Silicon photonic switching: from building block design to intelligent control


Yishen Huang


Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences


COLUMBIA UNIVERSITY


2020

# Abstract

Silicon photonic switching: from building block design to intelligent control

Yishen Huang

The rapid growth in data communication technologies is at the heart of enriching the digital experiences for people around the world. Encoding high bandwidth data to the optical domain has drastically changed the bandwidth-distance trade-off imposed by electrical media. Silicon photonics, sharing the technological maturity of the semiconductor industry, is a platform poised to make optical interconnect components more robust, manufacturable, and ubiquitous. One of the most prominent device classes enabled by the silicon photonics platform is photonic switching, which describes the direct routing of optical signal carriers without the optical-electrical-optical conversions. While theoretical designs and prototypes of monolithic silicon photonic switch devices have been studied, realizing high-performance and feasible switch systems requires explorations of all design aspects from basic building blocks to control systems. This thesis provides a holistic collection of studies on silicon photonic switching in topics of novel switching element designs, multi-stage switch architectures, device calibration, topology scalability, smart routing strategies, and performance-aware control plane.

First, component designs for assembling a silicon photonic switch device are presented. Structures that perform 2×2 optical switching functions are introduced. To realize switching granularities in both spatial and spectral domains, a resonator-assisted Mach-Zehnder interferometer design is demonstrated with high performance and design robustness. Next, multi-stage monolithic switching devices with microring resonator-based switching elements are investigated. An 8×8 switch device with dual-microring switching elements is presented with a well-balanced set of performance metrics in extinction ratio, crosstalk suppression, and optical bandwidth. Continued scaling in the switch port count requires both an

economic increase in the number of switching elements integrated in a device and the preservation of signal quality through the switch fabric. A highly scalable switch architecture based on Clos network with microring switch-and-select sub-switches is presented as a solution to reach high switch radices while addressing key factors of insertion loss, crosstalk, and optical passband to ensure end-to-end switching performance.

The thesis then explores calibration techniques to acquire and optimize system-wide control points for integrated silicon switch devices. Applicable to common rearrangeably non-blocking switch topologies, automated procedures are developed to calibrate entire switch devices without the need for built-in power monitors. Using Mach-Zehnder interferometer-based switching elements as a demonstration, calibration techniques for optimal control points are introduced to achieve balanced push-pull drive scheme and reduced crosstalk in switching operations. Furthermore, smart routing strategies are developed based on optical penalty estimations enabled by expedited lightpath characterization procedures. Leveraging configuration redundancies in the switch fabric, the routing strategies are capable of avoiding the worst penalty optical paths and effectively elevate the bottom-line performance of the switch device.

Additional works are also presented on enhancing optical system control planes with machine learning techniques to accurately characterize complex systems and identify critical control parameters. Using flexgrid networks as a case study, light-weight machine learning workflows are tailored to devise control strategies for improving spectral power stability during wavelength assignment and defragmentation. This work affirms the efficacy of intelligent control planes to predict system dynamics and drive performance optimizations for optical interconnect systems.

# Table of Contents

# List of Figures

iv

# List of Tables

# Glossary

| | | | | |
|---|---|---|---|---|
| ADC: | analog-to-digital converter | | MZM: | Mach-Zehnder modulator |
| AGC: | automatic gain control | | NRZ: | non-return-to-zero |
| ASIC: | application-specific integrated circuit | | OOK: | on-off keying |
| | | | OPM: | optical performance monitor |
| ASE: | amplified spontaneous emission | | OSA: | optical spectrum analyzer |
| B2B: | back-to-back | | OSNR: | optical signal-to-noise ratio |
| BER: | bit error rate | | PA: | power adjuster |
| CBR: | case-based reasoning | | PC: | polarization controller |
| CDF: | cumulative distribution function | | PCB: | printed circuit board |
| CMOS: | complementary metal–oxide–semiconductor | | PD: | photodiode |
| | | | PDK: | process design kit |
| CW: | continuous wave | | PIC: | photonic integrated circuit |
| DAC: | digital-to-analog converter | | PILOSS: | path-independent loss |
| DC: | directional coupler | | PLC: | planar lightwave circuit |
| DFB: | distributed feedback | | PMF: | polarization-maintaining fiber |
| DUT: | device under test | | PRBS: | pseudo-random bit sequence |
| DWDM: | dense wavelength-division multiplexing | | QoS: | quality of service |
| | | | RBF: | radial basis function |
| EDFA: | Erbium-doped fiber amplifier | | RMSE: | root mean square error |
| EO: | electro-optic | | RNB: | rearrangeably non-blocking |
| ER: | extinction ratio | | ROADM: | reconfigurable optical add/drop multiplexer |
| FC: | fiber coupler | | | |
| FoM: | figure of merit | | RSA: | routing and spectrum allocation |
| FPGA: | field programmable gate array | | RR: | ridge regression |
| FSR: | free-spectral range | | SE: | switching element |
| FWHM: | full width at half maximum | | SerDes: | serializer-deserializer |
| I/O: | input/output | | SNB: | strictly non-blocking |
| IPDR: | input power dynamic range | | SOA: | semiconductor optical amplifier |
| IL: | insertion loss | | SOI: | silicon-on-insulator |
| KBR: | kernelized Bayesian regression | | SSMF: | standard single-mode fiber |
| LCoS: | liquid crystal on silicon | | STDEV: | standard deviation |
| LR: | logistic regression | | TDM: | time-division multiplexing |
| LUT: | look-up table | | TE: | transverse electric |
| MbB: | Make-before-Break (defragmentation method) | | TIA: | transimpedance amplifier |
| | | | TLD: | tunable laser diode |
| MEMS: | micro-electromechanical system | | TO: | thermo-optic |
| ML: | machine learning | | TOF: | tunable optical filter |
| MMI: | multi-mode interferometer | | VOA: | variable optical attenuator |
| MPW: | multi-project-wafer | | WDM: | wavelength-division multiplexing |
| MRR: | microring resonator | | | |
| MSE: | mean square error | | WSNB: | wide-sense non-blocking |
| MZI: | Mach-Zehnder interferometer | | WSS: | wavelength-selective switch |

# Acknowledgments

First, I would like to thank my academic advisor, Prof. Keren Bergman, for her guidance and teaching throughout my studies and research at Columbia University.

To the members of my doctorate dissertation committee, Prof. Keren Bergman, Prof. Michal Lipson, Prof. Alexander Gaeta, Prof. Gil Zussman, and Prof. Christine Hendon, thank you for your time and feedback.

I would like to thank Dr. Qixiang Cheng, Dr. Alex Meng, Dr. Yuhan Hung, and Dr. Payman Samadi for their mentorship throughout my doctoral studies, and Nathan Abrams and Richard Dai, with whom I had the fortune and joy to share an office. Much appreciation to all my academic collaborators at the Lightwave Research Laboratory for your tremendous teamwork. I wish you all the best in your studies and endeavors.

The internships I did greatly enriched my studies with valuable industry perspectives. I would like to thank my mentors – Po Dong and Guilhem de Valicourt at Nokia Bell Labs, Brian Taylor, Gilad Goldfarb, Chris Berry, James Stewart, and Hans-Juergen Schmidtke at Facebook, for their guidance and help.

To my family, I thank you for your love, care, and support every step of the way.

To my parents, Qingmei and Yufei

# Chapter 1: Introduction and Background

## 1.1 An Optically Connected Globe

The ubiquitous connectivity to the internet has changed our lives in every way. Studies have projected that, by 2023, over 66% of the global population will have internet access [1], with 6% annual growth as shown in Figure 1-1. The accelerating increase in both the coverage and speed of mobile phones, broadband internet, and Wi-Fi hotspots has significantly enriched people's lives by instilling digital conveniences in areas such as communication, transportation, finance, and entertainment as necessities that we cannot live without. As the growing majority of the global population turns to mobile devices as the primary surface for digital interactions, the infrastructure of the internet that stores, processes, and transfers both the application and user data has centralized in hyperscale datacenters operated by cloud companies such as Google, Facebook, Microsoft, Amazon, Alibaba, Baidu, and Tencent.
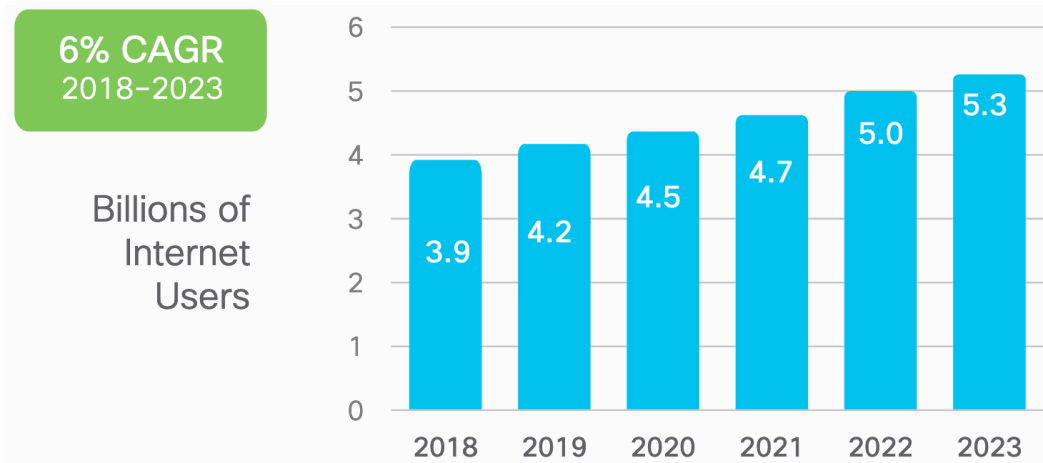


**Figure 1-1: Past and projected number of internet users from [1].**

While optical communication is commonly recognized as the main medium for continental and inter-continental connections, economic and power-efficient commodity optical transceivers are replacing electrical connections for short-distance links and becoming the dominant medium for intra-datacenter

interconnects. This inevitable transition is due to the limitation of resistive loss in the electrical lines at both high signal bandwidth and link length [2]. As intra-datacenter link rate grows beyond 100 Gbps and the physical link lengths extend beyond 2 km, fiber-based optical communications become the only viable candidate to support the continued scaling of the datacenter infrastructure. Nevertheless, the shift to optical transceivers is only half of the datacenter network transformation. Hyperscale datacenters are built on a high-availability, scalable, and flat network of Ethernet switches, which perform parsing, caching, processing, and forwarding of traffic packets. Each switch device consists of an electronic switch fabric and hundreds of electrical SerDes I/Os. The upgrade of the datacenter network speeds and architecture are typically bounded by the capabilities and functionalities of the switch devices, rather than the speed of the optics. For instance, in the new datacenter architecture, F16, deployed by Facebook [3], the company is able to drastically flatten the network with higher-radix switch ASICs, as shown in Figure 1-2.
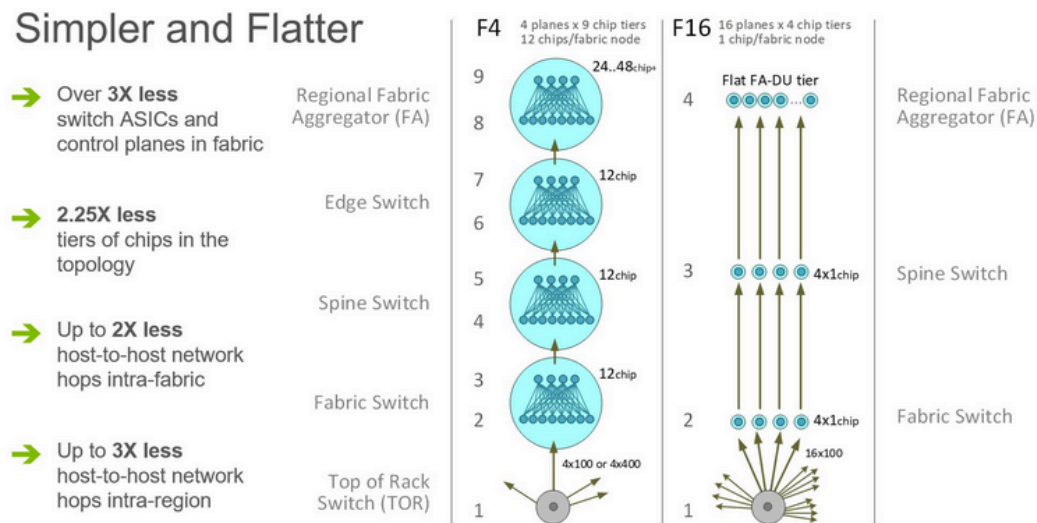


**Figure 1-2: The upgraded architecture of Facebook datacenter network, F16, and its comparison with the previous generation design, F4, from [3].**

Analogous to the transition from copper cables to optical fibers, the next evolution in high-speed communication for datacenter networks will be realized by closing the gap between the optics and the

switch ASICs. The communication and networking industry are increasingly focused on methods to integrate optics with the switch ASICs to significantly reduce electrical trace loss on the PCB and SerDes requirements, albeit at the cost of increasing complexities in design, fabrications, and packaging. Several proposed approaches for tightly integrated optics are shown in Figure 1-3. A recent demonstration [4] from Intel has also shown a reference design of a co-packaged switch ASIC with optical transceiver I/Os on the same package substrate.
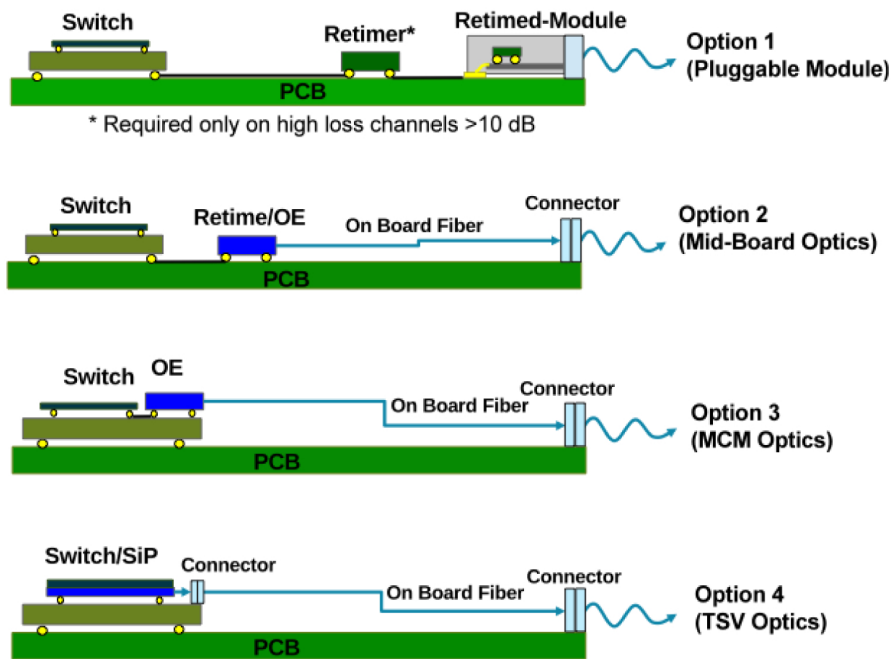


**Figure 1-3: Schematics for increasingly integrated optics and switch ASICs from [5].**

Photonic switching has the potential to further improve the efficiency and capability bottlenecks of the datacenter network by eliminating the pin and power constraints of the electronic switch ASICs, as well as the inefficiencies associated with opto-electronic conversions at every switch point of the network [6]. In addition, since the optical links are much more distance-agnostic than electrical links, performing switching in the optical domain can drive novel networking architectures and applications such as disaggregated hardware resources [7]. Disaggregation of the traditional server has been proposed as a

solution to improve utilization efficiency [8] by pooling similar resources, with the possibility of adaptively configuring and upgrading the resources for optimized performance. A disaggregated data center requires a robust interconnection fabric to handle the additional traffic between pooled hardware resources while offering high bandwidth and low latency for the system performance. Figure 1-4 shows a promising hardware disaggregation concept by using a photonically switched interconnection network to adaptively provision different computing resources [9] [10]. To realize these novel networking applications effectively and reliably, high-performance, economic, and scalable photonic switching solutions must first be developed. In this work, we comprehensively explore silicon photonic switching through building block designs, multi-stage switch architectures, device calibration, topology scalability, smart routing strategies, and performance-aware control plane.
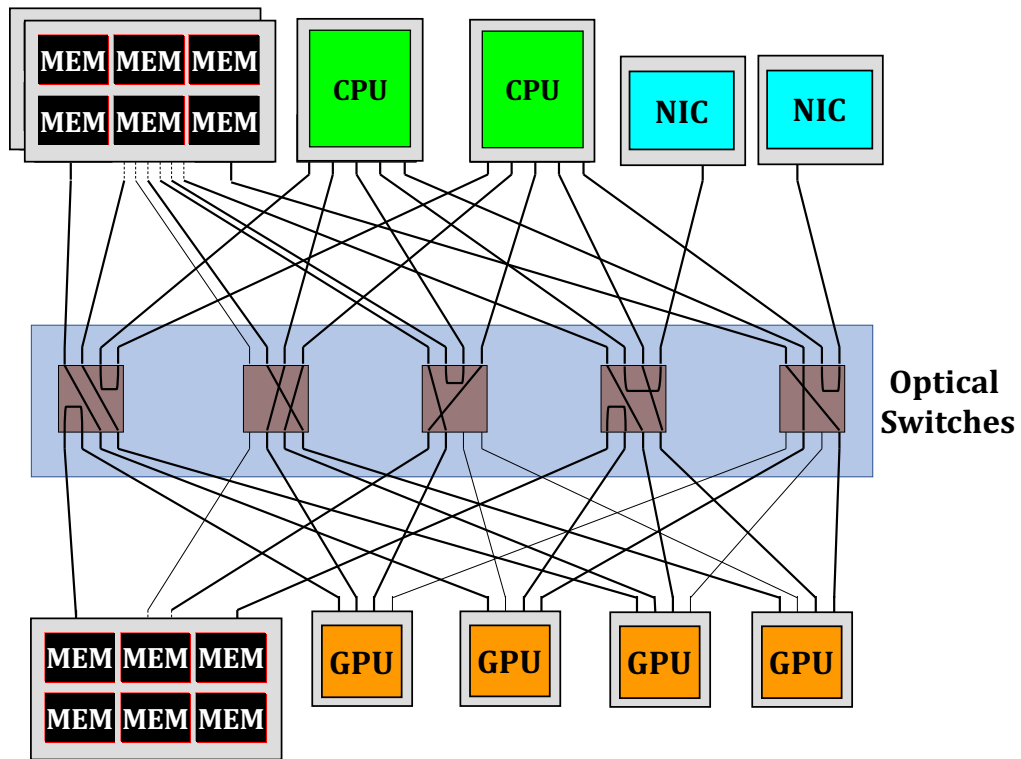


**Figure 1-4: A conceptual photonic integrated network for hardware resource disaggregation with photonic switches. Adopted from [8].**

## 1.2     Silicon Photonics

Silicon photonics is an emerging platform to design and fabricate nanoscale structures that guide and manipulate an optical field, and has experienced dramatic growth in adoption across a myriad of applications from data transmission, sensing, and photonic switching [11]. For datacenter interconnect applications, silicon photonics has the potential to enable highly energy efficient transceiver and networking systems [12]. The rise of silicon photonics can be attributed to three main factors: foundry manufacturability, efficient tuning mechanisms, and compact device footprints. On-going efforts on developing scalable testing and packaging techniques also contribute to fast iterations and optimizations in the silicon photonic design workflow.

*CMOS-compatible fabrication*

The silicon photonic platform shares many of the same processes and layer stack with the CMOS industry that has matured over the past few decades, and therefore promotes the reuse of past CMOS fabrication facilities. This not only significantly reduces the cost associated with design flow and device fabrication, but also ensures high scalability in manufacturing, as the CMOS industry is already positioned to supply the volume of the global computer hardware. This also means that innovative designs in silicon photonics can be realized in industrial scale, which significantly expedites the timeline for technology to market. Many of the major commercial CMOS foundries, such as Advanced Micro Foundry (AMF), GlobalFoundries, Interuniversity Microelectronics Centre (IMEC), Taiwan Semiconductor Manufacturing Company (TSMC), TowerJazz, and AIM Photonics, now offer MPW wafer runs as economic options for fabless companies and research groups to participate in silicon photonics innovations. An example of layer structures of the silicon photonic platform offered by AMF is shown in Figure 1-5.
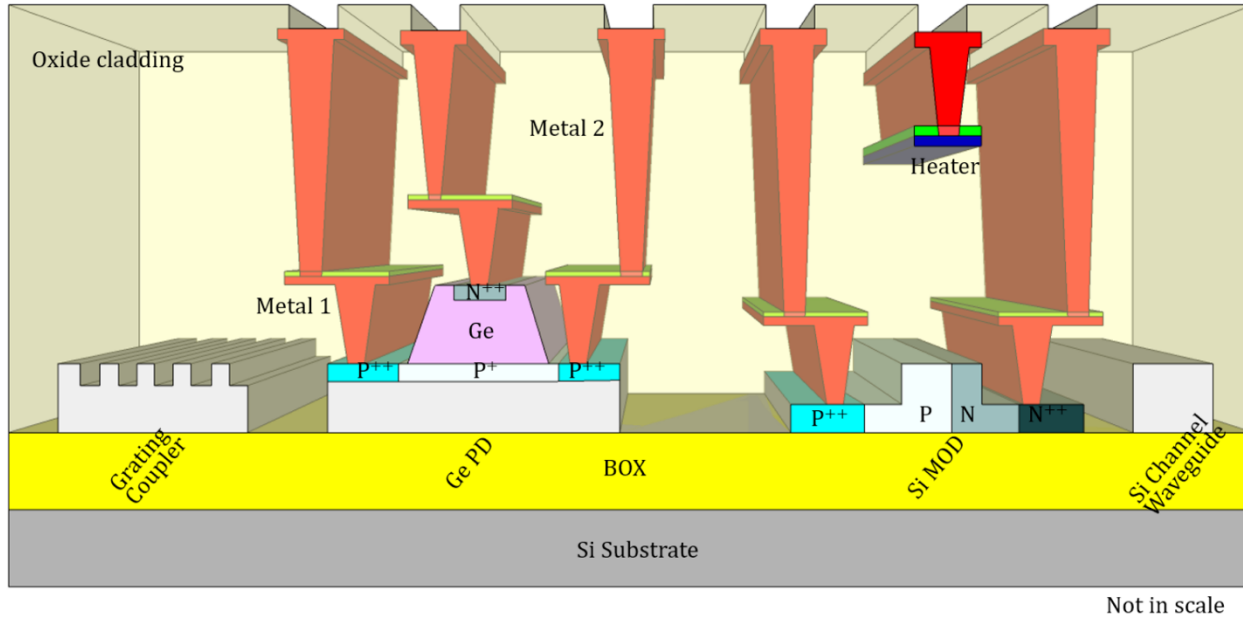
**Figure 1-5: Example layer stack of the silicon photonic platform from AMF [13].**

*Efficient tuneability of optical properties*

While silicon's indirect bandgap makes optical gain and emission infeasible, it provides very efficient mechanisms to manipulate the optical field by changing the refractive index of the material. Silicon's index has very strong temperature dependence at about 1.8e-4 $K^{-1}$ at room temperature [14], which is about 20 times greater than that of silica commonly used in PLC-based photonic circuits. Hence, the phase delay of the guided wave in a silicon PIC can be adjusted more efficiently by varying the temperature of the waveguide. Localized micro heaters in silicon can be achieved with resistive heater using a metal layer above the waveguide or through direct doping of the silicon itself. The former implementation induces no additional loss to the waveguide, but the latter can achieve higher tuning efficiency and faster tuning speed because the heat is generated directly at the waveguide core.

In addition, silicon's index has a strong carrier concentration dependence due to the plasma dispersion effect, making it feasible to modify the waveguide index through electro-doping. The free carrier induced index change, in both real and imaginary parts, is first empirically quantified in [15]. For common

telecom wavelengths around 1.55 µm, the electro-optic changes in the real and imaginary parts of the index are:

$$\Delta n = -\left[8.8 \times 10^{-22}\Delta N_e + 8.5 \times 10^{-18}(\Delta N_h)^{0.8}\right], \quad (1\text{-}1)$$

$$\Delta \alpha = 8.5 \times 10^{-18}\Delta N_e + 6.0 \times 10^{-18}\Delta N_h, \quad (1\text{-}2)$$

where $\Delta n$ and $\Delta \alpha$ are the changes in the real and imaginary parts of the material index, respectively; $\Delta N_e$ and $\Delta N_h$ are the changes in electron and hole concentrations in silicon, respectively. The EO effect on silicon index can realize sub-nanosecond speed phase shifter designs with p-n or p-i-n junctions that are integral to high-bandwidth optical modulators [16] [17] [18] and fast switching devices [19] [20].

*Highly compact photonic structures*

The silicon photonics platform achieves waveguiding with silicon (n = 3.5) as the core and silicon oxide (n = 1.5) as the cladding. The large index contrast between the core and the cladding results in tight confinement of the guided mode ($n_{eff}$ = 2.35 @ 1.55 µm, 450 nm waveguide width), allowing tight bending structures with <10 µm bend radius and overall reduced footprint. This ensures high integration density in PIC designs and compact chip dimensions. The miniaturized index-tuneable structures also result in low tuning power consumption. In addition, the transparency window of silicon photonic platform extends from 1.1 µm to 3.7 µm, bounded by the band edge of silicon at the lower end and the mid-infrared absorption of silicon oxide at the higher end. This makes 1 – 2 dB/cm waveguide propagation loss possible under industrial fabrication processes.

## 1.3    Photonic Switching

### 1.3.1    Historical Photonic Switch Designs

Photonic switching is a concept that entails on-demand connectivity change between a set of optical inputs and a set of optical outputs directly in the optical domain without conversion to electrical signals. To that end, there has been many commercial and research implementations to achieve viable photonic switch fabric designs. These historical efforts can be categorized based on their switching mechanisms:

*Free space switching*

Optical switching in free space typically relies on MEMS-actuated mirrors [21] [22] [23] or piezo-electric collimators [24] to deflect or steer collimated optical beams from an input fiber port to an output fiber port. Free space optical switches can offer <1 dB switching loss at high port count, and are insensitive to polarization and wavelength. However, due to the mechanical movements in the switching process, the switching speed of these systems are typically on the millisecond scale. In addition, the deflection and steering components need stringent alignment and calibration to ensure proper function, which results in bulky and expensive switching systems.

*Gain based switching*

Optical switching can be achieved by broadcasting the input signal and selective amplify or attenuate at the desired output port. SOA-based switch gates can be implemented for such switching mechanism [25] [26] [27]. Because gain is applied to the switched signal, a commonly adopted FoM is IPDR, which is defined as the range of input power levels within which error-free data transmission can be achieved. The SOAs can effectively compensate for the insertion loss of the switch device in gain mode, and eliminate crosstalk due to leakage power in attenuation mode. However, the high power consumption of the SOA gates prohibits large integrated device with high port counts. Later devices tend to utilize hybrid

designs that combine interferometric non-gain switching sections and selective placements of SOA gates to improve the switching extinction and crosstalk [28] [29] [30] [31].

*Phase based switching*

Steering of the optical signals can also be achieved by manipulating the phase of the optical field in an interferometer structure. Early demonstrations implemented in lithium niobate [32] [33] and silica PLC [34] [33] [35] have shown impressive results in achieving low loss, high extinction switching in integrated devices. In a different implementation, spatially and spectrally dispersed optical signals can be steered adeptly by phase changes actuated due to the birefringence in an LCoS array, enabling space-and-wavelength selective switching [36] [37]. LCoS-based switching has become a key technology in commercial WSS systems because of its high performance. Lithium niobate, silica, and LCoS switch devices tend to require bulky systems due to large component footprint, which can be adeptly addressed with a high index contrast material system such as the SOI platform.

### 1.3.2 Silicon Photonic Switch Designs

Because of silicon photonics' unique benefits in foundry manufacturability, efficient tuning mechanisms, and compact device footprints, switch fabrics designed with silicon PICs are viable options to address many of the issues present in previous photonic switching demonstrations. CMOS-compatible processes lower the manufacturing cost of silicon switch fabrics. Highly efficient index tuning mechanisms enable flexible and low power controls to achieve high switching performance. Compact device footprint allows more structures and components to be integrated in the same chip while reducing packaging cost and system scale. Silicon integrated photonic switches are typically built by a fabric of 2×2 SEs. Impressive demonstrations of silicon switch fabrics have been shown employing SEs realized with MZIs [19] [38] [39], MEMS couplers [40] [41] [42], MRRs [43] [44] [45] [46] [47] [48], and MRR-assisted MZIs [49] [50] [20] [51]. Figure 1-6 shows notable examples of each of these SE implementations.
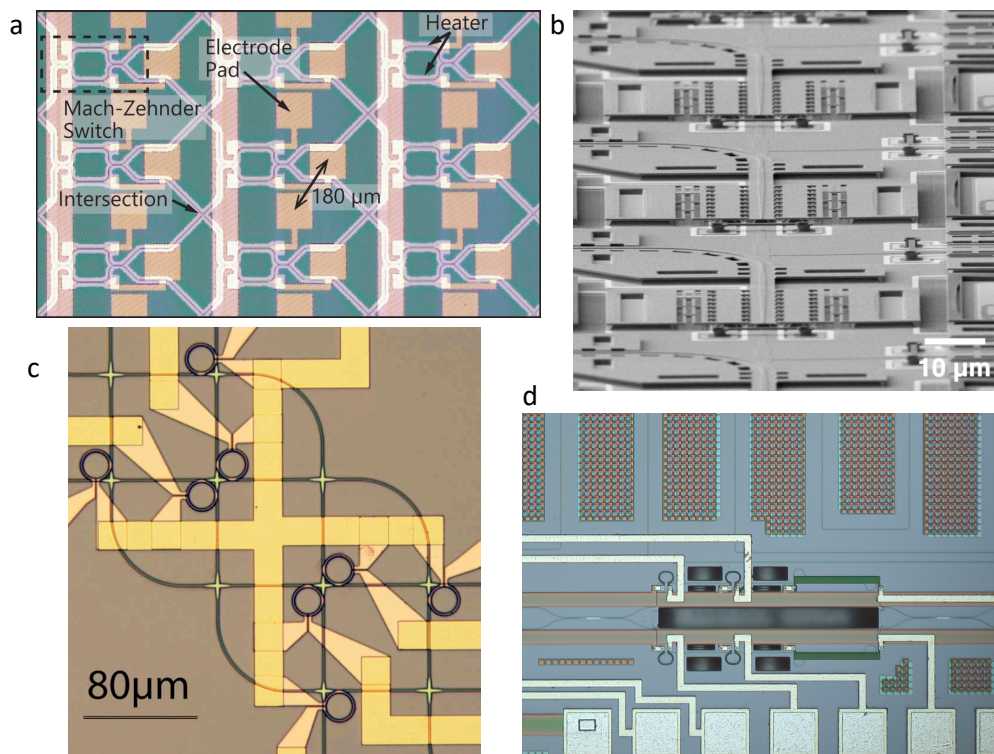


**Figure 1-6: Micrographs of silicon photonic SEs implemented with (a) MZI [39], (b) MEMS-actuated DCs [52], (c) MRRs [43], and (d) MRR-assisted MZIs [51].**

10

In comparison, MZI- and MEMS-based SEs can offer very high extinction and broadband switching, but they rely on couplers or phase shifters hundreds of microns in length, and the switch die area can reach about $50 - 150$ mm$^2$ [19] [38] [40]. In one demonstration of a high port count MEMS-based switch, the device exceeds the maximum reticle size for lithography and requires die-stitching [52]. In contrast, MRR-based SEs are usually tens of microns across and enable much higher integration density with greatly shrunk chip area. Leveraging traveling wave cavity dynamics, MRR SEs can perform narrowband through or drop operations around their resonance points [53]. Figure 1-7 shows a schematic of the through-drop configuration of the MRR SE, where the input optical field, when resonant with the MRR, can couple into the cavity, couple out into the top waveguide, and exit from the drop port. If the optical field is off resonant with the MRR, it bypasses the cavity and exits out of the through port. Figure 1-7 also identifies the self-coupling coefficients, $r_1$ and $r_2$, for the input and output waveguides respectively, as well as the roundtrip amplitude transmission, $a$, which describes the ratio of the preserved field amplitude after one roundtrip in the cavity. The spectral transfer functions for the field at the through and drop ports are described by Equations (1-3) and (1-4), respectively.

$$T_{through} = \frac{r_2{}^2 a^2 - 2r_1 r_2 a \cos\phi + r_1{}^2}{1 - 2r_1 r_2 a \cos\phi + (r_1 r_2 a)^2}, \qquad (1\text{-}3)$$

$$T_{drop} = \frac{(1 - r_1{}^2)(1 - r_2{}^2)a}{1 - 2r_1 r_2 a \cos\phi + (r_1 r_2 a)^2}, \qquad (1\text{-}4)$$

where $\phi$ is the phase detuning in radians from the center of the MRR's resonance and a full FSR is represented by $2\pi$.
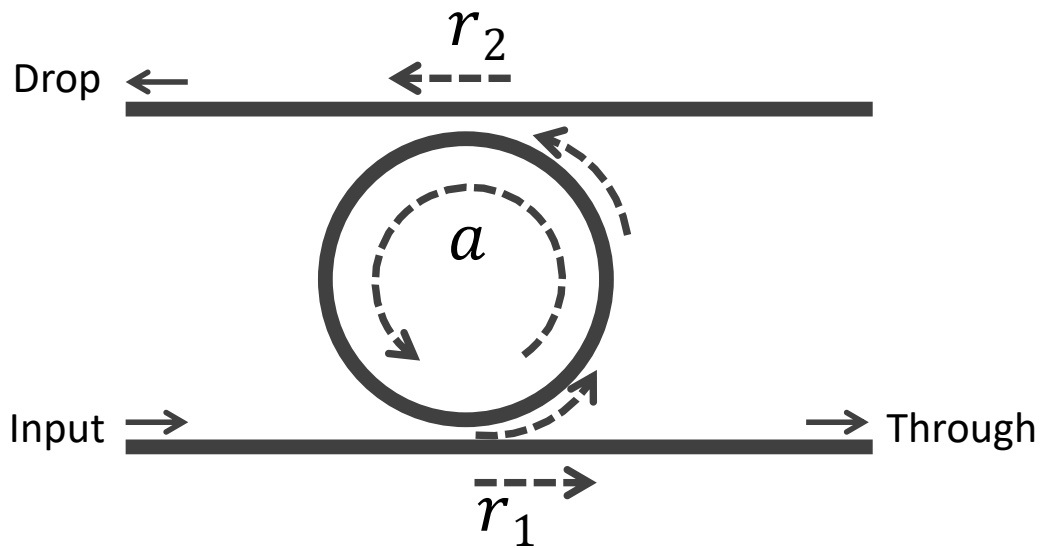
**Figure 1-7: Schematic of an MRR SE in the through-drop configuration.**
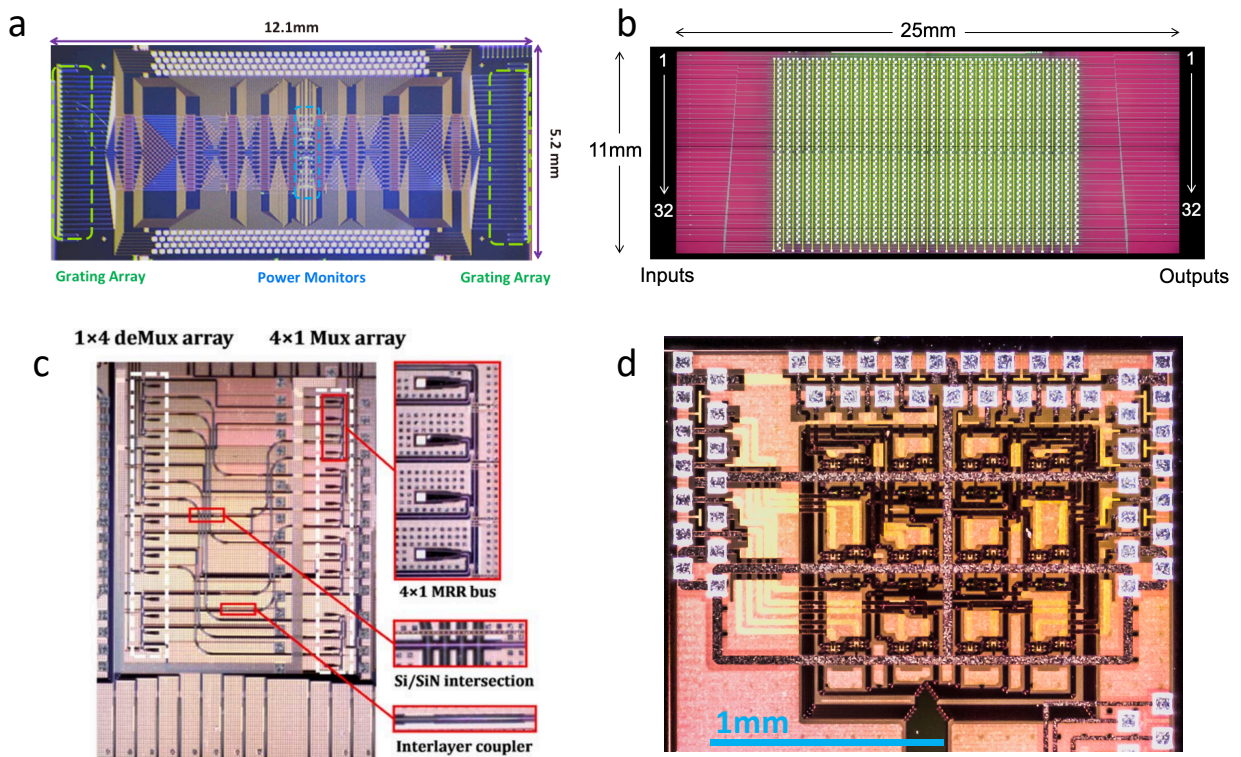


**Figure 1-8: Notable silicon photonic switch fabric demonstrations with (a) 32x32 MZI-based Beneš network [54], (b) 32x32 MZI-based PILOSS network [55], (c) 4x4 MRR-based switch-and-select network [44], and (d) 8x8 MRR-based Omega network [56].**

Larger switch networks are formed by arranging and interconnecting numerous 2×2 SEs in specific topologies to connect more inputs and outputs, and the network architecture dictates the type of connectivity and the routing control of the switch device. Figure 1-8 shows some notable demonstrations of silicon switch fabrics constructed with MZI- and MRR-based SEs with different network topologies. The wavelength-selective nature of the MRR SEs requires wavelength alignment across the switching circuit, which is typically achieved through schemes of fast and efficient wavelength locking of the MRRs as demonstrated in [57] [58] [59]. Higher-order MRR elements enable broadened passband, which relaxes the wavelength alignment requirement for the input optical signal [45], but is at the cost of higher insertion loss and fabrication complexity. Table 1-1 summarises a number of notable demonstrations of MRR-based photonic switch fabrics, with port counts ranging from 2 to 8 [45] [43] [60] [61] [62] [46] [63] [44] [48]. Key metrics, such as port count, on-chip loss, crosstalk, and optical power penalty are highlighted and compared.

**Table 1-1: Notable demonstrations of MRR-based silicon photonic switch fabrics**

| Port Count | Architecture | SE Type | On-chip loss [dB] | Crosstalk [dB] | Optical power penalty [dB] | Year and work |
|---|---|---|---|---|---|---|
| 4×4 | Hitless router | 1st-order MRR | - | >20 | - | 2008 [43] |
| 5×5 | Bidirectional optical router | 1st-order MRR | ~8 | >16 | < 1.75 @ 12.5 Gb/s | 2011 [61] |
| 8×7 | Crossbar | 5th-order MRR | 2 to 10 | 19.5 to 23.4 | < 1 @ 10 & 40 Gb/s | 2014 [45] |
| 8×8 | N-stage planar | 1st-order dual MRR | ~4 | <-10 | - | 2017 [62] |
| 8×8 | Crossbar | 2nd-order MRR | >5 | >-20 | - | 2017 [46] |
| 4×4 | Benes | 2nd-order MRR | ≤6.9 | <-13.6 | - | 2018 [63] |
| 4×4 | Switch-and-select | 1st-order MRR | >1.8 | -51.4 to -31.6 | -1 @ 12.5 Gb/s | 2019 [44] |
| 8×4 | Crossbar | 2nd-order MRR | 6 to 14 | <-32 | 0.2 @ 40Gb/s | 2019 [48] |
| 8×8 | Omega | 1st-order dual MRR | 4.4 to 8.4 | -18.8 to -14.7 | < 2 @ 32 Gb/s | 2019 [56] |

Another type of silicon photonic SEs is the MRR-assisted MZI, whose schematic is shown in Figure 1-9. This type of SEs resembles an MZI with its 4-port configuration and interferometer arms, but can be much smaller in footprint by replacing the linear waveguide phase shifters with compact, over-coupled MRRs. The sharp, narrowband phase delay induced by the MRR phase shifters across their resonances can be described by:

$$\varphi = \pi + \phi + \arctan\frac{r\sin\phi}{a - r\cos\phi} + \arctan\frac{ra\sin\phi}{1 - ra\cos\phi}, \qquad (\ 1\text{-}5\ )$$

where $\phi$ is the phase detuning in radians from the center of the MRR's resonance; $r$ is the self-coupling coefficient between the MRR and the single waveguide; $a$ is the roundtrip amplitude transmission of the MRR. The MRR-assisted MZI inherits many of the benefits from both MRR- and MZI-based SEs and will be explored in more details and implementations in Chapter 2.
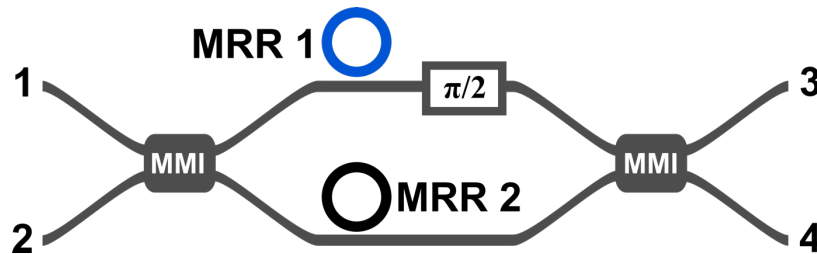


**Figure 1-9: Schematic of an MRR-assisted MZI SE.**

## 1.4    Scope of Thesis

The scope of this thesis focuses on a comprehensive exploration on silicon photonic switching components and devices, calibration and control techniques, topology and routing strategies, as well as intelligent control plane design. The chapters of this thesis are derived from the author's works published in peer reviewed journals and research conference proceedings.

This work adopts a bottom-up approach to structure the topics on silicon photonic switching devices and systems. Chapter 2 focuses on the building block components of a switch fabric and presents a novel space-and-wavelength selective SE using MRR-assisted MZI structures with high performance and design robustness. Chapter 3 explores integrated silicon photonic switch devices based on MRR SEs. The design and characterization of an Omega 8×8 switch device with dual-MRR SEs is presented with a well-balanced set of performance metrics in extinction ratio, crosstalk suppression, and optical bandwidth. For further increase in the switch port count, a Clos network of MRR switch-and-select sub-switches is proposed and analyzed to show optimal scalability and management of optical impairments. Chapter 4 discusses an approach for fast and efficient calibration techniques to address fabrication variations in switch devices and achieve precise and performance-optimal control points. Chapter 5 presents the analysis on the redundancies in switch configurations and derive routing strategies based on the redundancies to improve worst-case performance of the switch device. Chapter 6 examines how to best design ML-enabled control plane for optical systems to predict system dynamics and offer provisioning recommendations. As a case study, the ML approach introduced is applied to EDFA power excursion impairments and shows high efficacy in improving the spectral power stability during wavelength assignment and defragmentation.

# Chapter 2: Switching Elements for Spatial and Spectral Domains

## 2.1    Introduction

The capability and performance of a photonic switch system heavily depends on the design of the SEs, which are the fundamental building blocks of the switch fabric. Routing signal channels in both the spatial and spectral domains is one of the key functionalities to realize agile photonic switching in WDM applications. To perform space-and-wavelength selective switching in devices built with broadband DCs and MZIs, WDM channels need to be first distinguished spatially before being routed. Prior works [64] [65] [66] [67] of space-and-wavelength switch designs have introduced parallel switching planes bookended with wavelength (de)multiplexers, a generic schematic of which is shown in Figure 2-1a. In this design, the duplication of switching planes poses immense challenges to managing the complexity and footprint of the integrated systems. In contrast, space-and-wavelength selective switching can be significantly simplified by integrating independent controls for both spatial and spectral domains at the SE level, as shown in Figure 2-1b, with wavelength selective structures such as MRR [47], photonic crystal nanobeams [68], and waveguide Bragg gratings [69]. In particular, MRR-based SEs offer promising characteristics of large switching ER, high tuning efficiency, ultra-compact footprint (~100 $\mu m^2$), and commercial manufacturability, which are key factors for scaling to high-radix switch fabrics [47]. Previous works [48] [70] for MRR-based space-and-wavelength selective SEs have employed multiple MRRs with offset resonances as spectral add-drop filters. Achieving decent switching ER with add-drop MRRs requires careful examination of the critical-coupling conditions, which, however, are susceptible to fabrication variations. In contrast, over-coupled MRRs as phase shifters tend to have relaxed design constraints, and thus improved tolerance to variations.
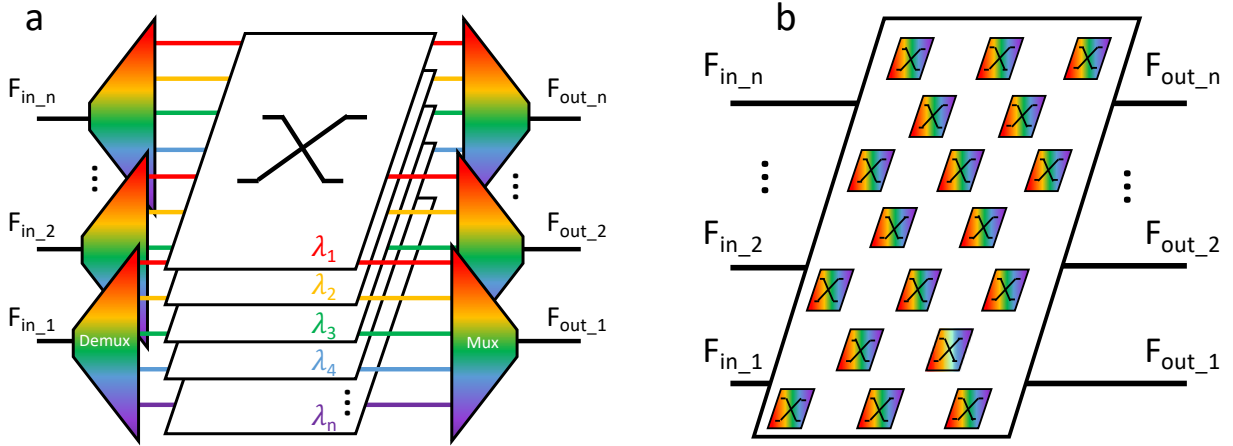
**Figure 2-1: (a) Schematic design for space-and-wavelength selective switch with parallel switching planes and wavelength (de)multiplexers. (b) Schematic design for space-and-wavelength selective switch built on SEs with spatial and spectral switching granularity.**

This chapter discusses a novel SE design that was first proposed in [51] with full space-and-wavelength switching using a symmetrical MZI assisted by pairs of over-coupled MRRs. The MRRs are operated as highly efficient narrowband phase shifters to enable independent switching of multiple wavelengths in the MZI structure. As a demonstration, we design and characterize a $2\times2\times2\lambda$ SE device, which shows high ER, highly suppressed crosstalk, and low signal penalty with push-pull control scheme when WDM signals are switched simultaneously.

