physical parameters that are used when designing each unit. We would like to explore the physical properties of the MRs themselves to improve the reliability. As we have previously said in the paper, making small changes to the shape, such as using racetracks [41] has led to an improvement in the reliability of MRs, and a reduction in the sensitivity to thermal variation. This means without changing the bending radius or waveguide thickness or material, they were able to improve reliability, and we would like to continue with such research.

Another item that would greatly aid the development of optical routers is the ability to buffer. Even more so the ability to read the data in multiple locations. Currently, splitting the data to be read will cause a large amount of insertion loss. Buffering is currently limited to causing a delay by creating optical coils, and can only delay it a very minimal amount of time, and cause a large amount of propagation loss [23]. Being able to read the data at multiple locations could allow for error correcting codes to not only fix some of the errors, but also aid in the fault diagnosis schemes.

This research mainly focused on improving the fault-tolerance of the network. We attempted to address the process variation problems of optical switches, but thermal variation is still a large problem. The temperature fluctuation can temporarily cause an MR to respond to an improper wavelength, which can result in larger problems.

References

- [1] Cisse Ahmadou Dit Adi, Hiroki Matsutani, Michihirio Koibuchi, Hidetsugu Irie, Takemitsu Miyoshi, and Tsutomu Yoshinaga. An efficient path setup for a photonic network-on-chip. In *2010 First International Conference on Networking and Computing*, pages 156–161, Nov 2010.
- [2] Mridul Agarwal, Bipul C Paul, Ming Zhang, and Subhasish Mitra. Circuit failure prediction and its application to transistor aging. In *25th IEEE VLSI Test Symposium (VTS'07)*, pages 277–286. IEEE, 2007.
- [3] Akram Ben Ahmed. *High-throughput Architecture and Routing Algorithms Towards the Design of Reliable Mesh-based Many-Core Network-on-Chip Systems*. PhD thesis, Graduate School of Computer Science and Engineering, University of Aizu, March 2015.
- [4] Jung Ho Ahn, Marco Fiorentino, Raymond G. Beausoleil, Nathan Binkert, Al Davis, D. Fattal, Norman P. Jouppi, Moray McLaren, C.M. Santori, Robert S. Schreiber, S.M. Spillane, Dana Vantrease, and Q. Xu. Devices and architectures for photonic chip-scale integration. *Applied Physics A*, 95(4):989–997.
- [5] Abderazek Ben Abdallah. *Multicore Systems-on-Chip: Practical Hardware/Software Design, 2nd Edition*. Atlantis, 2013.
- [6] Abderazek Ben Abdallah and Masahiro Sowa. Basic Network-on-Chip Interconnection for Future Gigascale MCSoc Applications: Communication and Computation Orthogonalization. In *JASSST2006*, 2006.
- [7] Achraf Ben Ahmed and Abderazek Ben Abdallah. Phenix: silicon photonic 3d-network-on-chip architecture for high-performance heterogeneous many-core system-on-chip. In *Sciences and Techniques of Automatic Control and Computer Engineering (STA), 2013 14th International Conference on*, pages 1–9, Dec 2013.
- [8] Achraf Ben Ahmed and Abderazek Ben Abdallah. Hybrid silicon-photonic network-on-chip for future generations of high-performance many-core systems. *The Journal of Supercomputing*, DOI: 10.1007/s11227-015-1539-0, 2015.
- [9] Achraf Ben Ahmed, Meyer Meyer, Yuichi Okuyama, and Abderazek Ben Abdallah. Hybrid photonic noc based on non-blocking photonic switch and light-weight electronic router. In *2015 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, October 2015.
- [10] Achraf Ben Ahmed, Michael Meyer, Yuichi Okuyama, and Abderazek Ben Abdallah. Efficient router architecture, design and performance exploration for many-core hybrid photonic network-on-chip (2d-phenix). In *Information Science and Control Engineering (ICISCE), 2015 2nd International Conference on*, pages 202–206, April 2015.
- [11] Achraf Ben Ahmed, Yuichi Okuyama, and Abderazek Ben Abdallah. Contention-free routing for hybrid photonic mesh-based network-on-chip systems. In *The 9th IEEE International Symposium on Embedded Multicore/Manycore SoCs (MCSoc)*, pages 235–242, September 2015.
- [12] Achraf Ben Ahmed, Yuichi Okuyama, and Abderazek Ben Abdallah. Non-blocking electro-optic network-on-chip router for high-throughput and low-power many-core systems. In *The World Congress on Information Technology and Computer Applications 2015*, June 2015.