18

## 2.2    Working Principles

Figure 2-2 shows a generic design of the proposed space-and-wavelength selective SE, independently switching M wavelength channels in an MZI structure with M-pairs of MRRs and a broadband $\pi/2$ phase difference between the arms. The SE maintains a compact footprint by using over-coupled MRRs as efficient phase shifters, which induce a sharp and continuous $2\pi$ phase change across their resonances. Each pair of identical MRRs aligns to a specific wavelength channel and operates differentially to create a symmetric MZI passband. By maintaining a 2-input and 2-output spatial configuration, the SE is a compatible building block for any multi-stage switch topologies based on $2\times2$ elementary cells [71] and can be interconnected to scale to larger $N\times N\times M\lambda$ fabrics. Since the MRRs have periodic resonances, the total number of channels supported in a single MZI structure, M, is limited by:

$$M \leq \left\lfloor \frac{\text{FSR}}{\Delta f + \text{FWHM}} \right\rfloor , \qquad (\,2\text{-}1\,)$$

where FSR and FWHM are respectively the free-spectral range and full width at half maximum transmission of the MRRs, and $\Delta f$ is the detuning between the pair of MRRs for a single channel.
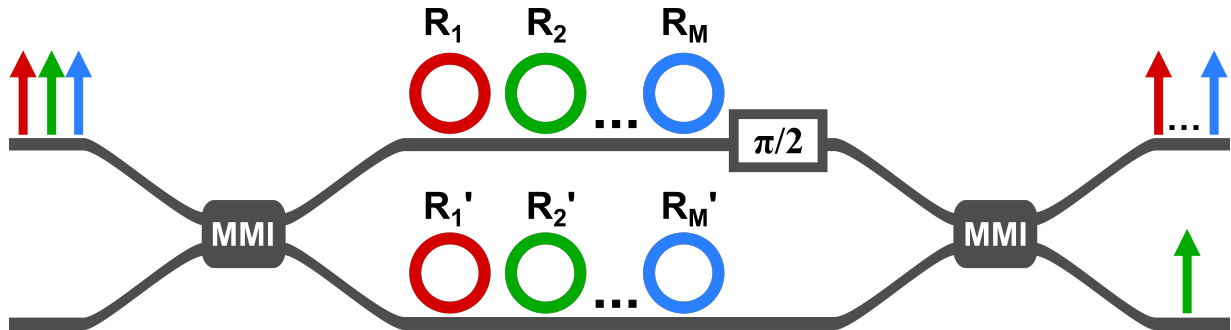


**Figure 2-2: Schematic of a 2×2×Mλ SE with MRR phase shifters. Each pair of MRRs aligns to the input wavelength channel of the same color.**

To illustrate the working principles of the MRR-assisted switching, we examine a single-channel wavelength-selective MZI with one pair of MRR phase shifters. This structure, as shown in Figure 2-3a, has been implemented in designs for efficient spatial SEs [49], modulators [72] [73], and wavelength interleavers [74]. With one MRR coupled to each arm, the MZI switching phase is achieved by slight

detuning of the MRRs' resonances. Figure 2-3b and Figure 2-3d visualize the phase delay induced by the MRRs, as well as the resultant phase differences between the MZI arms in the Bar and Cross switching states, respectively. The spectral transmission of the MZI outputs can be determined using the transfer matrix method:

$$\begin{bmatrix} E_{O1} \\ E_{O2} \end{bmatrix} = \begin{bmatrix} t & jr \\ jr & t \end{bmatrix} \begin{bmatrix} \alpha_{r1}e^{j(\varphi_{r1}+\phi_{bias}+\phi_{path})} & 0 \\ 0 & \alpha_{r2}e^{j(\varphi_{r2}+\phi_{path})} \end{bmatrix} \begin{bmatrix} t & jr \\ jr & t \end{bmatrix} \begin{bmatrix} E_{I1} \\ E_{I1} \end{bmatrix}, \qquad (2\text{-}2)$$

where $E_{I1}$, $E_{I2}$, $E_{O1}$, and $E_{O2}$ are the electric field amplitudes at Input Ports 1 and 2 and Output Ports 3 and 4, respectively; $t$ and $r$ are the field transmission and cross-coupling coefficients of the MMI structures; $\alpha_{r1}$ and $\alpha_{r2}$ are the field amplitude modulation factor by MRR1 and MRR2, respectively; $\varphi_{r1}$ and $\varphi_{r2}$ are the phase delay imposed by MRR1 and MRR2, respectively; $\phi_{bias}$ is the static phase bias difference between the two MZI arms; and $\phi_{path}$ is the phase accumulated in the straight waveguides, which are assumed to be equal between the two arms. Note that $\alpha_{r1}$, $\alpha_{r2}$, $\varphi_{r1}$, and $\varphi_{r2}$ are dependent on the detuning of the wavelength about the MRRs' resonances. In this work, we employ the push-pull control scheme for the pair of MRRs around the MZI's quadrature point – the resonances of the MRRs are driven in opposite directions about the switched channel, and a static and wideband phase bias of $\pi/2$ is set between the arms. Under this control scheme, both the Bar and Cross switching states have identical transmission spectra, as determined by Equation (2-2) and shown in Figure 2-3c and Figure 2-3e, and thus eliminating state-dependent performance variations. To route the wavelength channel from Port 1 to Port 3 (Figure 2-3b and Figure 2-3c), MRR 1 is red-shifted from the channel wavelength at zero detuning, while MRR 2 is blue-shifted. To route the wavelength channel from Port 1 to Port 4 (Figure 2-3d and Figure 2-3e), MRR 1 is blue-shifted from the channel, while MRR 2 is red-shifted.
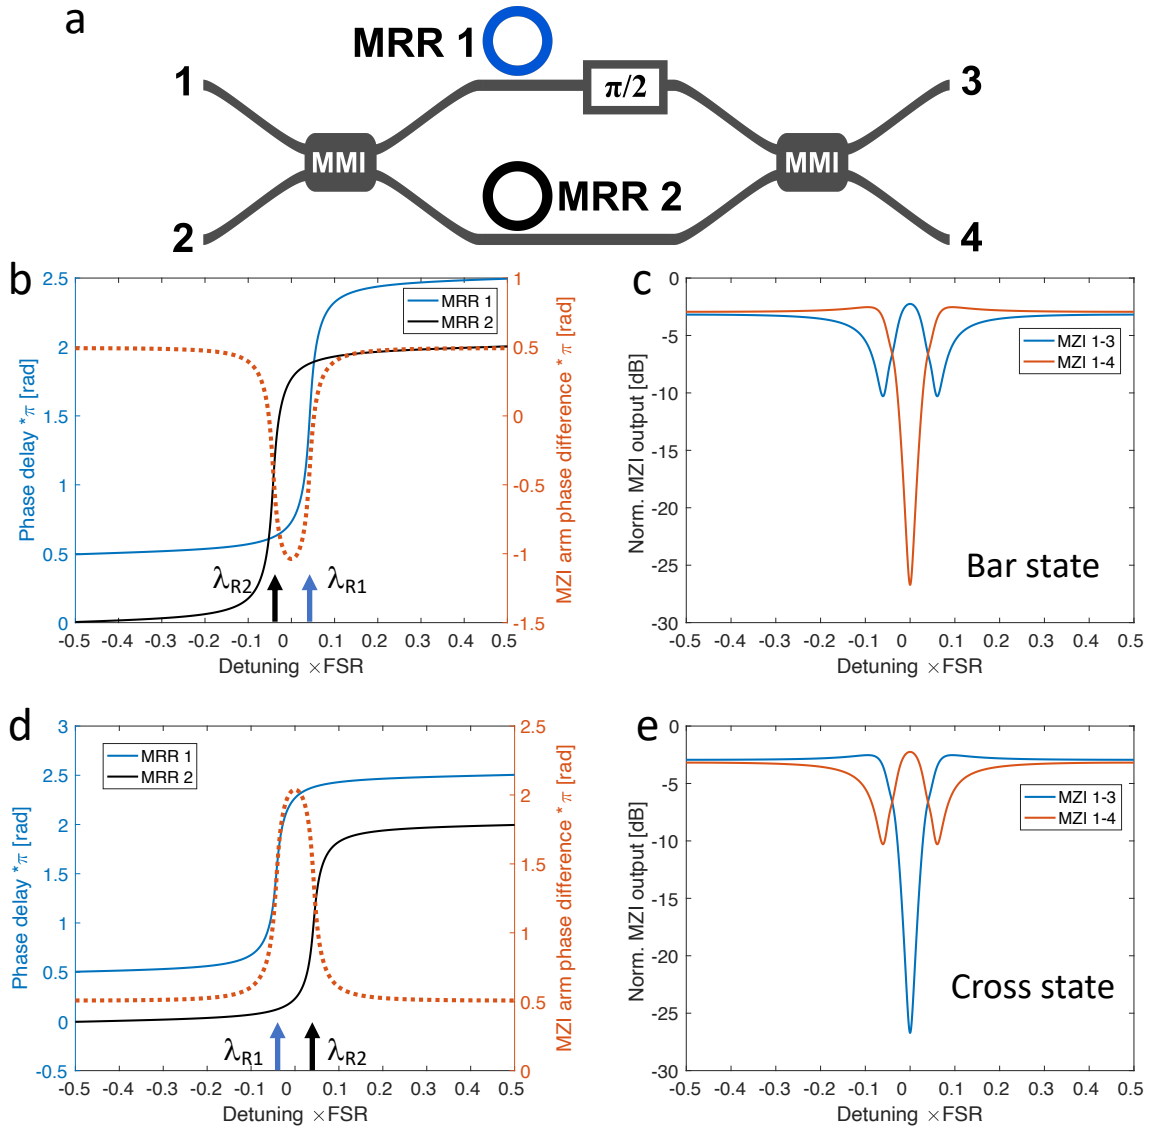
**Figure 2-3: (a) Schematic of the single channel MRR-assisted MZI. (b) Induced phase by MRR 1 and MRR 2 and the MZI arms' phase difference for Bar switching state, with center resonances of MRR 1 and MRR 2 labeled as $\lambda_{R1}$ and $\lambda_{R2}$, respectively. (c) Bar-state MZI output spectra for paths Port 1 to Port 3 and Port 1 to Port 4. (d) Induced phase by MRR 1 and MRR 2 and the MZI arms' phase difference for Cross switching state. (e) Cross-state MZI output spectra for paths Port 1 to Port 3 and Port 1 to Port 4.**

## 2.3 Device Design and Characterization

To demonstrate independent switching of multiple wavelength channels with the proposed switch design, we design and fabricate a $2\times2\times2\lambda$ SE device that incorporates two pairs of MRRs for simultaneous switching of two wavelength channels. Figure 2-4a shows the micrograph of the device, integrating 2 MMIs, 4 MRRs, and 2 MZI TO tuners under a footprint of 0.17 mm$^2$. The MRRs have a racetrack shape with an 8-μm bending radius and a 5-μm straight coupling section. The gap between the MRR and bus waveguide is 100 nm, imposing strong over-coupling conditions. Thermal isolation trenches are placed between the two MZI arms to reduce the thermal crosstalk between the phase shifters. The MZI tuners and the MRR phase shifters are implemented with TiN heaters to induce localized index change to the silicon waveguides without incurring additional IL. The switch chip is designed and fabricated in an MPW run through a commercial 200 mm SOI platform offered by Advanced Micro Foundry.
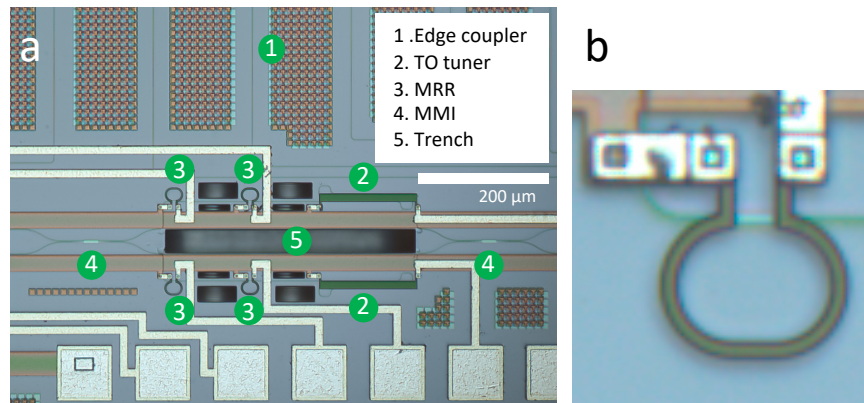


**Figure 2-4: (a) Micrograph of the $2\times2\times2\lambda$ SE. (b) MRR heater structure.**

The fabricated device is accessed optically with TE-polarized light via edge coupled fiber arrays and electrically via probes. The coupling loss is measured to be ~3 dB per facet with lensed fibers. The TO phase shifter performances are shown in Figure 2-5a and Figure 2-5b. The MZI tuner shows a tuning efficiency of 13.6 mW/π and results in an MZI ER over 32 dB. The MRR shifter, shown in Figure 2-4b, has a tuning efficiency of 0.262 nm/mW and is used to both align the resonance point and perform

switching. Figure 2-5c shows the transmission spectra of the MZI when a single pair of MRRs', R1 and R1', are far detuned. We measure the FSR of the MRRs to be 1.182 THz, and the 3dB-bandwidth to be 101 GHz, giving a finesse of 11.7. The switching speed of the switch is measured by applying a 2 V DC bias and a 15 kHz electrical square wave with 175 mV peak-to-peak amplitude to the MRR TO shifter. Observing the power levels of the resonant channel, the 0–90% optical rise and fall times of the MRR output are measured as 13 μs and 12 μs, respectively, as shown in Figure 2-5d.
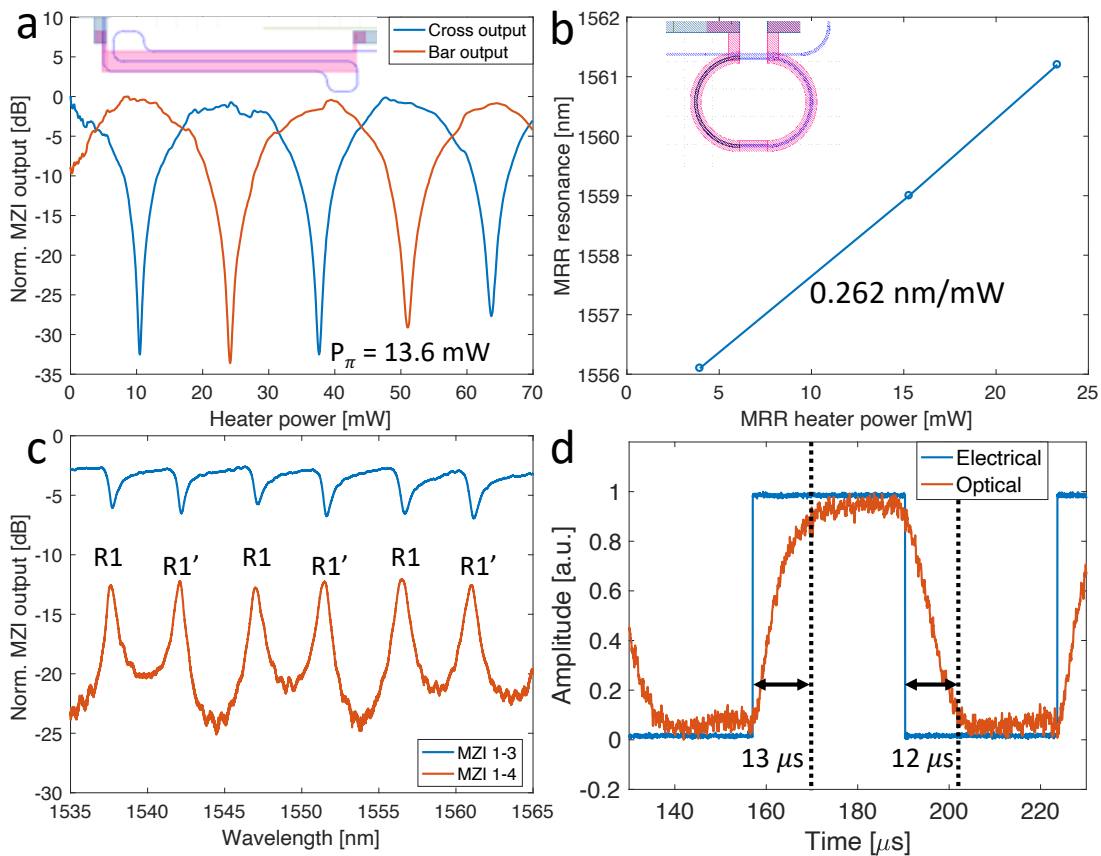


**Figure 2-5: (a) MZI switching output as a function of TO tuner power consumption; the inset shows the TO tuner design. (b) Resonance tuning power with the TO MRR shifter. (c) Offset MRR resonances while the MZI is biased to Bar state. (d) Optical rise and fall times when the MRR TO shifter is driven with a 15 kHz electrical square wave signal.**

For switching characterization of the 2×2×2λ SE, the TO tuner on the top MZI arm is set to a phase bias of π/2. Two CW laser signals at 1537.3 nm and 1541.3 nm, as Channel 1 and Channel 2 respectively,

are combined before inputting to Port 1 of the SE. We set the resonances of the first MRR pair, R1 and R1',

at Channel 1, and the resonances of the second MRR pair, R2 and R2', at Channel 2. In contrast to spatial

2×2 SE with binary states of Bar and Cross, the space-and-wavelength selective SE supports four switching

states - (Bar, Bar), (Bar, Cross), (Cross, Bar), and (Cross, Cross) – for (Channel 1, Channel 2) respectively,

as illustrated in Figure 2-6a-Figure 2-6d. Each pair of MRRs operates in push-pull for its corresponding

channel. To switch a channel to Bar state, the MRR corresponding to that channel on the top arm (R1 or

R2) red shifts its resonance with a slight increase of bias on its TO shifter, while the MRR on the bottom

arm (R1' or R2') blue shifts with a slight decrease in its TO bias, resulting in a $\pi$ phase difference between

the MZI arms at the channel wavelength. To switch a channel to Cross state, the top arm's MRR blue shifts

while the bottom arm's MRR red shifts, resulting in a $2\pi$ phase difference between the MZI arms.
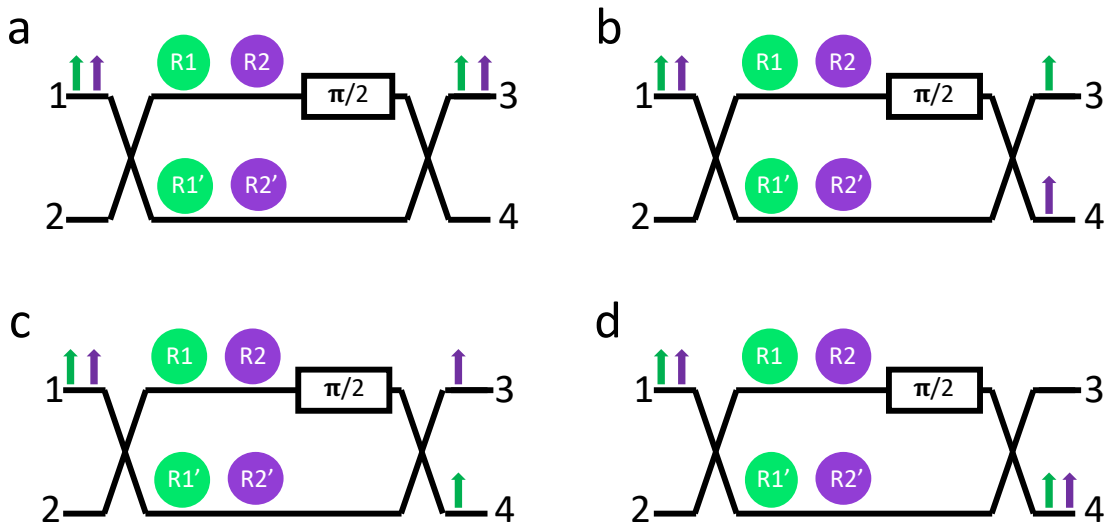


**Figure 2-6: (a)-(d) Illustrations of (Bar, Bar), (Bar, Cross), (Cross, Bar), and (Cross, Cross) states for the two independently switched channels.**

Switch outputs at Port 3 and Port 4 are monitored by an optical spectrum analyzer for the two

wavelengths under operation. Figure 2-7a – Figure 2-7d show the output channel power levels for Path Port

1 – Port 3 and Path Port 1 – Port 4 in the four switching states. We observe an average crosstalk suppression

ratio of 21.7 dB, as indicated in Figure 2-7a – Figure 2-7d. The on-chip loss for the switched signal averages to 5.1 dB, which is due to a combination of MMI loss and intensity modulation by the MRRs in the over-coupling regime. For a specific wavelength channel, we define the switching ER as the difference between the signal power of the channel in one state and the leakage power of the same channel switched to other states. For instance, the ER for the Channel 1 in <u>Bar</u> is determined by the average signal power at 1537.3 nm between States (<u>Bar</u>, Bar) and (<u>Bar</u>, Cross), subtracted by the average leakage power for Channel 1 between States (<u>Cross</u>, Bar) and (<u>Cross</u>, Cross). Between the two channels, we measure an average ER of 21.8 dB for Bar state, and 21.6 dB for Cross state. Figure 2-7e – Figure 2-7h show the output spectra of all switching states measured with a broadband source. The signal passbands of both channels in all switching states are fairly consistent and average to 75.1 GHz.
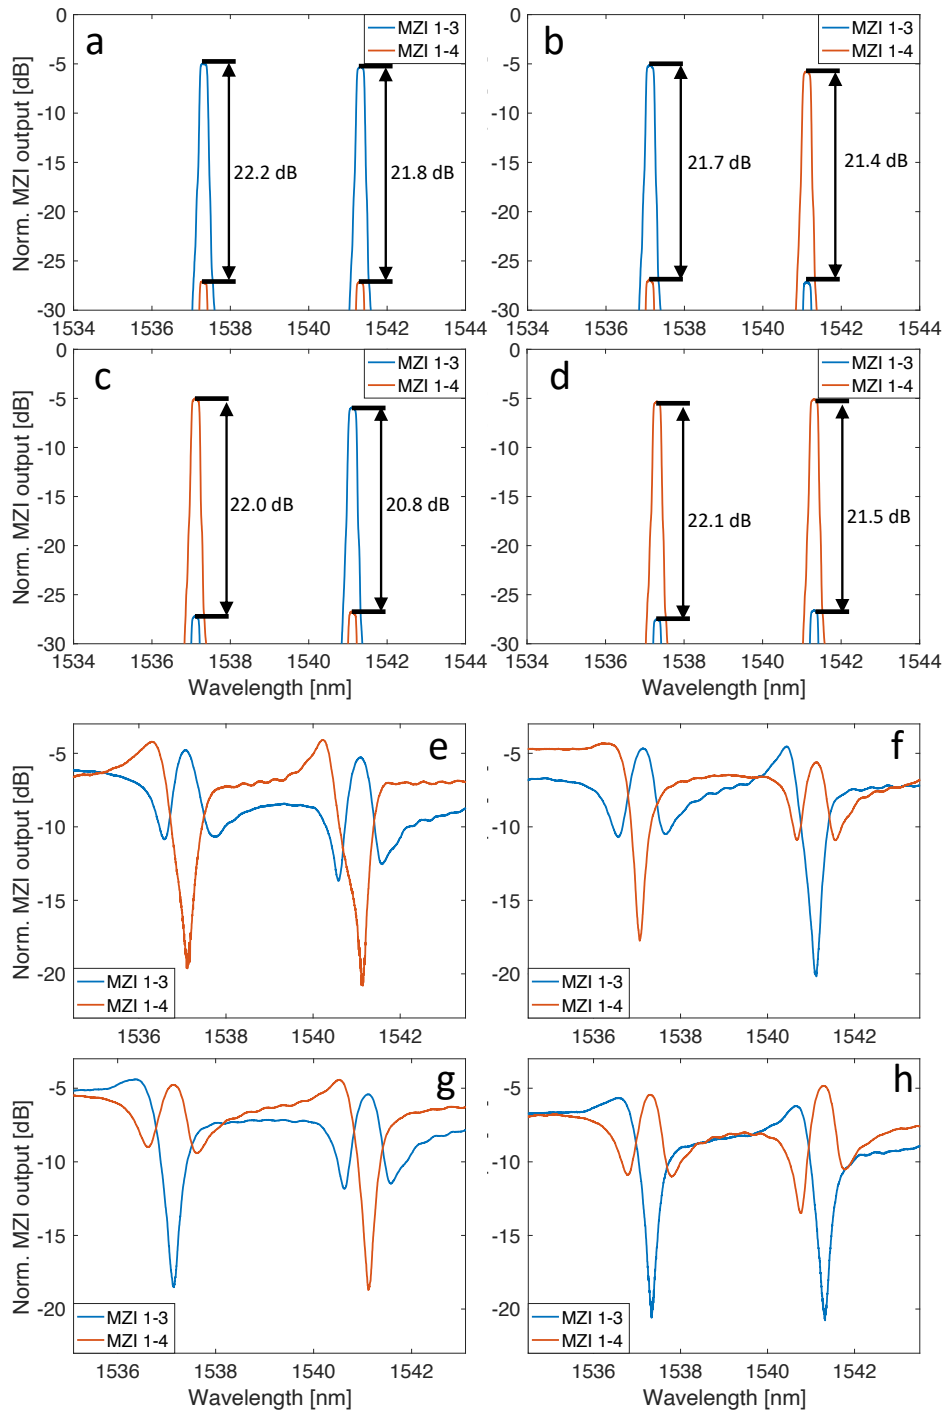
**Figure 2-7: (a)-(d) Signal and leakage power levels with two CW lasers in the four switching states. Crosstalk suppression levels are indicated. (e)-(h) Transmission spectra for both switch paths in the four switching states.**

We numerically simulate the 2×2×2λ SE using the transfer matrix method to optimize for the crosstalk suppression and IL during switching. In this analysis, the two wavelength channels are spaced at $\Delta f$ + FWHM – the closest spacing as defined by Equation (2-1). We maintain the experimentally extracted value for the MRR's finesse, $\mathcal{F}$, which is a function of the product between the MRRs' self-coupling coefficient, $r$, and roundtrip amplitude transmission, $a$:

$$\mathcal{F} = \frac{\text{FSR}}{\text{FWHM}} = \frac{\pi\sqrt{ra}}{1-ra}. \qquad (\ 2\text{-}3\ )$$

We examine the MRR detuning between 10% – 40% of their FSR and $a$ valued between 0.875 – 0.993, with corresponding $r$ values to maintain the same finesse value of 11.7. The sweep in $a$ can inform how to better design the intrinsic loss of the MRRs, while the sweep in MRR detuning can inform more precise push-pull control. From the results shown in Figure 2-8a and Figure 2-8b, the push-pull MRR-assisted switching is capable of achieving >25 dB crosstalk suppression and <1 dB IL by keeping $a$ greater than 0.96 and MRR detuning between 19% and 23% of the FSR, as indicated. The improved IL and crosstalk suppression can further enhance the SE design's scalability in multi-stage switch topologies [71].
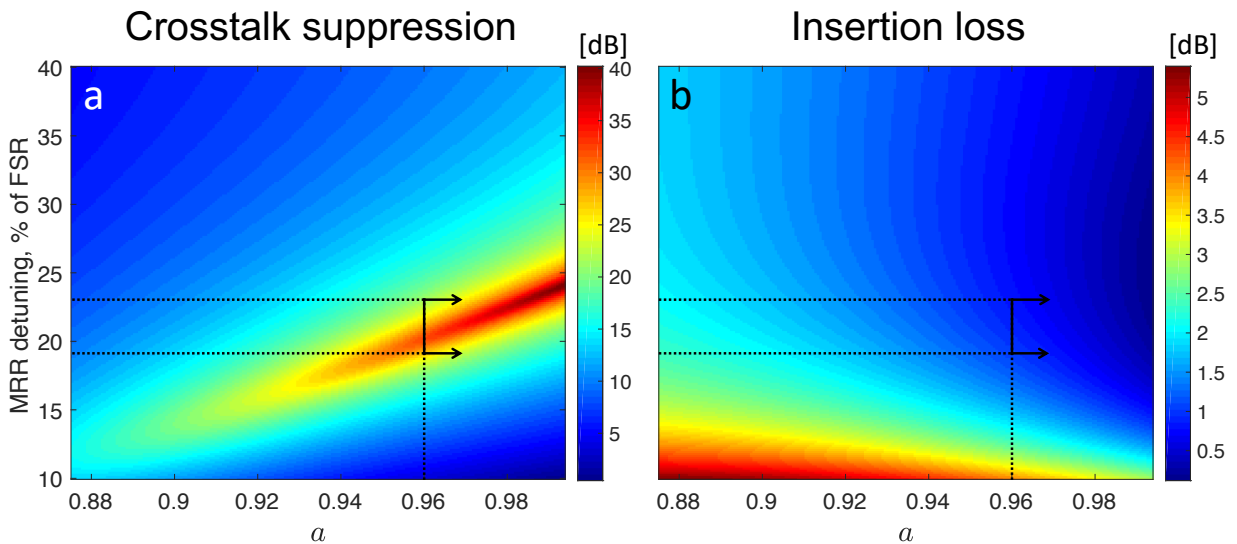


**Figure 2-8: Crosstalk suppression (a) and IL (b) over ranges of MRR detuning and roundtrip amplitude transmission.**

Data transmission was performed for switching states (Bar, Cross) and (Cross, Bar), where the two wavelength channels are launched into the same input port but switched to different output ports. A total of four switched signals are examined: (Bar, Cross)-Channel 1, (Bar, Cross)-Channel 2, (Cross, Bar)-Channel 1, (Cross, Bar)-Channel 2. We examine the switching states with one or both channels transmitting (single- or dual-channel cases) to study the impact of inter-channel crosstalk on the data routing performance. The data test schematic is shown in Figure 2-9a. An Anritsu MP1900A Signal Quality Analyzer generates two electrically decorrelated signals at 32 Gbps NRZ OOK using PRBS31. Two optical carriers at 1537.3 nm and 1541.3 nm are modulated by MZMs with 0 dBm output power and combined using a FC before entering the silicon photonic chip. A PA consisting of a VOA and an EDFA is used before the switch to compensate for coupling and propagation losses through the chip and ensure -10 dBm of optical power exits the chip. This PA is also used to replicate the device IL in the B2B reference case. A second set of EDFA and VOA adjusts the receiver optical power for the BER measurement. A TOF with a passband covering both wavelength channels is used to reject out-of-band ASE noise. The receiver consists of a Finisar XPDV3120 PD-TIA assembly, which performs the optical-to-electrical conversion and allows the data signal to be analyzed by the Anritsu error checker. For each switch path, the crosstalk channel is off for the single-channel case, and on for the dual-channel cases. Evident from Figure 2-9b – Figure 2-9e, all switch paths' dual-channel cases are within 1 dB power penalty at $10^{-9}$ BER compared to their single-channel cases without inter-channel crosstalk, and within 1.5 dB power penalty compared to the B2B reference case.
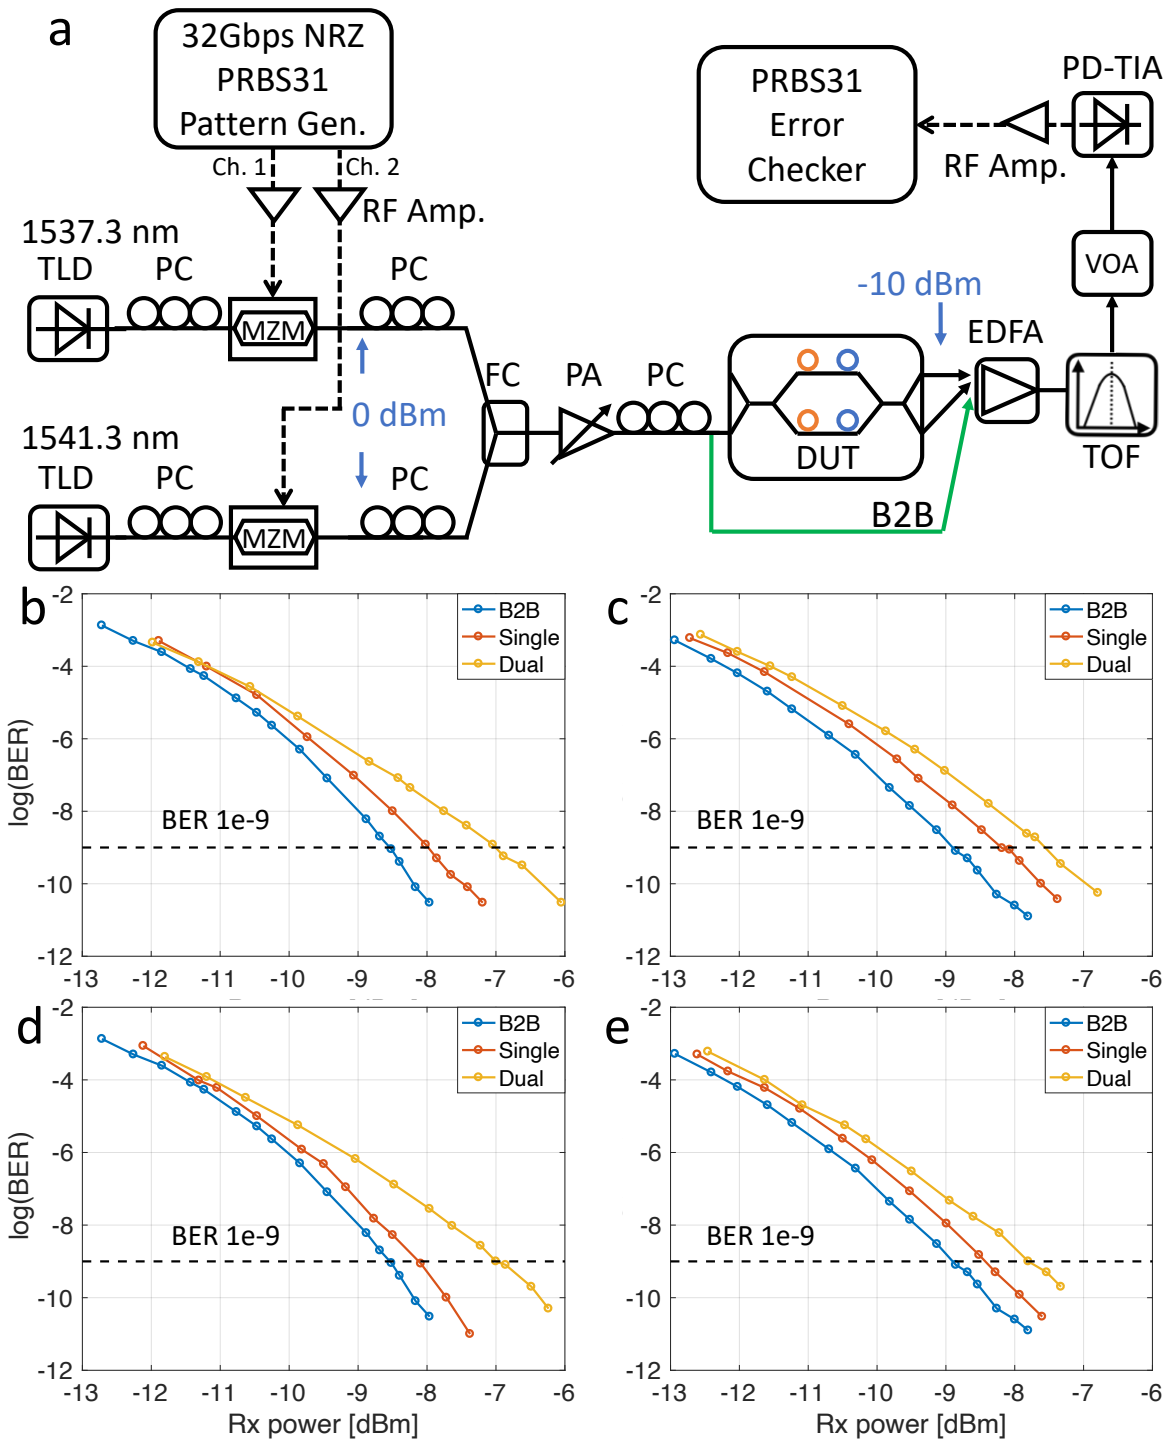
**Figure 2-9: (a)** Data experiment setup schematic. **(b)-(e)** BER curve for Switching States (Bar, Cross)-Channel 1, (Bar, Cross)-Channel 2, (Cross, Bar)-Channel 1, and (Cross, Bar)-Channel 2, respectively, compared among B2B, single- and dual-channel cases.

## 2.4　Chapter Summary

　　　　Space-and-wavelength selective switching can be simplified significantly by integrating spatial and spectral control granularity at the SE level. In this chapter, we introduce an MRR-assisted MZI SE that achieves independent switching of wavelength channels by using multiple pairs of MRRs as push-pull phase shifters. We discuss the working principles of the design and demonstrate a $2\times2\times2\lambda$ switch block experimentally. The SE switches two wavelength channels independently in a total of four switching states, with both crosstalk suppression and switching ER exceeding 21 dB. Less than 1 dB signal power penalty is observed from inter-channel crosstalk when transmitting 32 Gbps NRZ data signals. We further show a path to achieve >25 dB crosstalk suppression and <1 dB IL through optimizing the MRR intrinsic loss and more precise push-pull control. The SE's high performance, compact footprint, and efficient control make this design a promising building block for scaling to multi-stage space-and-wavelength selective optical switch fabrics. In the next chapter, we explore the unique architectures, key performance parameters, and design space and scalability analysis for multi-stage silicon photonic switch fabrics.

# Chapter 3: Multi-stage MRR-based Switch Fabrics

## 3.1 Introduction

By assembling basic SE blocks in an interconnected network, higher port-count switch fabrics can be realized. The MRRs are ideal structures for highly integrated switch devices due to their miniaturized footprint. Leverage traveling wave cavity dynamics, each MRR-based SE is about 50-100 times smaller compared to MZI- or MEMS-based SEs, and thus enabling much higher integration density with significantly shrunk chip area. In this chapter, we showcase the designs for two multi-stage switch architectures using MRR-based SEs that have the potential to enable high performance, high radix photonic switch devices – an 8×8 Omega switch using dual-MRR SEs and a Clos network architecture with switch-and-select sub-switches for scaling to higher port count.

Designing MRR-based multi-stage switch requires co-optimizing the SE performance with the selection of the architectures [47]. The scalability of MRR-based fabric can be limited on three fronts: power penalties due to insertion loss and crosstalk of the resonator SEs in on- and off-resonance states, successive passband narrowing of cascaded MRRs, and control complexity of the large number of resonators. Previous demonstrations of MRR-based optical switches to date achieve a record of 8×8 connectivity [45] [46], and both adopt the crossbar architecture that requires only one SE to be controlled to connect an input to an output. However, the number of SEs in a crossbar architecture scales poorly as $N^2$ for an N×N device, which poses a tremendous challenge in on-chip wiring and packaging complexity. In addition, a lightpath in a crossbar switch can traverse between 1 to $(2N - 1)$ SEs, which means both the worst-case insertion loss and the variations of path-dependent power penalties would grow quickly with increasing port count [45]. In contrast, multi-stage architectures, such as the Beneš and Omega topologies, can provide a balance in the trade-offs between the total number of SEs and the number of cascaded switching stages [9] [75]. Individual MRR SE needs to possess a balanced set of performance metrics that

meet the targets in both insertion loss and switching bandwidth, while the selection of architecture allows ease of control and limits the number of cascaded stages to preserve the end-to-end switch passband. Hence, the Omega architecture is a promising candidate for a modest-scale multi-stage MRR-based switch fabric. This design trades off the non-blocking connectivity for a much-reduced number of switch stages and simplified routing control compared to Beneš design. In the next section, we present an 8×8 demonstration of the Omega switch device designed and fabricated through Elenion's silicon photonic platform [76].

Scaling to switch fabrics with even higher radix, it is crucial to alleviate the impact of the first-order crosstalk, which can significantly deteriorate the signal quality and compromise the link budget. A recent demonstration of a monolithic Si/SiN MRR-based switch circuit [77] shows the switch-and-select topology can effectively cancel the first-order crosstalk. The switch-and-select topology, however, requires an $O(N^2)$ scaling of total number of MRR SEs for an N×N network. In addition, managing the required $N^2 \times N^2$ passive waveguide shuffle is increasingly difficult at high port counts. Instead, we devise an MRR-based Clos switch fabric architecture constructed with switch-and-select sub-switches. This design keeps the number of stages to the modest value of three while significantly reducing the required number of switching elements, in addition to the immunity to first-order crosstalk from the switch-and-select stages. We present performance and scalability analysis of the proposed switch architecture to verify a 16×16 device design. Key performance parameters used in this analysis is based on the 4×4 switch-and-select sub-switches [44], which was designed and fabricated through an AIM Photonics MPW run.

## 3.2 Dual-Microring-Based 8×8 Switch

In this work, we present the design and characterization of the first multi-stage silicon switch with 8×8 connectivity implementing dual add-drop MRRs, leveraging a commercial silicon photonics process and design flow at Elenion Technologies [76]. We highlight the combination of low on-chip loss, wide passband, and high tuning efficiency of the switch device as the key enablers to optically-switched datacenter network designs [75]. The operation of the dual-MRR SEs is discussed and comprehensive switching performance and usability analysis are presented. In the following sections, the switch fabric device and architecture are presented; the design, performance, and operation of the MRRs and dual-MRR SEs are discussed; and the end-to-end performance of the full switch fabric is reported.
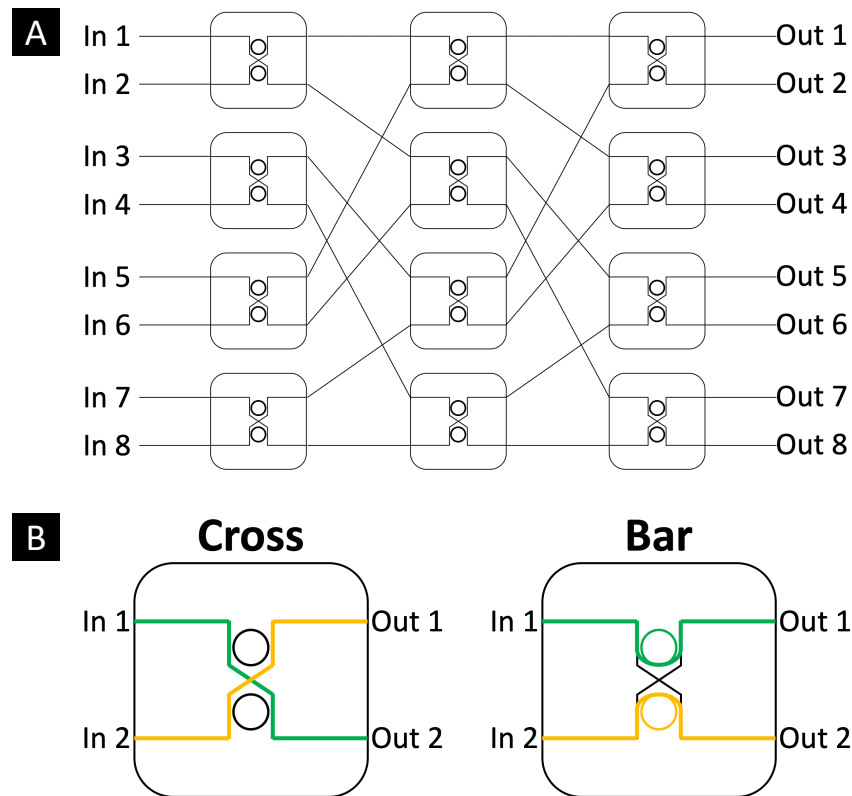


**Figure 3-1: (A) Schematic of the 8×8 Omega network implemented with 12 2×2 SEs. (B) Illustration of the SE's Cross and Bar states and their corresponding dual-MRR configurations.**

### 3.2.1 Switch Architecture and Design

The Omega architecture, as a Banyan-type network originally proposed for high-performance computer networks, is also attractive for high-speed electronic and optical switching applications. A Banyan-type network is defined as a class of multistage networks that have exactly one path from any input port to any output port. Generally, a Banyan switch fabric with N ports is constructed from $\frac{N}{d}(log_d N)$ d×d switching elements arranged in $log_d N$ stages, which is also referred to as d-nary switch [78]. In the optical domain, more attention has been focused on binary switch fabrics (d=2). In particular, the Omega architecture is defined by its perfect shuffle connection of SEs between adjacent stages, which interleaves each half of the previous stage's output ports. For a binary N×N Omega network, with N being a power of 2, the total number of 2×2 SEs is $\frac{N}{2}log_2 N$, and the number of cascaded switching stages is $log_2 N$.

To achieve connectivity between 8 inputs and 8 outputs, the switch arranges 12 SEs into an Omega network illustrated in Figure 3-1A. Each of the 2×2 SEs can be independently controlled as Bar state or Cross state, as shown in Figure 3-1B. The dual-MRR configuration differs from a single-MRR SE [43] [44] by operating two parallel-coupled resonators and achieving a broadened passband through mechanism discussed in Section 3.2.2. Figure 3-2 compares the worst-case insertion loss among three representative optical switch architectures – crossbar, Omega, and Beneš [79]. It is evident that the increment of insertion loss in multi-stage architectures is much slower comparing to the crossbar architecture as port counts increase. This is because the number of bypass MRRs, i.e. off-resonance rings, increases linearly with the switch port count in crossbar network and thus the accumulated through-MRR loss dominates over the drop-MRR loss. Omega network has a lower increment in loss than Beneš because it contains about half of the total stage counts. Reduced number of cascaded stages is also critical to preserving the lightpath passband [47].
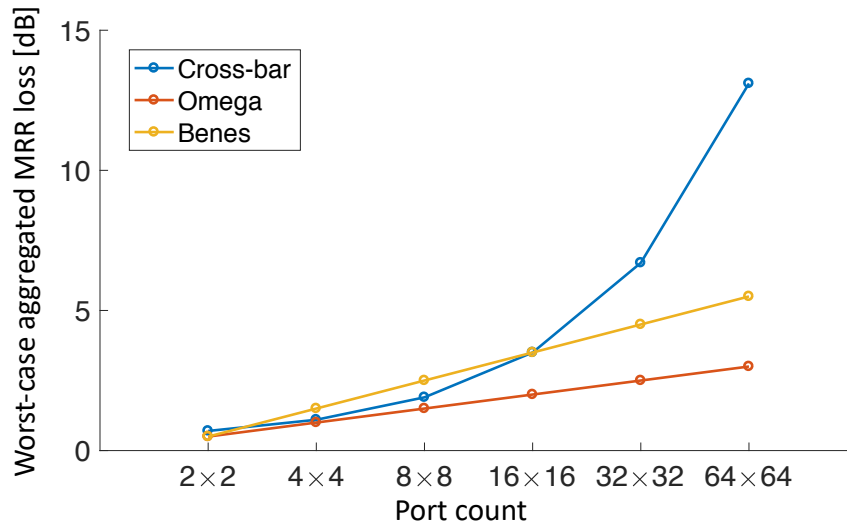
**Figure 3-2: Comparison of the worst path insertion loss on-chip for crossbar, Omega, and Beneš architectures with MRR SEs, based on the data of on- and off-resonance MRR loss values from [44].**

Since Omega networks have exactly one connecting path from any input to any output, they can take advantage of self-routing. Unlike Beneš, whose routing configurations need to be iteratively computed [80] or pre-determined as a look-up table (LUT) [81] [82], the Omega network can be controlled solely with destination-based routing – configuring each SE along the lightpath directly from the output label. Figure 3-3 illustrates an example of such routing procedure. The reduced routing control complexity of Omega network eschews the need for complicated routing logic and the associated control and computation overhead.

Step 1
(1ˢᵗ bit of destination tag)

Out 4(011)

In 7

Step 2
(2ⁿᵈ bit of destination tag)

Out 4(011)

In 7

Step 3
(3ʳᵈ bit of destination tag)

Out 4(011)

In 7
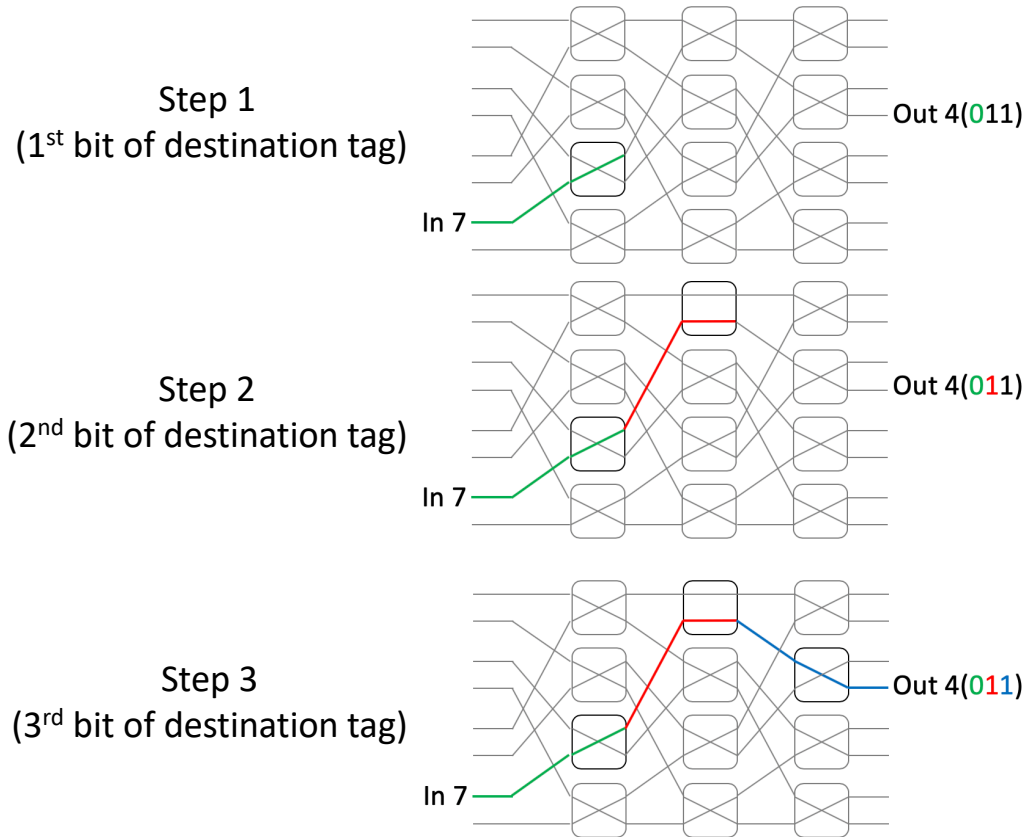
**Figure 3-3: An example of destination-based routing in Omega network. The output's binary label indicates the configuration states of the three SEs along the lightpath – in the sequence of the SEs traversed by the optical signal, the bit 0 selects the upper output of the corresponding SE, and the bit 1 selects the lower output. Hence, a label 011 indicates setting the first SE to Cross, second SE to Bar, and third SE to Cross, to connect a link from Input 7 to Output 4.**

Figure 3-4A shows the micrograph of the switch device. The entire switch chip, consisting of 24 TO MRRs, 8 co-integrated monitor PDs, and 46 electrical bonding pads, has a footprint of 4 mm². Shown in Figure 3-4B, the chip is die- and wire-bonded to a chip carrier placed on a custom PCB to allow electrical access of the TO phase shifters and PDs. An array of 18 fibers at 127-μm pitch are grating coupled to the chip to provide optical access to the switch inputs and outputs.

Electrical pads

A

1mm

16 waveguides

B

Fiber array

PCB

Wirebonds
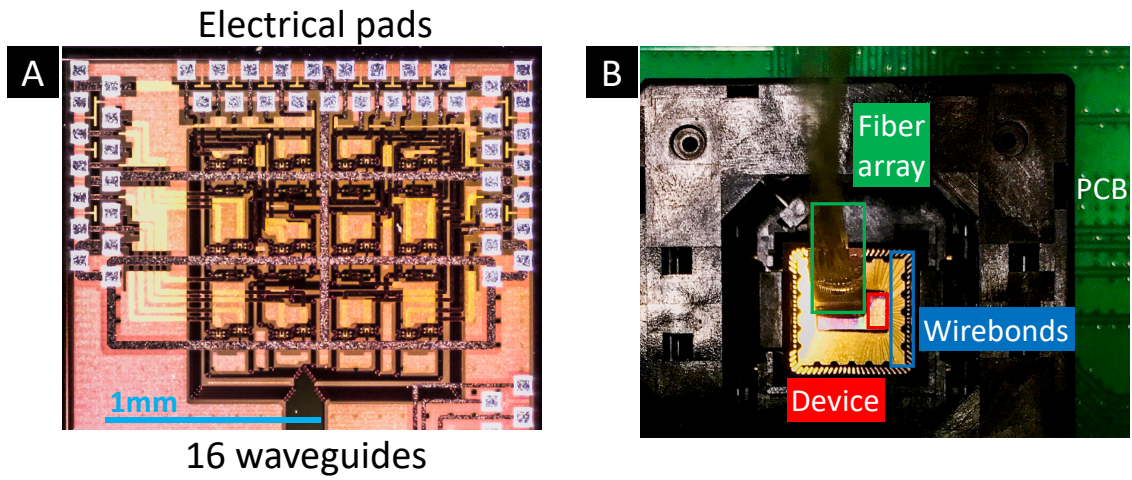
Device

**Figure 3-4: (A) Micrograph of the switch device showing a footprint of 4 mm² including electrical pads. (B) Photo of the switch device package.**

### 3.2.2　Switching Element Design and Characterization

Each of the 12 SEs integrated in the switch device contains 2 racetrack MRRs coupled to 2 parallel waveguides on either side. Figure 3-5A illustrates the arrangement of MRRs and the waveguide crossing in an SE. After each SE, 1% power is tapped on both SE output waveguides into an on-chip PD, which can be used to infer if optical power arrives at the corresponding SE. We characterize the through-MRR transmission spectra for the 8 MRRs in the last stage of the switch and demonstrate highly consistent resonance profile as shown in Figure 3-5B. Designed for operating with 120 GHz of passband, each single MRR shows an extinction ratio of about 9.5 dB. By operating in dual-MRR switching mode, as discussed later in this section, the extinction is extended to about 14.7 dB for Bar state and 18.8 dB in Cross state. To study the FSR of an MRR, we measure the transmission spectrum of the lightpath connecting Input 1 and Output 8, which traverses three SEs containing six MRRs. A bias voltage of 2.6 V is applied to the phase shifter of the last MRR while leaving the other five MRRs unbiased. It is evident from Figure 3-5C that the single biased MRR shows an FSR of about 1.831 THz (shallow troughs), and the unbiased MRRs show good alignment of resonances (deep troughs).
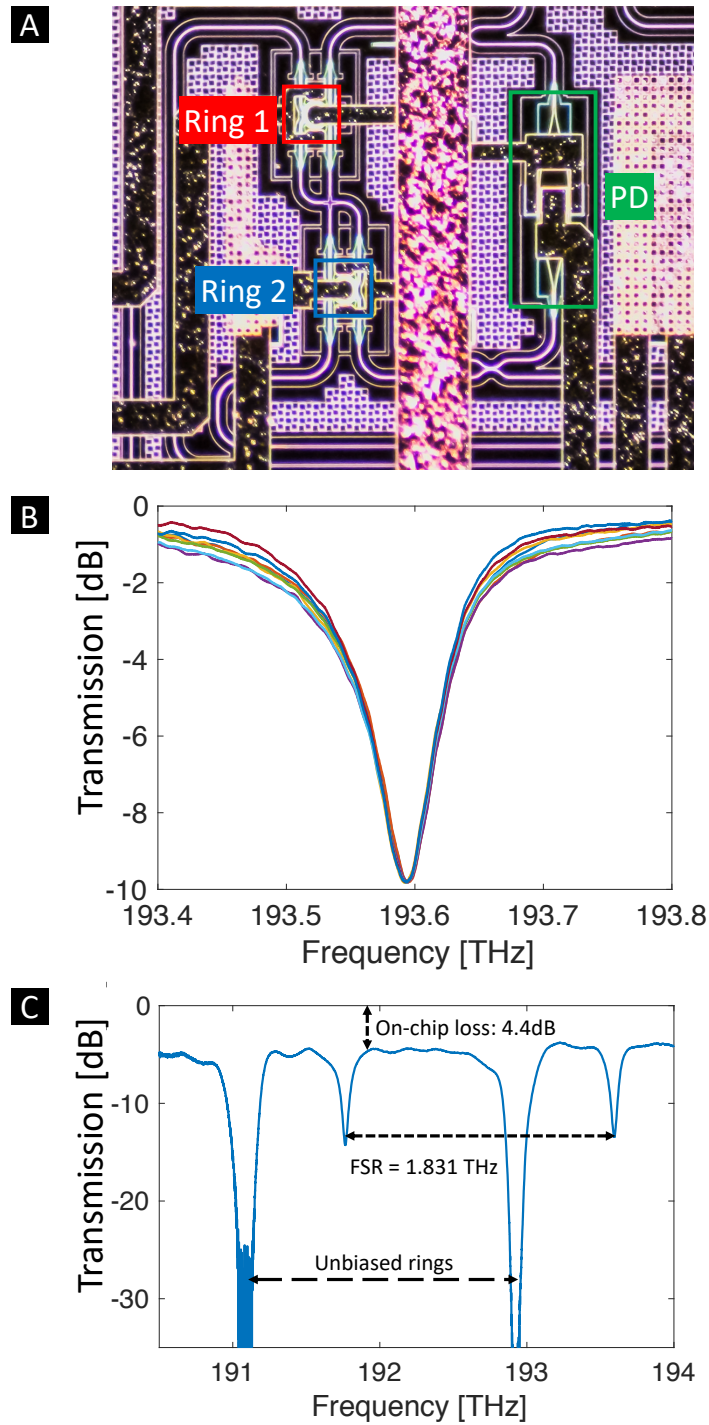
**Figure 3-5: (A)** Micrograph of a single SE showing configuration of both MRRs, waveguide crossing, and on-chip PD. **(B)** Bus waveguide transmission of 8 MRRs on the last switching stage, showing consistent filter profiles across MRRs. **(C)** Lightpath transmission spectrum through six MRRs, with one MRR biased at 2.6 V and the rest unbiased, showing an MRR FSR of 1.831 THz.

The resonance of each MRR can be adjusted by applying a DC voltage across the N-doped portions of the resonator, which behave as resistive heaters and thermo-optically change the roundtrip phase of the MRR. We show the resonance shift as a function of bias voltage between 2.6 V – 2.82 V for a single MRR in Figure 3-6A. The TO tuning efficiency can thus be extracted, as shown in Figure 3-6B, to be about 0.39 nm/mW or 48.85 GHz/mW, which corresponds to a $P_\pi$ of about 18.7 mW. We also characterize the switch reconfiguration speed by measuring the optical time-domain response of the TO switch. With a 150 KHz electrical square-wave signal at 50% duty cycle applied to the path from Input 7 to Output 1, we observe a switching rise time of 1.2 μs and fall time of 0.5 μs (0% to 100 %), as shown in Figure 3-7. For doped waveguide heaters, the heating process is typically longer than the cooling process. The longer rise time is due to the slower temperature increase of the phase-shifter relative to decreasing its temperature via dissipation.
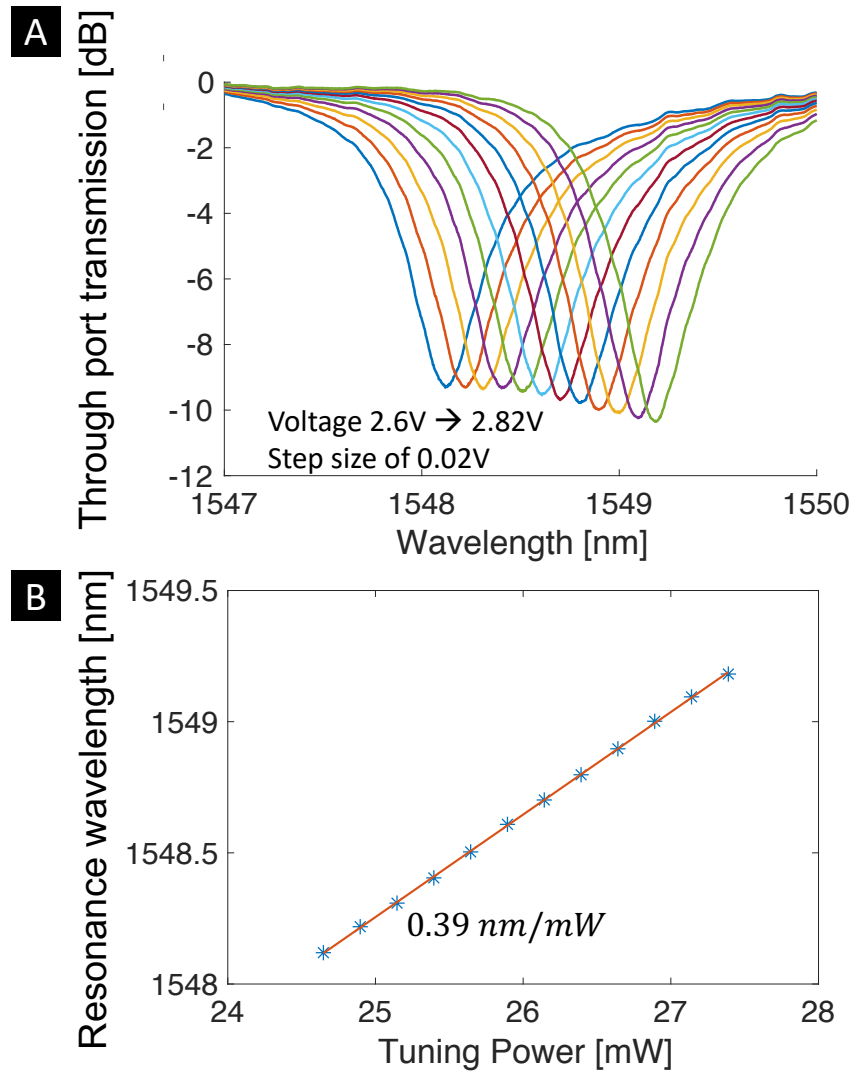
**Figure 3-6: (A)** Change in resonance wavelength of the MRR as heater bias is swept from 2.6 V to 2.82 V in 0.02 V step size. **(B)** Extracted resonance tuning power efficiency showing a linear trend around the wavelength range shown in Figure 3-6A.
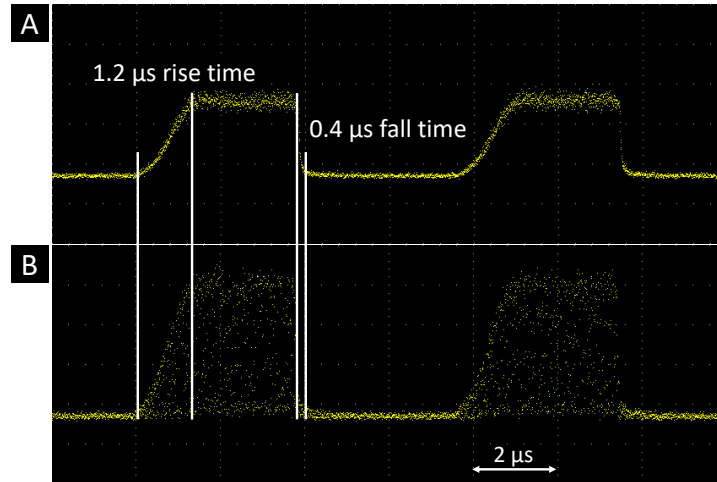
**Figure 3-7: (A) Rise and fall times of a continuous wave (CW) optical signal at 1548.5 nm through a single MRR. (B) Rise and fall time of an optical signal at 1548.5 nm modulated at 32 Gbps. Both are measured from 0% to 100%.**

The dual-MRR configuration of the SE resembles two MRRs simultaneously coupled to two parallel waveguides, as illustrated in Figure 3-8. The effective filter behavior is therefore a combined effect between both MRRs (Figure 3-8A) and the larger cavity formed by the MRRs and the waveguides (Figure 3-8B). When both MRRs in an SE are aligned in resonance, the dual-MRR switching widens the passband of the transmission. Figure 3-9 compares the transmission spectra of single- and dual-MRR switching mechanisms, as well as the crosstalk spectrum from the large cavity resonance when both MRRs are off-resonance. The single-MRR switching shows a passband of 120 GHz. By aligning the resonances of both MRRs with the resonance of the large cavity, the SE provides a boost in switching bandwidth to 165 GHz, as well as a modest improvement of 0.5 dB in peak switching power. The close agreement between bandwidths of the dual-MRR and the large cavity cases means that the cavity mode in dual-MRR switching largely exists in the large cavity. The large cavity, however, can still allow light to circulate even when both MRRs are at off-resonance, and therefore sets an extinction floor of about 15 dB. Future iterations of the device will address this issue by careful design of the waveguide section which detunes the peak of the larger cavity passband when the MRRs are off-resonance, or by inserting a phase shifter in the waveguide

section for additional tunability between switching states. In the context of current $25-50$ GHz per-channel baud rates for datacenter applications, the wide MRR passband can potentially eliminate the need for stabilization due to sufficient margin for thermal drifts in comparison to the signal bandwidth. In the following analysis, we evaluate the switch device's performance under dual-MRR switching mode because of both its extended passband and lower loss metrics.
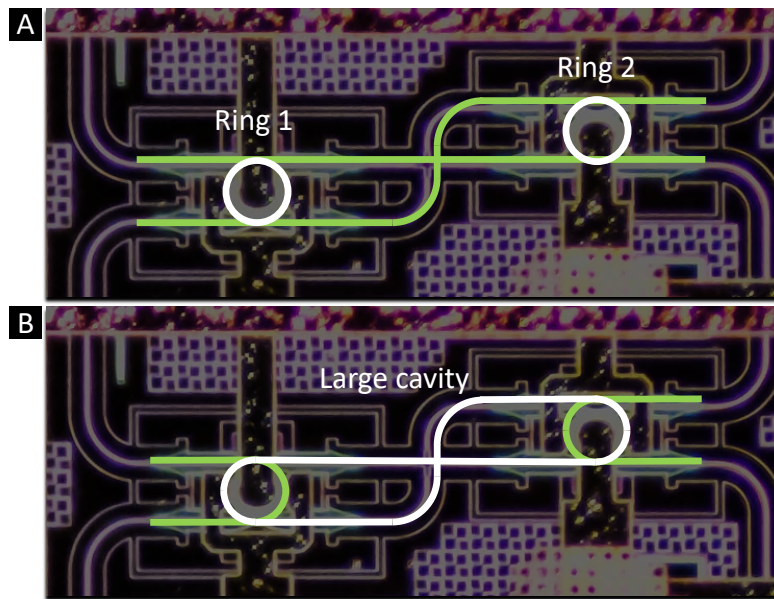


**Figure 3-8: Schematic of the parallel coupled resonators showing the two small cavities formed by the two MRRs (A), as well as the large cavity formed by halves of the two MRRs and the waveguides connecting them (B).**
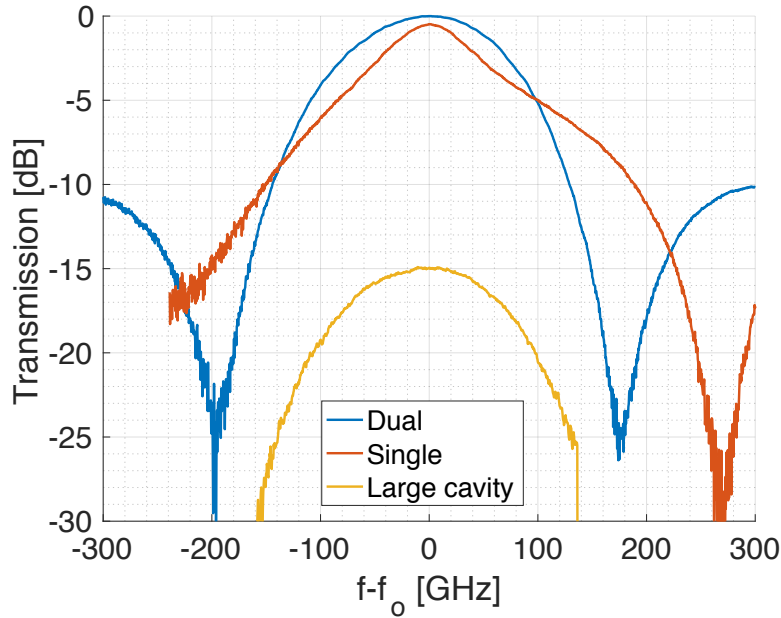
**Figure 3-9: Comparison of drop passbands and peak transmission for a dual-MRR SE under single- and dual-MRR switching. The large cavity transmission is taken with both MRRs far off-resonance. Transmission is normalized based on the dual-MRR case.**

To characterize individual SE performance, we examine 16 paths of the switch which differ pairwise by the state of one SE each. The paths studied connect Input 1 – Outputs 7/8, Input 2 – Outputs 3/4, Input 3 – Outputs 7/8, Input 4 – Outputs 3/4, Input 5 – Outputs 5/6, Input 6 – Outputs 1/2, Input 7 – Outputs 5/6, Input 8 – Outputs 1/2. For the same input, the SE in the last stage is toggled between Bar and Cross, allowing the routed signal and leakage power levels to be measured for both states, and their extinction ratio and crosstalk levels to be determined. Given a single input into a 2x2 SE in a particular switching state, the signal is the power level measured at the designated output, and the leakage is the power level measured at the undesignated output; their difference is the crosstalk level of that particular state. We define a state's extinction ratio as the difference between its signal power and the other state's leakage arriving at the same output. Figure 3-10 shows the signal and leakage power levels for the 16 paths; an average extinction ratio of 18.8 dB for the Cross state and 14.7 dB for the Bar state are observed. Crosstalk levels are similar between Cross and Bar states, averaging to -16.75 dB.
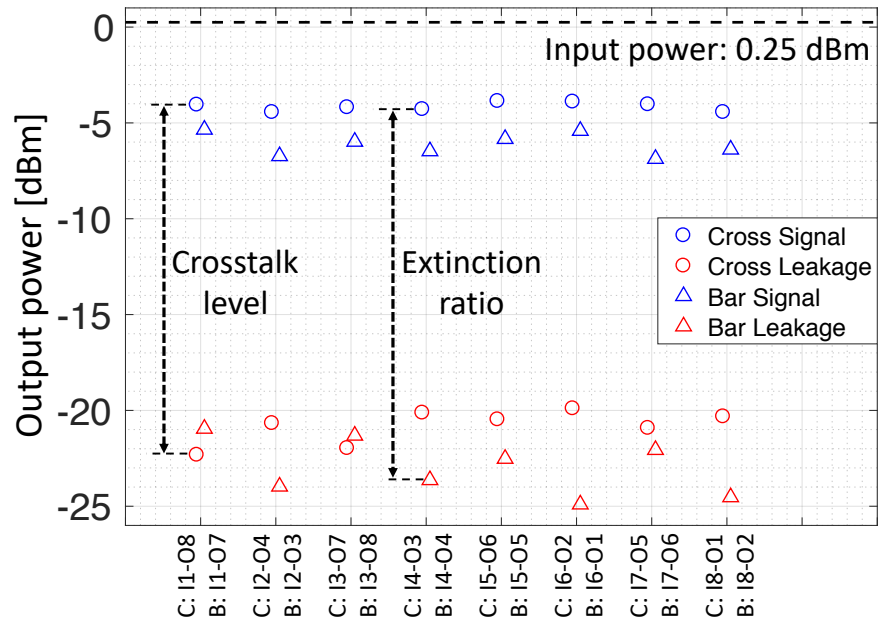
**Figure 3-10: Signal and leakage power levels for paths connecting 16 pairs of input-output as indicated, showing crosstalk and extinction ratio along each path. Each pair of the paths differ only in the state of one SEs, allowing per-SE extinction ratio and crosstalk levels to be extracted as indicated.**

### 3.2.3 Switch Fabric Performance

We first examine the switch's performance in two representative cases of operations – SEs configured in all-Cross and all-Bar states at 1548.5 nm, the peak transmission wavelength of the on-chip grating couplers. The switch fabric defaults to all-Cross with minimal loss along each lightpath with all MRRs unbiased and at off-resonance. In the all-Bar case, all MRRs are biased to resonate with the input wavelength, inducing the maximal attenuation on each lightpath due to loss associated with traversing in and out of the MRRs and attenuation in the doped waveguides. Shown in Figure 3-11, the average on-chip loss among all paths is 4.4 dB in all-Cross and 8.4 dB in all-Bar. The off-resonance and on-resonance losses are estimated at 0.67 dB and 2 dB per SE. A breakdown of the component loss contributions of the device is shown in Table 3-1. The performance of all switch paths from 8 inputs to 8 outputs is summarized in Figure 3-12, in which we show a mean path insertion loss of 6.7 dB. The worst-case crosstalk for each input-output connection, which is taken at the non-signal port with the highest leakage power, is -16 dB on average. While a majority of the worst-case crosstalk levels are within -13 dB and -23 dB, we attribute a few cases (Input 6 to Outputs 3 and 7, and Input 8 to Output 7), where high worst-case crosstalk levels are observed, to fabrication variations of the MRR elements. We further characterize the change in passband as the number of on-resonance SEs increases along a lightpath. In this case, a broadband signal is injected through Input 7 of the switch, and routed to Outputs 5, 1, 3, 4 via $0 – 3$ on-resonance SEs respectively, as illustrated in Figure 3-13A – Figure 3-13D. We show the spectra of the signal and three $1^{st}$ order crosstalks resulted from each of the 3 SEs traversed by the lightpath in each routing. The $2^{nd}$ order crosstalks are suppressed below -35 dB and therefore omitted for clarity. It is evident from Figure 3-13E – Figure 3-13H that, while cascaded MRR SEs increasingly narrows the switched passband, a lightpath traversing through $1 – 3$ on-resonance SEs still maintains 147 GHz, 96 GHz, and 55 GHz of bandwidth, respectively. The

power consumption per MRR is on average 0 mW and 25.6 mW for off- and on-resonance states, respectively.

**Table 3-1: Key component loss**

| Item | Loss |
|---|---|
| Waveguide propagation loss | 2 dB/cm |
| SE in Cross | 0.67 dB |
| SE in Bar | 2 dB |
| Grating coupler | 3.6 dB/facet |
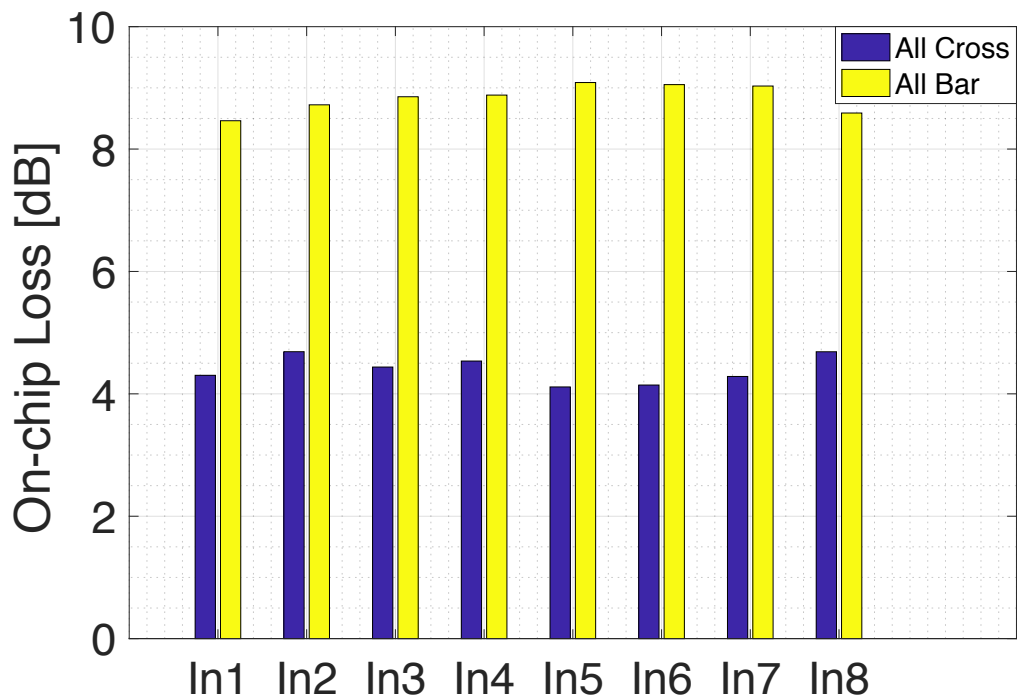


**Figure 3-11: On-chip loss of each input signal in all-Cross and all-Bar configurations.**
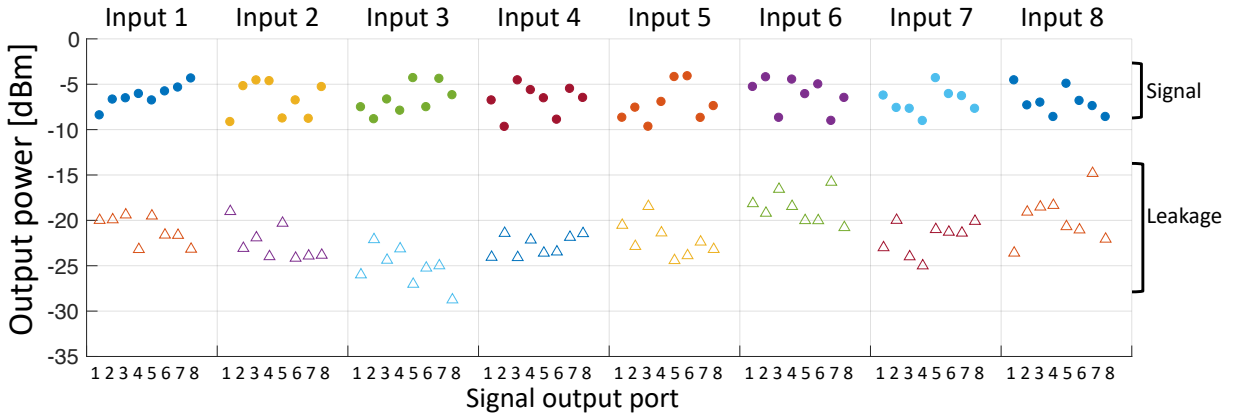
**Figure 3-12: On-chip signal power and worst-case leakage power levels for all switch paths connecting every input to every output. Grating coupler loss is compensated at injection. Worst-case crosstalk is taken as the difference between signal and leakage power for each connection.**
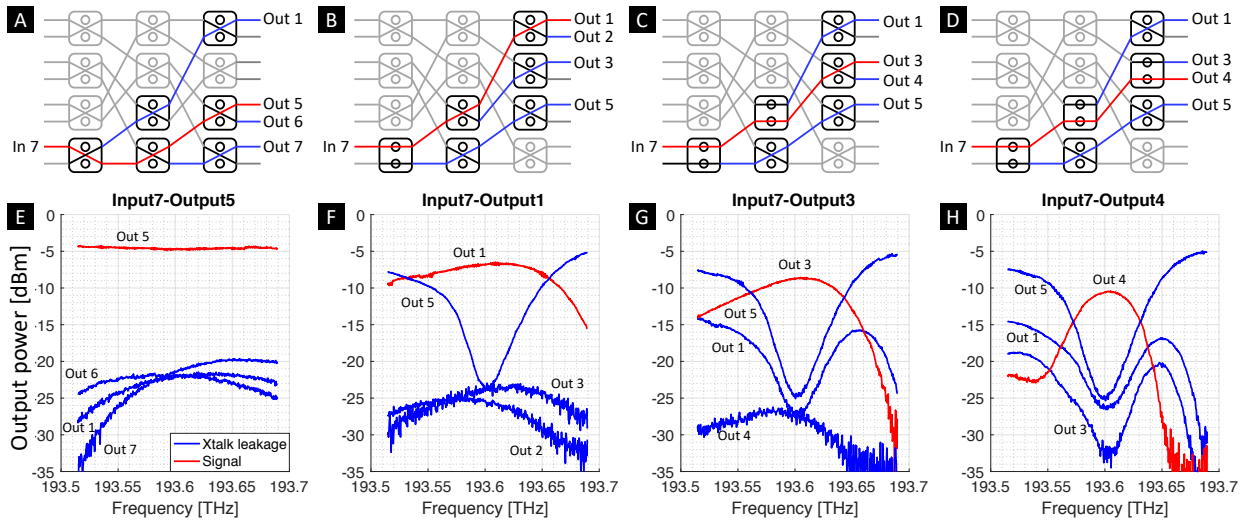


**Figure 3-13: (A-D) Illustration of the switch paths through 0 – 3 on-resonance SEs, respectively. (E-H) Signal and 1st order crosstalk spectra for the routing schemes shown in A-D.**

Data transmission was performed using an Anritsu MP1900A Signal Quality Analyzer and a Thorlabs MX35E Reference Transmitter at 32 Gbps NRZ-OOK using PRBS31. On-chip switch paths with varying number of on-resonance SEs all exhibit error free operations. The schematic of the data test is shown in Figure 3-14A, and the same four paths described in Figure 3-13 are tested. The optical carrier is launched from a TLD at 1548.5 nm and modulated by an MZM with 0 dBm output power. The intensity

48

modulated optical signal is then guided to the silicon photonic MRR-based switch via a PC. A PA consisting of a VOA and an EDFA is used before the switch to compensate for coupling and propagation losses through the chip, ensuring -10 dBm of optical power exits the chip. This PA is also used to replicate the device insertion loss in the B2B reference case. A second set of EDFA and VOA is used to adjust the receiver optical power for the BER measurement. A TOF with 180 GHz passband is used to reject out-of-band ASE noise. The receiver consists of a Finisar XPRV2022A PD-TIA assembly, which performs the optical-to-electrical conversion and allows the data signal to be analyzed by the Anritsu error checker. All lightpaths examined are within 2 dB power penalty comparing to the B2B reference at $10^{-9}$ BER, shown in Figure 3-14B. Despite going through $0 - 3$ on-resonance SEs with successively narrowed optical passband, all four switch paths' show clear eye-openings shown in Figure 3-14C and are within 1 dB penalty variations at 32 Gbps, which indicates negligible path dependence of power penalties due to consecutive filtering of the MRR switch device.
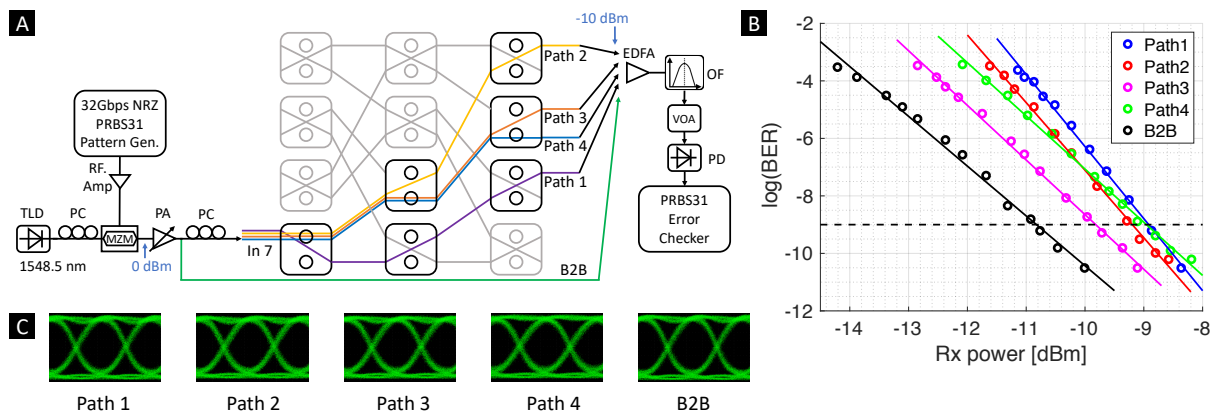


Figure 3-14: (A) Data transmission setup showing the switch paths and the B2B link examined. (B) BER-Rx power relationship for the switch paths and B2B reference path at 32G NRZ-OOK PRBS31. (B) Open eye diagrams of the switch paths and B2B reference path at 32G NRZ-OOK PRBS31.

### 3.2.4 Summary

An 8×8 silicon photonic switch is presented, implementing multi-stage Omega architecture and dual add-drop MRR SEs with a compact footprint of 4 mm². This device is taped out using a commercial silicon photonics design flow at Elenion Technologies. The switch demonstrates a well-balanced set of performance metrics, showing off- and on-resonance SEs losses to be 0.67 dB and 2 dB, respectively, and an end-to-end on-chip loss ranging between 4.4 dB and 9.6 dB. The worst-case first-order switching crosstalk levels have an average of -16 dB, mostly ranging between -13 dB and -23 dB. Component characterizations of the switch show a 120 GHz passband per MRR and a 165 GHz passband for dual-MRR switching. A minimum passband of 55 GHz is observed after three-stage SE filtering. The switching speed of the thermally driven SEs is measured as 1.2 μs rise time and 0.5 μs fall time, with a thermal tuning power efficiency of 48.85 GHz/mW. We perform data transmission tests with 32 Gbps NRZ-OOK, showing less than 2 dB power penalty incurred by the switch routing. The collective of appealing characteristics makes the device suitable for agile functionalities such as bandwidth steering and network reconfiguration for 200G and 400G datacenter applications and pave way for future designs of optically switched datacenter networks.

## 3.3 Scalable Microring-Based Clos Switch Fabric with Switch-and-Select Stages

To increase the port-count of photonic switch fabrics, careful examinations of signal penalties resultant from insertion loss and crosstalk should be conducted. As discussed in Section 3.2, the insertion loss accumulates based on total number of SEs traversed by an optical lightpath, and therefore limiting the total number of cascaded switching stages is an effective strategy when designing multi-stage switch fabrics. To further reduce signal power penalties, it would be beneficial to obtain lower crosstalk by topological modification to cancel the first-order crosstalk, as achieved typically in the dilated topology approach [28]. Modified switch-and-select architecture using MRRs has been shown to also be effective in eliminating first-order crosstalk along the lightpaths [77]. The MRR switch-and-select architecture first employs an array of MRR add-drop filters to perform an $1 \times N$ switching, and a second array of filters for perform an $N \times 1$ selection, which blocks the leakage optical power from the first array. However, the switch-and-select architecture imposes an $O(N^2)$ scaling of the number of MRRs for an $N \times N$ switch fabric, resulting in significant increase in design and control complexities even at moderate port counts.

In this chapter, we address the scalability challenge of the MRR-based switch-and-select switch by designing a Clos network interconnected with low-radix switch-and-select switching stages. This design allows for the construction of large-scale MRR-based switch fabrics while maintaining low crosstalk and low loss. The bounded number of required MRR SEs is significantly reduced while the number of cascaded SE stages is kept at a modest value of six. As a case study, the performance of a $16 \times 16$ Clos switch design is verified based on key parameters characterized from a silicon $4 \times 4$ switch-and-select sub-switch designed and fabricated with AIM Photonics.

### 3.3.1 Clos Architecture with Switch-and-Select Sub-Switches

*MRR-based switch-and-select sub-switch*

The MRR-based switch-and-select architecture is shown in Figure 3-15a. An N×N device consists two linear arrays of 1×N and N×1 spatial (de)multiplexers, connected by a perfect shuffle network. The MRR implementation requires a total of $2N^2$ MRR add-drop filters. This architecture provides SNB connectivity as each pair of input and output switching arrays dedicates a specific path. As the connectivity grows, the number of bypass MRRs increases, while the number of cascaded MRRs along a lightpath remains at 2, which contributes to significant preservation of optical signal bandwidth and power when scaling. Since every lightpath between all inputs and all outputs traverses through 1 switching filter array and 1 selection filter array, the optical path penalty remains largely uniform across all connections. as illustrated in Figure 3-15a, the first order crosstalk leakage (outlined by yellow arrows) from the input switching arrays will get dropped by the off-resonance MRRs at the output stage; hence only the second-order crosstalk occurs (represented by red arrows). In contrast, the crossbar topology suffers from the first-order crosstalk, as shown in Figure 3-15b. Figure 3-15c compares the simulated drop spectra of a crossbar switching device, which imposes a one-time MRR filtering to the optical signal, and the switch-and-select structure, which imposes the MRR filtering twice; the latter exhibits a sharper falloff due to the cascading of two narrowband add-drop elements.
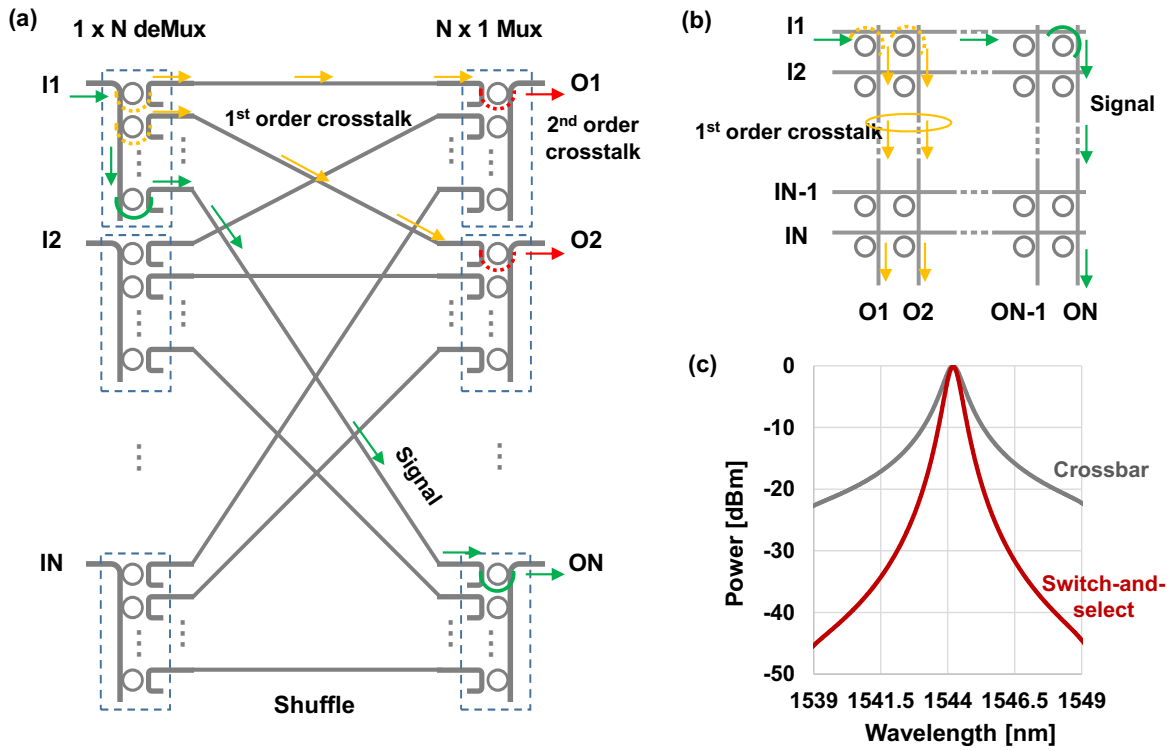
**Figure 3-15: (a) Schematic of the MRR-based switch-and-select topology. Signal (green lines) routes from input 1 to output N. (b) Schematic of the MRR-based crossbar topology. Signal (green lines) routes from input 1 to output N. (c) Comparison of simulated drop spectra of a crossbar switching device and the switch-and-select structure.**

*Clos switch topology*

The Clos architecture, a three-stage network that was first studied by Charles Clos in the early 1950s [83], presents a topology to efficiently interconnect small-radix sub-switches and realize large switch networks [84], therefore achieving a reduction in component complexity and cost. The same strategy applies to silicon photonic switching – switch device designs that show unique benefits at moderate radices but face challenges when scaling can be assembled into higher port-count devices. The Clos topology limits the total number of cascaded sub-switches to three, regardless of the overall device scale.

Figure 3-16a shows a generic Clos network, consisting of $r$ $n \times m$, $m$ $r \times r$, and $r$ $m \times n$ sub-switches, in the first, second, and third stage, respectively. The Clos network can be designed as either SNB or RNB, the conditions of which are dictated by [83]:

$$m \geq (2n - 1), \qquad (\text{3-1})$$

$$m \geq n. \qquad (\text{3-2})$$

Equations (3-1) and (3-2) respectively define SNB and RNB connectivity for the Clos architecture. In the following analysis, we select $m = (2n - 1)$ and $m = n$, which represents the cases with the least redundancy incurred while still satisfying the respective SNB and RNB conditions.

The MRR-based switch-and-select topology offers great flexibility in defining a sub-switch. As an example, a sub-switch with $n \times m$ connectivity by applying $n$ $1 \times m$ spatial de-multiplexer arrays and $m$ $n \times 1$ spatial multiplexer arrays, as shown in Figure 3-16b, with a total of $2nm$ MRRs.
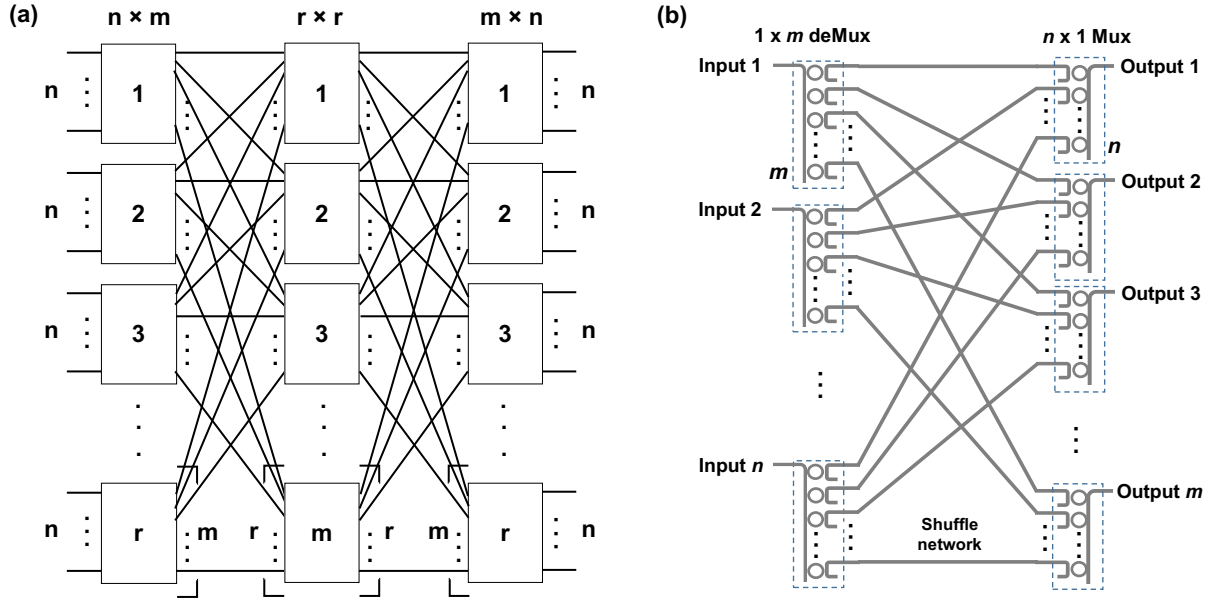
**Figure 3-16: (a) Layout of a generic three-stage Clos network built with $r$ $n{\times}m$, $m$ $r{\times}r$, and $r$ $m{\times}n$ sub-switches. (b) Schematic of an n×m MRR-based sub-switch in the switch-and-select topology.**

Given a total switch port count of $N$ and $r = N/n$, the numbers of required MRR SEs, $S$, for both

SNB Clos and RNB Clos are:

$$S_{SNB} = 4(2n-1)N + 2(\tfrac{2}{n} - \tfrac{1}{n^2})N^2 , \quad (\text{3-3})$$

$$S_{RNB} = 4nN + \tfrac{2}{n}N^2 , \quad\quad\quad (\text{3-4})$$

To compare the scalability between the Clos topology and the single-stage switch-and-select, we define $\alpha$

as the ratio between the total numbers of MRR switching elements in the Clos design and the single-stage

switch-and-select ($S = 2N^2$):

$$\alpha_{SNB} = \tfrac{2(2n-1)}{N} + (\tfrac{2}{n} - \tfrac{1}{n^2}) , \quad\quad (\text{3-5})$$

$$\alpha_{RNB} = \tfrac{2n}{N} + \tfrac{1}{n} , \quad\quad\quad (\text{3-6})$$

Figure 3-17a and Figure 3-17b examines the trend of $\alpha$ for SNB Clos and RNB Clos switch fabrics,

respectively, for various $n$ values. Figure 3-17a shows that for $N \leq 16$, the single-stage switch-and-select

is the preferred topology with the least number of MRRs for the SNB connections. The scaling advantage

of the Clos design becomes evident with $N \geq 64$, as the Clos topologies have less total number of MRRs than the single-stage switch-and-select. It is also evident that, for low port counts, smaller sub-switches result in lower number of MRRs, while the trend reverses for larger port counts. At $N \geq 128$, the Clos topology with $n = 8$ requires less than half the number of MRRs compared to the single-stage switch-and-select topology. For RNB Clos configurations, the reduction in total number of MRR cells is achieved at lower port counts, as shown in Figure 3-17b. At $n < 8$ and a port count of 16, the RNB Clos switch fabric saves over 25% of the MRRs compared to the single-stage switch-and-select design. At $n = 8$ and a port count of 128, the Clos only requires a quarter of the number of MRRs needed for single-stage switch-and-select. The total number of SEs, as a function of the switch port count N for both SNB and RNB Clos switch fabrics, is plotted in Figure 3-18, with $n$=2, 4 and 8, along with the numbers for the single-stage switch-and-select topology. We also identify the Clos RNB topology with $n = 4$ as the best configuration for the 16×16 switch for the rest of the analysis because of the low total number of MRR cells required.
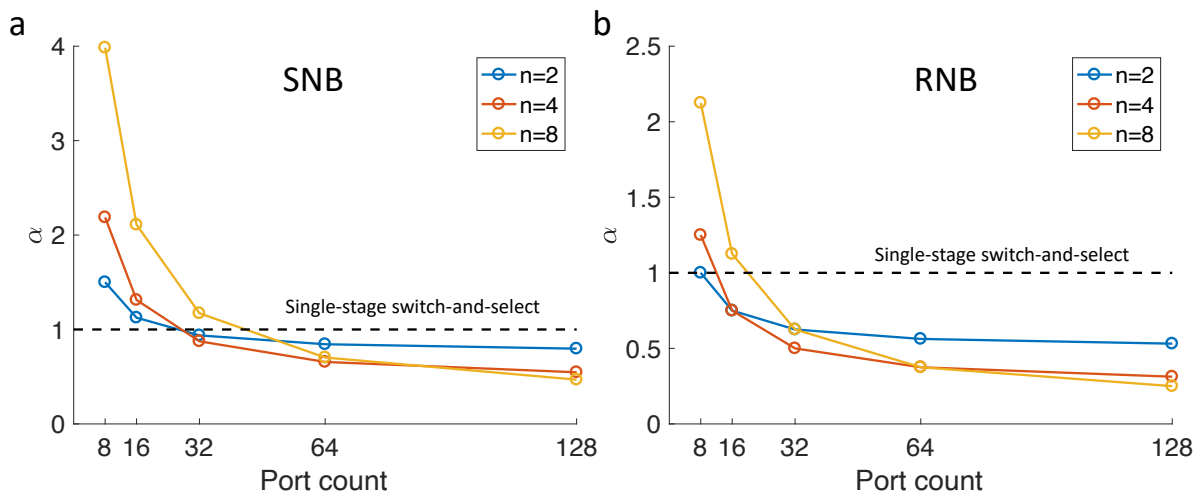


**Figure 3-17: Ratio between the number of switching elements in the (a) SNB Clos design and the single-stage switch-and-select, and (b) RNB Clos design and the single-stage switch-and-select, for n=2, 4, 8.**
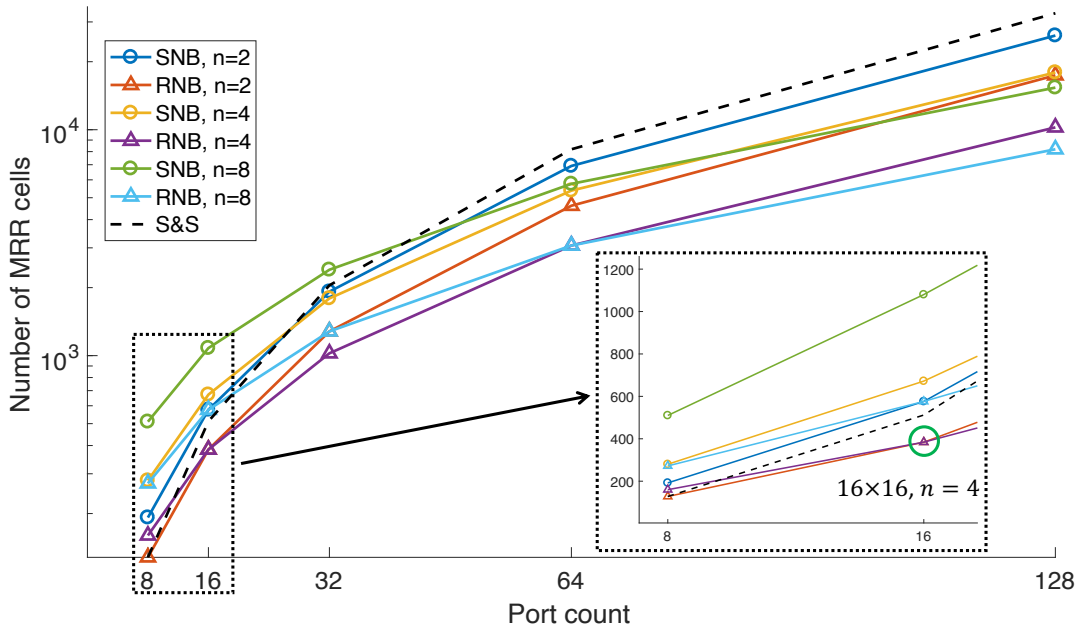
**Figure 3-18: Number of MRR cells as a function of the switch port count N for both SNB and RNB Clos networks, with n=2, 4, and 8, as well as for single-stage switch-and-select devices. The 16×16 switch configuration with n=4 is identified in the inset.**

The Clos topology also limits the number of crosstalk sources in the switch fabric to $(n + r + m - 3)$, which is a reduction from the single-stage switch-and-select value of $(N\text{-}1)$. Figure 3-19a compares the worst-case crosstalk ratio experienced by a lightpath between Clos SNB, Clos RNB, and single-stage switch-and-select, and shows Clos RNB consistently achieves lower crosstalk levels for all switch port counts. The Clos RNB also has the least number of maximum bypass MRRs – the off-resonant MRRs along a lightpath, which contributes to the aggregated IL of the switch fabric. The total MRR-induced IL is shown in Figure 3-19b and compared among the three topologies. It is clear that the single-stage switch-and-select topologies show a sharp increase in the loss due to the large number of bypass MRRs present; the IL discrepancy from the Clos designs is especially high at large port counts.
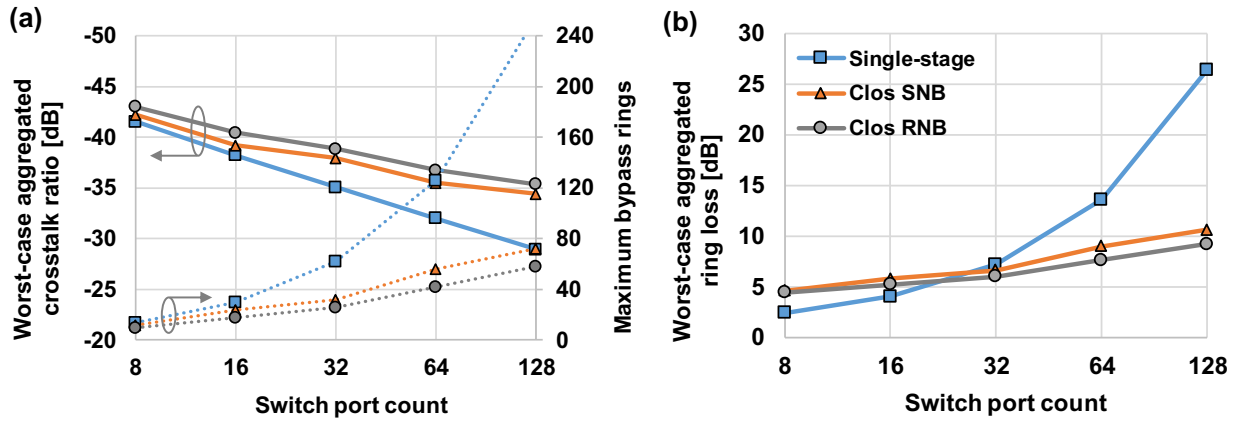
**Figure 3-19: (a) Worst-case aggregated crosstalk ratio, number of maximum bypass MRRs and (b) worst-case aggregated MRR loss, as a function of the switch port count $N$ for both SNB and RNB Clos networks, as well as single-stage switch-and-select devices. Calculations are based on the measured results reported in [44].**

58

*Bandwidth reduction in a multi-stage MRR switch*

The bandwidth narrow-down factor for MRR-based multi-stage switch is typically not addressed in previous works [63] [85] [86]. However, it is crucial to consider the reduction of optical bandwidth when designing large switch fabrics with cascaded narrowband filters. The drop spectrum of a first-order add-drop MRR can be described as [87]:

$$D(\lambda) = \frac{D_0}{1+\left(\frac{2\ FSR}{\pi\Delta\lambda_{3dB}}\sin\left(\pi\frac{(\lambda-\lambda_{res})}{FSR}\right)\right)^2} , \qquad (3\text{-}7)$$

where $D_0$ is the drop attenuation at the resonance, $\Delta\lambda_{3dB}$ is the 3 dB optical bandwidth, FSR is the free spectral range of the cavity, and $\lambda_{res}$ is the resonance of interest. For high finesse MRRs, the drop spectrum close to its resonance can be approximated as:

$$D(\lambda) \approx \frac{D_0}{1+\left(\frac{2}{\Delta\lambda_{3dB}}(\lambda-\lambda_{res})\right)^2} , \qquad (3\text{-}8)$$

If the optical signal is consecutively dropped by $k$ add-drop elements, the overall spectral response is $D_k(\lambda) = [D(\lambda)]^k$. The 3-dB optical bandwidth of the optical path is therefore given by:

$$\Delta\lambda_{3dB,k} \approx \Delta\lambda_{3dB} \times \sqrt{2^{\frac{1}{k}} - 1} , \qquad (3\text{-}9)$$

It should be noted that this equation also holds approximately true for the second-order add-drop elements with the *maximally flat* passband condition [88] imposed by:

$$\kappa_{RR} = \frac{\kappa_{RW}^2}{2-\kappa_{RW}^2} , \quad (3\text{-}10)$$

where $\kappa_{RW}$ is the MMR-bus field coupling and $\kappa_{RR}$ is the inter-MRR coupling coefficient. Figure 3-20 compares the bandwidth narrow-down factor for the crossbar, single-stage switch-and-select, and Clos of switch-and-select topologies, in which the lightpath traverses 1, 2, and 6 MRRs, respectively, regardless of the total number of port counts. When implementing the Clos topology based on switch-and-select sub-switches, the passband of the switch optical channel is only about 35% of the MRR filters employed. It is

important to co-design the MRRs and the high-speed optical communication applications and leave adequate margin to compensate for this bandwidth reduction.
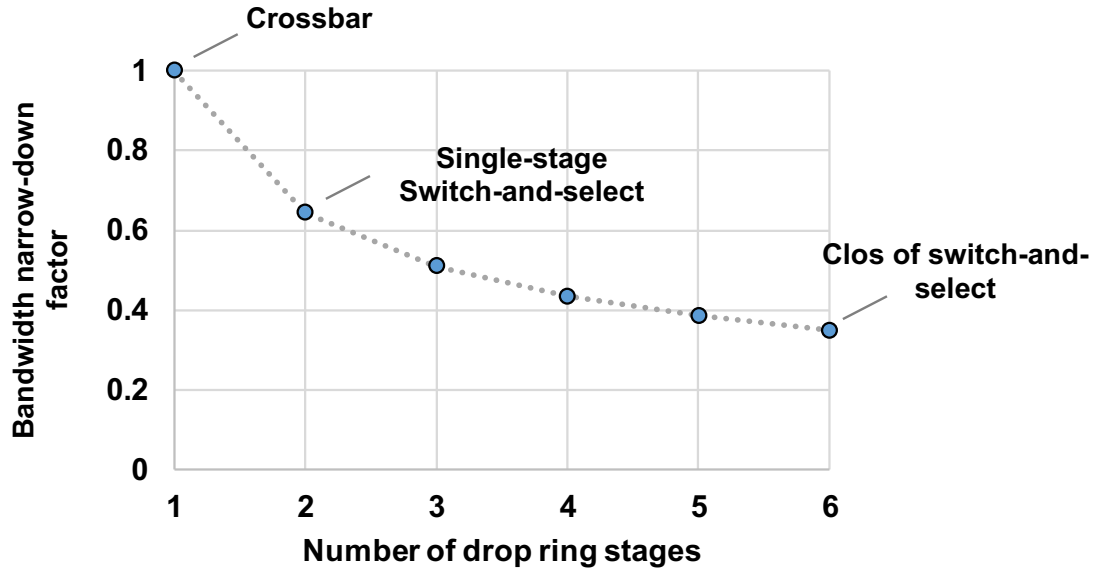


**Figure 3-20: Bandwidth narrow-down factor as a function of the number of drop MRRs per path.**

### 3.2.2 Silicon 4×4 Sub-Switch for a 16×16 Clos Switch Fabric

In this work, a 16×16 RNB Clos switch fabric is chosen as the initial demonstration for the Clos design. While both ($n=m=2$, $r=8$) and ($n=m=r=4$) configurations have the same number of MRRs (384), the latter consists of 12 identical 4×4 sub-switches and can be analyzed based on the performance metrics of a single 4×4 device. A detailed description of the design, fabrication and characterization of the 4×4 sub-switch is presented in this section, which is used to extrapolate the performance of a 16×16 Clos switch fabric.