- [13] Akram Ben Ahmed and Abderazek Ben Abdallah. Architecture and design of high-throughput, low-latency, and fault-tolerant routing algorithm for 3d-network-on-chip (3d-noc). *J. Supercomput.*, 66(3):1507–1532, dec 2013.
- [14] Akram Ben Ahmed and Abderazek Ben Abdallah. Graceful deadlock-free fault-tolerant routing algorithm for 3d network-on-chip architectures. *J. Parallel Distrib. Comput.*, 74(4):2229–2240, April 2014.
- [15] Wim Bogaerts, Peter De Heyn, Thomas Van Vaerenbergh, Katrien De Vos, Shankar Kumar Selvaraja, Tom Claes, Pieter Dumon, Peter Bienstman, Dries Van Thourhout, and Roel Baets. Silicon microring resonators. *Laser & Photonics Reviews*, 6(1):47–73, 2012.
- [16] Lars Brusberg, Henning Schrder, Marco Queisser, and Klaus-Dieter Lang. Single-mode glass waveguide platform for dwdm chip-to-chip interconnects. In *Electronic Components and Technology Conference (ECTC), 2012 IEEE 62nd*, pages 1532–1539, May 2012.
- [17] Johnnie Chan and Keren Bergman. Photonic interconnection network architectures using wavelength-selective spatial routing for chip-scale communications. *Optical Communications and Networking, IEEE/OSA Journal of*, 4(3), March 2012.
- [18] Johnnie Chan, Gilbert Hendry, Keren Bergman, and Luca P. Carloni. Physical-layer modeling and system-level design of chip-scale photonic interconnection networks. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 30(10):1507–1520, Oct 2011.
- [19] Johnnie Chan, Gilbert Hendry, Aleksandr Biberman, Keren Bergman, and Luca P Carloni. Phoenixsim: A simulator for physical-layer analysis of chip-scale photonic interconnection networks. In *Proceedings of the Conference on Design, Automation and Test in Europe*, pages 691–696. European Design and Automation Association, 2010.
- [20] Sai Vineel Reddy Chittamuru and Sudeep Pasricha. Crosstalk mitigation for high-radix and low-diameter photonic noc architectures. *Design Test, IEEE*, 32(3):29–39, June 2015.
- [21] Sai Vineel Reddy Chittamuru and Sudeep Pasricha. Improving crosstalk resilience with wavelength spacing in photonic crossbar-based network-on-chip architectures. In *Circuits and Systems (MWSCAS), 2015 IEEE 58th International Midwest Symposium on*, pages 1–4. IEEE, 2015.
- [22] Sai T. Chu, Wugen Pan, Shinya Sato, Taro Kaneko, Brent E. Little, and Yasuo Kokubun. Wavelength trimming of a microring resonator filter by means of a uv sensitive polymer overlay. *Photonics Technology Letters, IEEE*, 11(6):688–690, June 1999.
- [23] Sasan Fathpour and Nabeel A. Riza. Silicon-photonics-based wideband radar beam-forming: basic design. *Optical Engineering*, 49(1):018201–018201–7, 2010.
- [24] Sheng guang Yang, Li Li, Yu ang Zhang, Bing Zhang, and Yi Xu. A power-aware adaptive routing scheme for network on a chip. In *ASIC, 2007. ASICON '07. 7th International Conference on*, pages 1301–1304, Oct 2007.
- [25] Gilbert Hendry, Eric Robinson, Vitaliy Gleyzer, Johnnie Chan, Luca P. Carloni, Nadya Bliss, and Keren Bergman. Circuit-switched memory access in photonic interconnection networks for high-performance embedded computing. In *High Performance Computing, Networking, Storage and Analysis (SC), 2010 International Conference for*, pages 1–12, Nov 2010.