*Device design, fabrication, and packaging*

The silicon 4×4 MRRs-based switch-and-select layout is shown schematically in Figure 3-21a, similar to the one reported in [77]. It has 32 MRRs tuneable with TO phase shifters with optical terminations placed at each through port to eliminate reflections. A perfect shuffle is used to connect the 16 (4×4 arrays) MRRs at the input stage to the 16 MRRs at the output stage. Low-loss and low-crosstalk waveguide crossings are realized with MMIs. The input and output MRR switching arrays are staggered by an offset of ~130 μm to reduce the number of waveguide bending structures. An edge coupler array at 127 μm pitch is used to couple TE-polarized light in and out of the chip, with two pairs of looped couplers to facilitate the coupling process.
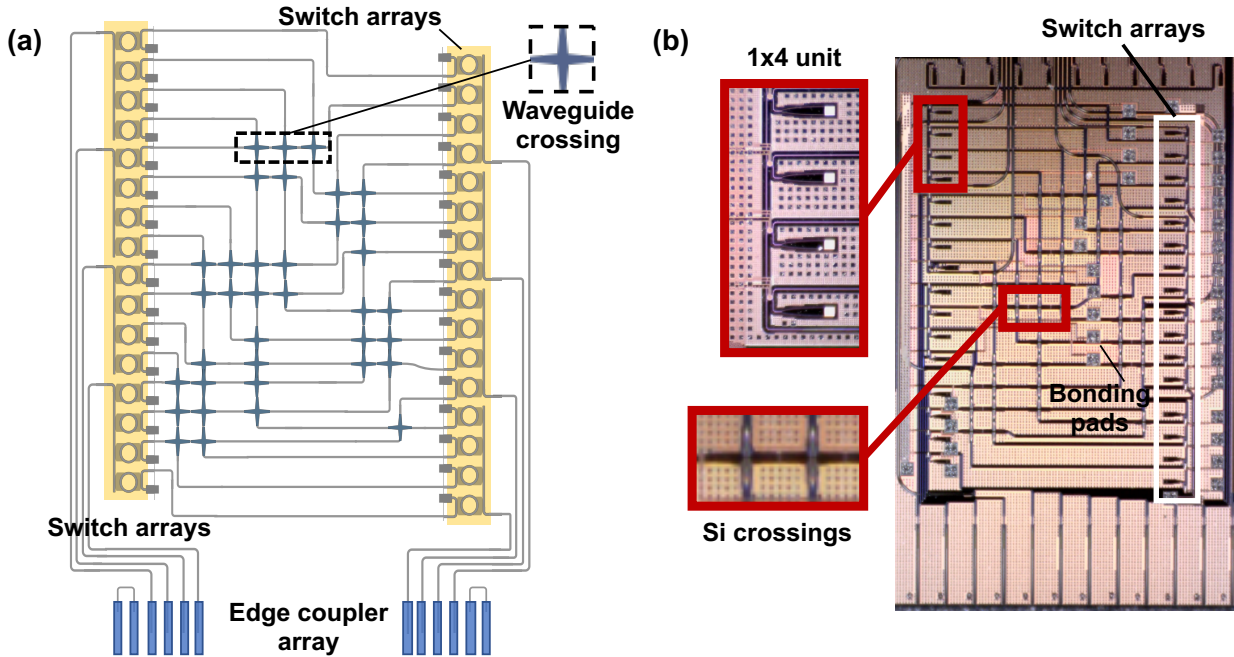
**Figure 3-21: (a) Schematic of the 4×4 MRR-based switch-and-select layout with inset showing the silicon waveguide crossing. (b) Microscope photo of the fabricated device with insets showing the 1×4 spatial de-multiplexer and the silicon waveguide crossing.**

The device was taped out using standard PDK elements through the AIM Photonics MPW run [89]. A microscope photo of the fabricated device is shown in Figure 3-21b. The MRRs are placed at a pitch of 100 μm to minimize the thermal crosstalk. The measured resonance shift shows a TO tuning efficiency of 1 nm/mW [90]. The switch fabric has a compact footprint of 1.6×2.5 mm$^2$ with 34 electrical bonding pads. The fabricated device was wire-bonded to a custom designed PCB for electrical fan-out, as shown in Figure 3-22a. The silicon chip was first die-bonded onto a custom copper sub-mount with one overhung edge. The facet of the chip was polished together with the copper mount to remove the chip ledge residual from the dicing trench (Figure 3-22b). This allows a standard PMF array to be attached using a UV curable epoxy. MMCX connectors are used for a compact PCB footprint. The packaged device was mounted with an aluminium fiber holder and the completed package is shown in Figure 3-22c.
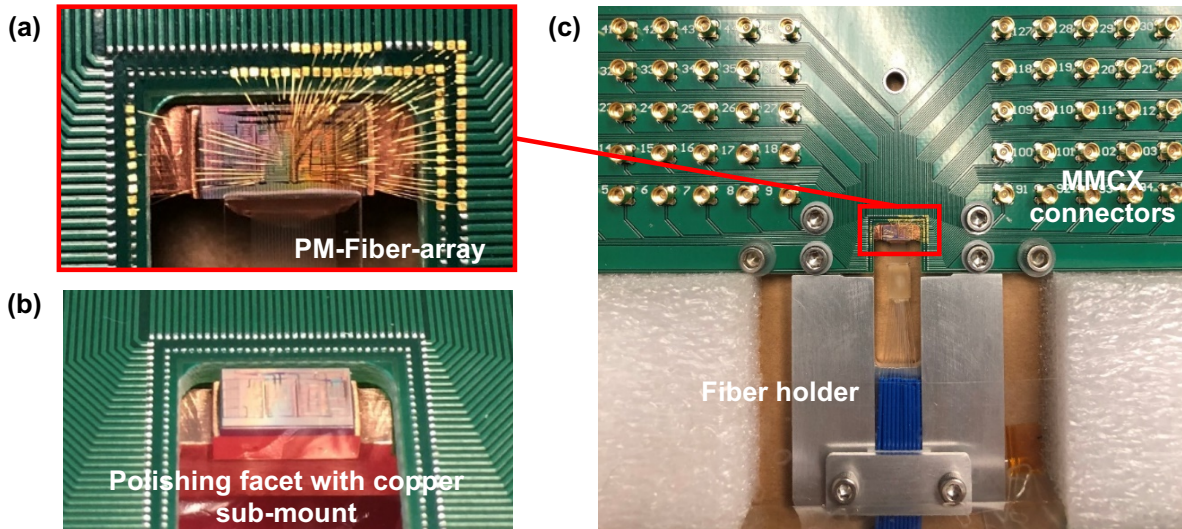
**Figure 3-22: (a) Photo of the wire-bonded chip with UV-curved PM fiber array. (b) Photo of the copper sub-mount, which was polished together with the chip facet for fiber attachment. (c) Photo of the packaged 4×4 switch-and-select device.**

The device evaluation is performed for Inputs 1 and 4, to all four output ports. The measurement covers the shortest path and the longest path cases, which will be used in the three-stage Clos switch fabric analysis. The on-state and off-state are verified by searching for the highest and lowest output power, respectively. An OSA is used to record transmission spectrum. Figure 3-23 summarizes the switch on-chip insertion loss and crosstalk ratio. The on-state and off-state bias lies in the range of 2.2 to 2.7 V and 0.2 to 0.7 V, respectively. The power consumption per path (including both on-state and off-state tuning) is estimated at 20 mW for the two switching arrays. The measured on-chip loss ranges between 2.3 to 4.8 dB. The path-dependent loss mainly comes from the central shuffle network. The passive loss of each path is a linear combination of different loss sources: MRRs drop and through loss, waveguide propagation and crossing loss, and the chip coupling loss, which are summarized in Table 3-2. The coupling loss is determined to be 6 dB/facet by measuring the looped pair of edge couplings. As shown in Figure 3-23, the measured switch crosstalk ratio is in the range of -57 dB to -48.5 dB, benefitting from the first-order crosstalk cancellation in the switch-and-select architecture. The worst-case crosstalk ratio occurs in path 1-

4 as marked in red. The slight degradation is believed to come from the higher path insertion loss. Figure 3-24a illustrates the elimination of first-order crosstalk by the two MRR filters along a lightpath, showing that the on-off extinction of single MRR elements is more than 23 dB. A detailed measurement on the extinction ratio of path 1-3 is also presented in Figure 3-24b as the MRR TO bias changes, showing a maximal optical switching extinction ratio of 49 dB. The normalized optical path spectra are shown in Figure 3-25a and have consistent passbands of 22.5 GHz, which reflects individual MRR filters' bandwidth as 35 GHz, based on Equation (3-9).
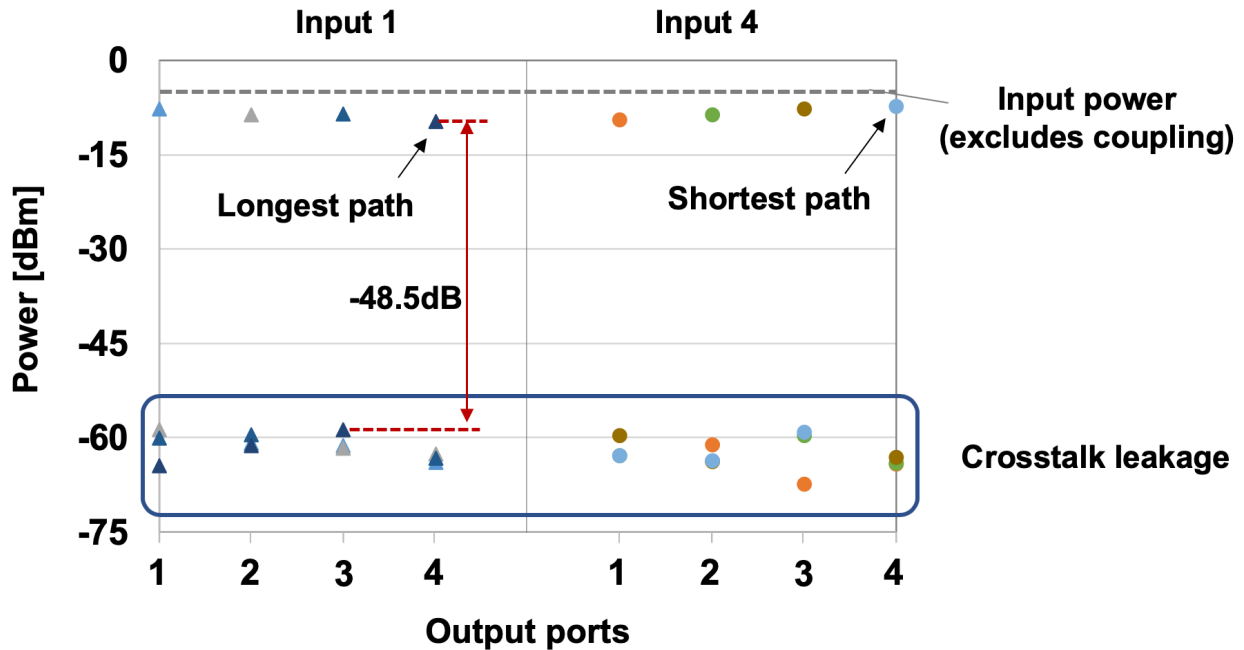


**Figure 3-23: Measured insertion loss and crosstalk for paths from input 1 and 4 to all four outputs.**

**Table 3-2: Estimated component loss**

| Item | Loss |
|---|---|
| Waveguide propagation loss | 2 dB/cm |
| TO MRR drop | 0.5 dB |
| TO MRR through | 0.1 dB |
| Waveguide crossing | 0.2 dB |
| Edge coupler | 6 dB/facet |
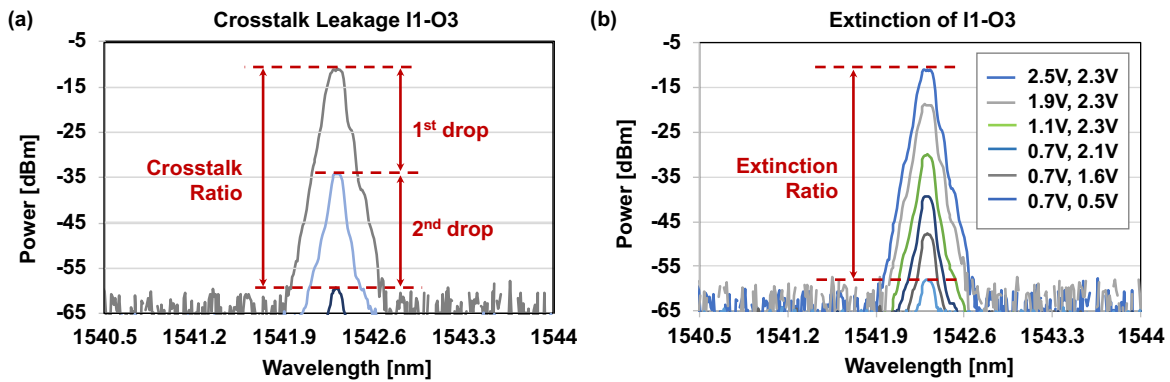


**Figure 3-24: (a) Breakdown measurement on the worst-case crosstalk. Data routed from input 1 to output 4 and the crosstalk leakage to output 3. (b) Power tuning for path 1-3 with various bias voltages.**
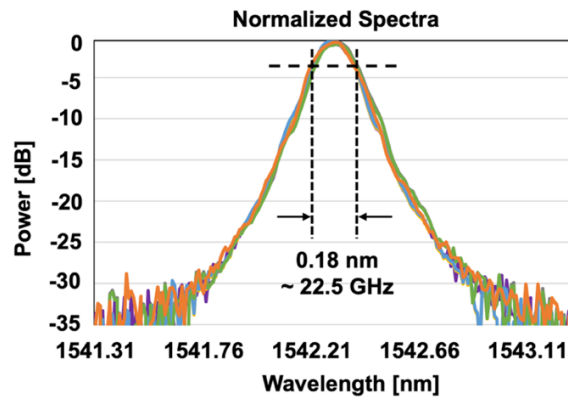


**Figure 3-25: (a) Normalized spectra with resolution at 0.1 nm. (b) BER as a function of received optical power at 12.5 Gb/s for path 1-4 and 4-4. Insets show the measured eye diagrams.**

*16×16 Clos switch fabric design*

A case study for a 16×16 switch is performed to validate the Clos switch concept with the switch-and-select sub-switches. The schematic of the switch fabric is shown in Figure 3-26. It consists of 12 4×4 sub-switches in three stages, connected via two inter-stage shuffle networks. The footprint of the shuffle network is mostly limited by the number of cascaded waveguide crossings, as seen in Figure 3-21a. Since the inter-stage shuffle connects the 16 outputs of one stage to the 16 inputs of the following stage, its width is roughly the same as that of the inner shuffle within the 4×4 sub-switch (~0.8 mm). For each sub-switch, electrical pads can be placed right next to the MRR unit and for flip-chip packaging with a PCB or an interposer as demonstrated in [25]. The footprint of the 16×16 Clos switch fabric with the fabricated 4×4 sub-switches is estimated as 10 mm × 6.4 mm. The device dimensions can be further reduced if thermal isolation trenches are placed around the MRR components to allow less spacing between the MRRs.
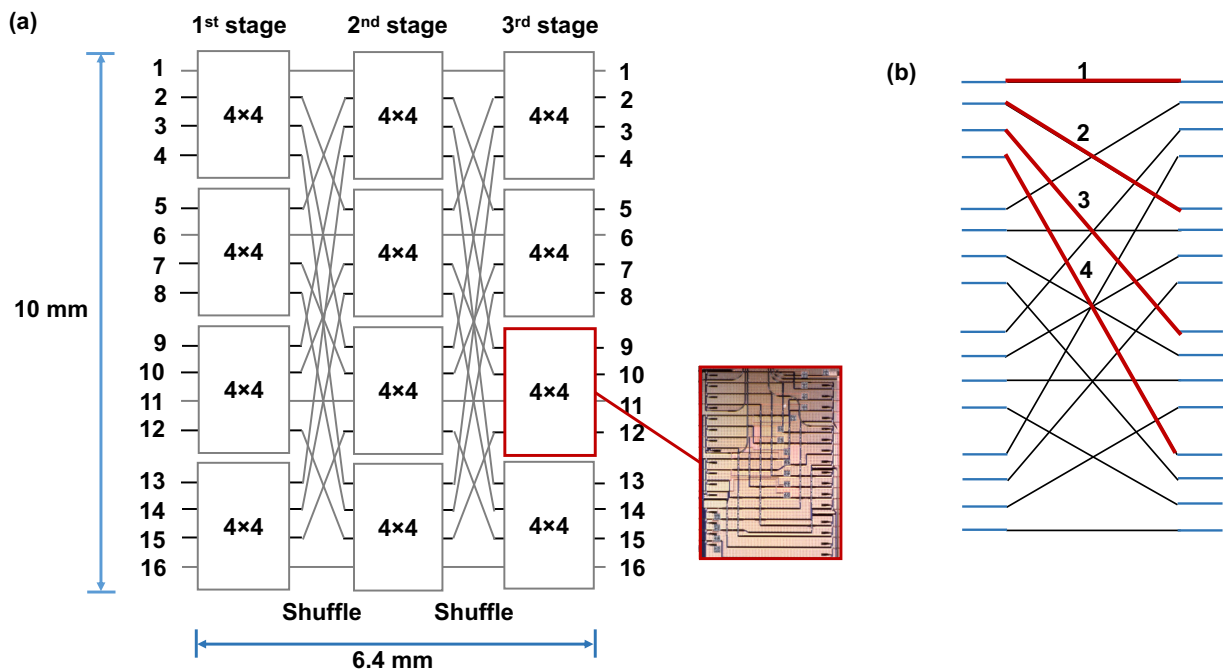


**Figure 3-26: Schematic of (a) the 16×16 Clos switch fabric and (b) the inter-stage shuffle network. Path loss of the shuffle is estimated in Table 3-3.**

66

The inter-stage shuffle loss is estimated based on the experiment-validated component-level loss listed in Table 3-2. By verifying the detailed shuffle parameters, the shuffle path loss is calculated and shown in Table 3-3. The shortest path travels directly from input 1 to output 1, while the longest path travels from input 1 to output 16 via two inter-stage shuffles. The on-chip insertion loss for the 16×16 Clos switch fabric can thus be summed, in the range of 4.7 dB to 15.58 dB, as shown in Table 3-4. The path-dependent loss is primarily a result of the loss variations in the shuffles. Under conservative and realistic improvements in the component loss values of the waveguide propagation, crossing, and MRR drop, it is possible to refine the total on-chip loss to the range of 2.5 dB to 6.44 dB. Taking into account the measured worst-case crosstalk ratio of -48.5 dB for the 4×4 sub-switch, a worst-case crosstalk ratio of -39 dB for the 16×16 Clos switch fabric is determined accounting for all crosstalk sources under a full load.

**Table 3-3: Estimated shuffle loss**

| Path | Waveguide propagation length [mm] | Number of waveguide crossings | Loss [dB] |
|------|-----------------------------------|-------------------------------|-----------|
| 1 | 0.8 | 0 | 0.2 |
| 2 | 2.7 | 3 | 1.1 |
| 3 | 4.6 | 6 | 2.1 |
| 4 | 6.5 | 9 | 3.1 |

Waveguide and crossing loss values are based on the measurements shown in Table 3-2.

Based on the measured optical passband of the MRRs, the end-to-end 3 dB bandwidth of all optical paths through the 16×16 is about 12.25 GHz, which would be limiting for applications at rates beyond 25 GBaud. This can be resolved through a redesign of the MRR elements in the sub-switches. Experimentally, we are able to verify the bandwidth narrow-down factor with a test structure with 1, 3, and 6 cascaded MRRs in one path, each with a passband of 125 GHz. Figure 3-27b illustrates the normalized spectra of the cascaded MRRs are shown in Figure 3-27a, with the 6-cascaded-MRR case resulting in an overall passband 42 GHz. The measured bandwidths of cascaded MRRs also agree well with the bandwidth narrow-down factor relation determined by Equation (3-9). The power consumption of the Clos switch fabric is the sum

of each switch-and-select sub-switch. This results in ~60 mW per path for a total power consumption of 960 mW at a full load.
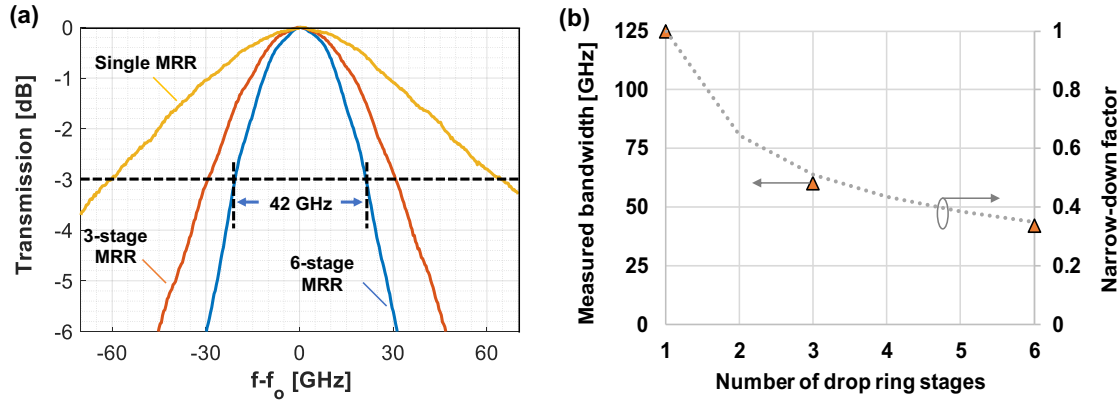


Figure 3-27: (a) Measured optical spectra for single, 3-stage, 6-stage add-drop MRRs. (b) Measured optical bandwidth with fitted narrow-down factor from Figure 3-20.

Table 3-4: Estimated total on-chip insertion loss of the 16×16 Clos switch fabric

| Source | Unit loss | Number of units | Subtotal loss [dB] | Refined unit loss | Subtotal loss [dB] |
|---|---|---|---|---|---|
| Propagation in sub-switch | 2 dB/cm | (0.8 − 2.8 mm)×3 | 0.48 − 1.68 | 1 dB/cm | 0.24 − 0.84 |
| Propagation in inter-stage shuffle | 2 dB/cm | (0.8 − 6.5 mm)×2 | 0.32 − 2.6 | 1 dB/cm | 0.16 − 1.3 |
| Crossings in sub-switch | 0.2 dB | (0 − 9)×3 | 0 − 5.4 | 0.05 dB | 0 − 1.35 |
| Crossings in inter-stage shuffle | 0.2 dB | (0 − 9)×2 | 0 − 3.6 | 0.05 dB | 0 − 0.9 |
| MRR drop | 0.5 dB | 6 | 3 | 0.2 dB | 1.2 |
| MRR through | 0.1 dB | 0 − 18 | 0 − 1.8 | 0.1 dB | 0 − 1.8 |
| Total [dB] | | | 4.7 − 15.58* | | 2.5 − 6.44* |

* Realistic range of the optical paths is less than the sum of absolute range.

### 3.3.3   Clos Switch Fabric Performance Exploration

To further understand the scalability of the proposed Clos topology with switch-and-select sub-switches, we perform a design space analysis on individual MRR SEs and identify optimal values of key design parameters. Using foundry-validated data and experimentally extracted component performance, we examine the impact of MRR radius and coupling gap on the optical power penalty and the scalability of the three-stage Clos switch fabric. We define the cross-section of the waveguide studied in this analysis as 400 nm × 220 nm. The bending loss of the waveguide is extract from past AIM Photonics MPW devices as a function of the bend radius and plotted in Figure 3-28a. As the bend radius increases, the bending loss asymptotically approaches the propagation loss of straight waveguide at ~2 dB/cm. Both the coupling gap between the MRR and the straight bus waveguide and the radius of the MRRs affect the coupling coefficient, as shown by the contour plot in Figure 3-28b based on the compact models designed in [87]. The component loss, crosstalk, and ER of the MRR cells are then aggregated into the Clos switch fabric of various radices, which allows us to determine the optical signal penalty through the fabric. Further, to maintain an end-to-end optical passband of 25 GHz after 6 cascaded MRR filters, as defined by Equation (3-9), we set the lower bound of individual MRR's 3-dB bandwidth at 70 GHz. We also show that, by employing second-order MRR filters instead of first-order MRR filters and adopting racetrack-shaped MRRs, the optimal design space of the MRR SEs can significantly increase for building moderate to high radix Clos switch fabrics.
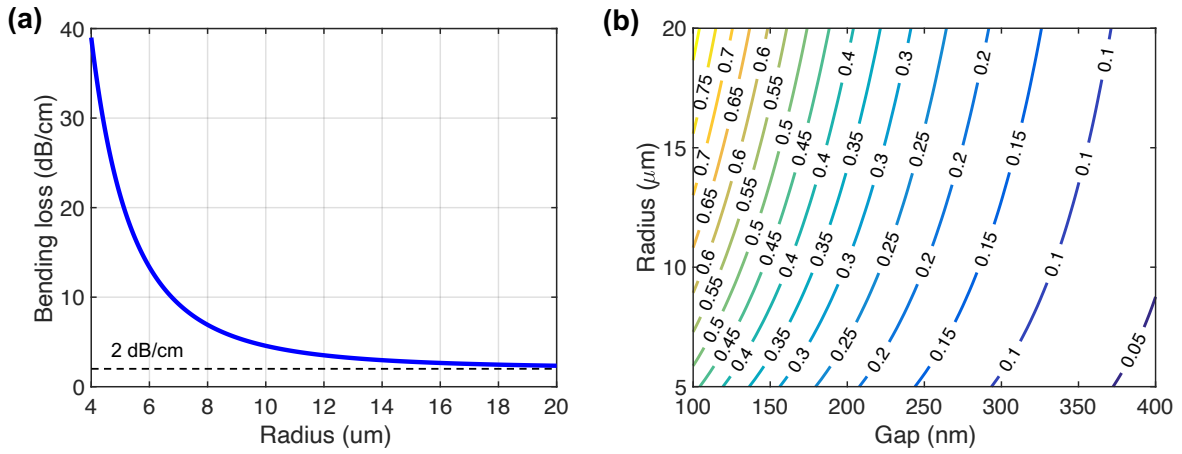
**Figure 3-28: (a) Bending loss model for 400 nm × 220 nm silicon strip waveguides. (b) Contours of electric field coupling coefficient between the bus waveguide and the MRR.**

Detailed performance contours of a first-order add-drop MRR as a function of the MRR radius and MRR-bus gap are plotted in Figure 3-29a, Figure 3-29b and Figure 3-29c, respectively, to explore the drop loss, out-of-band ER, and 3-dB bandwidth. The MRR is assumed to be critically coupled. It can be seen that the requirement on bandwidth of over 70 GHz calls for small gap sizes for strong coupling (Figure 3-29c). We target a drop loss of 0.2 dB (Figure 3-29a), while maintaining a reasonable out-of-band extinction ratio of >20 dB (Figure 3-29b). These performance bounds outline an optimal design space as well as the design tolerance as shown in Figure 3-29d, yielding a drop loss of 0.15 – 0.2 dB, out-of-band extinction ratio of 20 – 22 dB, and 3-dB bandwidth of 70 – 100 GHz.
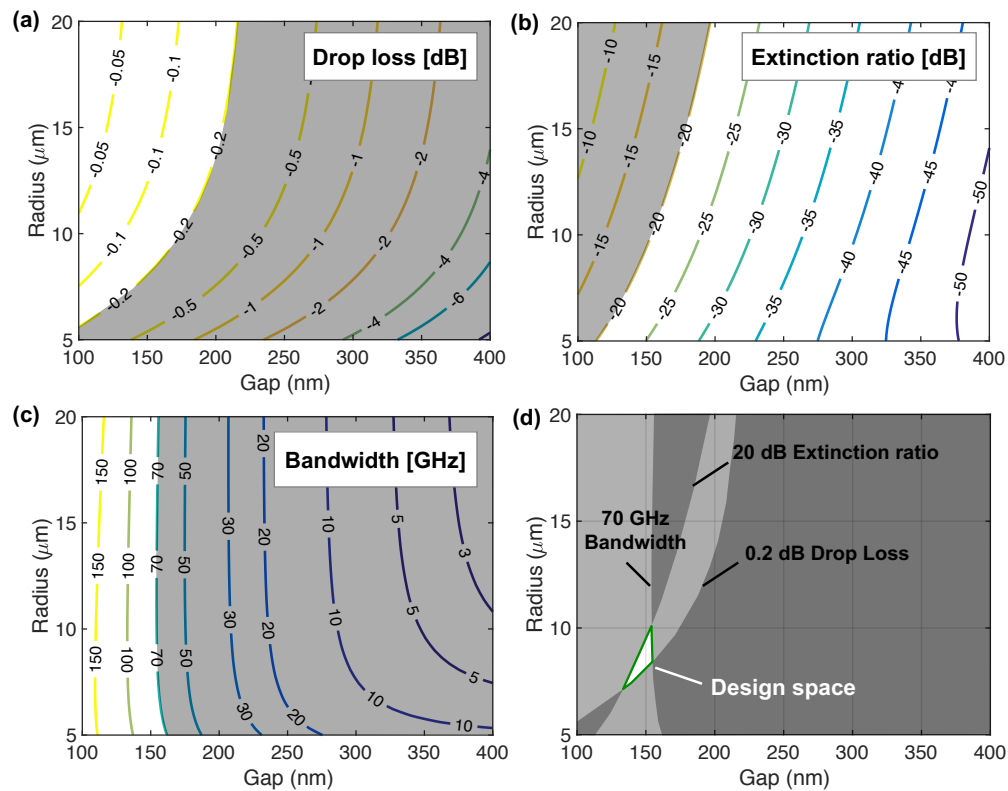


**Figure 3-29: Design space exploration of first-order add-drop MRRs based on 400 nm × 220 nm silicon strip waveguides for (a) drop-port insertion loss, (b) drop-port out-of-band ER, (c) 3-dB optical bandwidth, and (d) overall design space.**

We apply the component performance values obtained in the MRR design space (drop-port loss of 0.2 dB, ER of 20 dB) to analyze the Clos switch fabric's performance. The through-port loss of the MRR is obtained through simulation as 0.05 dB, and the waveguide propagation loss is obtained based on Table 3-2. A loss of 0.05 dB per waveguide crossing is chosen based on previously demonstrations using CMOS-compatible fabrication [91] [92] [93]. A detailed breakdown of the power penalty for the worst performing lightpath is summarized in Figure 3-30 for Clos networks of port counts 8×8, 16×16, 32×32, 64×64, and 128×128, with (n, r, m) values (2, 4, 2), (4, 4, 4), (4, 8, 4), (8, 8, 8), and (8, 16, 8), respectively, as they represent the lowest number of MRR cells in each network.

The total optical signal penalty consists of contributions from MRR drop and through loss, crosstalk-induced data signal penalty [94], waveguide propagation loss, and waveguide crossing loss. The MRR drop loss is unchanged with increasing port count because the Clos switch fabric maintains a constant of six cascaded MRR cells regardless of fabric size. The MRR through cells add a 0.05 dB per-MRR penalty and therefore scale slowly with increasing fabric size. In addition, the highly suppressed first-order crosstalk by the switch-and-select sub-switches means that the aggregate penalty due to crosstalk contributes modestly. The limiting factor on the switch fabric performance and scalability is the insertion loss of the shuffle networks due to long waveguide propagation loss and waveguide crossing loss. A potential solution to improving the waveguide shuffle penalties is using a Si/SiN multi-layer platform to avoid waveguide crossings altogether, as studied in [44] [95].
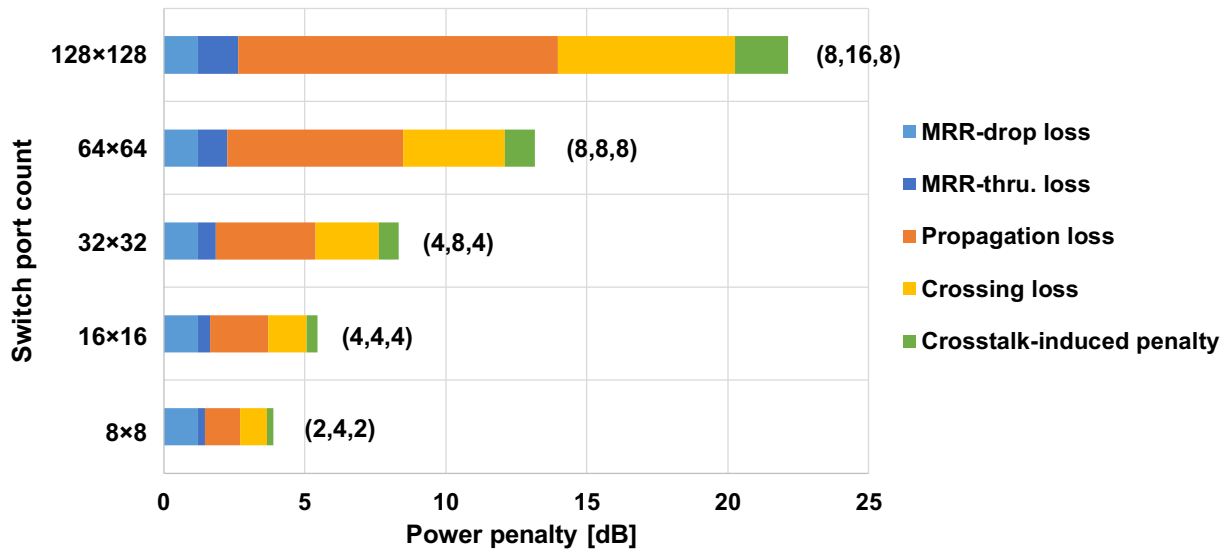
**Figure 3-30: Power penalty breakdown for the longest path in various Clos switch fabrics using first-order MRR SEs. The (*n*, *r*, *m*) parameter values of each network are shown next to the corresponding bar.**

*Second-order MRR filter cell*

The strong trade-off between ER and bandwidth of the MRR can be largely relaxed with higher-order MRR filters. Here we examine the second-order MRR cells operating under the *maximally flat* conditions [88] defined by Equation (3-10). The design spaces for the second-order MRR cell's drop loss, out-of-band ER, and 3 dB optical bandwidth in relation to the MRR-bus coupling gap and MRR radius are shown in Figure 3-31a, Figure 3-31b and Figure 3-31c, respectively. Figure 3-31b shows that, at ER greater than 20 dB, the lower bound on the MRR-bus gap is more relaxed comparing to the first-order MRR case in Figure 3-29b. In this configuration, we define the design space for drop loss $\leq$ 0.2 dB, ER $\geq$ 30 dB, and 3 dB bandwidth $\geq$ 70 GHz, as illustrated in Figure 3-31d. We restrict the MRR-bus gap's lower bound at 100 nm due to common fabrication constraints. The design space identified in Figure 3-31d yields a drop loss of 0.1-0.2 dB, out-of-band ER of 30-40 dB, and 3 dB bandwidth of 70-110 GHz.
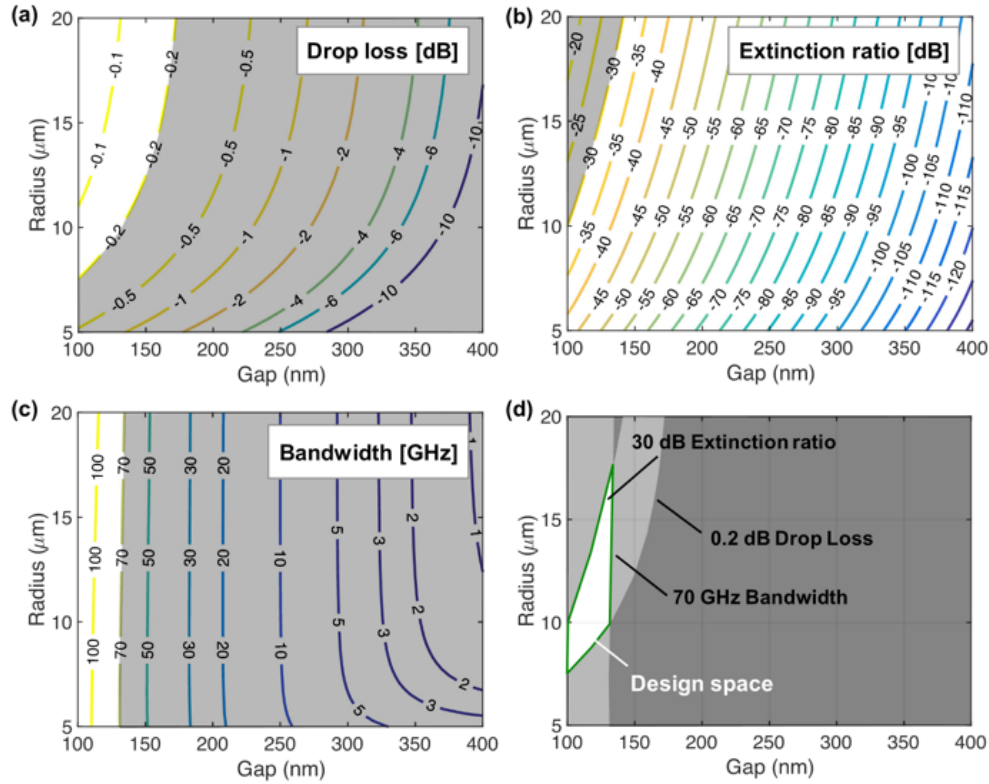
**Figure 3-31: Design space exploration of the second-order add-drop MRR on 400 nm × 220 nm silicon strip waveguides for (a) drop-port insertion loss, (b) drop-port out-of-band extinction ratio, (c) 3 dB optical bandwidth, and (d) overall design space.**

While Figure 3-31c shows the MRR filter passband improves with stronger coupling between the MRR and the bus waveguides, foundry fabrication processes typically restrict minimal feature spacing at 100 nm. To enable larger passband, the necessary stronger MRR-waveguide coupling can be achieved by utilizing the racetrack structure or a curved bus [87] to lengthen the coupling section. The schematic of a racetrack structure is shown in Figure 3-32a, with $L$ as the straight section length, $d$ as the gap, and $R$ as the radius. Figure 3-32b plots the MRR-bus coupling coefficient as a function of ring-bus gap for $L$=0, 1, 2, 3, and 4 μm. The coupling coefficient can be flexibly set by adjusting the length of the straight coupling section. We set $L$=2 μm and examine the performance contours of the second-order racetrack MRR cell's drop loss, out-of-band ER, and 3-dB optical bandwidth in Figure 3-33a, Figure 3-33b, and Figure 3-33c,

respectively. The defined design space shown in Figure 3-33d indicates a drop loss of 0.1-0.2 dB, out-of-band ER of 30-40 dB, and a significantly improved 3-dB bandwidth of 70-200 GHz.
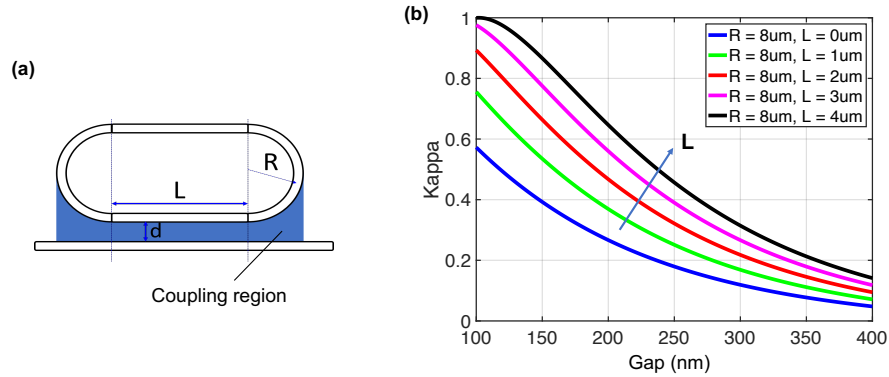


**Figure 3-32: (a) Schematic of a racetrack MRR with straight section length of *L*, gap of d, and radius of R. (b) Comparison of ring-bus coupling coefficient between *L*=0, 1, 2, 3, and 4 μm, as a function of MRR-bus gap.**
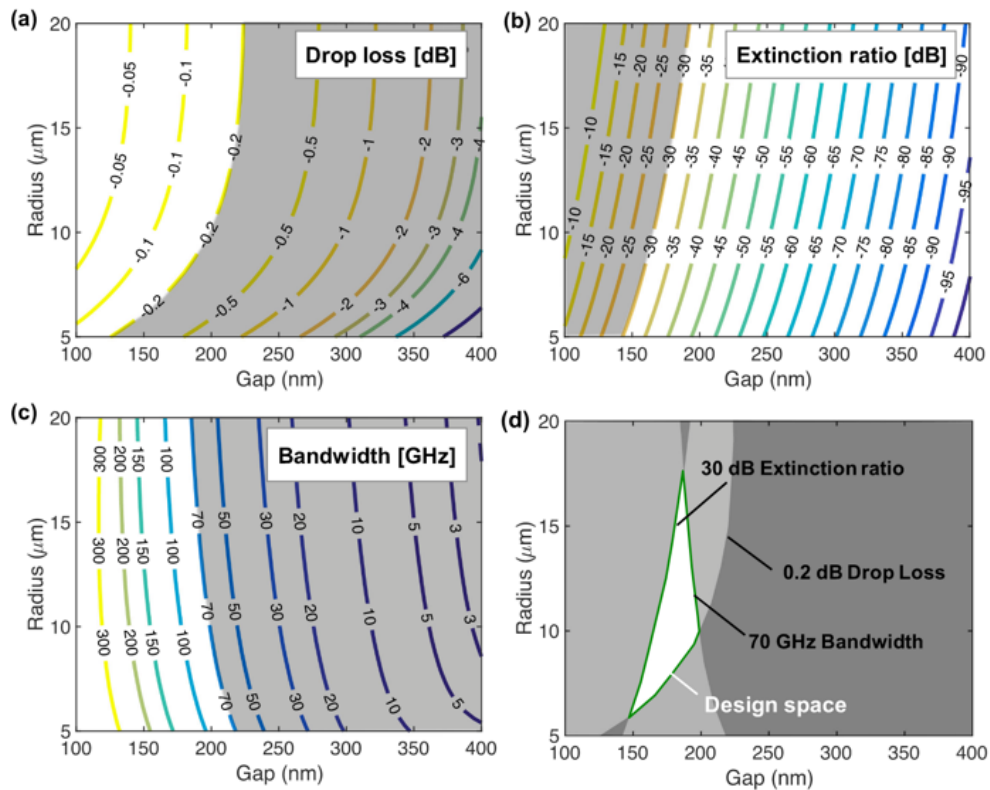


**Figure 3-33: Design space exploration of second-order add-drop MRR in racetrack structure with *L*=2 μm for (a) drop-state insertion loss, (b) out-of-band ER, (c) 3 dB optical bandwidth, and (d) overall design space.**

The through state loss is determined by simulation as 0.07 dB. We employ the same values as the first-order MRR analysis for the waveguide propagation and crossing loss. A detailed breakdown of the power penalty in relation to switch scales is shown in Figure 3-34. The 30 dB out-of-band extinction offers higher suppression on the switch crosstalk, and therefore reducing the crosstalk-induced penalty compared to Figure 3-30. While the penalty due to waveguide insertion loss remains unchanged compared to first-order designs, the MRR design space is much more relaxed for the second-order elements with higher design tolerance. Moreover, the second-order racetrack MRRs can offer broadened passbands that support an end-to-end 3-dB bandwidth of up to 70 GHz in the Clos switch fabrics.
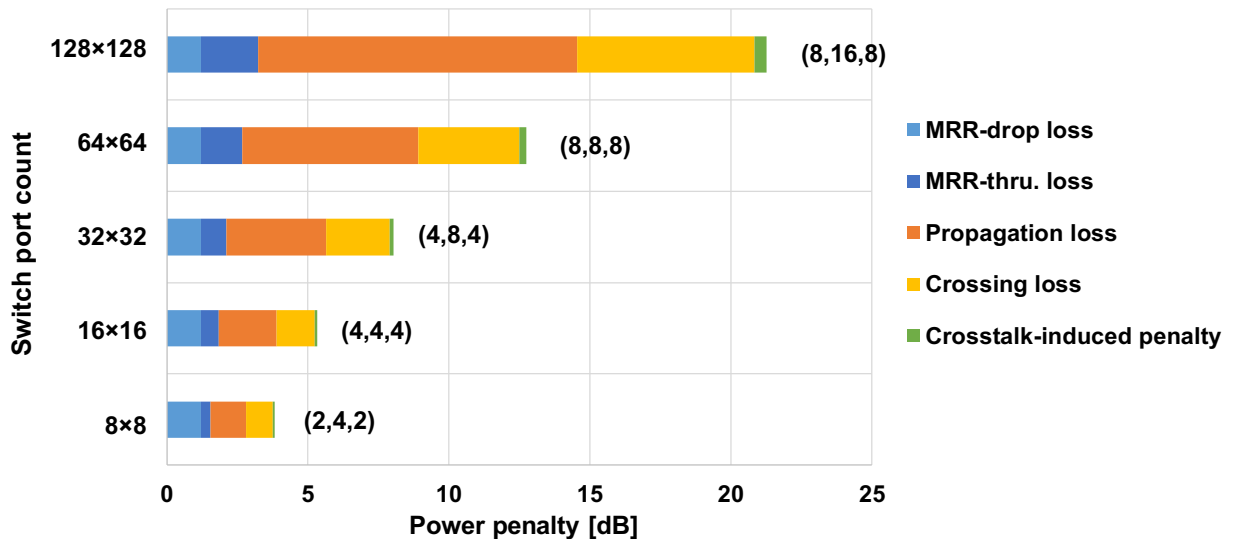


**Figure 3-34: Power penalty breakdown for the longest path in various Clos switch scales using second-order racetrack MRR units. The (*n*, *r*, *m*) parameter values of each network are shown next to the corresponding bar.**

### 3.3.4 Summary

A scalable MRR-based Clos switch fabric architecture constructed with switch-and-select sub-switches is proposed. We analyse the scalability of both SNB and RNB Clos implementations when designing moderate- and high-radix photonic switch fabrics. The Clos switch design inherits the high suppression of first-order crosstalk with the switch-and-select sub-switches, while limiting the number of cascaded SEs in a lightpath to six. Feasibility and performance analysis are performed for a 16×16 port count RNB Clos switch assembled with MRR-based 4×4 silicon switch-and-select sub-switches. Furthermore, using the foundry-validated data, a detailed design space exploration is conducted on the insertion loss, out-of-band ER, and 3-dB bandwidth for MRR-based SEs to show the feasibility of building large-scale MRR-based switch fabrics in the Clos architecture. We also highlight the higher design tolerance of the second-order racetrack MRR SEs when compared to the first-order designs.

## 3.4    Chapter Summary

In this chapter, we demonstrate key considerations and unique designs to achieve moderate- and high-radix silicon photonic switch networks with MRRs as the basic building blocks. The MRR SEs are interconnected in multi-stage networks to realize higher port count switch fabrics. To achieve high end-to-end switching performance, co-design and optimization of both the basic SEs and the switch fabric topology are necessary. We first demonstrate an 8×8 Omega switch device based on dual-MRR 2×2 SEs, which offer a combination of high switching ER, wide optical passband, and low switching crosstalk that contribute to the overall high performance of the switch. Furthermore, we analyze the design of using Clos network to interconnect silicon MRR-based switch-and-select sub-switches and achieve high-radix, high-performance switch fabrics. The switch-and-select sub-switches have the unique advantage of eliminating first-order crosstalk, while the Clos network limits the number of cascaded MRR cells to six, regardless of the total switch port count. The combination of these two benefits make possible a highly scalable architecture for silicon MRR-based switch devices. Based on experimental characterization of the switch-and-select sub-switches, we extrapolate the end-to-end performance for switch fabrics with port counts of 16×16 and above. The analysis helps use identify major contributions to the lightpath signal penalty, as well as corresponding component level optimizations to further improve the scalability for larger switch fabrics. In the next chapter, we discuss functionalization and calibration techniques for photonic switch devices.

# Chapter 4: Calibration Techniques for Photonic Switches

## 4.1    Introduction

Silicon photonic platform offers unique advantages to design highly compact, highly tuneable integrated photonic structures. These benefits, however, also contribute to reduced design tolerance in silicon photonic devices. For instance, the high index contrast of the silicon platform makes the interferometers susceptible to phase errors resulted from fabrication variations. In addition, the high TO coefficient of silicon, which offers an efficient index tuning mechanism, also means the phase-shifter-based SEs need to be actively bias-controlled to alleviate environmental drifts. Consequently, it is desirable to have a simpler switching network topology that limits both the total number of SEs on-chip and the number of cascaded SEs traversed by a lightpath in order to reduce the complexity of the monitoring and control system overhead. RNB networks, such as the Beneš topology [79], scale economically as $O(N \log_2 N)$ for $N$ inputs and need much fewer SEs compared to the crossbar topology, thus reducing the associated cost in design, fabrication, and packaging. However, RNB topologies need to coordinate multiple SEs in cascade to perform switching, and each individual SE's configuration is difficult to access due to other SEs in the same lightpath. Since fabrication variations and local phase errors can be pervasive in silicon photonic integrated circuits and lead to non-uniform switching voltages across identical switching devices [19], switch fabric calibration that achieves precise control points of each SE is particularly important to functionalize the switch devices.

Previous studies of large-scale silicon photonic switches have mostly employed built-in power monitors for calibration. For instance, about 900 germanium PDs are integrated on-chip for the switch device in [38]. While this implementation presents direct access and monitoring of each SE's switching state, the hundreds of additional components significantly increase the design complexity, packaging cost, and risk to the fabrication yield. Other switch designs [19] [54] achieve calibration by using less number of

chip-integrated PDs with additional control complexity. In [19], $N/2$ bi-directional PDs for an $N$ input device are placed in the middle stage of the fabric; by injecting optical power from either inputs and outputs of the fabric, the two halves of the switch can be calibrated respectively. This approach largely reduces the number of on-chip PDs required, but the need for injecting optical power from the switch outputs may not be compatible with applications that require unidirectional propagation. In [96], non-invasive, contactless integrated photonic probes are used for calibration, leveraging small optically induced changes in the waveguide conductance and capacitive sensing. This technique results in non-invasive photonic circuit design but requires advanced sensing circuitry to detect the minute changes in conductivity. Hence, it is important to devise a calibration methodology that eliminates the need for built-in power monitors, while achieving fast, scalable, and precise calibration of all SEs in an integrated switch fabric.

In this chapter, we introduce an approach to precisely calibrate and characterize every SE in common RNB switch fabrics with only PDs at the fabric outputs. This scheme significantly simplifies the device and system design. We review the techniques of simultaneous input activation and differential output power monitoring for switch calibration first introduced in [97] [98] [99].

## 4.2 Automated Calibration Techniques for MZI-based Switch Fabrics

### 4.2.1 Rearrangeably Non-blocking Switch Calibration without Built-in Monitors

RNB architectures such as Beneš are appealing candidates for efficient scaling of optical switch fabrics because of the reduced number of SEs needed for all-to-all connectivity. By utilizing multiple cascaded 2×2 SEs to steer a lightpath, the number of SEs in RNB architectures scales at $O(N\log_2 N)$ with $N$ inputs, rather than with $O(N^2)$ in the crossbar architecture. The scaling efficiency, however, results in a few key challenges to calibrating individual SEs. First, cascaded SEs have degenerate switching configurations that result in the same input-output mapping; hence, unlike in crossbar architecture, individual SEs' switching states cannot be determined from an input-output connection alone. Second, to distinguish Bar and Cross states of an SE, only one of its inputs can be excited, which is often difficult for SEs in the middle stages when the states of the proceeding SE layer are unknown. Third, random SE states result in power fanout in the multi-stage fabric, and therefore high-sensitivity PDs are needed to detect optical power variations at the fabric output. Hence, the design criteria of a calibration methodology include the abilities to (A) precisely direct optical power to a single input of each SE under calibration, (B) deliver sufficient optical power into the fabric to relax output PD sensitivity requirements, and (C) extend the same procedures to higher radix fabrics.

To address these challenges, we propose a method that first addresses the central stage of the RNB Beneš switch fabric and utilize multiple fabric inputs simultaneously. The switching states of individual SEs can be inferred from the power difference between subsets of the fabric outputs, completely eliminating the need for built-in power taps within the fabric. The scalability of the calibration methodology is achieved based on the recursive scaling of Beneš architecture. Figure 4-1 illustrates the working principle of the methodology. A Beneš switch with N inputs and N outputs encapsulates two N/2 × N/2 Beneš sub-switches in the middle stages. The upper N/2 fabric inputs in the input layer are redirected to the top inputs of the

sub-Beneš switches, as shown by the green arrows. Similarly, the lower N/2 fabric inputs are redirected to the bottom inputs of the sub-Beneš switches, as shown by the black arrows. If we follow the green paths in Figure 4-1B, which activates the top inputs of the sub-Beneš switches, the sub-sub-Beneš switches' top inputs will also be activated because of the recursion, which continues until the central layer of 2×2 SEs. This means that, by specifically exciting the top half of the fabric inputs, we can selectively excite only the top inputs of every SE in the central layer without any knowledge or control of the switching states of other SEs. Because of the symmetry of Beneš topology, the top outputs of the central layer SEs are connected to N/2 fabric outputs in the upper half, while their bottom outputs are connected to the N/2 fabric outputs in the lower half, regardless of the configurations of SEs in layers after the central layer. Hence, by observing the power difference between the upper N/2 outputs and the lower N/2 outputs, we can determine the switching state and the corresponding control points of each SE in the central layer. In Figure 4-1, the lightpath follows the red routes if the central layer SEs are set to Bar, or follows the blue routes if they are set to Cross. This procedure can be automated by searching for the maximum difference between upper and lower fabric output power levels, which indicates the Bar control points for the central layer SEs, while a search of minimum difference can determine the Cross control points. Since the fabric output power difference is monotonic to individual SE's output power difference between the red and blue routes, all SEs in the middle layer can be tuned together to Bar or Cross in a single optimization step.
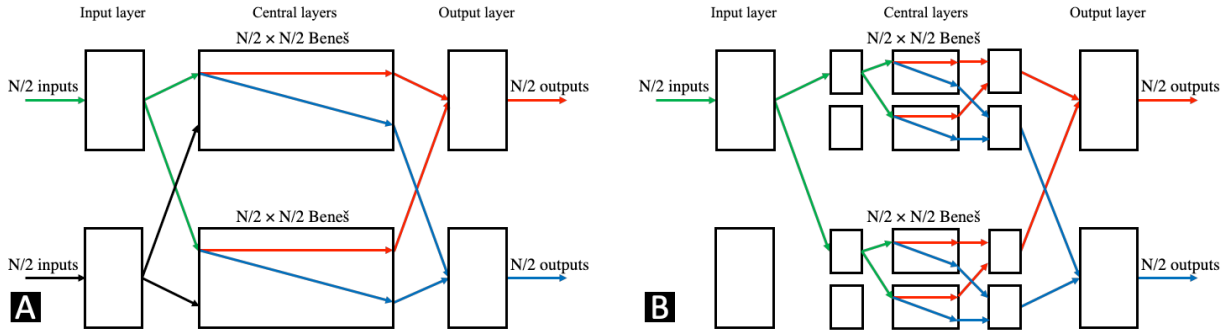
**Figure 4-1: (A)** Illustration of the top half of inputs routed to top (red path) or bottom (blue path) set of outputs by the sub-Beneš. **(B)** Recursive structure of Beneš architecture allows upper fabric inputs to directly address the top input of SEs in the central layer.

After the central layer SEs' control points are determined, they can be used to excited single inputs of SEs in the succeeding layers, as illustrated in Figure 4-2A. The difference in output power levels is used again to determine the switch control points of these SEs. Once the output layer SEs are calibrated, the second half of the fabric can be controlled to steer lightpaths and enable calibration of the SEs in the input layer of the fabric, as shown in Figure 4-2B. Individual input of the input layer SEs can be excited and routed to the fabric outputs for measurements. Power level difference between the red and blue fabric outputs in Figure 4-2B is used to infer switching states of the input layer SEs. The specific implementation of the calibration procedure for an 8×8 switch is shown in Figure 4-3, in which the inactivated SEs are grey; SEs with undetermined switching states are blank; and the red and blue lines show the Bar and Cross paths, respectively, of the SEs being calibrated in each step. The activated fabric inputs are indicated by black arrows. Power difference is determined between the sum of red output power levels and the sum of the blue output power levels.
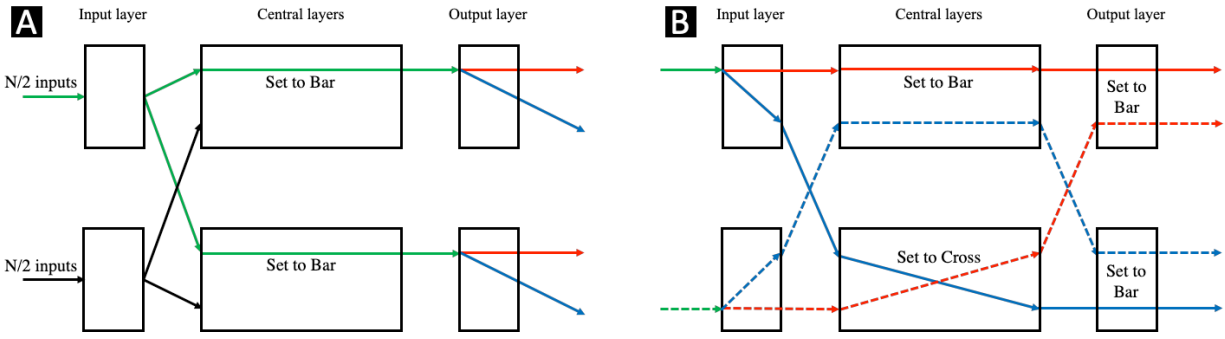
Figure 4-2: (A) Calibration of the output layer SEs while controlling the central layer. (B) Calibration of the input layer SEs while controlling the central layer and the output layer.
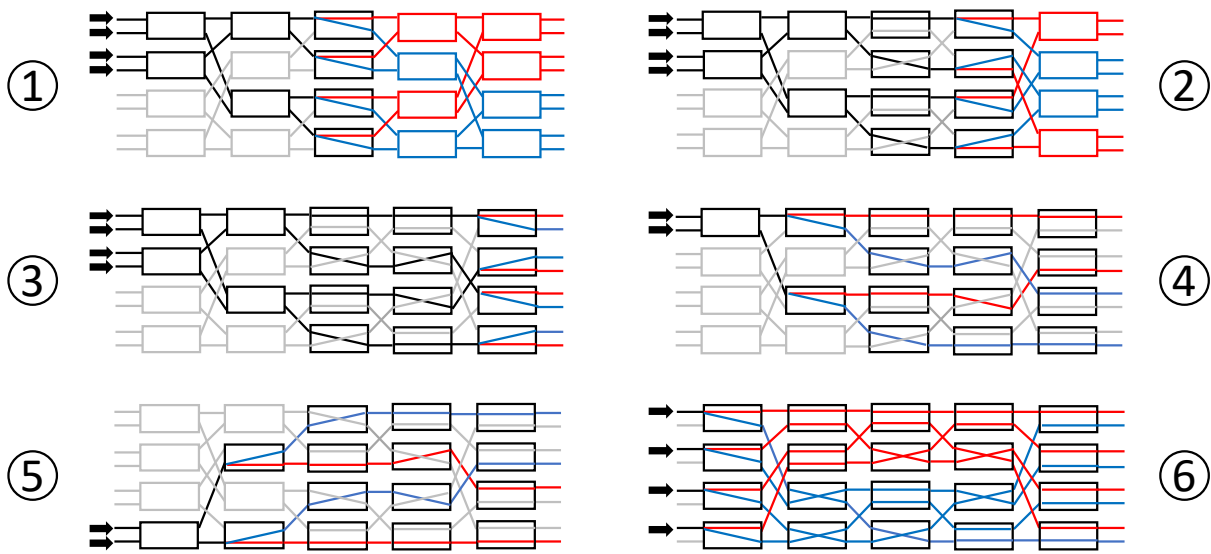


Figure 4-3: Step-by-step calibration procedure for an 8×8 Beneš switch. 4 SEs in the central layer are calibrated in Step 1; 4 SEs in the central+1 layer are calibrated in Step 2; 4 SEs in the output layer are calibrated in Step 3; 4 SEs in the central-1 layer are calibrated in Step 4-5; 4 SEs in the input layer are calibrated in Step 6.

The calibration of each SE requires a sweep of its control voltage to identify Bar and Cross points indicated by differential power levels at the output. Parallelism enabled by the architecture allows multiple SEs to be calibrated simultaneously – if the set of SEs in the same layer require distinct outputs to monitor their respective differential power levels, then they can be calibrated in parallel. For instance, in Step 6 shown in Figure 4-3, all 4 SEs in the input layers can be tuned simultaneously using different outputs in order to expedite the process. Switch architectures derived from Beneš can also implement the calibration

methodology with minor modifications. These architectures resemble a partial Beneš to provide

connectivity with blocking constraints. Figure 4-4 and Figure 4-5 show how the calibration methodology

can be applied to 8×8 Banyan and 8×8 Butterfly architectures. By exciting one single input at a time for

every SE in the fabric and inferring switching states without built-in PDs, the proposed calibration

methodology is an efficient procedure to determine all operating points of the switch device. It is also

important to note that the calibration procedures introduced here apply to 2×2 SEs regardless of their

strcutures and switching mechanisms. In the following section, we show how the calibration methodology

is implemented to enable specific controls and driving schemes of the photonic switch.
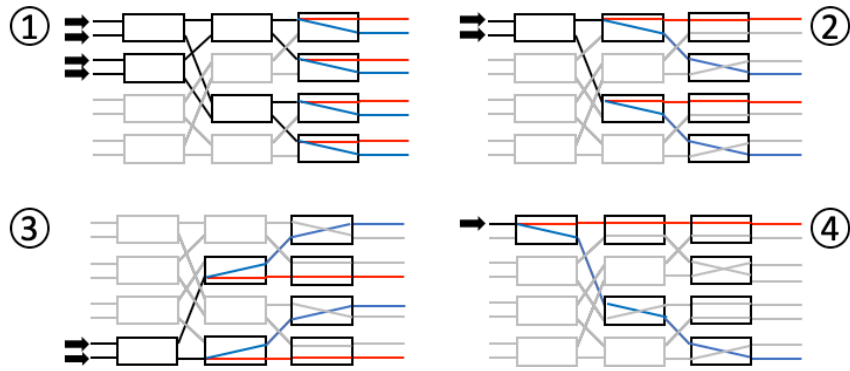


**Figure 4-4: Step-by-step calibration procedure for an 8×8 Banyan switch. Step 4 is generalized to all SEs in the input layer for brevity.**
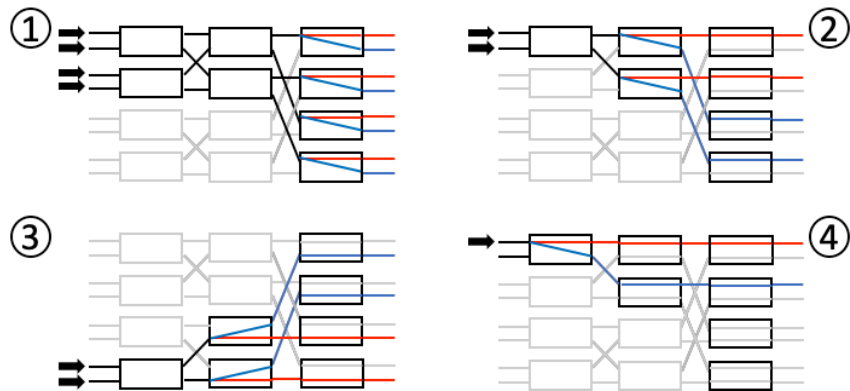


**Figure 4-5: Step-by-step calibration procedure for an 8×8 Butterfly switch. Step 4 is generalized to all SEs in the input layer for brevity.**

### 4.2.2   Demonstration of Rapid Switch Calibration

We experimentally demonstrate the methodology applied to a 4×4 silicon photonic switch designed and fabricated on the OpSIS MPW platform. The packaged switch device is shown in Figure 4-6A. Six MZI SEs are arranged in the Beneš architecture to provide all-to-all connectivity between 4 inputs and 4 outputs. The switch device is driven directly by FPGA-controlled DACs. Figure 4-6B illustrates the experimental setup consisting of the FPGA system programmed with the calibration logic, 18 DAC channels applying voltages to the phase shifters of all SEs, 4 Thorlabs PDA10CS PDs at the fabric outputs, and 4 ADC channels to provide real-time feedback of the output power levels. An x86 Linux system provides high-level programming interfaces to the underlying hardware in the FGPA and the DAC and ADC channels. Shown in Figure 4-6B, each of the SEs has two EO phase shifters and two TO shifters – one of each on either MZI arm. Balanced push-pull drive scheme can be realized on this device by applying a quadrature biasing for each MZI using one TO shifter and performing switching by turning on the EO shifters alternatively. Push-pull drive scheme at the quadrature point reduces the switching phase swing of both states to $\pi/2$. The EO shifter on one arm is used to switch to Bar, and the EO shifter on the other arm is used to switch to Cross. In the following analysis, we designate the MZI arms as Bar arm and Cross arm based on their respective EO control functions.

87
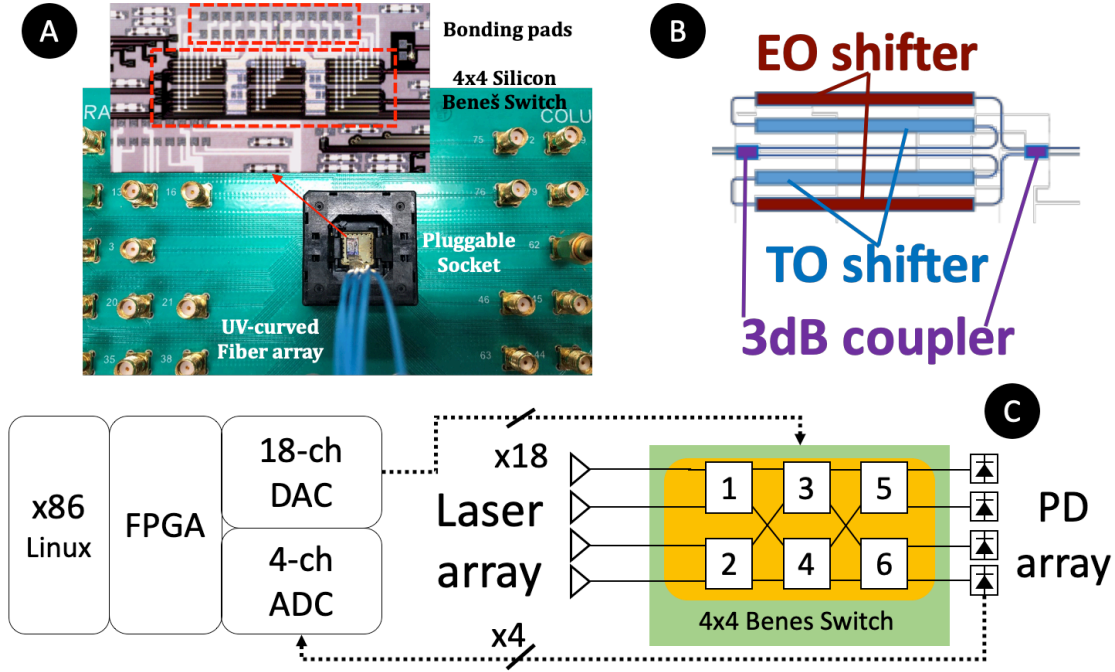
**Figure 4-6: (A)** The packaged silicon photonic switch chip wire-bonded to a PCB breakout board and grating coupled. **(B)** Schematic of a single MZI SE showing placements of EO and TO phase shifters. **(C)** Schematic of SE arrangement in Beneš and calibration setup showing the FPGA control system, laser and PD array.

**Algorithm 4-1: Initial Control Point Calibration Procedure**

---

1: <u>Notation:</u>
2: $se$(L/R): the left or right arm bias of the element $se$;
3: (Out1−Out2): the difference between Out1 and Out2 power in [mW].
4: <u>Procedure:</u>
5: **for** $se$ in [SE3, SE4, SE5, SE6, SE1, SE2] **do**
6:     **if** $se$ == SE3 or SE4 **then**
7:         turn on In1 and In2;                 ▷ Set laser input
8:         $diff = $ (Out1+Out2−Out3−Out4)   ▷ Define differential power
9:     **else if** $se$ == SE5 or SE6 **then**
10:         set SE3 to Bar, SE4 to Cross;           ▷ Set known SE states
11:         $diff = $ (Out1+Out4−Out2−Out3)
12:     **else if** $se$ == SE1 **then**
13:         turn off In 2, set SE5, SE6 to Bar;
14:         $diff = $ (Out1−Out4)
15:     **else if** $se$ == SE2 **then**
16:         turn on In 4;
17:         $diff = $ (Out2−Out3)
18:     **end if**
19:     sweep $se$(L) bias; find $V_L^{bar}$ at $max(diff)$ and $V_L^{cross}$ at $min(diff)$; set $se$(L) to $zero$ bias;
20:     sweep $se$(R) bias; find $V_R^{bar}$ at $max(diff)$ and $V_R^{cross}$ at $min(diff)$; set $se$(R) to $zero$ bias;
21:     $V_{se}^{bar} = \min(V_L^{bar}, V_R^{bar})$
22:     $V_{se}^{cross} = \min(V_L^{cross}, V_R^{cross})$
23: **end for**

---

Algorithm 4-1 shows the calibration procedure of the 4×4 Beneš device. For every SE following the methodology sequence, an initial pair of push-pull switching control points are determined by performing a voltage sweep on either arms' EO bias between 0.5-2 V in 2 mV steps, while keeping the TO bias and the other arm's EO bias at 0 V. The EO sweep range is bracketed between the p-i-n junction turn-on voltage and $V_\pi$, thus guaranteeing reduced scanning range and non-cyclic solutions. The output power differences are monitored by the PDs at the fabric outputs and computed in real-time to identify the Cross and Bar points. The arm with the higher maximum differential power is designated as the Bar arm, and the arm with the lower minimum differential power is designated as the Cross arm. Once the initial switching points without TO bias are found, the quadrature point can be determined by linearly ramping the TO bias while performing fast local search on the EO bias – this ensures that as the TO bias changes, new EO switching points are updated. As described by Algorithm 4-2, quadrature bias is found when the Bar and Cross EO switching voltages coincide [98]. Overall, the calibration process completes in one minute. Figure 4-7 compares the initial switching point and the balanced push-pull points for all SEs after the calibration. The process of finding the quadrature inherently accounts for MZI arms' phase variations and environmental drift at the time, and therefore can enable in-situ recalibration of the switch device. We characterize the effect of balanced calibration by showing the power levels at all outputs as Input 1 signal is switched to Outputs 1–4 respectively, as shown in Figure 4-8. Comparing to unbalanced operations, the balanced operations of the switch improve the end-to-end loss of the switched signal by 0.6 dB on average and reduce crosstalk on 10 out of the 12 paths. The significant crosstalk increase at Output 1 when Input 1 is switched to Output 2 is likely due to the abnormally high carrier absorption loss in one arm of SE5, which becomes more apparent after the balancing.

## Algorithm 4-2: Balanced Push-Pull Calibration Procedure of a Single SE

---

1: <u>Notation:</u>
2: $V_E^{se}$(B/C): EO voltage corresponding to $se$'s Bar arm or Cross arm in push-pull;
3: $V_T^{se}$: TO bias applied to the arm with lower initial EO Bar or Cross voltage
4: $V_T^{max}$: maximum range for $V_T^{se}$
5: $V_T^{step}$: ramping step for $V_T^{se}$
6: $V_E^{step}$: dither step for $V_E^{se}$(B/C)
7: $V_E^{diff}$: Equalized EO bias tolerance
8: $P_{out}(V_T^{se})$: differential power output associated with $se$ at a given ($V_T^{se}$), with maximum at Bar and minimum at Cross
9: <u>Procedure:</u>
10: **for** $V_T^{se} = 0 : V_T^{step} : V_T^{max}$ **do**
11:      fast dither $V_E^{se}$(B) in range $[V_E^{se}(B) - V_E^{step}, V_E^{se}(B) + V_E^{step}]$
12:      update $V_E^{se}$(B) $s.t.$ $P_{out}(V_T^{se})$ is maximized
13:      fast dither $V_E^{se}$(C) in range $[V_E^{se}(C) - V_E^{step}, V_E^{se}(C) + V_E^{step}]$
14:      update $V_E^{se}$(C) $s.t.$ $P_{out}(V_T^{se})$ is minimized
15:      **if** $abs(V_E^{se}(B) - V_E^{se}(C)) \leq V_E^{diff}$ **then**
16:          break
17:      **end if**
18: **end for**

---
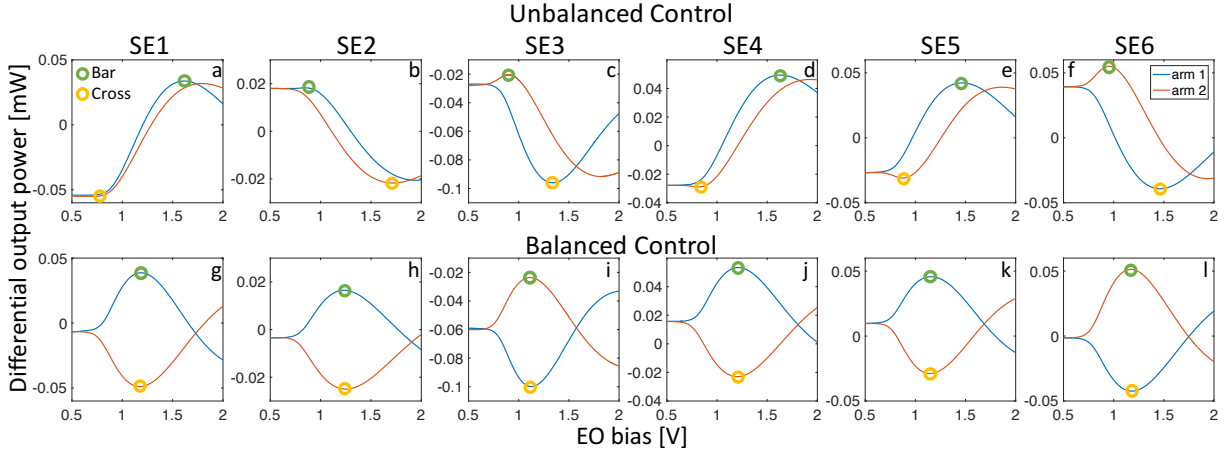


**Figure 4-7: (a-f) Differential power level of the 6 SEs in the 4×4 switch without TO bias, showing Bar and Cross points as maxima and minima, respectively. (g-l) Differential power level of the 6 SEs in the 4×4 switch biased to quadrature to enable balanced push-pull points.**
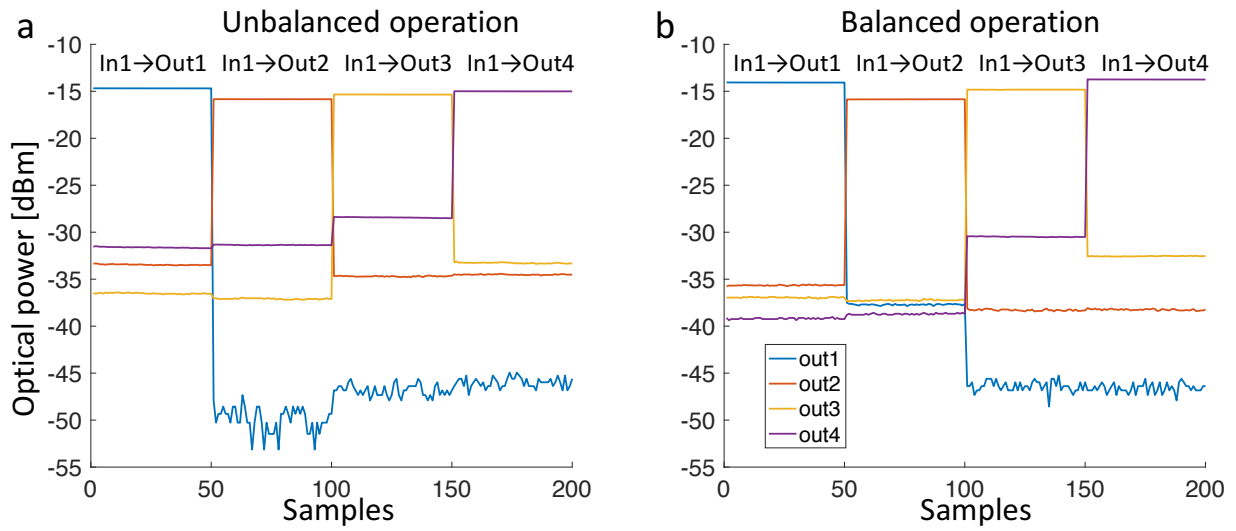
**Figure 4-8: Optical power from Input 1 is switched to Outputs 1-4 sequentially under unbalanced operations (a) and balanced operations (b).**

## 4.3    Crosstalk-aware Calibration Techniques for Switch Fabrics

It is possible to further reduce the worst-case crosstalk levels in MZI SEs using advanced switch calibration with fast characterization [99]. Typically, the crosstalk leakage power in MZI SEs is the result of phase error and intensity imbalance between the MZI arms. Through the expedient calibration technique introduced in the previous section, the phase differences between MZI arms required for switching can be determined precisely. The intensity imbalance between the two MZI arms, however, still remains as the main cause of crosstalk – the uneven field intensities between the two arms of the MZI lead to incomplete destructive interference and power leakage. The intensity of the crosstalk can be denoted by Equations (4-1) – (4-4) using a transfer matrix method for the MZI:

The Bar state crosstalk power at outputs 1 and 2 are, respectively:

$$\left[\left(\kappa_2 - \kappa_1\kappa_2\right)\alpha_U{}^2 + \left(\kappa_1 - \kappa_1\kappa_2\right)\alpha_L{}^2 - 2\sqrt{\left(\kappa_1 - \kappa_1{}^2\right)\left(\kappa_2 - \kappa_2{}^2\right)}\alpha_U\alpha_L\right]\left|I_2\right|^2 \ , \quad (\textbf{4-1})$$

$$\left[\left(\kappa_1 - \kappa_1\kappa_2\right)\alpha_U{}^2 + \left(\kappa_2 - \kappa_1\kappa_2\right)\alpha_L{}^2 - 2\sqrt{\left(\kappa_1 - \kappa_1{}^2\right)\left(\kappa_2 - \kappa_2{}^2\right)}\alpha_U\alpha_L\right]\left|I_1\right|^2 \ ; \quad (\textbf{4-2})$$

and the Cross state crosstalk power at outputs 1 and 2 are, respectively:

$$\left[\kappa_1\kappa_2\alpha_U{}^2 + \left(1 - \kappa_1 - \kappa_2 + \kappa_1\kappa_2\right)\alpha_L{}^2 - 2\sqrt{\left(\kappa_1 - \kappa_1{}^2\right)\left(\kappa_2 - \kappa_2{}^2\right)}\alpha_U\alpha_L\right]\left|I_1\right|^2 \ , \quad (\textbf{4-3})$$

$$\left[\left(1 - \kappa_1 - \kappa_2 + \kappa_1\kappa_2\right)\alpha_U{}^2 + \kappa_1\kappa_2\alpha_L{}^2 - 2\sqrt{\left(\kappa_1 - \kappa_1{}^2\right)\left(\kappa_2 - \kappa_2{}^2\right)}\alpha_U\alpha_L\right]\left|I_2\right|^2 \ , \quad (\textbf{4-4})$$

where $\kappa_1$ and $\kappa_2$ are the power coupling coefficients of the in- and out-coupler; $I_1$ and $I_2$ are input fields at inputs 1 and 2; $\alpha_U$ and $\alpha_L$ are the intensity scaling factors of the upper and lower MZI arms accounting for the loss. In fast silicon photonic MZI with EO phase shifters, we assume the propagation loss is consistent between the arms, and the only difference in field attenuation is the result of tuning the EO shifter on and off. Hence, the intensity imbalance in the MZI originates from two interacting mechanisms – coupler variations [100] and EO shifter attenuation. The coupler variations result in static imbalance in the MZI and are due to imperfect 50/50 splitting and combining due to design and fabrication variations [101]. The EO

shifter attenuation, however, results in a dynamic imbalance between the SE states. In a push-pull drive scheme for the MZI, the EO phase shifter on each arm alternatively turns on for its respective switching state. The arm with its EO shifter turned on experiences an additional free-carrier-induced loss compared to the other arm. This means there is a nonzero crosstalk floor for EO MZI SEs [102] due to switching dynamics even if all other component structures are perfectly fabricated.

The interaction between the dynamic intensity variations due to the shifters and the static intensity imbalance due to the couplers means that Cross and Bar states have different crosstalk levels when operated under push-pull. In one switching state, the worst-case SE crosstalk occurs when the static imbalance is exacerbated by the shifter loss; conversely, in the other switching state, the shifter loss counteracts the static intensity imbalance and reduces the switching crosstalk, as illustrated in Figure 4-9A. In the case study shown in Figure 4-9B, it is evident that the SE's worst-case crosstalk is reduced when the Bar state crosstalk and Cross state crosstalk coincide, which can be achieved by tuning the TO bias slightly away from quadrature. Intuitively, this precise TO biasing effectively modifies EO induced loss levels for both switching states, and even out the static imbalance of the couplers. Figure 4-9C and Figure 4-9D compare the worst-case crosstalk over a range of coupler coupling coefficient values for quadrature biasing and crosstalk-aware TO biasing, and show that the latter results in a tremendous relaxation of the performance margins of the in-coupler, which would be particularly effective in optimizing SE crosstalk performance post-fabrication.
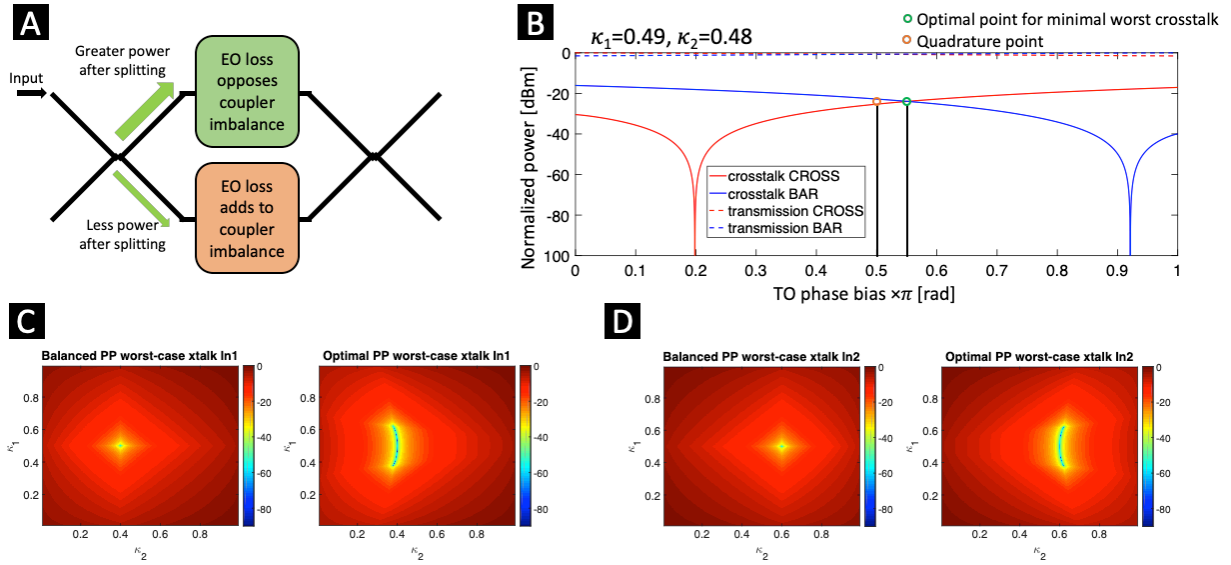
**Figure 4-9: (A)** Illustration of the interactions between uneven splitting of in-coupler and the switching loss of EO shifters in both MZI arms. **(B)** Case study illustrating crosstalk levels of an MZI SE under quadrature bias and crosstalk-aware bias, given a pair of in-coupler and out-coupler power coupling coefficients $\kappa_1$ and $\kappa_2$. **(C-D)** Comparison of the worst-case crosstalk heatmaps for a signal from Input 1 **(C)** and Input 2 **(D)** between quadrature biased balanced push-pull switching and crosstalk-optimal operation at wide range of coupler coupling coefficient values, showing the proposed methods can drastically increase tolerance margins of $\kappa_1$ while maintaining low worst-case crosstalk.

The optimization for TO biasing can be efficiently incorporated into the calibration methodology by exciting a single input of every SE and inferring signal and crosstalk power levels based on subsets of the fabric outputs – the crosstalk-optimal bias is achieved when the crosstalk levels of both Cross and Bar states converge in a close-loop control procedure. We validate the crosstalk optimal tuning with the experimental system shown in Figure 4-6 and observe an averaged reduction of worst-case crosstalk across all six SEs by 1.72 dB over balanced push-pull with quadrature TO bias. Figure 4-10 compares the crosstalk levels for the SE4 with the largest reduction of 3.8 dB over balanced push-pull. We also highlight the speed of the fully automated calibration process, which determines both the EO switching voltages and the crosstalk optimal TO bias for all SEs in the experimental switch under 2 minutes.
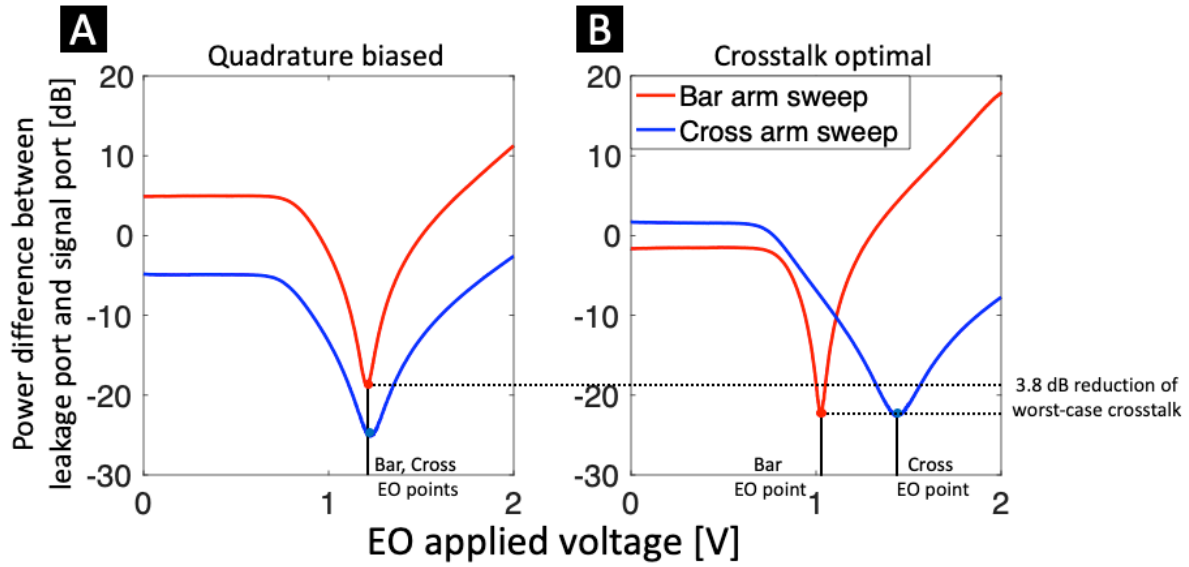
**Figure 4-10: Comparison among power level differences between the signal output and leakage output observed for SE4 when (A) quadrature TO bias is applied to achieve balanced EO operating voltages, and (B) TO bias is optimized to reduce the worst-case crosstalk levels.**

## 4.4    Chapter Summary

Silicon photonic switch fabrics present a highly integrated, low cost, and small footprint solution with high affinity to scalable CMOS manufacturing. As the number of switching elements increase with the port count and connectivity of the fabric, post-fabrication calibration and characterization of the switch are crucial to ensure the performance and usability of the device. In this chapter, we present an efficient, highly scalable calibration methodology to determine the exact switching points of every MZI SE without the need for intermediate power taps, which drastically reduces complexities in both design and testing. Leveraging the redundancy in the Beneš and other derivative RNB architectures, we emphasize key design aspects of the calibration methodology with central-layer-first sequence, simultaneous fabric inputs, and differential power monitoring at the fabric outputs. We demonstrate experimentally the application of the methodology to determine the control points of a 4×4 silicon photonic Beneš switch, achieving both quadrature biasing for balanced push-pull drive and crosstalk-optimal biasing to minimize worst-case crosstalk at each SE. The systematic methodology introduced here can significantly simplify photonic switch fabric PIC design, expedite the functionalization process of the switch devices, and optimize component level performance post-fabrication. In the next chapter, we discuss how routing strategies can be devised for RNB switching topologies to avoid worst-case performance and improve end-to-end power penalties for integrated silicon switch fabrics.

# Chapter 5: Optical Switching Topologies and Smart Routing

## 5.1    Introduction

As discussed in Chapter 4, expedited calibration and characterization of SEs in a silicon photonic switch fabric can determine precise and performance-optimal control points for the switch building blocks. The insight gained in the SE operation is key to devise system-level control and routing for the switch network. Performance-aware routing is necessary for photonic switch devices because, in contrast to electronic switches, where data signal packets are stored, retimed, and error-corrected, optical signals traversing a photonic switch experience the aggregated optical impairments along the optical path. The level of impairments depend on the type of SEs implemented – thermally or electrically actuated MZIs [39], MRRs [44], or MEMS [41], the number of SEs traversed by the lightpath, and the specific states of SEs as defined by the routing control of the switch network. In previous chapters, we have explored methods to design high performance SEs and scalable switch topologies. In this chapter, we examine how to design routing strategies tailored for the switch network topologies in order to elevate the end-to-end performance of the integrated switch device when scaling up the switch port count.

While a number of calibration techniques have been studied to facilitate fast and accurate characterization of optical switching circuits [103] [96] [104] [97] [105] [106] [98] [99] [107], research on photonic switch routing control has been sparsely reported. The routing algorithms in electronic switches are primarily developed to resolve viable paths with least contentions [80] [108] [109]. However, the optical integrated switches, as repeaterless fabrics where signal's amplitude and timing margins are not restored, generally have path-dependent performance. Such path-dependent variation normally lies in imperfect fabrication and design limitations, which could cause a discrepancy in the different switching states of the elementary cells, limit the switch performance, and lead to excess loss in waveguide shuffling. A photonic routing strategy should take these factors into consideration. Previous works on this topic typically consider

device-specific metrics when performing routing, such as loss and IPDR improvement in Clos network [110] and crosstalk reduction in dilated Banyan topology [111]. Here, we present a generic methodology to for photonic routing strategies that addresses key considerations in both the switch topologies and fabric-wide power penalty optimization [112]. We review the commonly applied optical switching architectures and examine the redundancies in the switch permutations that can be exploit for photonic switch routing control. Via both simulation and experimental demonstrations, we analyze how performance metrics can applied to devise routing strategies that reduce the performance variations of switch lightpaths and avoid worst-case switch configurations due to fabrication variation.