- [26] Zhan-Shuo Hu, Fei-Yi Hung, Kuan-Jen Chen, Shoou-Jinn Chang, Wei-Kang Hsieh, and Tsai-Yu Liao. Improvement in thermal degradation of zno photodetector by embedding silver oxide nanoparticles. *Functional Materials Letters*, 6(01):1350001, 2013.
- [27] Andrew B. Kahng, Bin Li, Li-Shiuan Peh, and Kambiz Samadi. Orion 2.0: A power-area simulator for interconnection networks. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 20(1):191–196, Jan 2012.
- [28] Pradheep Khanna Kaliraj. Reliability-performance trade-offs in photonic noc architectures. 2013.
- [29] Roman Kappeler. Radiation testing of micro photonic components. Stagiaire Project Report. ESA/ESTEC. September 29, 2004. Ref. No.: EWP 2263.
- [30] John Keane and Chris H Kim. An odometer for cpus: Microprocessors don't normally show wear and tear, but wear they do. *IEEE SPECTRUM*, 48(5):26–31, 2011.
- [31] John Keane, Tae-Hyoung Kim, and Chris H Kim. An on-chip nbti sensor for measuring pmos threshold voltage degradation. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 18(6):947–956, 2010.
- [32] Brian R. Koch, Alexander W. Fang, Oded Cohen, and John E. Bowers. Mode-locked silicon evanescent lasers. *Optics Express*, 18(15), 2007.
- [33] Kelin Kuhn, Chris Kenyon, Avner Kornfeld, Mark Liu, Atul Maheshwari, Wei-kai Shih, Sam Sivakumar, Greg Taylor, Peter VanDerVoorn, and Keith Zawadzki. Managing process variation in intel's 45nm cmos technology. *Intel Technology Journal*, 12(2), 2008.
- [34] Hui Li, Alain Fourmigue, Sébastien Le Beux, Xavier Letartre, Ian O'Connor, and Gabriela Nicolescu. Thermal aware design method for vcsel-based on-chip optical interconnect. In *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, pages 1120–1125. EDA Consortium, 2015.
- [35] Zheng Li, Moustafa Mohamed, Xi Chen, Eric Dudley, Ke Meng, Li Shang, Alan R Mickelson, Russ Joseph, Manish Vachharajani, Brian Schwartz, et al. Reliability modeling and management of nanophotonic on-chip networks. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 20(1):98–111, 2012.
- [36] Peter KK Loh and Wen-Jing Hsu. Design of a viable fault-tolerant routing strategy for optical-based grids. In *Parallel and Distributed Processing and Applications*, pages 112–126. Springer, 2003.
- [37] Serge Luryi, Jimmy Xu, and Alex Zaslavsky. *Future trends in microelectronics: up the nano creek*. John Wiley & Sons, 2007.
- [38] Moray McLaren, Nathan Lorenzo Binkert, Alan Lynn Davis, and Marco Florentino. Energy-efficient and fault-tolerant resonator-based modulation and wavelength division multiplexing systems, 22 2014. US Patent 8,705,972.
- [39] Michael C. Meyer, Akram Ben Ahmed, Yuichi Okuyama, and Aabderazek Ben Abdallah. Fttdd: Microring fault-resilient optical router for reliable optical network-on-chip systems. In *Embedded Multicore/Many-core Systems-on-Chip (MCSoc), 2015 IEEE 9th International Symposium on*, pages 227–234, Sept 2015.
- [40] Evelyn Mintarno, Joëlle Skaf, Rui Zheng, Jyothi Bhaskar Velamala, Yu Cao, Stephen Boyd, Robert W Dutton, and Subhasish Mitra. Self-tuning for maximized lifetime energy-efficiency in the presence of circuit aging. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 30(5):760–773, 2011.