## 5.2     Switch Topologies and Switching States

A switching device's performance depends critically on the selection of topology, which dictates a switch's blocking characteristics, crosstalk suppression, total number of SEs, and number of cascaded stages [9]. Some of the classical switch architectures, such as crossbar, Banyan, Clos, and Beneš networks, are adopted from electronic switch network designs, while the others are made by pioneers in optical switch fabrics to offset the limitations of photonic integration technologies. For instance, N-stage planar architecture was proposed to eliminate waveguide crossings; PILOSS network was designed to achieve a loss uniformity across all paths; and dilated networks were used to completely cancel the first order crosstalk. Some of the most commonly applied optical switch architectures based on 2×2 SEs, such as Banyan, N-stage planar, crossbar, Beneš, PILOSS, and switch-and-select are shown in Figure 5-1a – Figure 5-1f, and compared in Table 5-1 for their respective blocking characteristics, order of crosstalk, total number of SE, and the number of cascaded switching stages. We define the global switching states as the total number of switch configurations that lead to valid connection between an input and an output, and switch permutations as the unique connection mappings between all inputs and all outputs. The number of global switching states is dictated by the number of SEs and the way in which they are interconnected; the number of switch permutations is dictated by the port count of the switch fabric and its blocking characteristics. In this section, we investigate the global switching states for various switch topologies in relation to the number of switch permutations.
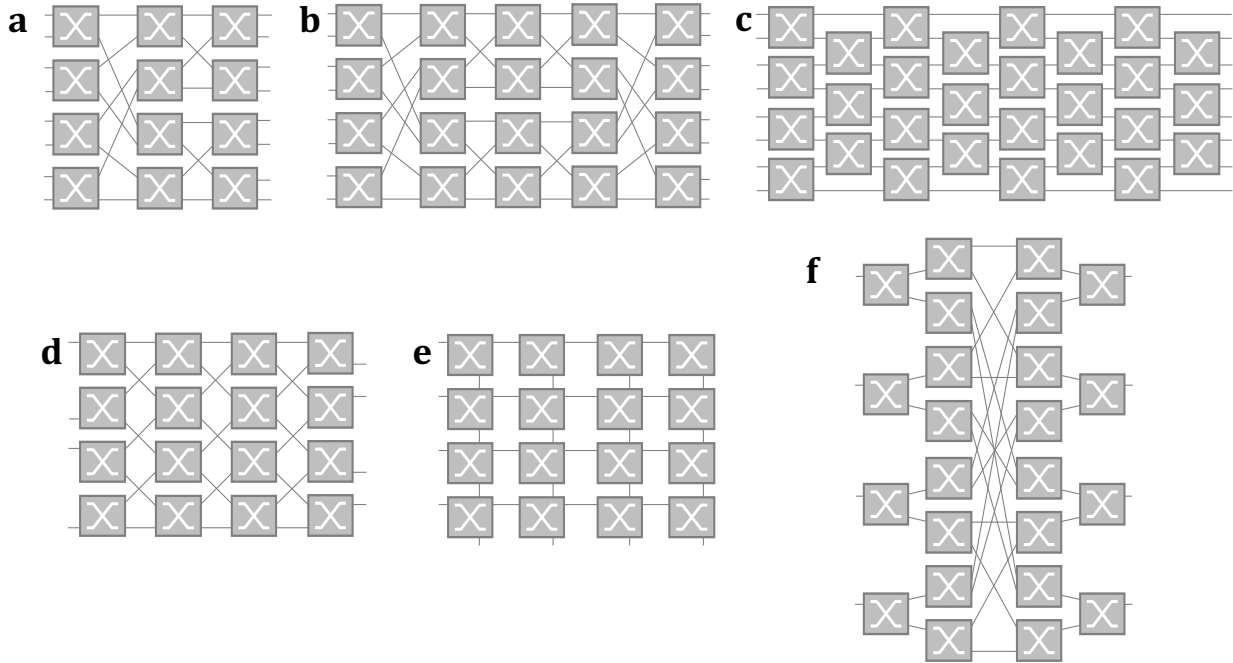
**Figure 5-1: Schematic of switch architectures: (a) Banyan, (b) Beneš, (c) N-stage planar, (d) PILOSS, (e) crossbar, and (f) switch-and-select. Note that (a), (b), and (c) have 8×8 port count, while (d), (e), and (f) have 4×4 port count.**

**Table 5-1: Common optical switch architectures**

| Blocking Characteristics | Architecture | Order of Crosstalk | Total Number of SEs | Max Number of Switch Stages | Global Switching States |
|---|---|---|---|---|---|
| Blocking | Banyan | First | $\frac{N}{2}log_2N$ | $log_2N$ | $2^{\frac{N}{2}log_2N}$ |
| RNB | Beneš | First | $\frac{N}{2}(2log_2N-1)$ | $2log_2N$-1 | $2^{\frac{N}{2}(2log_2N-1)}$ |
| RNB | N-stage planar | First | $\frac{N}{2}(N-1)$ | $N$ | $2^{\frac{N}{2}(N-1)}$ |
| WSNB | PILOSS | First | $N^2$ | $N$ | N! |
| WSNB | Crossbar | First | $N^2$ | $2N-1$ | N! |
| SNB | Switch-and-Select | Second | $2N(N-1)$ | $2log_2N$ | N! |

*Banyan network*

Banyan switch fabric, which was first proposed for computer networks, is an attractive candidate for optical switching applications. A Banyan network is defined as a class of multistage networks that have no path diversity and the minimum diameter for a fully connected network. Banyan network also requires the least number of SEs to provide all-to-all connectivity, which translates to simplified switch fabric design with minimum footprint. For an N×N Banyan switch (with N being powers of 2), the total number of 2×2 SEs needed is $\frac{N}{2}log_2N$, which leads to $2^{\frac{N}{2}log_2N}$ global switching states. Since the number of global switching states is less than the maximum number of the switch permutations for an N×N network, N!, Banyan is a blocking network, meaning that while any input port can connect to any output port, multiple connections can be in conflict. Due to a lack of path redundancy between an input port and an output port, each switch path maps to just a single routing configuration that can be pre-determined.

*Rearrangeably non-blocking networks*

RNB switch fabrics improve upon the Banyan topology by offering non-blocking connectivity, while still leveraging multi-stage structures to reduce the total number of SEs and fabric complexity. By definition, RNB networks require rerouting of existing connections to accommodate new input-output connections. The Beneš topology is a popular topology of choice for large scale silicon photonic switch devices, because it requires the minimum number of SEs to construct a non-blocking N×N network while retaining efficient logarithmic scaling with increasing switch port count. The number of cascaded stages in a Beneš topology is $2log_2N$-1, with a total of $\frac{N}{2}(2log_2N-1)$ cells and $2^{\frac{N}{2}(2log_2N-1)}$ global switching states. Since the number of global switching states is greater than the maximum number of switch permutations, N!, in a non-blocking network, path redundancy is guaranteed and routing control is needed.

Another example of the RNB topology is the N-stage planar network, which is the topology of choice for some works in modest-scale switch fabrics [113] [114]. The main benefit for N-stage planar

topology is the removal of waveguide crossings in the switch fabric, which helps simplify photonic circuit design and eliminate optical impairments from waveguide crossings [115]. The N-stage planar architecture has a total of $\frac{N}{2}(N-1)$ 2×2 SEs and can lead to a much larger switch device footprint than Beneš. In addition, optical paths traversing a N-stage planar fabric can see different numbers of SEs in the range of N/2 to N, which leads to highly path-dependent optical impairments. N-stage planar has a total of $2^{\frac{N}{2}(N-1)}$ global switching states, which is significantly higher than the maximum switch permutations and offer high path redundancy in the fabric.

*Wide-sense and strictly non-blocking networks*

WSNB and SNB networks can set up paths between any idle inputs to any idle outputs without any conflict with existing connections. The difference between WSNB and SNB is that the former requires specific rules when establishing the connections, while the latter does not. Typically, WSNB and SNB networks are much larger in both SE number and fabric size but retain much simplified routing control. Crossbar and PILOSS are two typical WSNB networks, both of which require $N^2$ SEs. While a lightpath through the crossbar network can traverse between 1 to (*2N-1*) SEs, the PILOSS network uniformly has *N* cascaded SEs in a lightpath, which results in low path-dependence in the lightpaths [116]. The routing rules of these two WSNB networks dictate that the number of global switching states matches exactly the number of switch permutations, offering full non-blocking connectivity with no path redundancy.

The switch-and-select topology explored in Chapter 4 is an SNB network that follows the binary-tree logic with two mirrored switching arrays to demultiplex and multiplex the input signals. A switch-and-select network has a total of $2N(N-1)$ 2×2 SEs arranged in $2log_2 N$ stages, and gives rise to the largest footprint among the networks listed in Table 5-1. The fabric shows one-to-one mapping between an input-output path and a switch configuration, which leads to a global switching states of N! and a lack of path redundancy.
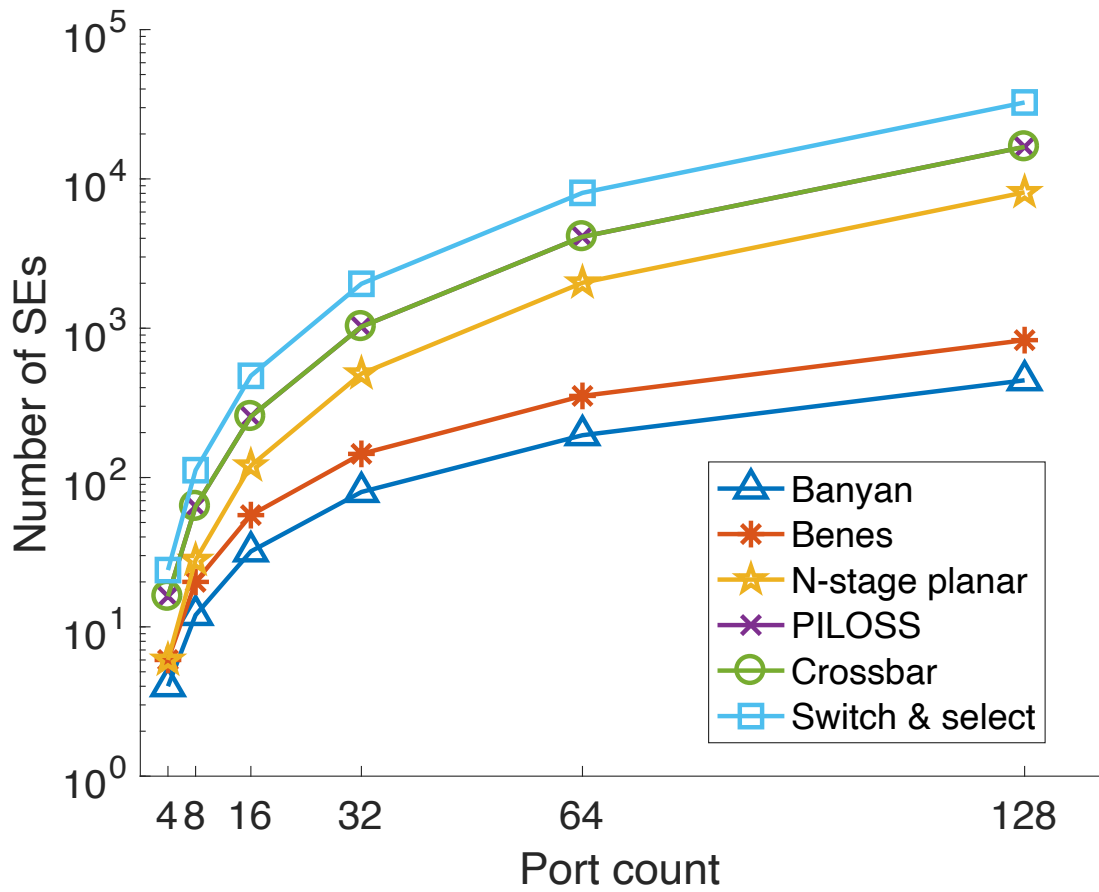
**Figure 5-2: Comparison of the total number of SEs among various switch topologies as a function of port counts in an N×N network.**
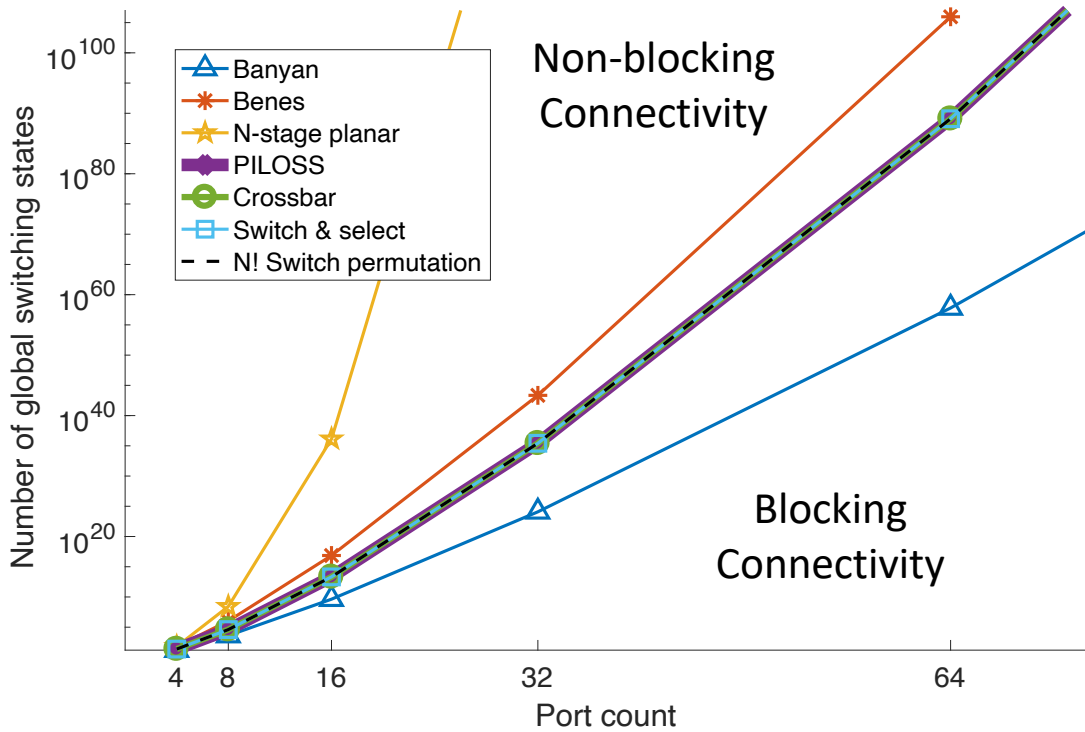
**Figure 5-3: Number of global switching states for Banyan, Beneš, N-stage planar, PILOSS, crossbar, and switch-and-select networks with different port counts.**

To further illustrate the scaling characteristics for each of the switch topologies explored, we show the changes in the number of SEs and the number of global switching states as the fabric port count increases in Figure 5-2 and Figure 5-3, respectively. In particular, for RNB topologies that have global switching states greater than the maximum switch permutations, numerous switch configurations can lead to the same switch permutation. In Figure 5-4, we show two examples of routing redundancies in Beneš and N-stage planar switch fabrics. In the following section, we exploit such routing redundancy to develop a smart routing strategy based on switch component performance data.
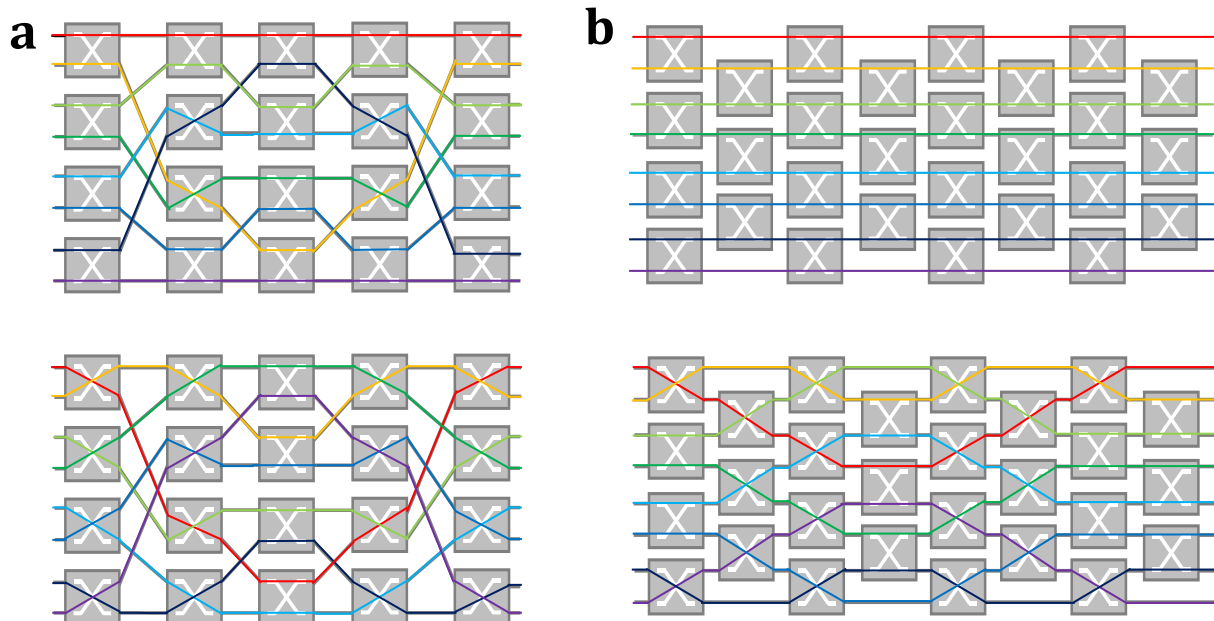
**Figure 5-4: Redundant routing paths with same input-output permutations for (a) Beneš, (b) N-stage planar RNB switch fabrics.**

## 5.3    Fabric-wide, Penalty-Optimal Switch Routing Strategies

In this section, we present the proposed generic routing strategy that is capable of optimizing the fabric-wide switch performance quantified in optical power penalties. The routing strategy is developed based on fast characterization of lightpath performance based on the calibration techniques discussed in Chapter 4 and can be incorporated into the switch control plane [117]. Given redundant routing options as introduced in the previous section, different lightpaths connecting the same input-output pairs can see different waveguide lengths, numbers of waveguide crossings, and SE switching states, all of which contribute to path-dependent power penalties [118]. Quantifying the path power penalties requires a full characterization of the switch device, which can leverage the fast and automated calibration techniques shown in [103] [96] [104] [97] [105] [106] [98] [99] [107]. The characterization summarizes device-specific performance measurements of loss and crosstalk that are converted into an optical power penalty metric for each feasible fabric path, which is used to generate routing tables based on physical-layer performance.

### 5.3.1 Switch Path Characterization

The expedited SE calibration and precise performance-optimal biasing provide a starting point to determine the best routing strategies of the switch fabric. To precisely determine the loss and crosstalk performance of a single switching state, the calibrated switch can be characterized by activating its inputs one at a time with input power, $\varrho_i$. By monitoring the power levels at all outputs, we can distinguish the signal power and the leakage power, thus determining loss and crosstalk levels of the fabric for each input under that configuration. For a generic N×N switch fabric, the path loss and path crosstalk results are determined for all available global switching states, which requires iterating over the valid combinations of all SEs in both Bar and Cross states, with one input injected with optical signal power at a time. We define a switching transfer function $S$, which maps an input $i$ to an output $j$ for a single state:

$$S(i) = j , \qquad (5\text{-}1)$$

$$S^{-1}(j) = i , \qquad (5\text{-}2)$$

We let the signal power measured at the designated output be $\rho_{i,j}$, and the leakage power levels at the other N-1 output port be $\sigma_{i,j,k}$, where $k$ denotes the leakage port and $k \neq j$. Hence, the path insertion loss ratio ($l_{i,j}$) and crosstalk ratio ($\kappa_{i,j,k}$) can be determined as:

$$l_{i,j} = \rho_{i,j}/\varrho_i , \qquad (5\text{-}3)$$

$$\kappa_{i,j,k} = \sigma_{i,j,k}/\rho_{i,j} , \qquad (5\text{-}4)$$

For a single switch state, the aggregated crosstalk power to output $k$ at full switch load with all input ports activated is:

$$\mu_k = \sum_{i=1, S(i) \neq k}^{N} \sigma_{i,S(i),k} , \qquad (5\text{-}5)$$

and the extinction ratio for the path connecting input $\boldsymbol{i}$ and output $k$ is:

$$\epsilon_{i,k} = \rho_{i,k}/\mu_k . \qquad (5\text{-}6)$$

Based on the model studied in [119], the aggregated crosstalk-induced power penalty is estimated as:

$$\delta_k = -10log\left(1 - 2\sqrt{1/\epsilon_k}\right). \quad (\text{5-7})$$

We can then have the total path power penalty in dB as:

$$PP_{i,j} = -10\log(l_{i,j}) + \delta_k. \quad (\text{5-8})$$

Equation (5-8) allows us to conveniently combine both IL and crosstalk into a single power penalty metric as the overarching performance parameter for a switch lightpath.

As indicated in Equation (5-5), to have an accurate measurement of the crosstalk levels, switch fabric inputs need to be activated one at a time. It is possible to aggregate single-input results to infer the performance of a fully-loaded switch state, as shown in Figure 5-5 using a 4×4 Beneš switch fabric as an example. In this case, the switch state connects Inputs 1, 2, 3, 4 to Outputs 4, 2, 1, 3, respectively.
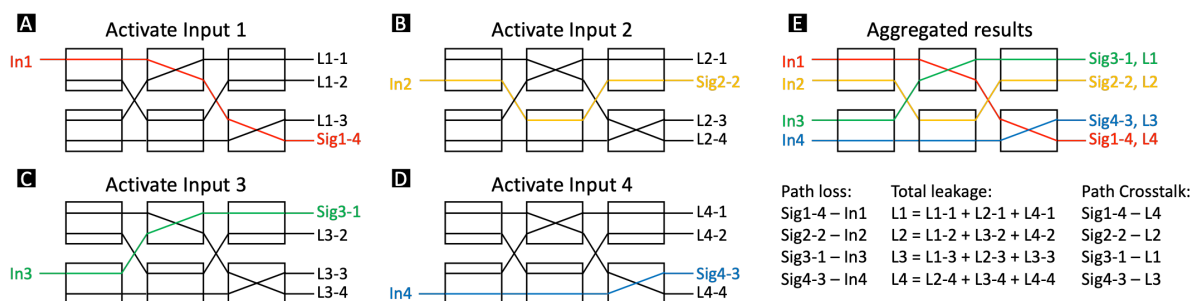


**Figure 5-5: (A-D) Single input results of the same switching state, showing signal power (Sig) and leakage power (L). (E) Aggregated results showing loss and crosstalk of all paths for the same switching state.**

Each switching state is characterized, but instead of iterating through $2^C$ configurations $N$ times for a fabric with $N$ inputs and $C$ SEs, the only SEs that need to be examined are the ones reachable by the optical signal from the specific input tested. For instance, in Figure 5-5a, the optical signal from Input 1 and its crosstalk would not reach the bottom left SE. This reduces the total number of measurements to $N \cdot 2^{C-S}$, where $S$ is the number of SEs unreachable by an input and its crosstalk. For a Beneš network, for instance, $S = N/2log_2N - N + 1$. $S$ can also vary among lightpaths and switching states for networks such as crossbar and N-stage planar. The full characterization of a calibrated switch fabric can be performed in an

automated fashion based on the Characterization Procedure in Algorithm 5-1. For each switching state, the aggregated performance characterization results are converted into switch fabric lightpath power penalties according to Equation (5-8), and its input-output mapping is determined using the block-diagonal transfer matrix method [120], as indicated in the Aggregation Procedure in Algorithm 5-1. For a specific input-output mapping, or switch permutation, all switching states with the same mapping are grouped together and become valid routing options for that switch permutation.

**Algorithm 5-1: Characterization and routing generation for switch**

1: Definition:
2: $I$: Set of switch inputs;
3: $O$: Set of switch outputs;
4: $M$: Set of switch input-output mapping;
5: $S$: Set of switching elements;
6: $S_i$: Set of switching elements reachable by Input $i$;
7: $C$: Set of switch configurations; element has the format of binary arrays in the order of switching element indexing, with 0 for Bar, 1 for Cross.
8: $C_i$: Set of switch element configurations for $S_i$; element has the format of binary arrays in the order of switching element indexing, with 0 bit for Bar, 1 bit for Cross, and x (*don't care*) as placeholder for switching elements $s \notin S_i$.
9: Characterization procedure:
10: **for** $i \in I$ **do**
11:     Find corresponding $S_i$, $C_i$
12:     Excite Input $i$ with CW laser
13:     **for** $c_i \in C_i$ **do**
14:         Set appropriate control signal for $s \in S_i$ based on $c_i$
15:         Record power levels at Outputs $o \in O$ corresponding to $c_i$
16:     **end for**
17: **end for**
18: Aggregation procedure:
19: **for** $c \in C$ **do**
20:     Find corresponding mapping $m_c \in M$
21:     Combine signal and leakage power for $c_i \equiv c$, for $i \in I$, to determine fully loaded power penalties for $o \in O$
22: **end for**
23: Routing generation procedure:
24: **for** $m \in M$ **do**
25:     rank $\{c\}|m_c = m$ based on metric $\mathcal{M}$ and select best configuration $c'$
26:     Store in routing table with format $\{m{:}c'\}$
27: **end for**

### 5.3.2 Performance-Aware Switch Routing Schemes

The power penalties of all lightpaths in each switching state conveniently allow a ranking system to be designed for evaluating routing options for each switch permutations. The ranking metric can be designed with different weighting factors to emphasize specific routing optimizations. In this section, we study the implementations of two of these ranking schemes. The first scheme compares the switching states that satisfy a switch permutation based on the penalty of the worst-performing path. By selecting the switching state with the lowest worst-path power penalty as the routing candidate, we eliminate other poorly performance switching states from the switch's operation, and effective raise the bottom-line performance for that switch permutation. This scheme is summarized in Algorithm 5-2, in which $\{S\}$ is the set of switching states with the same switch permutation $m_i$, $s$ is the switching state selected, and $max\ (PP_s)$ is the worst path's penalty for $s$, and $max(PP_{\{S\}})$ is the set of worst path's penalties for all states in $\{S\}$.

**Algorithm 5-2: Worst-path-optimal routing selection**

for $\{S\}$ with mapping $m_i$

select $s$ s.t.

$s \in \{S\}$ and

$$max(PP_s) \leq\ min\ (max(PP_{\{S\}}))$$

The second performance ranking scheme considers the routing options based on the range of power penalties amongst the switch paths, and select the switching state with the minimal discrepancy between the best and worst performing paths. This metric has the benefit of reducing the variations of switch output signal penalties, but does not guarantee to avoid the worst-case penalty paths. This scheme is summarized in Algorithm 5-3.

**Algorithm 5-3: Range-optimal routing selection**

for $\{S\}$ with mapping $m_i$

select $s$  s.t.

$s \in \{S\}$ and

$$range(PP_s) \leq min\,(range(PP_{\{S\}}))$$

For each switch permutation, the highest ranked routing option based on the chosen metric is recorded into a LUT, of which the switch permutation is the key of the entry and the best routing options with detailed SE configurations are the associated values. This ensures that, during routing operation, the switch control plane can directly query the LUT to actuate configurations with high confidence of the performance.

### 5.3.3 Routing Control for a Beneš Switch Device

*4×4 Beneš switch routing demonstration*

We perform a fully automated characterization and routing control on a 4×4 electro-optic Beneš switch device, as shown by the micrograph in Figure 5-6a. The switch device is fabricated at the Institute of Microelectronics via an OpSIS MPW, and consists of six 2×2 MZI SE. Each MZI arm is equipped with a TO tuner for device calibration and a fast EO p-i-n phase shifter for switching. The 4-port Beneš switch can be configured in $2^6$ global switching states that map to 4! switch permutations. The device is die-bonded onto a chip carrier and fixated on a custom PCB fan-out board, as shown in Figure 5-6a. A fiber-array is UV-cured and couples into the silicon chip via grating couplers. The SEs are calibrated in balanced push-pull scheme with a TO tuner. Switching between the Cross and Bar states is achieved by alternatively applying a $\pi/2$ phase shift with the EO p-i-n shifter.

The switch device is subsequently characterized for all 64 switching states using the methodology discussed in Section 5.3.1. The 4-port device exhibits on-chip IL levels ranging between 2.5 dB and 6.5 dB, and a crosstalk ratio below -15 dB. The loss variations mainly lie in the optical shuffling with different propagation distances and numbers of crossovers, and the loss difference in Bar and Cross states due to free-carrier absorption. Figure 5-6b and Figure 5-6c compare the optimal routing option and the worst routing option for the worst-path-optimal and range-optimal routing scheme, respectively. When using the minimal worst-case penalty metric, the optimal routings can reduce power penalties by up to 3.2 dB from the worst routing options. When minimal penalty range metric is applied, the optimal routing options can reduce the discrepancy of penalties by up to 5.9 dB comparing to the worst routing options. We also indicate in Figure 5-6c the different routing configurations that are consider optimal by the two optimization metrics applied.
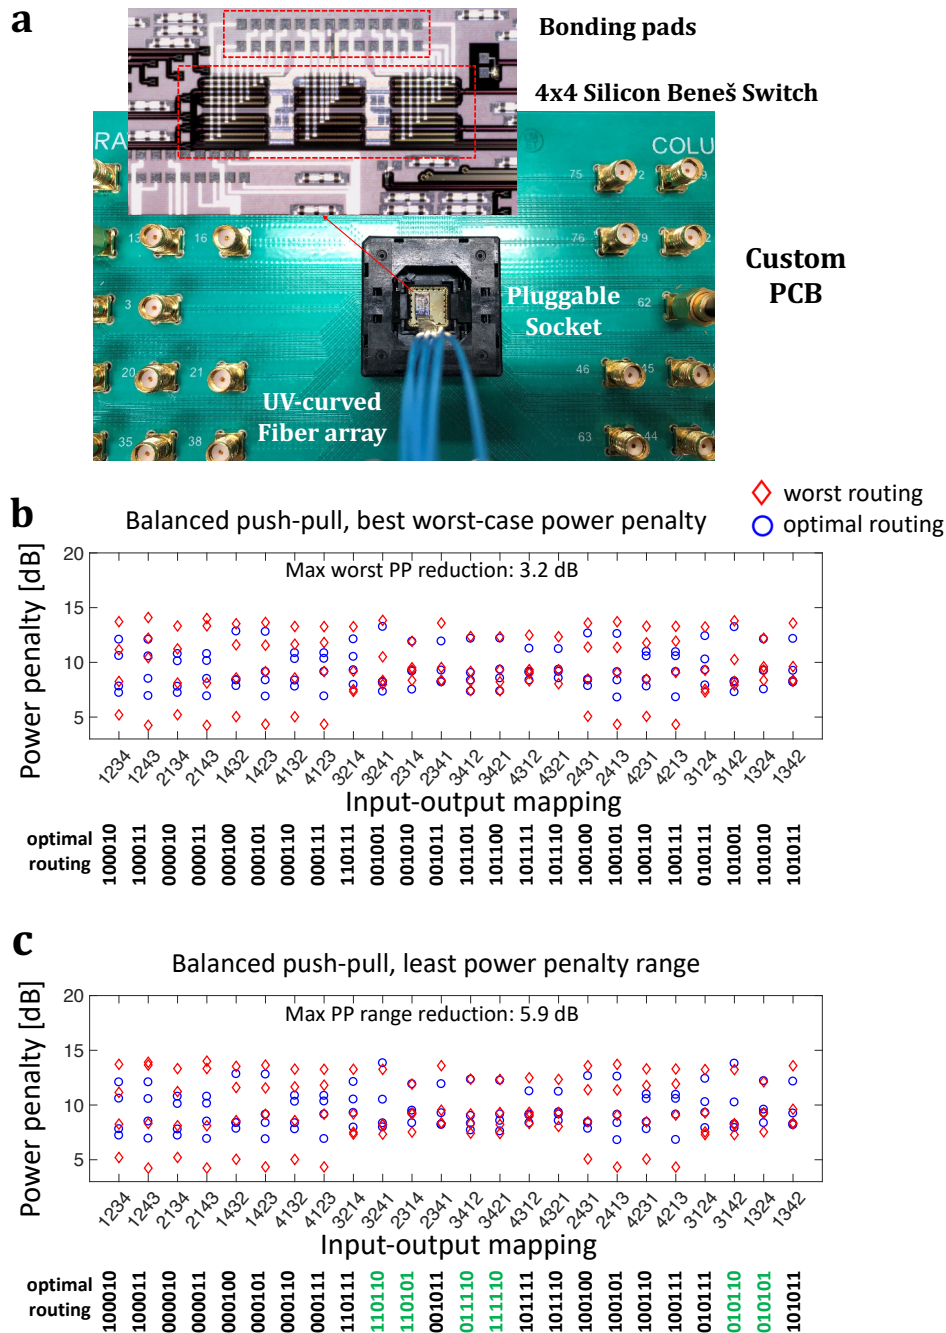
**Figure 5-6: (a)** Switch test-bed with micrograph inset showing the OPSIS 4×4 silicon MZI-based Beneš switch photo. Comparison of the on-chip penalties between optimal and worst routing options under **(b)** worst-path-optimal routing and **(c)** range-optimal routing. Input-output mappings are denoted by how Inputs 1,2,3,4 are reordered at the outputs. Routing configurations are indicated as binary states of 0 (Bar) and 1 (Cross) in the sequence for SEs 1 to 6 for the 4×4 Beneš switch. Green routing states indicate differences in optimal routing as defined by the metrics.

*Simulated 8×8 Beneš switch fabric routing scheme*

To further study the benefit of the proposed routing strategy in a larger switch fabric, an 8×8 Beneš switch model is analyzed using the cross-layer simulation platform, *PhoenixSim* [121], which is a simulation tool that enables integrated and interactive design space exploration over the physical, networking and application layers. System-level metrics such as loss and crosstalk can be extracted from physical-layer compact models of silicon photonic components, including waveguides, bends, crossings, directional couplers, PIN phase shifters, and thermal tuners.

Physical-layer parameters based on the measurements from the 4×4 Beneš switch device are used to define the 8×8 Beneš switch fabric with five stages of MZI SEs with realistic waveguide shuffles that reflect the discrepancies in waveguide propagation lengths and crossing numbers among different paths. The switch fabric contains $2^{20}$ switching states, 8! switch permutations, and 8·8! individual switch paths. We characterize the routing options for each switch permutations in procedures defined in Section 5.3.1, and perform the RMSE-optimal routing scheme. Equation (5-9) defines the RMSE, η, which is a better metric to reflect the variance of the output penalties than the range.

$$\eta = \sqrt{(\textstyle\sum_{i=1}^{N}(PP_{i,j} - \overline{PP}))^2/N} \,. \qquad (\textbf{5-9})$$

We visualize the routing performance metric by examining a specific permutation connecting Inputs [1,2,3,4,5,6,7,8] to Outputs [1,2,3,4,5,6,7,8]. Figure 5-7a illustrates the power penalties of the most RMSE-optimal switching state the out of the 256 states that satisfy the switch permutation, and Figure 5-7b illustrates the worst switching state in the available routing options. In comparison, not only was the worst-case power penalty reduced by up to 6 dB by selecting the RMSE-optimal routing option, the variance among the switch paths are also much lower.
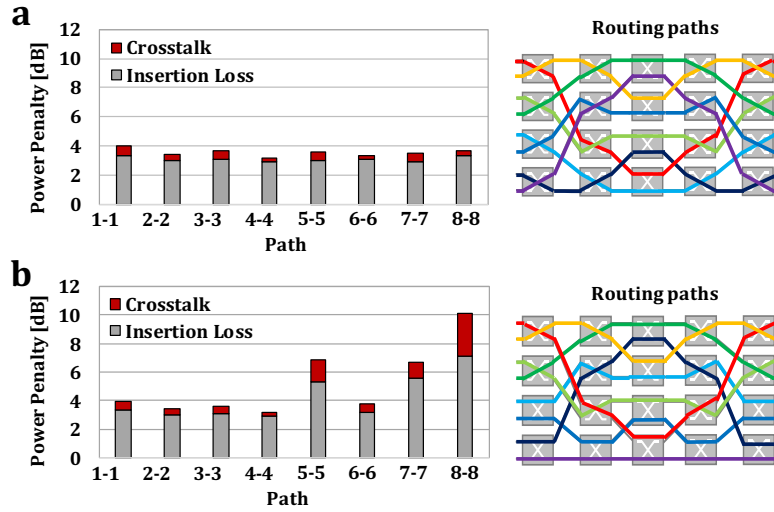
**Figure 5-7: Path power penalties and the switch state for (a) the RMSE-optimal routing option and (b) the least RMSE-optimal routing option.**

An overview of power penalties for all switch paths in the 8! switch permutations is shown in Figure 5-8 to illustrate the difference in the power penalty distributions between the least optimal and the most optimal routing options. It can be seen that the distribution of the most optimal routing paths (Figure 5-8b) has a significantly lower variance than the least optimal routing paths (Figure 5-8a), as well as a much-improved worst-case power penalty by up to 17 dB.
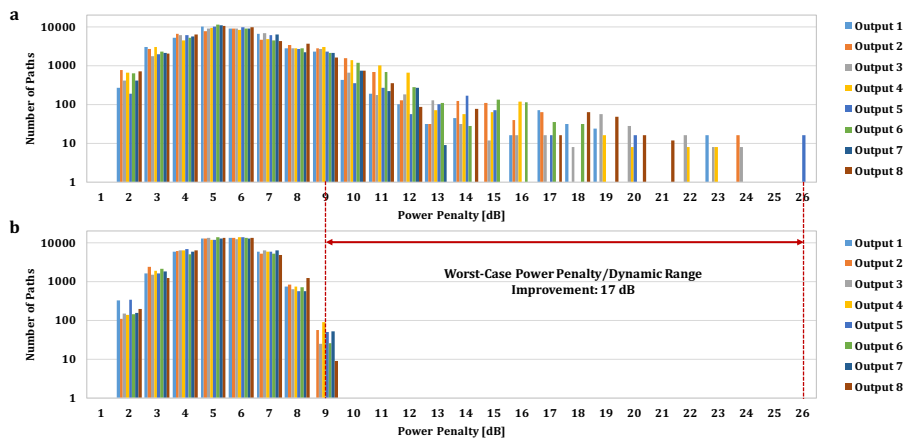


**Figure 5-8: Histogram of path power penalties for (a) worst-case routing and (b) optimized routing for all 40320 (8!) switch permutations, indicating an improvement of ~17 dB for both the worst-case path power penalty and dynamic range.**

## 5.4    Chapter Summary

In this chapter, we analyze the importance and implementations of routing strategies for silicon photonic switch fabrics. By defining and quantifying the number of global switching states in various switching topologies, we reveal how redundancies in switching states can be exploited to optimize fabric-wide path power penalties. Significant power penalty improvements with low variance among switched lightpaths are demonstrated experimentally and in simulation, verifying that performance-aware routing strategies offer great potentials to compensate for device fabrication variations.

Increasing switch port count can result in greater variance in the path power penalties among switching states satisfying the same switch permutation, as shown in the 8×8 Beneš switch fabric analysis when compared to the 4×4 fabric. The significant increases in the number of SEs, global switching states, waveguide shuffle configurations, and SE states all contribute to large discrepancies in redundant switching states that, in some cases, lead to poorly performing light paths. As silicon switch fabrics scale to even higher radices, such as 16×16 and 32×32, routing strategies are increasingly important for improving the performance and consistency of lightpaths through the switch devices. Hence, the continuing improvement for silicon photonic switch designs require advances from component design, device integration, post-fabrication calibration, and system optimal control scheme. In the next chapter, we explore how statistical and machine learning methods can help generate additional insights when designed as part of the system control plane.

# Chapter 6: Intelligent Control Plane with Machine Learning

## 6.1    Introduction

As we have established in the previous chapters, control plane design for optical devices and systems can be tailored to integrate specific functionalities and optimizations to enhance the physical layer performance. As optical hardware begins to evolve into open and interoperable designs, advanced and intelligent control plane design is becoming an emergent research and development area. In particular, control planes that incorporate ML capabilities have shown tremendous potential to enable the underlying systems to autonomously monitor, optimize, and adapt from historical operation data [122]. Works in [123] and [124] have shown informative summaries of the techniques, applications, and results in building intelligent optical communication systems with ML-enabled toolkits. It is also demonstrated in [125] that ML techniques can be applied to not only networking systems, but also to component performance monitoring by building statistical, rather than analytical, models of the operated component. Nevertheless, ML-informed strategic operations of optical components to improve physical layer stability in dynamic networking have been largely overlooked. It is especially beneficial to design a generalized cognitive workflow that widely applies to unique systems. The learning process of the cognitive workflow allows each system to characterize its responses from historical data and offer tailored best practices that mitigate power excursions.

To that end, we introduce an ML engine to address the challenge of EDFA power excursion in WDM networks, and present two applications of this control plane design. AGC-enabled EDFA tends to induce channel-dependent power excursions as WDM channels are added or dropped and can undesirably increase the discrepancy among amplified channel power levels. In the first application, we use ML techniques to statistically characterize the EDFA response and predict best practices for channel addition and removal. We show that the ML techniques can accurately predict the system's power response to

channel changes and avoid channel configurations that would trigger significant power excursions. In the second application, we adapt the ML engine to further address more complicated EDFA power excursions for flexgrid networks and the wavelength defragmentation process. We show how the ML engine can recommend precise power pre-adjustments that will reduce the post-EDFA power variance for various spectrum defragmentation methods. Throughout this chapter, we want to highlight our design philosophy in building ML-enhanced control plane, and show how ML as a tool can unearth, rather than obscure, insights valuable to system designers and operators.

## 6.2　EDFA Power Excursions

Broadband optical amplifier systems, such as the EDFA, are conventionally designed to operate with slow-varying channel usage [126]. Despite EDFAs' ability to achieve economic regeneration of WDM channels, they face an unsolved challenge of wavelength-dependent power excursions during rapidly changing channel configurations [127]. The fast power transients in EDFAs occur on the time scale of 10-100μs [128], and have been largely addressed by localized feedback and feedforward controls of individual EDFA instruments [129]. However, steady-state power excursions that occur across multiple EDFAs can adversely impact channel power levels during dynamic channel add/drop operations, and therefore still demand a reliable and efficient solution. Resolving the critical issue of power excursion in EDFA systems would eradicate a major obstacle to achieving power stability in dynamic optical networking.

The cause of steady-state EDFA power excursions is the interaction between the non-flat gain tilt and the gain control mechanism of the amplifier [127]. Modern EDFA systems employ AGC to maintain the post-EDFA power levels within a tolerance window [130], with additional controls also in place to reduce power transients in response to changing input power [131]. If a channel with high gain is added, AGC responds to an increase in the mean output power by reducing the gain on all channels. This response leads to the high-gain channel effectively stealing power from lower-gain channels [132]. Conversely, adding a low-gain channel feeds power to higher-gain channels [127]. Excursions up to 2dB have been demonstrated experimentally in as few as three cascaded EDFAs through haphazard channel additions [127]. Power excursion is undesirable because it increases the discrepancy among post-EDFA channel power levels, which may be further exacerbated by downstream amplifiers.

Power excursions are also a critical challenge to the power stability in emergent flexgrid networks, in which the bandwidth of each channel is unfixed. The added variability of channel bandwidth in flexgrid

networks means that EDFA systems need to respond to a variety of spectral power changes [126]. In particular, adding, dropping, and shuffling flexgrid super-channels often involve provisioning multiple WDM channels at the same time [133], and therefore require more robust strategies to retain power stability.

To further complicate the matter, the exact power excursion response of an EDFA depends on its specific gain tilt, gain control mechanisms, and the number of EDFAs in a lightpath. While there are attempts to mitigate EDFA power excursions [127] [130] [132] [134] [135] [136], these approaches are limited when generalized to other systems. In [111], a compact model is developed for EDFA power excursions based on mean power levels and gain differences among the channels to enable single-step channel provisioning. In [132], different pre-emphasized power levels are set for the channels to compensate for the excursion in gains. In [127], pairs of TDM channels are added through fast TLDs to achieve a balanced average power, preventing the EDFA controls from adjusting gain levels on existing channels. In [130], the pumping level of the Raman/EDFA hybrid amplifier is adjusted to reduce the power transient variations and steady-state excursions. In [135], optimized wavelength assignment has been shown experimentally to reduce the excursions, but is limited to 5-15% reduction of gain deviations. In [136], CBR has also been applied to make heuristic guesses on EDFA tuning, but is limited in its prediction capability for unknown network scenarios. These techniques, while effective for the specific systems analyzed, are not necessarily transferrable to different networks. In addition, they also rely on deterministic models and full knowledge of the gain profile, details difficult to obtain for live-network equipment that cannot be disrupted. It is critical for EDFA power excursion solutions to be transferrable among heterogeneous systems, and especially beneficial to design a generalized workflow that characterizes and responds to different system. ML techniques are powerful tools for model building based on historical power response data and offer tailored best practices to channel provisioning and spectral defragmentation.

## 6.3 Mitigation of EDFA Power Excursions during Dynamic Channel Provisioning

In this section, we introduce a ML engine to statistically characterize EDFA systems and predict best practices in fixgrid channel addition and removal. In contrast to a fixed, closed-form analytical model of the system, the ML approach flexibly adapts to the historical responses of the amplifiers by performing regression on the selection of channels and the discrepancy of their post-EDFA power levels. We show that the ML techniques can accurately predict the system's power response to channel changes, avoid channel configurations that would trigger significant power excursions, and predict best practices of channel add/drop to mitigate undesired excursions. We also demonstrate that the ML approach is directly transferrable among EDFA systems with different equipment and scales.

Two ML models are presented and implemented in the ML engine to characterize EDFA power responses. We further show an extension of the ML-based approach to provision WDM channels with variable spectral bandwidth and super-channels to support flexgrid networks. By associating the channel selection process with the knowledge of the system's overall power response, our approach shows more optimal heuristics for fast channel provisioning decisions. When performing super-channel provisioning, our ML approach shows significant mitigation of power excursions when compared to the conventional first-fit algorithm.

### 6.3.1 Testbed Design

To emulate the complexity of a multi-span EDFA system, we develop an experimental testbed consisting of multiple cascaded EDFAs of different brands and models. The multi-span AGC-enabled EDFA system constructed is shown in Figure 6-1. We vary the number of EDFAs between experiments to examine the transferability of the ML approach among different system scales. Each EDFA demonstrates a different gain tilt and responds differently to input power changes when channels are added or dropped. The channel dependence of the system's power excursion is determined by training a regression model that determines the contribution of ON/OFF channels to the overall discrepancy among the channel power levels after the final EDFA.
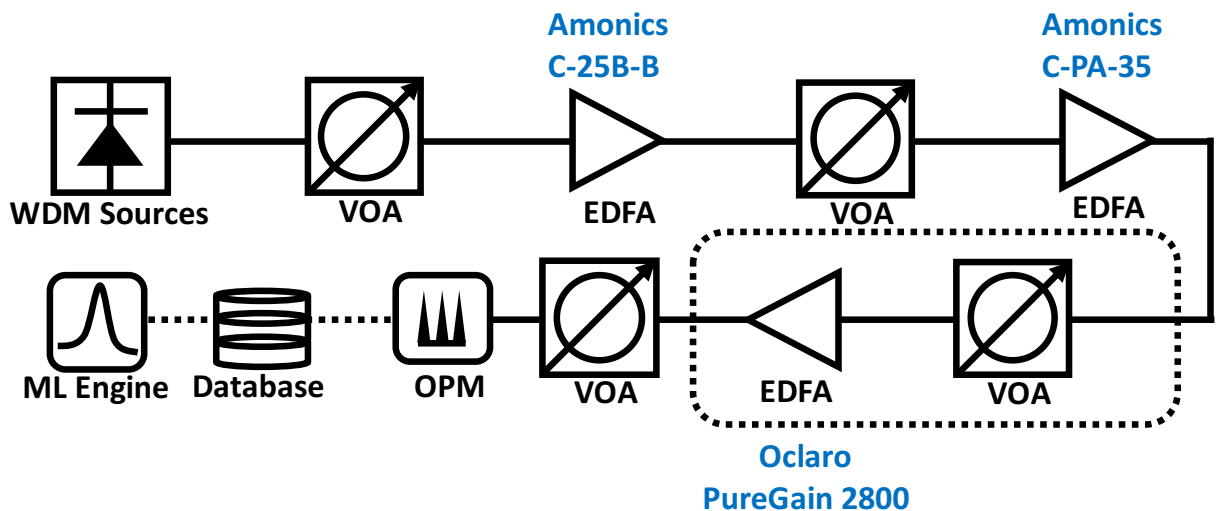


**Figure 6-1: Setup of the multi-span EDFA system; the additional EDFA and VOA in the dashed box are included for the 3-span system.**

The WDM sources transmit 24 WDM channels from ITU-T grid 194.40THz to 192.10THz with 100 GHz spacing at a uniform laser launch power level of 13dBm per channel, which are combined via a WSS. VOAs are placed before each span's EDFA to emulate a 20dB per-span transmission loss of SSMF. We adjust the EDFA pumping levels to achieve an overall gain tilt with a maximum disparity of 10dB between the highest and lowest gains; this is to ensure the receiver dynamic range is sufficient to identify

channels after the EDFA cascade. Figure 6-2 shows the post-EDFA channel power spectra of the 2-span and 3-span systems. While the gain tilt is exaggerated compared to telecom conventions, it presents an extreme case that produces sufficient EDFA power excursions to examine the ML engine's capabilities. To ensure an adequate number of available channels for add/drop, we maintain 10-20 ON channels at any given time, which corresponds to a spectral utilization between 42% and 83%. The post-EDFA power levels are recorded with a C-band OPM, and stored in a database for analysis. The ML engine is trained on the post-EDFA power levels and channel ON/OFF states and constructs a regression model for future predictions.
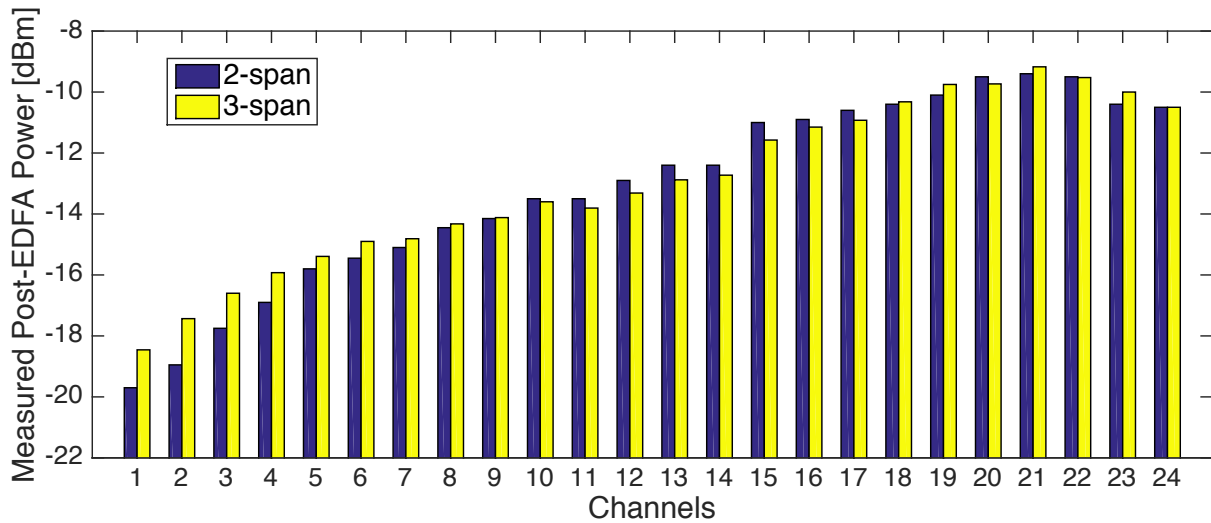


**Figure 6-2: Measured post-EDFA power spectra with 24 ON channels for both systems. Channels 1 to 24 correspond to ITU-T C-band 194.40THz to 192.10THz with 100GHz spacing. Channels are launched with uniform power.**

### 6.3.2 Machine Learning Models and Analysis

We define a regression problem with supervised ML to statistically model the channel dependence of EDFA power excursions. The predictor variable of the regression model is a 24-bit array, with each bit indicating an ON channel as 1, or an OFF channel as 0. This can be scaled up to longer arrays to accommodate more WDM channels. The response variable of the regression model is a measurement of the post-EDFA power discrepancy. In the following analysis, we use the STDEV of the channel power levels as the measurement of discrepancy. We select STDEV of post-EDFA power levels as the optimal metric instead of other FoMs, such as the OSNR, because changes in STDEV directly infer the extent of undesired power excursions induced by cascaded EDFAs. STDEV is also industrially practical because it necessitates the transceivers' dynamic range, which impacts system complexity and cost. Alternatively, the power excursion measured for each ON channel can be aggregated to reflect the overall impact on the system; this has the ability to reflect the channels' power stability instead of power discrepancy. To accentuate the mitigation of the power discrepancy among channels due to EDFA power excursions, we focus on the power STDEV metric in the following analysis.