- [41] Moustafa Mohamed. *Silicon Nanophotonics for Many-Core On-Chip Networks*. PhD thesis, University of Colorado, 2013.
- [42] Mahdi Nikdast and Jiang Xu. On the impact of crosstalk noise in optical networks-on-chip. In *Design Automation Conference (DAC)*, 2014.
- [43] Mahdi Nikdast, Jiang Xu, Xiaowen Wu, Wei Zhang, Yaoyao Ye, Xuan Wang, Zhehui Wang, and Zhe Wang. Systematic analysis of crosstalk noise in folded-torus-based optical networks-on-chip. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 33(3):437–450, 2014.
- [44] 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, pages 95–104, New York, NY, USA, 2011. ACM.
- [45] Yan Pan, Prabhat Kumar, John Kim, Gokhan Memik, Yu Zhang, and Alok Choudhary. Firefly: illuminating future network-on-chip with nanophotonics. In *ACM SIGARCH Computer Architecture News*, volume 37, pages 429–440. ACM, 2009.
- [46] Kyle Preston, Nicolas Sherwood-Droz, Jacob S. Levy, and Michal Lipson. Performance guidelines for wdm interconnects based on silicon microring resonators. In *Lasers and Electro-Optics (CLEO), 2011 Conference on*, pages 1–2, May 2011.
- [47] Martin Radetzki, Chaochao Feng, Xueqian Zhao, and Axel Jantsch. Methods for fault tolerance in networks-on-chip. *ACM Computing Surveys (CSUR)*, 46(1):8, 2013.
- [48] D. Rafizadeh, J.P. Zhang, S.C. Hagness, A. Taflove, K.A. Stair, S.T. Ho, and R.C. Tiberio. Temperature tuning of microcavity ring and disk resonators at 1.5- μm . In *Lasers and Electro-Optics Society Annual Meeting, 1997. LEOS '97 10th Annual Meeting. Conference Proceedings., IEEE*, volume 2, pages 162–163 vol.2, Nov 1997.
- [49] Gayatri Ramesh and S SundaraVadivelu. A reliable and fault tolerant routing for optical wdm networks. *arXiv preprint arXiv:0912.0602*, 2009.
- [50] Samar K. Saha. Modeling process variability in scaled cmos technology. *IEEE Design & Test of Computers*, 27(2):0008–16, 2010.
- [51] Laxman Sahasrabuddhe, Senthil Ramamurthy, and Biswanath Mukherjee. Fault management in ip-over-wdm networks: Wdm protection versus ip restoration. *Selected Areas in Communications, IEEE Journal on*, 20(1):21–33, 2002.
- [52] Assaf Shacham, Keren Bergman, and Luca P. Carloni. On the design of a photonic network-on-chip. In *Networks-on-Chip, 2007. NOCS 2007. First International Symposium on*, pages 53–64, May 2007.
- [53] Assaf Shacham, Keren Bergman, and Luca P. Carloni. Photonic networks-on-chip for future generations of chip multiprocessors. *Computers, IEEE Transactions on*, 57(9):1246–1260, Sept 2008.
- [54] Zhijuan Tu, Zhiping Zhou, and Xingjun Wang. Reliability considerations of high speed germanium waveguide photodetectors. In *SPIE OPTO*, pages 89820W–89820W. International Society for Optics and Photonics, 2014.
- [55] Dana Vantrease, Robert Schreiber, Matteo Monchiero, Moray McLaren, Norman P Jouppi, Marco Fiorentino, Al Davis, Nathan Binkert, Raymond G Beausoleil, and Jung Ho Ahn. Corona: System implications of emerging nanophotonic technology. In *ACM SIGARCH Computer Architecture News*, volume 36, pages 153–164. IEEE Computer Society, 2008.

- [56] Dong Xiang, Yan Zhang, ShuChang Shan, and Yi Xu. A fault-tolerant routing algorithm design for on-chip optical networks. In *Reliable Distributed Systems (SRDS), 2013 IEEE 32nd International Symposium on*, pages 1–9, Sept 2013.
- [57] Qi Xingyun, Feng Quanyou, Chen Yongran, Dou Qiang, and Dou Wenhua. A fault tolerant bufferless optical interconnection network. In *Computer and Information Science, 2009. ICIS 2009. Eighth IEEE/ACIS International Conference on*, pages 249–254. IEEE, 2009.
- [58] Yi Xu, Jun Yang, and Rami Melhem. Tolerating process variations in nanophotonic on-chip networks. In *ACM SIGARCH Computer Architecture News*, volume 40, pages 142–152. IEEE Computer Society, 2012.
- [59] Yaoyao Ye, Xiaowen Wu, Jiang Xu, Wei Zhang, Mahdi Nikdast, and Xuan Wang. Holistic comparison of optical routers for chip multiprocessors. In *Anti-Counterfeiting, Security and Identification (ASID), 2012 International Conference on*, pages 1–5. IEEE, 2012.
- [60] Jing Zhang and B Mukherjee. A review of fault management in wdm mesh networks: basic concepts and research challenges. *Network, IEEE*, 18(2):41–48, 2004.
- [61] Shiyang Zhu and Guo-Qiang Lo. Vertically-stacked multilayer photonics on bulk silicon toward three-dimensional integration. *Lightwave Technology, Journal of*, PP(99):1–1, 2015.