We collect 870 historical channel ON/OFF states and power STDEV values from each experimental system, which are split into a training set of size 600 and a testing set of size 270. During training with a system's training set, the ML model's parameter values are optimized for that specific system in order to achieve accurate prediction capability. We design the ML engine to streamline the training process by using historical channel ON/OFF states and post-EDFA power levels, which are operation data that can be easily collected. Because existing operation data of the system can also be used to train, the ML engine can continue to improve its accuracy with increasing size of available training data. For a larger system with complex effects, the ML engine can implement an online, continuously evolving training process to gradually capture the system's power dynamics. The ML model is evaluated with the

testing set by two metrics: A) the MSE between the predicted and the measured STDEV of the testing set, and B) the correctness of the best channel provisioning recommended based on the predictions.

The predictor and response variable values are preprocessed before training and testing according to Equations (6-1) and (6-2). The per-dimension mean is removed from the predictor variable $x$ and the response variable $y$. Each predictor dimension is also standardized with a variance of 1. When used for prediction, the model takes in standardized inputs and returns offset outputs. These commonly practiced preprocessing techniques prevent dimensions with large variances or means from masking the contribution of other dimensions.

$$x_{ij}^{prep} = \frac{x_{ij} - \bar{x}_j}{\sigma_j}, \qquad (\textbf{6-1})$$

$$y_{ij}^{prep} = y_{ij} - \bar{y}_j. \qquad (\textbf{6-2})$$

In Equations (6-1) and (6-2), $i = 1..n$ is for $n$ total data points; $j = 1..d$ labels a dimension of the predictor ($d = 24$) or the response ($d = 1$); $\sigma_j$ is the per-dimension STDEV of the predictor, and $\bar{x}_j, \bar{y}_j$ are the respective per-dimension means. We examine the efficacy of two regression models – RR and KBR – in the context of predicting power STDEV solely from channel ON/OFF states. The rationale for RR is to examine a low-complexity model that can be trained and applied quickly. The rationale for KBR is to perform an efficient statistical implementation of CBR with improved prediction capabilities. We explain in detail the formalism and implementation of each model below.

*Ridge regression model*

RR is a linear regression model with an additional $\ell_2$ penalty. The weights associated with the channels are defined by Equation (6-3) and determined by Equation (6-4):

$$w_{RR} = \arg\min_w \left( \|y - Xw\|^2 + \lambda \|w\|^2 \right), \quad \textbf{( 6-3 )}$$

$$w_{RR} = \left( \lambda I + X^T X \right)^{-1} X^T y, \quad \textbf{( 6-4 )}$$

where $w_{RR}$ is the array of weights learned by RR; $y$ is the columnized response values of the training set of shape $n \times 1$; $X$ is the vertically stacked predictor values of the training set of shape $n \times 25$. The dimension in addition to the 24 channels is used for the learned bias. $I$ is the identity matrix of shape $25 \times 25$. $\lambda$ is the complexity parameter. Using the cross-validation technique, in which the ML model is repeatedly validated with randomized subsets of the training set and different values of $\lambda$, the parameter is evaluated as 2.6 and 2.7 for the 2-span and 3-span systems, respectively. RR implements the $\lambda$ penalty term to encourage a small variance across the determined weights. This means the model prefers to characterize contributions from every channel, instead of attributing the cause of power excursions to a few isolated channels. As a result, RR in general can better prevent over-fitting the training data than linear regression.

Figure 6-3 shows the weights associated with each channel in its contribution to the post-EDFA power discrepancy after training with 600 historical samples for the 2-span and 3-span EDFA systems. The magnitude of each weight indicates the corresponding channel's influence on the post-EDFA power STDEV. The sign of each weight indicates whether the addition of the corresponding channel will increase (positive weight) or decrease (negative weight) the power STDEV. In the case of our experiment, channels close to the ends of the spectrum exacerbate the power discrepancy, while channels at the center of the spectrum mitigate the power discrepancy. We can also note that, comparing the 2-span system with the 3-span system, the same channel has different contribution weights, indicating that the two systems, while sharing some common equipment, have different power excursion responses. Channel 24, whose weights

have different signs, contributes opposite effects to post-EDFA power discrepancy in the two systems, even if the magnitudes of its contributions are relatively small.



**Figure 6-3: Associated weights assigned to each channel by RR for the 2-span and 3-span EDFA systems, indicating each channel's contribution to the post-EDFA power discrepancy in respective systems.**

*Kernelized Bayesian regression model*

KBR relies on mapping the training inputs to a kernel, which is equivalent to an expansion in the input feature dimensions. The regression model is trained using the kernel to determine the posterior probability distribution of the response variable. Specifically, we construct a kernel called the RBF for every pair of predictors as shown in Equation (6-5):

$$K(x,x') = \alpha \exp\left( -\frac{1}{b}\|x - x'\|^2 \right), \qquad \textbf{( 6-5 )}$$

where $a$ and $b$ are parameters that adjust the strength of the kernel, which we set as 0.0001 and 3.5 respectively for our analysis using cross-validation, consistent for both the 2-span and 3-span systems. The variables $x$ and $x'$ are two 24-dimensional predictor values; the value of the kernel function decreases exponentially with the magnitude square of their difference. The RBF kernel indicates that two network scenarios with similar ON/OFF channels are assumed to have similar extents of power excursions [136], which is leveraged to efficiently emulate the use of a database in CBR. Given a new network scenario, we can infer its predicted power STDEV based on the individual contribution of each channel, as well as how similar it is to the known network scenarios. The predictions are obtained from the linear combinations of the training outputs weighted by the kernel function values [137]. In contrary to learning the linear weights of each channel's contribution, KBR with RBF models the system as a Gaussian process by learning the mean and variance of the power STDEV given the ON/OFF channels. One advantage of KBR is its ability to pinch the variance of the Gaussian process near known data points. Hence, for a future scenario that has appeared in the training set, our KBR model can accurately recall the associated post-EDFA power STDEV, thus achieving the functionality of CBR.

We characterize the adaptability of the models on two different systems – with two and three EDFA amplified spans respectively. Figure 6-4 compares the prediction MSE between KBR and RR as a function of the training set size for both systems. For each training size, the models are trained 5 times with random subsamples of the training set, and the prediction MSEs of the testing set are averaged. In both EDFA systems, RR outperforms KBR between the training sizes of 30 and 600. In addition, the ML models perform better on the 2-span system than the 3-span system. It is also noticeable that KBR's performance improves more significantly than RR with increasing training set size, despite its poorer performance when trained with a small data set. For RR on both systems, there is no significant performance improvement beyond 120 training data points.
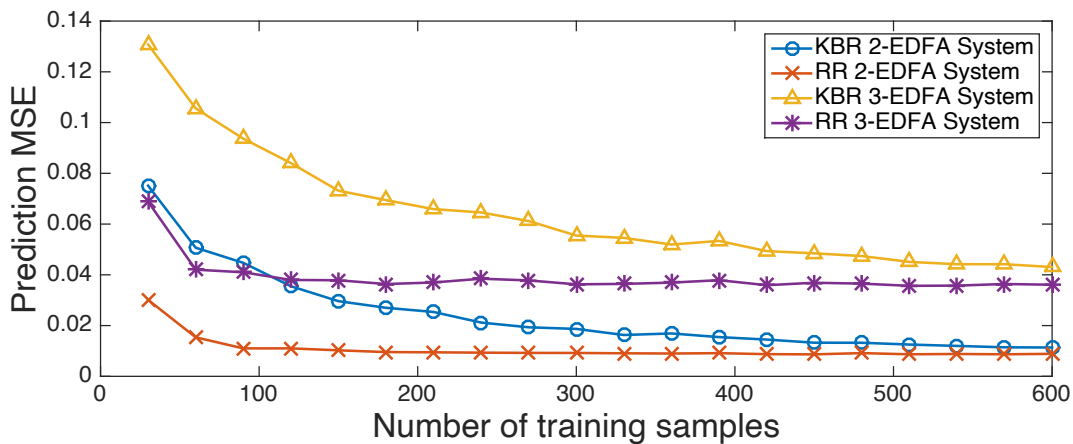


**Figure 6-4: Reduction of the ML models' prediction MSEs of power STDEV with increasing training data size.**

RR and KBR demonstrate different training and prediction mechanisms. For RR, the training data is used to determine a set of weights associated with the input dimensions. After training, the training data can be discarded; only the set of weights is necessary to make predictions. For KBR, the training process involves building the RBF kernel; the kernel and the training data are both used when making predictions. RR has a training computation complexity of $O(N^2M)$ and single-scenario prediction computation

complexity of $O(N)$, while KBR with RBF kernel is $O(NM^2)$ for training and $O(M^2)$ for single-scenario prediction, where $N$ is the number of predictor values (number of channels $+ 1$) and $M$ is the number of training samples. Consequently, RR scales more optimally with increasing number of training samples, and KBR scales more optimally with increasing number of channels. If $M \gg N$, RR would be more efficient in both training and prediction. Both models are implemented in Python 3.5 and executed on a personal computer. An example of their respective training and prediction times are shown in Table 6-1.

**Table 6-1: Time consumption of training and prediction for RR and KBR**

| Model | RR | KBR |
|---|---|---|
| Time to train with 600 data points [ms] | 82.2 | 2300 |
| Time to predict for a single scenario [ms] | 0.068 | 467 |

### 6.3.3 Machine Learning Assisted Channel Provisioning

*Single-channel add/drop*

The trained ML engine is applied to predict the induced power excursions given an unseen channel configuration, and can be used to effectively predict the power changes during adding and dropping fixgrid single channels. Figure 6-5 illustrates the ML engine's capability when used to identify channel add/drop locations that induce the least undesired post-EDFA power excursions. The examples are retrieved from 3-span EDFA system with the ML engine employing KBR and trained with 600 training data points. In both cases shown in Figure 6-5, we start with a randomly initialized scenario of ON/OFF channels. For adding a channel, as shown in Figure 6-5(a), Figure 6-5(c), and Figure 6-5(e), the model predicts the power STDEV if a hypothetical channel is added to each available slot. Then the model recommends the best slots to add a channel that will result in the lowest power STDEV and the least undesired power excursions. Recommending up to four slot options provides flexibility for network operators. This test is repeated for dropping a channel, shown in Figure 6-5(b), Figure 6-5(d), and Figure 6-5(f). In the two tests shown, the slots recommended by the ML engine correctly align with the best slots from the actual measurements.

In Figure 6-6, we examine the recommendation accuracy of the ML engine with 100 tests of single-channel addition and 100 tests of single-channel removal on both EDFA systems with randomly initialized starting conditions. We report the ML engine's ability to identify the top one and top four channels to add/drop after training with various training sizes between 30 to 600, as well as the corresponding prediction MSE at these training sizes. Overall, the ML engine achieves a recommendation accuracy of 81% and 84% respectively, for identifying the correct top one and top four candidates on the 2-span system, and 75% and 89% respectively on the 3-span system. Both RR and KBR demonstrate strong correlation between the prediction MSE and the recommendation accuracy. Despite showing higher prediction MSE on both

systems, KBR demonstrates better recommendation accuracy in identifying the best channels to add/drop in general.
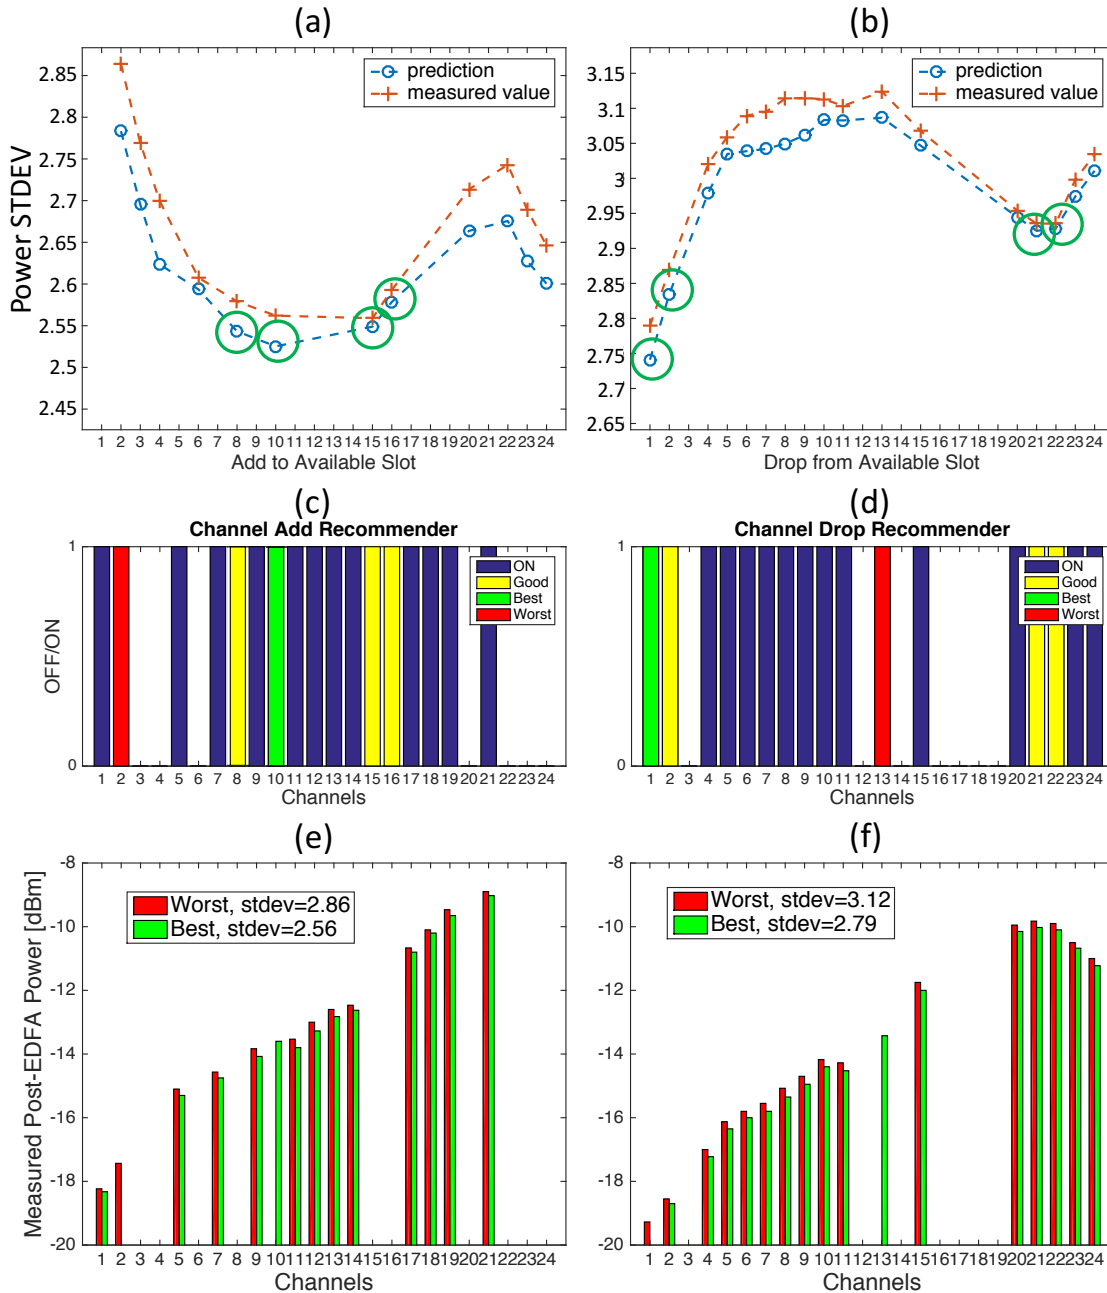


**Figure 6-5: Comparisons between predictions and measurements of post-EDFA power discrepancy for single-channel add (a) and drop (b). The top four slot options with the lowest predicted power STDEV are circled. (c) and (d) illustrate the good channels to add/drop, and the worst channels to avoid. (e) and (f) compare the measured post-EDFA power spectra of ON channels between the best and worst cases in channel add and drop, respectively.**
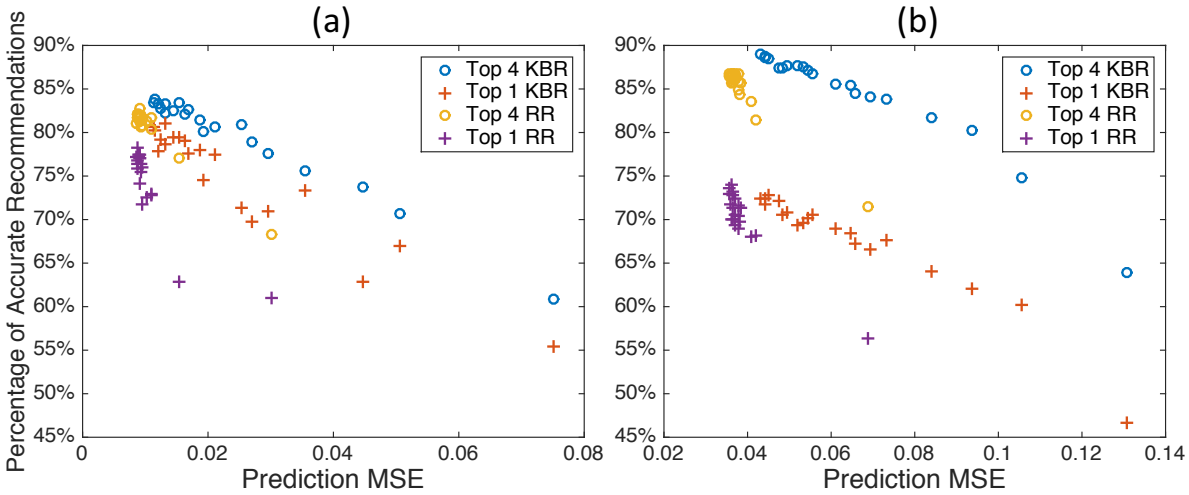
**Figure 6-6: Correlation between the prediction MSE (horizontal axis) and the percentage of accurate recommendations (vertical axis) for KBR and RR for the (a) 2-span and (b) 3-span EDFA systems. The percentage of accurate recommendations reflects the number of tests in which the ML engine correctly recommends the top channels, out of the 200 tests performed.**

Even when the ML engine does not recommend the actual best channel, the recommendations would still result in post-EDFA power discrepancy comparable to the lowest achievable power discrepancy, known as the "ground truth". This makes it feasible to provision channels solely based on the ML engine's recommendations. Figure 6-7 shows the cumulative distribution of the percent error between the ML recommendations' power discrepancy and the ground truths. Each CDF plot is generated from 200 single-channel add/drop cases with randomly initialized ON/OFF channels. The vertical axis represents the fraction of the channel recommendations in the 200 tests that fall within a certain percent error from the lowest possible post-EDFA power discrepancy. In the 2-span system, over 95% of the ML recommendations are within 1% error of the lowest achievable power discrepancy for both RR and KBR. In the 3-span system, over 94% of the ML recommendations are within 1% error for both RR and KBR. This means that even if the ML engine does not give the exact best channel for add/drop, its recommendation would still be almost as good as the best outcome. If we were to randomly pick a channel to add/drop, the choice would be within 1% error of the best provisioning only 16% of the time. Figure 6-8

shows a systematic analysis of the training set size and recommendation accuracy for RR and KBR on the

2-span and 3-span systems. KBR shows monotonic improvement in recommendation accuracy with larger

training set, while RR shows less correlation between recommendation accuracy and different training set

sizes from 100 to 600, with the best result given by a training set size of 200; these results verify the training

convergence trend shown in Figure 6-4.
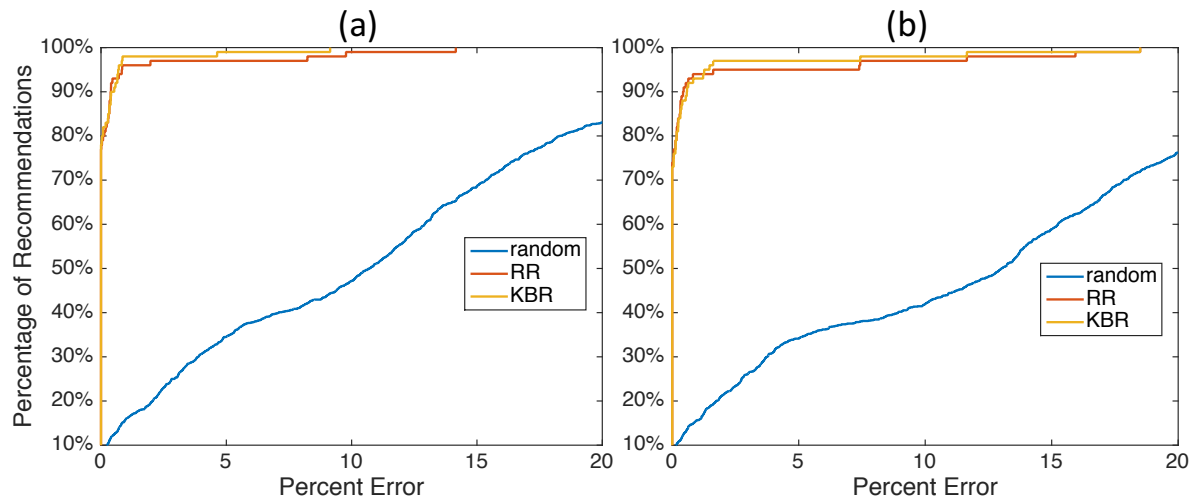


**Figure 6-7: CDF of the percent error between the power discrepancy of recommended options and the actual lowest power discrepancy, plotted for single channel candidates based on KBR, RR, and random selections of channels for the (a) 2-span and (b) 3-span EDFA systems. Compared to random selections, ML recommendations with KBR and RR show significant improvement in approaching the lowest possible post-EDFA power discrepancy.**
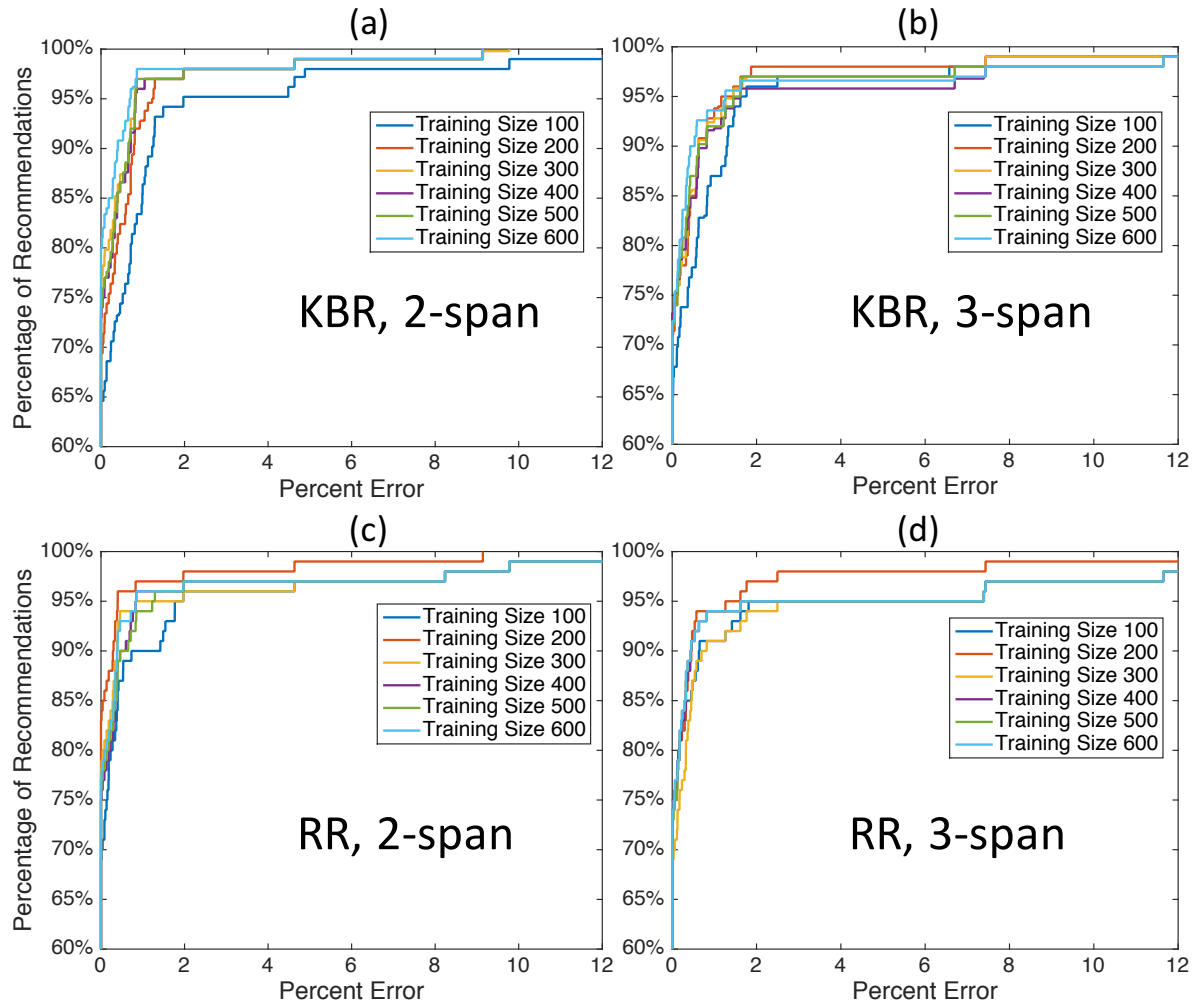
**Figure 6-8: CDF of the percent error between the power discrepancy of recommended options and the actual lowest power discrepancy, at different training data sizes. (a) and (b) illustrate for the KBR model on the 2-span and 3-span EDFA systems respectively. (c) and (d) illustrate for the RR model on the 2-span and 3-span EDFA systems respectively.**

In order to further support the dynamicity of traffic demands, flexgrid optical networks employ variable channel bandwidth to accommodate diverse modulation formats and data rates [138]. One implementation of flexgrid networks employs super-channels consisting of multiple contiguous sub-channels to enable higher data bandwidth and greater spectral efficiency [139]. In the context of EDFA-induced power excursions, the addition of a super-channel results in a greater input power change than the addition of a regular channel. Applying the ML engine to predict and avoid undesired power excursion induced by a super-channel would greatly improve the power stability of the EDFA system.

We experimentally demonstrate that the ML engine can assign two to three contiguous sub-channels to form an optimal super-channel while mitigating the undesired power excursions. Figure 6-9 illustrates the workflow of the ML engine for super-channel provisioning. Figure 6-10 shows two examples of ML recommendations for super-channel addition on the 3-span EDFA system. Figure 6-10(a) and Figure 6-10(b) demonstrate the close fit between the predicted and measured post-EDFA power STDEV values of the super-channels at all available locations. Figure 6-10(c) and Figure 6-10(d) illustrate the recommendation view for the best and worst locations. The ML engine is trained with KBR and the same 600 training datasets used in the previous single-channel provisioning.
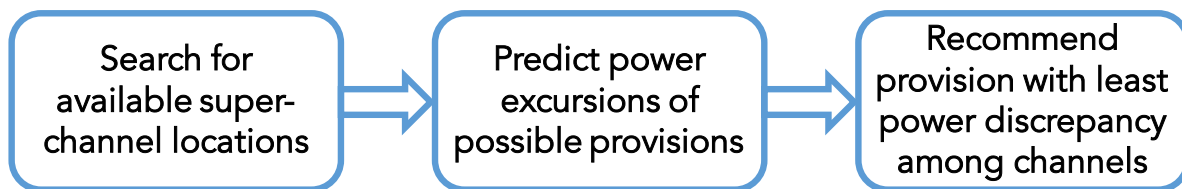


**Figure 6-9: Functional workflow of the ML engine for super-channel addition.**
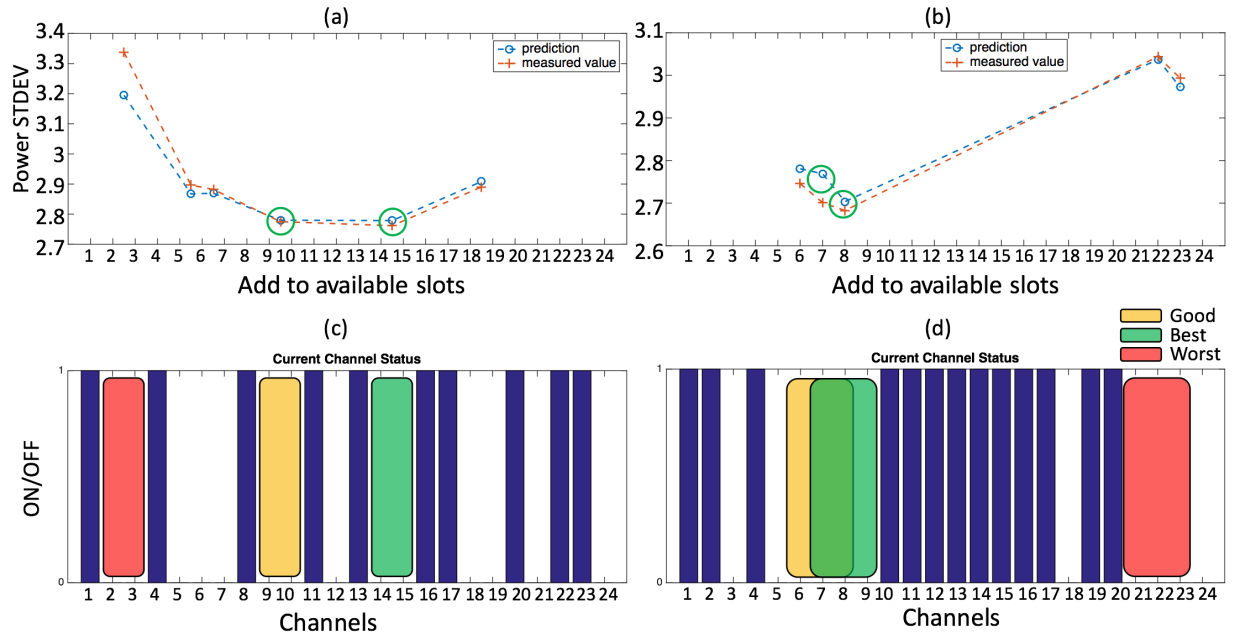
**Figure 6-10: Comparisons between predictions and measurements of post-EDFA power discrepancy for super-channel addition consisting of (a) two contiguous sub-channels and (b) three contiguous sub-channels. The top two super-channel candidates with the lowest predicted power STDEV are circled. (c) and (d) illustrate the best, good, and worst super-channel candidates.**

We see a significant mitigation in post-EDFA power discrepancy when using the ML engine to determine the optimal locations for the super-channels. Figure 6-11 shows a comparison of the post-EDFA power STDEV among the best provisioning options, the ML engine recommendations, and super-channel allocation based on the first-fit algorithm [140], with which the first available location starting from the lower end of the spectrum is selected. Overall, the ML engine recommendations agree closely with the best provisioning options based on actual measurements. For cases where the first-fit algorithm results deviate from the best options, the ML recommendations demonstrate clear improvements over the first-fit algorithm results for mitigating the post-EDFA power discrepancy.
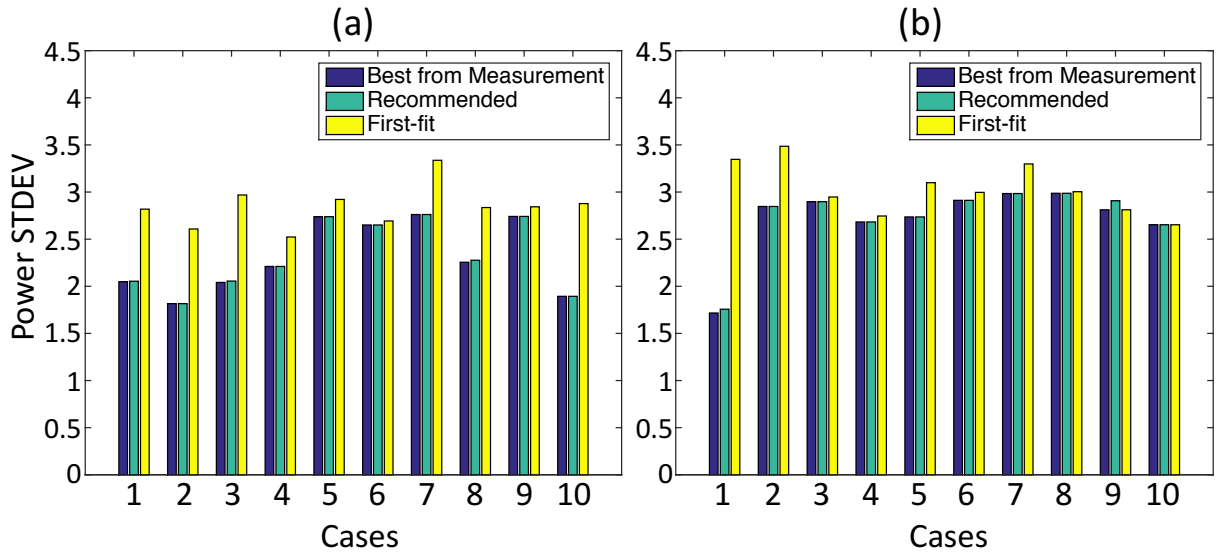
**Figure 6-11: Twenty randomly initialized scenarios to compare the resultant power STDEV values among super-channel additions based on actual measurements, ML predictions, and the first-fit allocation algorithm for (a) two-channel-wide and (b) three-channel-wide super-channels in the 3-span EDFA system.**

### 6.3.4  Scalability of the ML Engine

While the ML engine's efficacy in mitigating post-EDFA power discrepancy is demonstrated experimentally for single- and super-channel provisioning, it is important to discuss how the ML engine would scale with increasing dimensions and complexity in an optical network. For a greater number of amplifiers and spans of fibers in a lightpath, the ML engine's training and prediction processes are unaltered. This is because the ML models are trained on the overall channel ON/OFF states, as well as the power discrepancy after all cascaded amplifiers, which are independent from the number of EDFAs and fiber spans in a lightpath.

A more complex optical network may employ wavelength add/drop capability at intermediate nodes along a lightpath, as shown in Figure 6-12. In this case, the set of channels entering the first EDFA, labelled A, are different from the set of channels exiting an intermediate EDFA, labelled B. To operate the ML engine in this scenario, the predictor variable to the ML models would capture the ON/OFF states of all 4 channels. The response variable of the ML models is determined as the post-EDFA power STDEV calculated from the power levels of Channels 1 and 2 after EDFA C, and the power levels of Channels 3 and 4 after EDFA B. The recommendation process of the ML engine would take into account the channel assignment constraint – only Channels 1 and 2 are available from EDFA A to EDFA C, and Channels 3 and 4 are available for add/drop at EDFA B. This represents a limited search space for the ML engine to make the channel add/drop recommendations.

For more complex networks implementing ROADMs, network edges that connect pairs of ROADMs may carry different wavelength channels. Each edge may also contain multiple EDFAs for optical power regeneration. In this case, multiple ML engines can be deployed to individual edges and operate in a distributed fashion. This allows individual ML engines to mitigate the power excursions in each edge of the network – consistency of channel power levels on one edge would be optimized before

handing off to the next edge. However, the distributed ML engines will need to coordinate for lightpaths that traverse multiple network edges to take into account wavelength continuity and system-wide QoS. These additional functionalities are out of the scope of this paper and would be suitable for future explorations to augment the ML engine's capabilities.
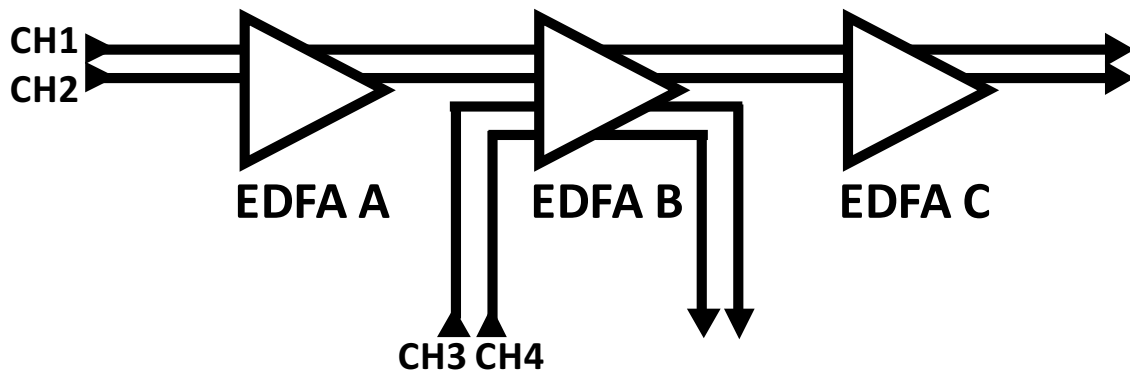


**Figure 6-12: Illustration of channel add/drop operations at the intermediate EDFA node.**

### 6.3.5   Summary

Channel dependent power excursions in gain controlled EDFAs present critical challenges to system performance and agility in dynamic optical networking. The strong dependence of the excursion on the gain profile of an EDFA makes it infeasible to transfer analytical solutions to different systems. We introduce an ML engine based on a regression approach that characterizes the channel dependence of power excursions in a WDM network with multiple cascaded EDFAs. Two machine learning models – RR and KBR with RBF kernel – are trained with 600 historical network usage data points and the discrepancy of post-EDFA channel power levels. Channel provisioning options recommended by the trained ML engine achieve within 1% error of the lowest possible post-EDFA power discrepancy in over 94% of randomized network scenarios. We also demonstrate that the ML engine's workflow can be transferred to systems of different EDFAs and achieve similar performance, and is applicable to both single- and super-channel provisioning to support flexgrid optical networking. By using ML to accurately predict the power excursions in WDM channel provisioning, network operators can make quick and precise decisions to address network demands and optimize EDFA power dynamics. For future explorations, the capabilities of the proposed ML engine can be further augmented with optimizations of OSNR and additional QoS metrics. For larger networks, distributed and coordinated operations of multiple ML engines can be explored to optimize power consistency for individual network edges, while ensuring wavelength continuity and overall system level QoS.

## 6.4    Power Excursion Mitigation for Flexgrid Defragmentation with ML

The second application we examine in this chapter is the implementation of ML techniques to perform optimized spectrum defragmentation for flexgrid networks to improve the agility and efficiency of WDM systems. By implementing channels with variable allocated bandwidths, flexgrid networks support workloads at different data rates, modulation formats, and QoS requirements. Nevertheless, flexgrid networking faces numerous stringent challenges when designing for dynamic network scenarios and rapidly changing traffic demands. First, the arrival and departure of channels with different bandwidths result in spectrum fragmentation that increases the blocking probability of the network. As a result, defragmentation is performed to re-assign channels wavelengths – the groomed spectrum would then be able to accommodate future channel provisioning. Second, as discussed in the previous sections, broadband optical amplifiers such as the EDFA express wavelength-dependent power excursions when channels are dynamically added or dropped [126] and can increase the post-amplifier channel power variance. EDFA power excursions may greatly impact flexgrid networks because of the implementation of super-channels, which occupy multiple contiguous channel bandwidths and induce greater power changes to the EDFAs than individual fixgrid channels. As demonstrated in the previous section, a change in a super-channel's spectral location can trigger power excursions undesirable to the power stability of the system. Defragmentation, if uncompensated, can increase post-EDFA power variance both during and after its process. In this section, we analyze three main defragmentation method's mechanisms that trigger power excursions, and propose and implement an ML engine that performs power pre-adjustments to reduce the post-EDFA power variance for the defragmentation methods.

### 6.4.1 Defragmentation Methods

Defragmentation is defined as the process to re-assign channels wavelengths in order to groom the spectrum of a flexgrid network to reduce blocking probability when accommodating future channel provisioning. Recent studies [141] [142] have shown impressive models in network RSA that implement one or more wavelength defragmentation methods. These models coordinate the optimal allocation of routing and assignment of channel wavelengths to enable efficient flexgrid networking, while considering their respective impact on data traffic. Three main defragmentation methods – Hop [143], MbB [144], and Sweep [145] – have been demonstrated to effectively re-arrange the spectral usage in a flexgrid system. In Hop, a new channel is established while dropping the original channel simultaneously. In MbB, a new channel is created, allowing the traffic to be re-established before dropping the original channel. In Sweep, the channel wavelength is shifted gradually at the spectral granularity of the equipment without disruption to the traffic. Evidently, the spectral usage and the total number of channels present at a given instance differ for each method, inducing a unique power impact on the EDFAs. Hop can be treated as a single-step process in which a channel relocates instantly; the new channel location may experience a different EDFA gain and therefore trigger EDFA power excursions under AGC. MbB is a two-step process, in which a new channel co-exists with the original channel for a short time; the change in both the number and locations of channels interact with the EDFA gain spectrum and may result in power excursions in both steps. Sweep is a multi-step process as the channel shifts gradually across the spectrum, during which the channel may experience a varying range of EDFA gain and consequently result in the amplifier power excursions throughout the process. We demonstrate that the ML engine is capable of alleviating post-EDFA power discrepancy through predicted power pre-adjustments for all three defragmentation methods, showcasing the wide applicability of our proposed design.

143

## 6.4.2 Methodology

In this section, we describe the design philosophy and logical workflow of the ML-based approach, and present the implementation of the methodology as a lightweight control plane application for flexgrid networks with cascaded EDFAs.

*Machine Learning Models*

To accurately predict power pre-adjustments throughout the defragmentation process, it is important to determine two factors: 1) to what extent a channel's spectral location exacerbates the post-EDFA discrepancy, and 2) whether to adjust a channel's power prior to the EDFAs *up* or *down* in order to minimize the post-EDFA discrepancy. We denote the first factor as the *magnitude* of impact, and the second factor as the *correlation* of impact in the following sections. Since channels experiencing high and low gain levels can both trigger undesired excursions, we implement two low-complexity ML models to simultaneously learn the magnitude and correlation of a channel's impact on the post-EDFA power level discrepancy.

RR is performed, similar to the technique introduced in Section 6.3, to examine to what extent a specific channel contributes to the post-EDFA power discrepancy. RR determines a set of weights, $w_{RR}$, from which the variance in post-EDFA power levels, $y_{var}$, is predicted as a weighted linear combination of channel usage:

$$y_{var} = x w_{RR}, \qquad (6\text{-}6)$$

where $x$ is an $1 \times (n + 1)$ array indicating the ON/OFF channel states, as 1 or 0 respectively, of an $n$-channel system. The one additional dimension is used for the learned bias. $w_{RR}$ has shape $(n + 1) \times 1$ and is determined by

$$w_{RR} = arg_w^{min}\left(\|Y - Xw\|^2 + \lambda\|w\|^2\right), \qquad (6\text{-}7)$$

$$w_{RR} = \left(\lambda I + X^T X\right)^{-1} X^T Y, \qquad (6\text{-}8)$$

where $X$ is the set of training data for channel ON/OFF states, consisting of $m$ data points arranged in $m \times (n + 1)$, and $Y$ is the set of training data for post-EDFA power level variances arranged in $m \times 1$. $I$ is the identity matrix, and $\lambda$ is the complexity parameter that encourages a small distribution amongst the dimensions of the learned weight. It is used to avoid heavily attributing the cause of power discrepancy to a few specific channels, and thus preventing overfitting of the training data. Through cross-validations, $\lambda$ is set to 2.8 to achieve best prediction accuracy through repeated training across subsets of the training data. Figure 6-13 shows the learned RR weights corresponding to each of the 24 channels in the experiment system, with further details discussed in Section 6.4.3. The RR model is trained with 800 data points with randomized channel usage. In the previous section, we document the prediction accuracy with respect to the training set size is for a similar system, showing marginal improvements of prediction accuracy beyond more than 600 randomized training data points, although this number is expected to increase with increasing system size. Positive and negative weights indicate that a channel would increase or decrease, respectively, the post-EDFA power discrepancy when turned ON. The magnitude of each weight shows relatively each channel's contribution to the increase or decrease of the power discrepancy.
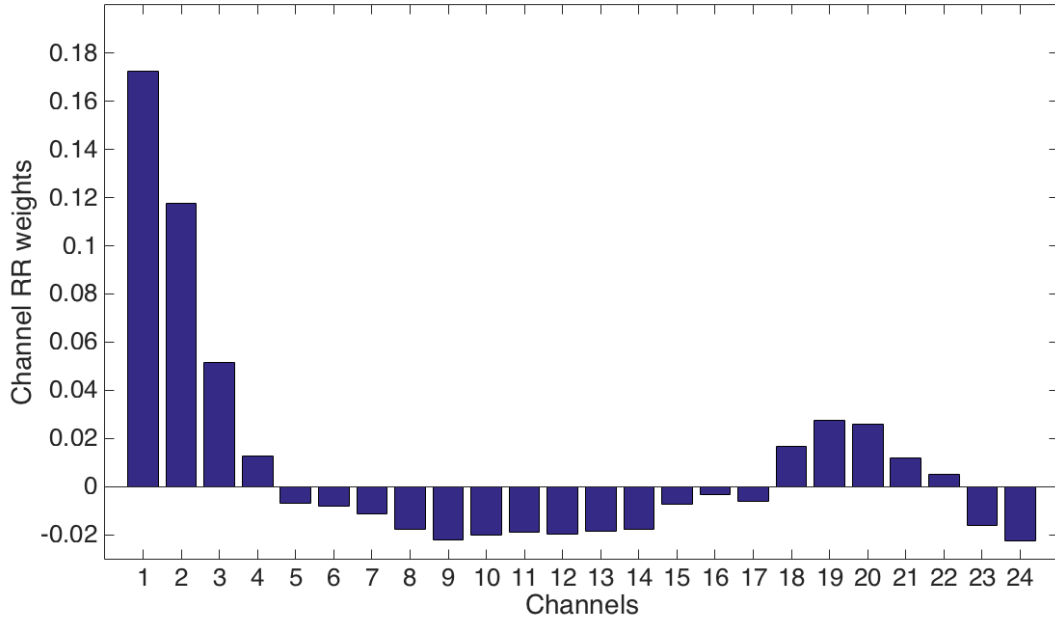
**Figure 6-13: Values of the learned RR weights corresponding to the 24 channels of the experiment system.**

A rise in a channel's pre-EDFA power may increase (correlation = +1) or decrease (correlation = -1) the post-EDFA power discrepancy, conditional upon the other ON channels in the system. Hence we perform LR as a classification technique to determine whether the specific channel's pre-EDFA power level needs to be increased or decreased to reduce the power variance given the channel usage. During training, LR constructs a distribution $P(s_{ch}|\{x^-\})$, where $s_{ch}$ is the correlation specific to a channel ch and the set $\{x^-\}$ indicates the ON/OFF states of all other channels in the same lightpath. We treat $s_{ch}$ in the form of a sigmoid function, whose sign conveniently reflects the correlation factor:

$$s_{ch} \sim \sigma\left(x w_{LR}^{ch}\right), \qquad (\ 6\text{-}9\ )$$

where $w_{LR}^{ch}$ is a channel-specific set of weights that is learned by LR during training. Since the derivatives of a sigmoid function are well-defined, optimization methods such as gradient descent or Newton method can be used to determine a set of weights that achieve the desired training accuracy. $w_{LR}^{ch}$ is channel-specific because it captures how a channel's power level compares to other ON channels in the system. Note that $x$

labels an ON channel as 1 and an OFF channel as 0, which conveniently counts the contribution of ON channels only. Figure 6-14 illustrates the trained LR weights based on the experiment system. Each row in Figure 6-14 represents how the gain levels on other channels (vertical axis) compare to the gain level on the corresponding channel on the horizontal axis – a negative weight means the vertical axis channel has lower gain compared to the horizontal axis channel, and a positive weight means the vertical axis channel has higher gain. This illustration is significant because it reflects the fact that the dynamics of the EDFA system depend both on the channel being added or dropped, as well as the configuration of existing channels currently in the lightpath. This intuition forms the basis of our control strategy.
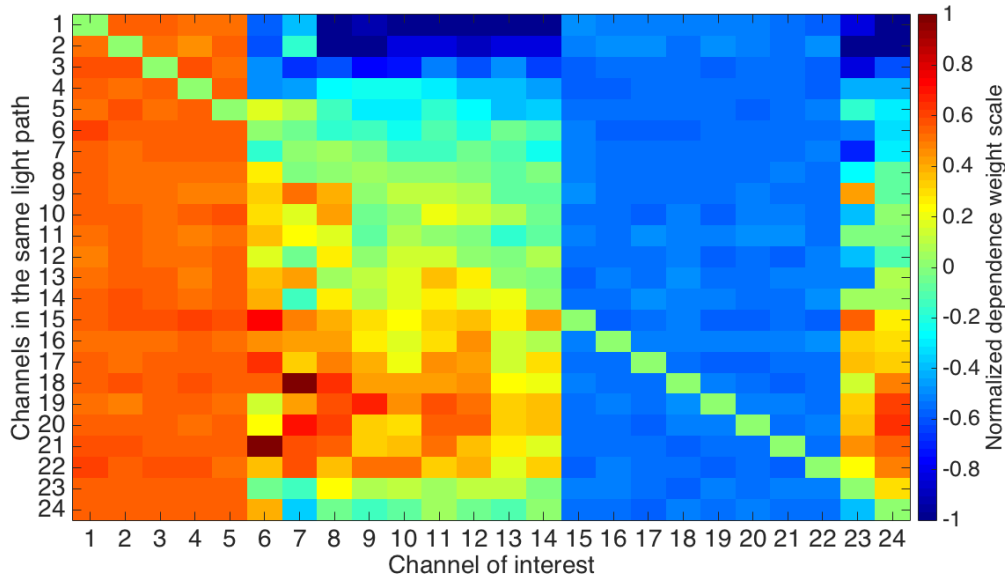


**Figure 6-14: Values of the learned channel-specific LR weights corresponding to the 24 channels of the experiment system.**

The two ML models learn from the same corpus of data consisting of the ON/OFF states of all channels in the system, and the variance of their post-EDFA power levels – this ensures the two models can operate in parallel, expediting the system training process. RR has a training computation complexity of $O(N^2 M)$, where $N$ is the number of predictor values (number of channels + 1) and $M$ is the number of training samples. The complexity of LR training may vary with the approaches to solve the optimization

147

problem. Using gradient descent, each step has complexity $O(N)$, but the overall execution time depends

on the number of iterations and stopping conditions. Both models have prediction computation complexity

of $O(N)$. In addition to the efficient training and prediction processes, the ML model complexities have no

dependence on the number of EDFAs in the system, and therefore are capable of scaling up with increasing

number of amplifiers.

Figure 6-15 illustrates the logical workflow of the ML engine in training and operating on a flexgrid system during the defragmentation process. We present this workflow to be applicable to different defragmentation methods, and discuss the detailed implementations in the following section. The ML engine, including the RR and LR models, is trained with historical channel ON/OFF states and post-EDFA power discrepancy. The same corpus of data is used to train both ML models in parallel, whose results are stored as magnitude and correlation metrics of impact, respectively. These two metrics together determine 1) whether a channel will adversely trigger EDFA excursions to increase the power discrepancy, and if so, 2) how to adjust its pre-EDFA power to reduce the adverse effect. Once trained, the ML engine can make a single-step prediction to adjust the power of the relocated channel by monitoring the ON/OFF channel usage alone, eliminating the need for iterative power tuning processes that require numerous measurements of post-EDFA power. Since the training process is independent from the defragmentation methods, the ML engine is expected to function on the same system when different defragmentation methods are applied. If the system setup or equipment changes, the ML engine can be conveniently re-trained with operational channel ON/OFF states and post-EDFA power levels.
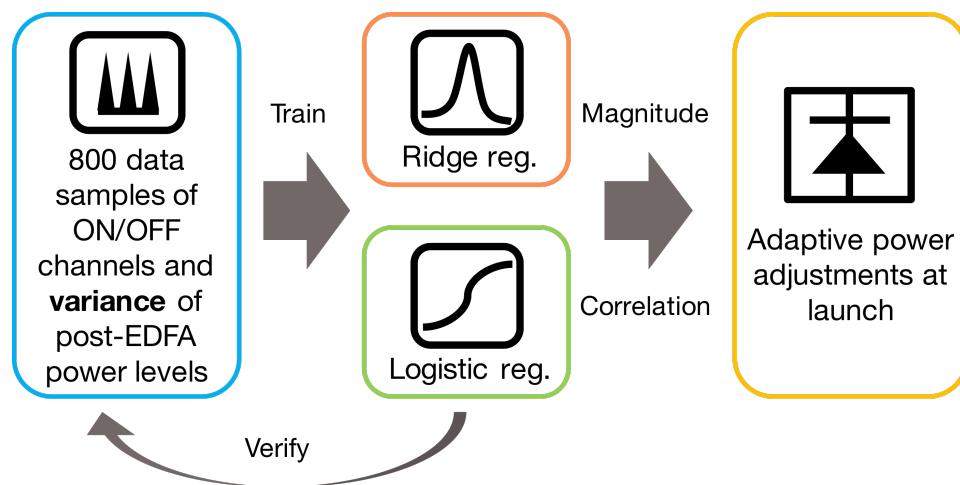


**Figure 6-15: Workflow schematic of the ML engine showing training of RR and LR models, whose results are used to determine power adjustments.**

*Implementation for Defragmentation Processes*

The trained RR and LR models yield two sets of weights, $w_{RR}$ and $w_{LR}^{\{ch\}}$, for the set of channels $\{ch\}$ in the system directly from channel ON/OFF states and variance of post-EDFA power levels. These two metrics summarize each channel's effect on the power discrepancy, and its power level relative to other ON channels in the system. Both the training data and training process are agnostic to the defragmentation methods implemented, and therefore can be transferred flexibly among them. For each defragmentation step, the RR weights are used to predict whether the new channel arrangement would result in a higher post-EDFA power variance than the old arrangement. If the power variance is predicted to increase, the ML engine then uses the LR weights to determine whether the new channel's power is too high or too low comparing to the other ON channels currently in the system, and perform a compensating adjustment to the new channel power prior to EDFAs.

Hop can be treated as a single step process in which a new channel is added simultaneous to a channel being taken down. The power adjustment is performed by the ML engine once the new channel location is determined. MbB is treated as a two-step process, including an intermediate step when both the new and old channels co-exist. Hence, the ML engine performs the power adjustment process *twice* to the new channel – once when the system contains both the old and the new channel, and again after the old channel is taken down. In Sweep, a channel's spectral location is continuously changed, and therefore for every step the channel is shifted, the ML engine determines a power adjustment, if necessary, given the channel's new location and the overall spectrum usage at that instance during the defragmentation. For all three implementations, the ML engine only adjusts the power levels when necessary for the newly provisioned channels, and thus minimizing the interference on other channels in the system.

### 6.4.3 Experimental Demonstration

*Experiment setup*

A total of three C-band EDFAs are cascaded together, in the arrangement shown in Figure 6-1, to concatenate their power spectra over 24 WDM channels from ITU-T grid 194.40 THz to 192.10 THz with 100 GHz spacing, launched by 24 Thorlabs Pro8 dense WDM DFB laser modules. The power of individual lasers at channel wavelengths are tuneable between 7 dBm and 13 dBm at 0.1 dBm steps. Note that in this work, the channel launch power levels are adjusted by the lasers, but they can also be controlled by other means such as using a WSS or VOA. The EDFAs cascaded are of different brands and models to emulate a more complex combination of gain tilts. Wideband VOAs are used to simulate fiber propagation loss and match the per-span average amplifier gain of around 15 dB per channel. A fourth VOA with -25 dB attenuation per channel is used to reduce the optical power due to the input power limitation of the OPM, which is used to record channel power levels after the EDFA-VOA spans and communicate with the computer system implementing the database and ML engine. The database records channel usage and power variance data, and the ML Engine trains on the data collected and controls the WDM sources with computed channel power adjustments. Figure 6-16 shows the widely discrepant channel power levels measured by the OPM when all channels are launched at a uniform power of 10 dBm.
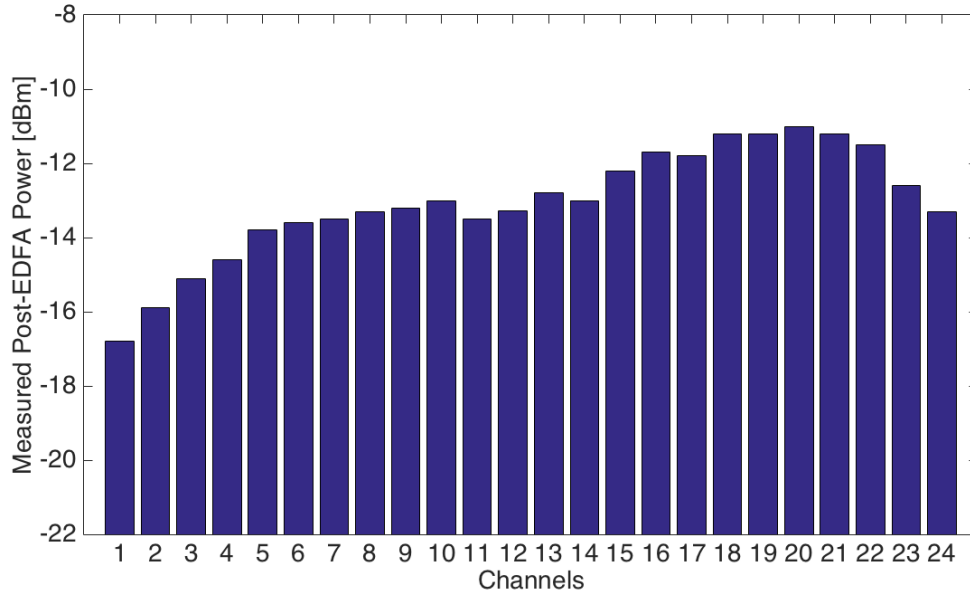
**Figure 6-16: Discrepant post-EDFA power levels from 24 channels launched at uniform power.**

We can deduce a correlation between Figure 6-13 and Figure 6-16 – channels with post-EDFA power much higher or lower than the mean both contribute positively to the post-EDFA power variance, while channels with post-EDFA power near the mean contribute negatively to the variance. Information about the post-EDFA power spectrum is also captured by the LR model, shown in Figure 6-14, whose values range from -1 to 1 and indicate how the power levels of other channels compare to a specific channel. Hence, the dot product between $w_{LR}^{ch}$ and the 24-bit ON/OFF vector gives an estimate on whether Channel $ch$ power is high or low relative to the other ON channels. A power adjustment, $\Delta P_{ch}$, can then be determined as:

$$\Delta P_{ch} = sgn\left(\sigma\left(xw_{LR}^{ch}\right) - 0.5\right) \cdot \frac{w_{RR}(ch)}{\max\left(|w_{RR}|\right)} \cdot P_{step}, \qquad (\,6\text{-}10\,)$$

where $P_{step}$ is a power tuning unit that we predefine as 3 dBm, and $\Delta P_{ch}$ is scaled by the ratio between the channel's RR weight and the maximum RR weight among all channels. It is also evident in Figure 6-14 that the LR model was unable to capture finer details for Ch. 1-5, and 15-22, which are respectively at the low and high ends of the gain levels. We expect the finer details can be extracted by the LR model if more

system data is available, specifically with cases containing only Ch. 1-5 or Ch. 15-22 as ON. In our demonstration, the lack of details in the extreme low and high ends of the spectrum by the LR model does not impact the ML engine's performance significantly, since the current set of LR weights still results in positive adjustments for Ch. 1-5 and negative adjustments for Ch. 15-22.

*Hop defragmentation*

We implement a wavelength grooming scenario for a super-channel in two experiments using the Hop defragmentation method. The super-channel is emulated by assigning three spectrally adjacent channels into a single channel entity that is relocated together, with each sub channel's power adjusted individually by the ML engine. In the first experiment, as illustrated in Figure 6-17, a super-channel is relocated from Ch. 9-11 to Ch. 1-3, while Ch. 13-24 are in use. This is a representative case since the new location of the super-channel promotes a significant increase in the post-EDFA power discrepancy, according to Figure 6-13. The Hop method is a single step procedure in our implementation, and Figure 6-19(a) shows immediate change in the variance of post-EDFA channel power levels. We show that, by implementing the ML engine to perform a single-step power adjustment on the relocated super-channel, the post-EDFA power variance is improved by 63%, without re-measuring channel power levels during the defragmentation.



**Figure 6-17: Illustration of the first defragmentation experiment – relocated super-channel is shown in green; other ON channels are shown in blue.**

In the second experiment, we evaluate a more complex spectrum usage shown in Figure 6-18, in which a super-channel is relocated from Ch. 18-20 to Ch. 1-3. This defragmentation procedure opens up an available band at Ch. 17-20, which can accommodate a super-channel of four channel spacing that is

154

previously unavailable. In addition, the original location of the super-channel experiences the higher end of the gain tilt, while the new location of the super-channel experiences the lower end of the gain tilt, which examines how the ML engine performs with channels shifting between extreme ends of the gain spectrum. Figure 6-19(b) illustrates that with power adjustments by the ML engine, the post-EDFA variance is improved by 62%, effectively maintaining the same post-EDFA power discrepancy before and after the defragmentation.



**Figure 6-18: Illustration of the second defragmentation experiment for Hop and MbB – relocated super-channel is shown in green; other ON channels are shown in blue.**
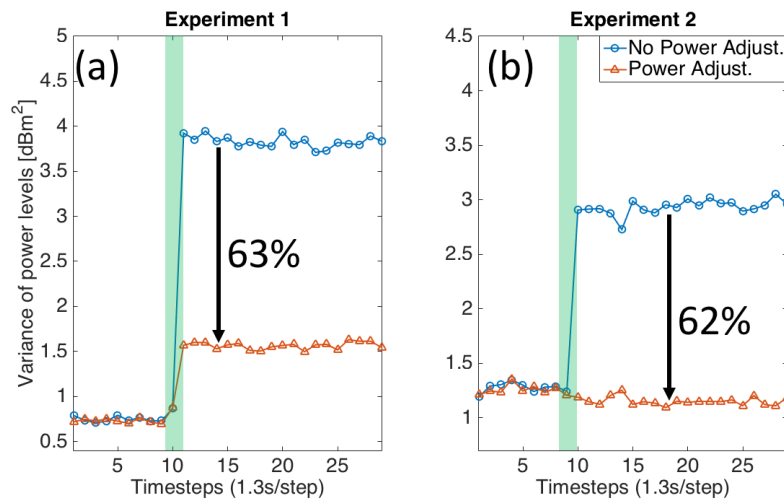


**Figure 6-19: Comparison of post-EDFA power variance with and without ML-enabled power adjustments in two experiments of Hop defragmentation, the duration of which is shaded in green.**

*MbB defragmentation*

We repeat the two experiments with the MbB defragmentation method, which is treated as a two-step process – a new super-channel is turned on, allowing the network traffic to transfer over before turning off the original super-channel. The intermediate stage of MbB captures when both the original and new channels co-exist in the system. The ML engine performs two power adjustments on the new super-channel based on the intermediate and the final channel usage. As shown in Figure 6-20(a) and Figure 6-20(b), we observe improvements in post-EDFA power variance both at the intermediate stage and after the defragmentation for both experiments. In both cases, the final post-EDFA power variances after the MbB defragmentation improve by 62% and match well with the Hop defragmentation method, showing consistency of the ML engine across different defragmentation methods.
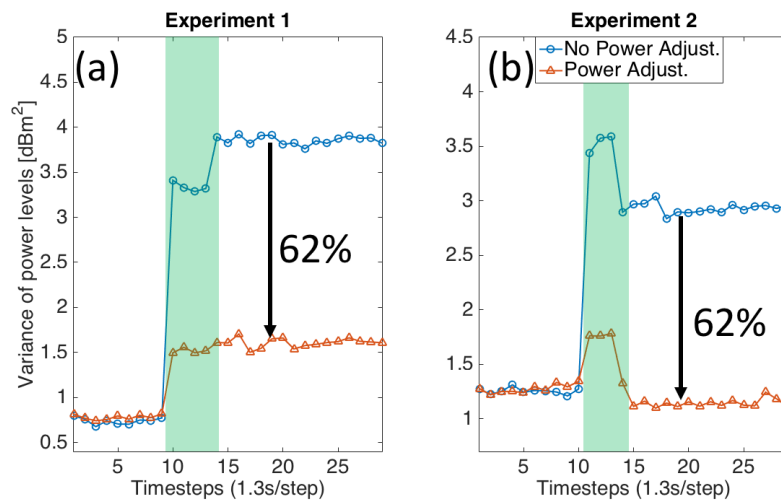


**Figure 6-20: Comparison of post-EDFA power variance with and without ML-enabled power adjustments in two experiments of MbB defragmentation, the duration of which is shaded in green.**

The two defragmentation experiments are again repeated with the Sweep defragmentation method, in which the super-channel is shifted at the spectral granularity of the system until it reaches its new spectral location. In the first case, a super-channel is swept from Ch. 9-11 to Ch. 1-3, while Ch. 13-24 are in use. The super-channel is moved directly since the continuous spectrum between the start and end locations is available. During the experiment, the central frequency of super-channel is shifted every 1.3 seconds at 100 GHz steps, limited by the measurement sampling frequency of the OPM and the spectral granularity of the system. The ML engine performs a power adjustment on the super-channel at every step based on the spectrum usage at that instance. Figure 6-22(a) shows that, with ML-enabled power adjustments, the change in post-EDFA power variance is greatly suppressed and optimized. After the defragmentation completes, the ML engine helps to achieve a 75% reduction in post-EDFA power variance. The greater improvement over Hop and MbB is due to the multi-step process of Sweep, which allows the ML engine to perform more adjustments throughout the process.

In the second experiment, in which a super-channel is relocated from Ch. 18-20 to Ch. 1-3, existing channels prevent the super-channel from being directly swept across the spectrum. Hence, we perform a sequential sweep of every channel in the system to the lower wavelength, effectively maximizing the available continuous bandwidth of the spectrum, as shown in Figure 6-21. We illustrate the evolution of the post-EDFA power variance throughout this process with and without ML-enabled power adjustments in Figure 6-22(b), and show that the ML engine drastically suppresses the change in power variance. At the end of the defragmentation, the spectrum variance is improved by 89%, resulting in a set of much less dispersed post-EDFA channel power levels.
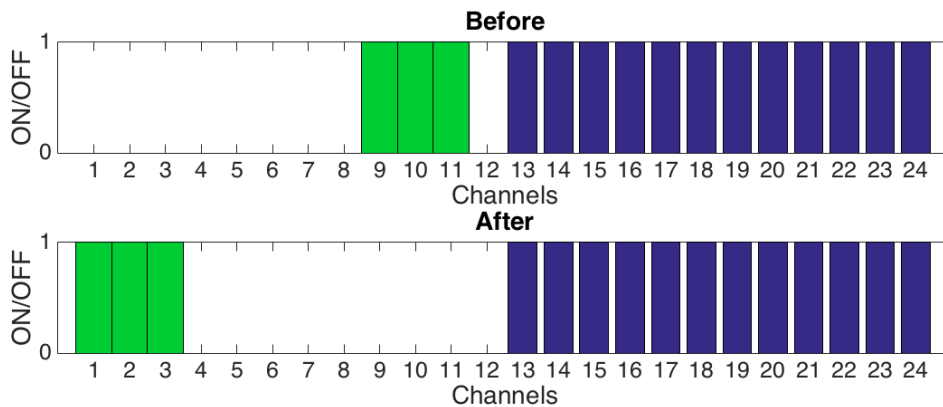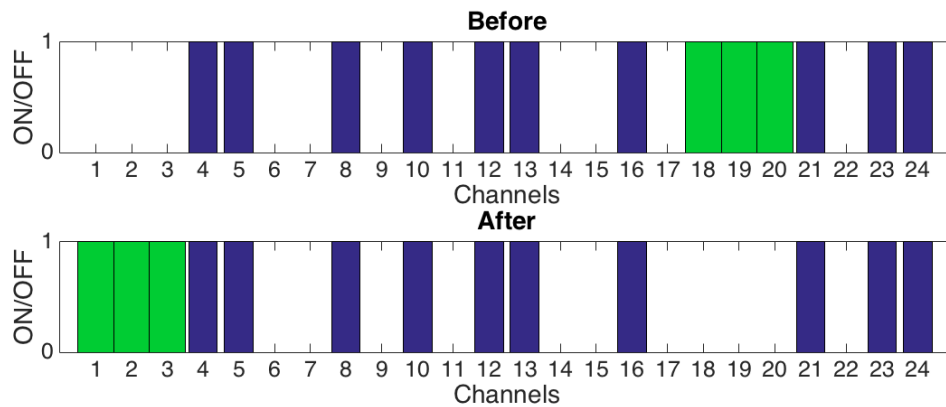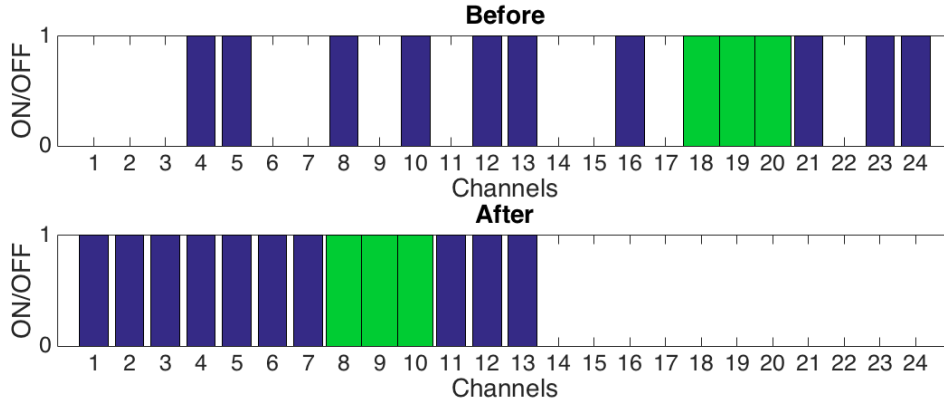
**Figure 6-21: Illustration of the second defragmentation experiment for Sweep – relocated super-channel is shown in green; other ON channels are shown in blue.**
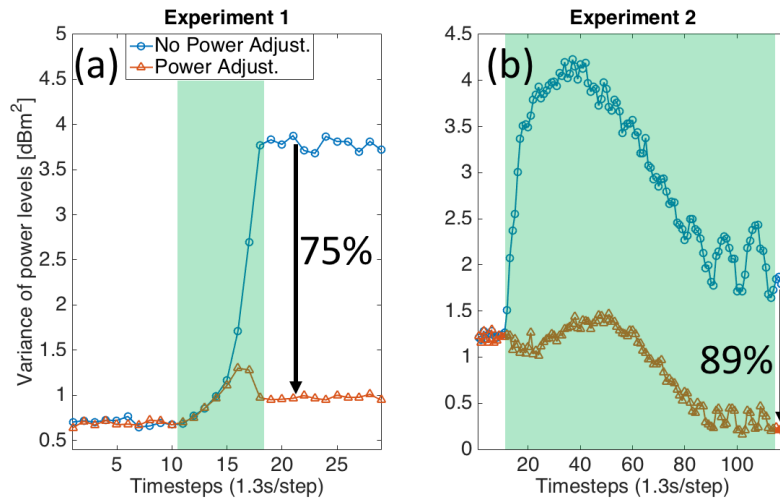


**Figure 6-22: Comparison of post-EDFA power variance with and without ML-enabled power adjustments in two experiments of Sweep defragmentation, the duration of which is shaded in green.**

### 6.4.4   Scalability of the Approach

The ML engine presented effectively reduces the post-EDFA power discrepancy during and after the defragmentation process and is applicable to all three main defragmentation methods. Here we discuss the potential to scale the ML engine for larger and more complex networks. One limitation of using variance as the optimized metric of the flexgrid optical network is that the discrepancy amongst the channel power levels, instead of their mean, is captured. This may overlook crucial aspects such as lightpath power penalty as the ML engine promotes equalizing rather than maximizing channel power levels. A potential solution to this concern can be to set appropriate bounds within which the ML engine can modify the channel power, and thus avoiding adjusting a channel's power too low or too high. Another solution can consider a joint metric to achieve both high mean and low variance among post-EDFA power levels.

For optical networks with a variety of channel injection and termination points, such as a mesh network, a distributed implementation of the ML engine is possible. Individual ML engines can be trained on each network edge and focus on power discrepancy optimization of the specific edges. Power adjustments can be performed at the network nodes where channels enter the edge. The low-complexity and small-footprint operations of the RR and LR models, in addition to their memory efficient weights, encourage parallel and distributed operations of multiple ML engines in a scaled-up and complex network.

### 6.4.5 Summary

Maintaining the channel power stability during fast changing spectrum utilizations is crucial to ensuring the QoS of flexgrid optical networking. We introduce an ML engine to preserve channel power consistency during defragmentation process of a flexgrid network experiencing EDFA power excursions. The proposed ML engine employees low-complexity ML models in a fully automated workflow, which extracts EDFA power dynamics and performs power adjustments without iterative power measurements. Experimentally we demonstrate the effectiveness of the ML engine in diverse spectral usage scenarios, and show consistent performance and applicability among three main defragmentation methods – Hop, MbB, and Sweep. In addition, we explain possible improvements and scalable implementations of the ML engine for larger, more complex networks.

## 6.5   Chapter Summary

In this chapter, we examine how ML-based approaches can be designed and integrated into optical system control plane and show real efficacy in extracting and predicting system dynamics and further drive performance optimizations for optical networking applications. We introduce an ML engine to address the challenge of EDFA power excursion in WDM networks, and show that the ML-enabled control plane can accurately predict the EDFA system's power response to channel changes and avoid channel configurations that would trigger significant power excursions. In addition, the ML engine is adapted to address more complicated EDFA power excursions for flexgrid networks and the wavelength defragmentation process. The ML engine is shown to recommend precise and single-step power pre-adjustments that can reduce the post-EDFA power variance for main spectrum defragmentation methods.

We want to highlight key aspects of the design philosophy when applying ML techniques to the optical system control plane. First, the ML models implemented are light weight and scalable in deployment with larger systems, which ensures economic integration with the control plane without incurring high cost or system upgrades. Second, the ML techniques are employed to recommend single step adjustments to the system, which represents significantly speed-up from conventional feedback-based monitoring and control systems. Third, we avoid using hidden or internalized features in the ML approach and make sure the learned ML feature weights can be interpreted directly as characteristics of the underlying system, in order to provide additional insights to optical systems that are hard to describe analytically.

# Chapter 7: Final Remarks

## 7.1    Summary of Contributions

As optical connectivity becomes the physical medium of choice for datacenter networking, interconnect systems based on photonic switching have the potential to deliver high-bandwidth and modulation indifferent routing that can enable future networking applications and emergent datacenter designs. Silicon photonic switch fabrics offer unique advantages in tunability and manufacturability, but still face unsolved challenges in functionality, scalability, and system control. In this work, we holistically explore all segments of building and functionalizing silicon photonic switch designs. We introduce a novel design of space-and-wavelength selective SE with MRR-assisted MZI structure in Chapter 2, which allows spatial and spectral switching granularities to be integrated at the building block level for a switch fabric. The high performance and improved switching agility of the design pave way for highly compact and simplified large-scale space-and-wavelength selective switching systems without the need for (de)multiplexers and parallel switching planes. We further demonstrate the first design and characterization of an 8×8 multi-stage silicon switch fabric in Chapter 3 based on MRR SEs, which enable about 100× reduction in the footprint compared to MZI-based SEs and therefore much greater integration density. The analysis of the moderate-radix switch fabrics helps us emphasize key performance metrics for MRR-based SEs and their impact on the scalability of silicon switch design. To that end, we introduce a Clos switch topology built with MRR-based switch-and-select sub-switches. By eliminating first-order crosstalk and controlling the number of cascaded MRR SEs, we show that the switch architecture is feasible for moderate- to high-radix switch designs.

Furthermore, we address key challenges in calibration, control, and functionalization techniques for silicon switch devices in RNB topologies in Chapter 4. Calibration algorithms are developed to expeditiously determine precise control points for each individual SE in a switch fabric without on-chip

power monitors and the associated PIC design overhead. We further develop calibration functionalities that address fabrication variations of the photonic structures, and experimentally verify their capability in improving crosstalk suppression in MZI-based SEs. To improve end-to-end performance of the switch fabric and address SE- and path-dependent impairments, we explore common switch network topologies and identify conditions for routing redundancies in Chapter 5. Routing strategies are developed to exploit these redundancies in RNB switch fabrics and are shown to be highly effective in avoiding poorly performing switch lightpaths and elevate the bottom-line performance of the switch device. In Chapter 6, we discuss control plane architectures that incorporate ML techniques for system characterization, prediction, and recommendation, and validate our approach in experimental applications addressing EDFA power excursions during WDM channel provisioning and flexgrid network channel defragmentation. We emphasize the light overhead, single-step provisioning capability, and actionable insights as key design benefits for applying ML techniques to monitor and optimize optical systems.

## 7.2 Recommendations for Future Work

While this work presents a compressive study on designing high performance silicon photonic switch building blocks, architectures, control techniques, and routing strategies, there are tremendous opportunities to continue evolving silicon switch designs in research and development. The future works recommended here can be largely categorized into five areas:

### *Component design*

As discussed in Chapters 2 – 4 of this work, the end-to-end performance of a silicon switch device heavily depends on the accumulated impairments along a lightpath. With thousands of optical components monolithically integrated in a high-radix switch fabric, improvements in the performance of individual components, including optical coupler, waveguide, crossing, MRR SE, and MZI SE, can translate to significant improvement in the end-to-end switching penalty. State-of-the-art silicon components have shown impressive performance in intrinsic loss, with waveguide propagation loss <0.5 dB/cm [146], waveguide crossing loss <0.01 dB [92], Y-junction coupler loss <0.1 dB [147], thermo-optic MZI element loss <0.2 dB [39], and coupling loss <1 dB [148]. In addition, among the available layers for silicon photonics design flow, SiN shows great potential in realizing both ultra-low loss waveguide [149] and switch waveguide shuffle without crossing structures [77]. Integration of low-loss silicon photonic components and SiN structures can continue to enable novel and scalable silicon switch architectures.

### *Device design*

New features can be incorporated into silicon photonic switch devices to enhance the designs' functionalities. In particular, polarization diversity [150], space-and-wavelength selectivity [65] [51], and lossless hybrid switching [29] are promising ideas. Because the silicon waveguides are typically suited for TE polarization, polarization diversity silicon switch can gain affinity with polarization multiplexed applications. Space-and-wavelength selectivity, as discussed in Chapter 2, is crucial for WDM applications.

Hybrid switch devices, which combine the benefits of both III-V gain blocks and silicon switch fabric, can perhaps achieve a good balance among footprint, cost, power consumption, and lightpath penalty when performing optical switching.

*Packaging*

Silicon switch devices can easily scale to hundreds of optical ports and thousands of electrical control signals, and therefore the switch PIC needs to be designed for packaging. On the optical side, grating coupled fiber arrays can achieve low coupling loss, but are wavelength sensitive and take up large real estate on the PIC, while edge coupled fiber arrays offer broadband optical access and minimal on-chip footprint, but can only achieve decent mode matching with the PIC waveguides with expensive lensed tips. New optical coupling techniques leveraging high-$\Delta$ silica coupler [39], polymer alignment structure [151], or spot-size converter [4] can offer economic, low-loss, and broadband coupling options. On the electrical side, wire bonding to PCB is an option for small to medium size devices, but for large-radix switch devices, electrical pads can exceed the immediate sections near the PIC perimeter and make wire bonding difficult. Flip-chip bonding configurations such as chip-on-PCB or chip-on-interposer-on-PCB are promising options for electrical access for a variety of PIC sizes and electrical pad densities, but the sheer number of electrical pads in larger devices would require development of reliable and high yield process.

*Control plane*

The switch control plane is an integral part of a photonic switch system. The control plane incorporates features such as performance monitoring, calibration, actuation, and routing computation, as we have discussed in this work, as well as integration with higher layer protocols and applications [12][152]. A robust and flexible control plane design can motivate further adoption of optical switching as a viable candidate for networking applications.

*Applications*

While a commonly proposed application for photonic switching is in communication networks to circumvent the pin and power limitations imposed by electronic switches, integrated silicon photonic switch fabrics can play important roles in a myriad of areas such as quantum computing [153], optical computing [154], and opto-genetics [155]. The requirements for these non-communication applications generally align with the ones we have explored in this work – low lightpath penalty, consistent performance among lightpaths, reduced system footprint, and robust control. Likewise, the techniques and studies performed in this work are extensible to address the design challenges of silicon photonic switch fabrics in other applications.

# Bibliography

[1] "Cisco Annual Internet Report," Cisco, 2020. [Online]. Available: https://www.cisco.com/c/en/us/solutions/executive-perspectives/annual-internet-report/index.html.

[2] D. A. B. Miller, "Device requirements for optical interconnects to silicon chips," *Proceedings of the IEEE,* vol. 97, no. 7, 2009.

[3] "Reinventing Facebook's data center network," Facebook, 2019. [Online]. Available: https://engineering.fb.com/data-center-engineering/f16-minipack/.

[4] S. Fathololoumi, K. Nguyen, H. Mahalingam, M. Sakib, Z. Li, C. Seibert, M. Montazeri, J. Chen, J. K. Doylend, H. Jayatilleka, C. Jan, J. Heck, R. Venables, H. Frish, R. A. Defrees, R. S. Appleton, S. Hollingsworth, S. McCargar, R. Jones, D. Zhu, Y. Akulova, and L. Liao, "1.6Tbps Silicon photonics integrated circuit for co-packaged optical-IO switch applications," in *Optical Fiber Communication Conference*, 2020.

[5] A. Ghiasi, "Large data centers interconnect bottlenecks," *Optics Express,* vol. 23, no. 3, 2015.

[6] B. G. Lee and N. Dupuis, "Silicon photonic switch fabrics: technology and architecture," *Journal of Lightwave Technology,* vol. 37, no. 1, 2018.

[7] Z. Zhu, Y. Shen, Y. Huang, A. Gazman, M. Hattink and K. Bergman, "Flexible resource allocation using photonic switched interconnects for disaggregated system architectures," in *Optical Fiber Communication Conference*, 2019.

[8] P. X. Gao, A. Narayan, S. Karandikar, J. Carreira, S. Han, R. Agarwal, S. Ratnasamy and S. Shenker, "Network requirements for resource disaggregation," in *USENIX Conference on Operating Systems Design and Implementation*, 2016.

[9] Q. Cheng, M. Bahadori, M. Glick, S. Rumley and K. Bergman, "Recent advances in optical technologies for data centers: a review," *Optica,* vol. 5, no. 11, 2018.

[10] K. Wen, P. Samadi, S. Rumley, C. P. Chen, Y. Shen, M. Bahadori, J. Wilke and K. Bergman, "Flexfly: enabling a reconfigurable dragonfly through silicon photonics," in *International Conference for High Performance Computing, Networking, Storage and Analysis*, 2016.

[11] M. Hochberg, N. C. Harris, R. Ding, Y. Zhang, A. Novack, Z. Xuan and T. Baehr-Jones, "Silicon photonics: the next fabless semiconductor industry," *IEEE Solid-State Circuits Magazine,* vol. 5, no. 1, 2013.

[12] Y. Shen, X. Meng, Q. Cheng, S. Rumley, N. Abrams, A. Gazman, E. Manzhosov, M. S. Glick and K. Bergman, "Silicon photonics for extreme scale systems," *Journal of Lightwave Technology,* vol. 37, no. 2, 2019.

[13] "Silicon Photonics Multi-Project Wafer (MPW) Service," Advanced Micro Foundry, 2018. [Online]. Available: http://www.advmf.com/wp-content/uploads/2018/09/MPW-Brochure-2018.pdf.

[14] J. Komma, C. Schwarz, G. Hofmann, D. Heinert and R. Nawrodt, "Thermo-optic coefficient of silicon at 1550nm and cryogenic temperatures," *Applied Physics Letters,* vol. 101, 2012.

[15] R. Soref and B. Bennett, "Electrooptical effects in silicon," *IEEE Journal of Quantum Electronics ,* vol. 23, no. 1, 1987.

[16] J. Sun, R. Kumar, M. Sakib, J. B. Driscoll, H. Jayatilleka and H. Rong, "A 128 Gb/s PAM4 silicon microring modulator with integrated thermo-optic resonance tuning," *Journal of Lightwave Technology,* vol. 37, no. 1, 2019.

[17] G. T. Reed and C. E. J. Png, "Silicon optical modulators," *Nature Photonics,* vol. 4, 2010.

[18] J. Zhou, J. Wang, L. Zhu and Q. Zhang, "High baud rate all-silicon photonics carrier depletion modulators," *Journal of Lightwave Technology,* vol. 38, no. 2, 2020.

[19] L. Qiao, W. Tang and T. Chu, "32 × 32 silicon electro-optic switch with built-in monitors and balanced-status units," *Scientific Reports,* vol. 7, 2017.

[20] Z. Guo, L. Lu, L. Zhou, L. Shen and J. Chen, "16 × 16 Silicon Optical Switch Based on Dual-Ring-Assisted Mach–Zehnder Interferometers," *Journal of Lightwave Technology,* vol. 36, no. 2, 2018.

[21] P. D. Dobbelaere, K. Falta, L. Fan, S. Gloeckner and S. Patra, "Digital MEMS for optical switching," *IEEE Communications Magazine ,* vol. 40, no. 3, 2002.

[22] "Optical Circuit Switch," Calient Technologies, 2020. [Online]. Available: https://www.calient.net/products/s-series-photonic-switch/.

[23] J. Kim, C. Nuzman, B. S. Kumar, D. Lieuwen, J. Kraus, A. Weiss, C. Lichtenwalner, A. Papazian, R. Frahm, N. Basavanhally, D. Ramsey, V. Aksyuk, F. Pardo, M. Simon, V. Lifton, H. B. Chan, M. Haueis, A. Gasparyan, H. R. Shea, and J. V. Gates, "1100 x 1100

port MEMS-based optical crossconnect with 4-dB maximum loss," *IEEE Photonics Technology Letters,* vol. 15, no. 11, 2003.

[24]    "Polatis Optical Switches," Huber+Suhner, 2020. [Online]. Available: https://www.hubersuhner.com/en/products/fiber-optics/optical-switches/polatis-optical-switches.

[25]    Q. Cheng, A. Wonfor, J. L. Wei, R. V. Penty and I. H. White, "Low-energy, high-performance lossless 8×8 SOA switch," in *Optical Fiber Communication Conference*, 2015.

[26]    S. Tanaka, S.-H. Jeong, S. Yamazaki, A. Uetake, S. Tomabechi, M. Ekawa and K. Morito, "Monolithically integrated 8:1 SOA gate switch with large extinction ratio and wide input power dynamic range," *IEEE Journal of Quantum Electronics,* vol. 45, no. 9, 2009.

[27]    R. Stabile, A. Albores-Mejia, A. Rohit and K. A. Williams, "Integrated optical switch matrices for packet data networks," *Microsystems & Nanoengineering,* vol. 2, 2016.

[28]    Q. Cheng, A. Wonfor, R. V. Penty and I. H. White, "Scalable, low-energy hybrid photonic space switch," *Journal of Lightwave Technology,* vol. 31, no. 18, 2013.

[29]    N. Dupuis, F. Doany, R. A. Budd, L. Schares, C. W. Baks, D. M. Kuchta, T. Hirokawa and B. G. Lee, "A 4×4 electrooptic silicon photonic switch fabric with net neutral insertion loss," *Journal of Lightwave Technology,* vol. 38, no. 2, 2020.

[30]    Q. Cheng, A. Wonfor, J. L. Wei, R. V. Penty and I. H. White, "Monolithic MZI-SOA hybrid switch for low-power and low-penalty operation," *Optics Letters,* vol. 39, no. 6, 2014.

[31]    M. Ding, A. Wonfor, Q. Cheng, R. V. Penty and I. H. White, "Hybrid MZI-SOA InGaAs/InP photonic integrated switches," *IEEE Journal of Selected Topics in Quantum Electronics,* vol. 24, no. 1, 2018.

[32]    E. Murphy, T. Murphy, A. Ambrose, R. Irvin, B. Lee, P. Peng, G. Richards and A. Yorinks, "16x16 strictly nonblocking guided-wave optical switching system," *Journal of Lightwave Technology,* vol. 14, no. 3, 1996.

[33]    P. J. Duthie and M. J. Wale, "16*16 single chip optical switch array in lithium niobate," *Electronics Letters,* vol. 27, no. 14, 1991.

[34]    T. Goh, M. Yasu, K. Hattori, A. Himeno, M. Okuno and Y. Ohmori, "Low loss and high extinction ratio strictly nonblocking 16x16 thermooptic matrix switch on 6-in wafer using

silica-based planar lightwave circuit technology," *Journal of Lightwave Technology,* vol. 19, no. 3, 2001.

[35]   S. Sohma, T. Watanabe, N. Ooba, M. Itoh, T. Shibata and H. Takahashi, "Silica-based PLC type 32 x 32 optical matrix switch," in *European Conference on Optical Communications*, 2006.

[36]   G. Baxter, S. Frisken, D. Abakoumov, H. Zhou, I. Clarke, A. Bartos and S. Poole, "Highly programmable wavelength selective switch based on liquid crystal on silicon switching elements," in *Optical Fiber Communication Conference*, 2006.

[37]   B. Robertson, H. Yang, M. M. Redmond, N. Collings, J. R. Moore, J. Liu, A. M. Jeziorska-Chapman, M. Pivnenko, S. Lee, A. Wonfor, I. H. White, W. A. Crossland and D. P. Chu, "Demonstration of multi-casting in a 1 × 9 LCOS wavelength selective switch," *Journal of Lightwave Technology,* vol. 32, no. 3, 2014.

[38]   D. Celo, D. J. Goodwill, J. Jiang, P. Dumais, C. Zhang, F. Zhao, X. Tu, C. Zhang, S. Yan, J. He, M. Li, W. Liu, Y. Wei, D. Geng, H. Mehrvar and E. Bernier, "32×32 silicon photonic switch," in *OptoElectronics and Communications Conference*, 2016.

[39]   K. Suzuki, R. Konoike, J. Hasegawa, S. Suda, H. Matsuura, K. Ikeda, S. Namiki and H. Kawashima, "Low-insertion-loss and power-efficient 32 × 32 silicon photonics switch with extremely high-Δ silica PLC connector," *Journal of Lightwave Technology,* vol. 37, no. 1, 2019.

[40]   S. Han, T. J. Seok, N. Quack, B.-W. Yoo and M. C. Wu, "Large-scale silicon photonic switches with movable directional couplers," *Optica,* vol. 2, no. 4, 2015.

[41]   T. J. Seok, K. Kwon, J. Henriksson, J. Luo and M. C. Wu, "240×240 wafer-scale silicon photonic switches," in *Optical Fiber Communication Conference*, 2019.

[42]   T. J. Seok, N. Quack, S. Han, W. Zhang, R. S. Muller and M. C. Wu, "64×64 Low-loss and broadband digital silicon photonic MEMS switches," in *European Conference on Optical Communication*, 2015.

[43]   N. Sherwood-Droz, H. Wang, L. Chen, B. G. Lee, A. Biberman, K. Bergman and M. Lipson, "Optical 4x4 hitless silicon router for optical networks-on-chip (NoC)," *Optics Express,* vol. 16, no. 20, 2008.

[44]   Q. Cheng, L. Y. Dai, N. C. Abrams, Y.-H. Hung, P. E. Morrissey, M. Glick, P. O'Brien and K. Bergman, "Ultralow-crosstalk, strictly non-blocking microring-based optical switch," *Photonics Research,* vol. 7, no. 2, 2019.

[45] P. Dasmahapatra, R. Stabile, A. Rohit and K. Williams, "Optical crosspoint matrix using broadband resonant switches," *IEEE Journal of Selected Topics in Quantum Electronics,* vol. 20, no. 4, 2014.

[46] A. S. P. Khope, A. M. Netherton, T. Hirokawa, N. Volet, E. J. Stanton, C. Schow, R. Helkey, A. A. M. Saleh, J. E. Bowers and R. C. Alferness, "Elastic WDM optoelectronic crossbar switch with on-chip wavelength control," in *Photonics in Switching*, 2017.

[47] Q. Cheng, M. Bahadori, Y.-H. Hung, Y. Huang, N. Abrams and K. Bergman, "Scalable microring-based silicon Clos switch fabric with switch-and-select stages," *IEEE Journal of Selected Topics in Quantum Electronics,* vol. 25, no. 5, 2019.

[48] A. S. P. Khope, M. Saeidi, R. Yu, X. Wu, A. M. Netherton, Y. Liu, Z. Zhang, Y. Xia, G. Fleeman, A. Spott, S. Pinna, C. Schow, R. Helkey, L. Theogarajan, R. C. Alferness, A. A. M. Saleh and J. E. Bowers, "Multi-wavelength selective crossbar switch," *Optics Express,* vol. 27, no. 4, 2019.

[49] L. Lu, L. Zhou, X. Li and J. Chen, "Low-power 2×2 silicon electro-optic switches based on double-ring assisted Mach–Zehnder interferometers," *Optics Letters,* vol. 39, no. 6, 2014.

[50] L. Lu, L. Zhou, Z. Li, D. Li, S. Zhao, X. Li and J. Chen, "4 × 4 silicon optical switches based on double-ring-assisted Mach–Zehnder interferometers," *IEEE Photonics Technology Letters,* vol. 27, no. 23, 2015.

[51] Y. Huang, Q. Cheng, A. Rizzo and K. Bergman, "High-performance microring-assisted space-and-wavelength selective switch," in *Optical Fiber Communication Conference*, 2020.

[52] T. J. Seok, K. Kwon, J. Henriksson, J. Luo and M. C. Wu, "Wafer-scale silicon photonic switches beyond die size limit," *Optica,* vol. 6, no. 4, 2019.

[53] W. Bogaerts, P. DeHeyn, T. V. Vaerenbergh, K. DeVos, S. KumarSelvaraja, T. Claes, P. Dumon, P. Bienstman, D. VanThourhout and R. Baets, "Silicon microring resonators," *Laser & Photonics Reviews,* vol. 6, no. 1, 2012.

[54] T. Chu, L. Qiao, W. Tang, D. Guo and W. Wu, "Fast, high-radix silicon photonic switches," in *Optical Fiber Communication Conference*, 2018.

[55] K. Tanizawa, K. Suzuki, M. Toyama, M. Ohtsuka, N. Yokoyama, K. Matsumaro, M. Seki, K. Koshino, T. Sugaya, S. Suda, G. Cong, T. Kimura, K. Ikeda, S. Namiki and H. Kawashima, "Ultra-compact 32 × 32 strictly non-blocking Si-wire optical switch with fan-out LGA interposer," *Optics Express,* vol. 23, no. 13, 2015.

[56] Y. Huang, Q. Cheng, Y.-H. Hung, H. Guan, X. Meng, A. Novack, M. Streshinsky, M. Hochberg and K. Bergman, "Multi-stage 8×8 silicon photonic switch based on dual-microring switching elements," *Journal of Lightwave Technology,* vol. 38, no. 2, 2020.

[57] A. S. P. Khope, T. Hirokawa, A. M. Netherton, M. Saeidi, Y. Xia, N. Volet, C. Schow, R. Helkey, L. Theogarajan, A. A. M. Saleh, J. E. Bowers and R. C. Alferness, "On-chip wavelength locking for photonic switches," *Optics Letters,* vol. 42, no. 23, 2017.

[58] K. Padmaraju, D. F. Logan, T. Shiraishi, J. J. Ackert, A. P. Knights and K. Bergman, "Wavelength locking and thermally stabilizing microring resonators using dithering signals," *Journal of Lightwave Technology,* vol. 32, no. 3, 2014.

[59] C. Sun, M. Wade, M. Georgas, S. Lin, L. Alloatti, B. Moss, R. Kumar, A. H. Atabaki, F. Pavanello, J. M. Shainline, J. S. Orcutt, R. J. Ram, M. Popović and V. Stojanović, "A 45 nm CMOS-SOI monolithic photonics platform with bit-statistics-based resonant microring thermal tuning," *IEEE Journal of Solid-State Circuits,* vol. 51, no. 4, 2016.

[60] Y. Huang, Q. Cheng, Y.-H. Hung, H. Guan, A. Novack, M. Streshinsky, M. Hochberg and K. Bergman, "Dual-microring resonator based 8×8 silicon photonic switch," in *Optical Fiber Communication Conference*, 2019.

[61] R. Ji, L. Yang, L. Zhang, Y. Tian, J. Ding, H. Chen, Y. Lu, P. Zhou and W. Zhu, "Five-port optical router for photonic networks-on-chip," *Optics Express,* vol. 19, no. 21, 2011.

[62] G. Fan, R. Orobtchouk, B. Han, Y. Li and H. Li, "8 × 8 wavelength router of optical network on chip," *Optics Express,* vol. 25, no. 20, 2017.

[63] Q. Zhu, X. Jiang, Y. Yu, R. Cao, H. Zhang, D. Li, Y. Li, L. Zeng, X. Guo, Y. Zhang and C. Qiu, "Automated wavelength alignment in a 4 × 4 silicon thermo-optic switch based on dual-ring resonators," *IEEE Photonics Journal,* vol. 1, 10.

[64] R. Stabile, A. Rohit and K. A. Williams, "Monolithically integrated 8 × 8 space and wavelength selective cross-connect," *Journal of Lightwave Technology,* vol. 32, no. 2, 2013.

[65] T. J. Seok, J. Luo, Z. Huang, K. Kwon, J. Henriksson, J. Jacobs, L. Ochikubo, R. S. Muller and M. C. Wu, "MEMS-actuated 8×8 silicon photonic wavelength-selective switches with 8 wavelength channels," in *Conference on Lasers and Electro-Optics*, 2018.

[66] Q. Cheng, M. Bahadori, M. Glick and K. Bergman, "Scalable space-and-wavelength selective switch architecture using microring resonators," in *Conference on Lasers and Electro-Optics*, 2019.

172

[67] A. Rohit, J. Bolk, X. J. M. Leijtens and K. A. Williams, "Monolithic nanosecond-reconfigurable 4×4 space and wavelength selective cross-connect," *Journal of Lightwave Technology,* vol. 30, no. 17, 2012.

[68] H. Zhou, C. Qiu, X. Jiang, Q. Zhu, Y. He, Y. Zhang, Y. Su and R. Soref, "Compact, submilliwatt, $2 \times 2$ silicon thermo-optic switch based on photonic crystal nanobeam cavities," *Photonics Research,* vol. 5, no. 2, 2017.

[69] R. A. Soref, F. D. Leonardis and V. M. N. Passaro, ""Mach-Zehnder crossbar switching and tunable filtering using N-coupled waveguide Bragg resonators," *Optics Express,* vol. 26, no. 12, 2018.

[70] T. Dai, G. Wang, X. Guo, C. Bei, J. Jiang, W. Chen, Y. Wang, H. Yu and J. Yang, "Scalable bandwidth-tunable micro-ring filter based on multi-channel-spectrum combination," *IEEE Photonics Technology Letters,* vol. 30, no. 11, 2018.

[71] Q. Cheng, Y. Huang, H. Yang, M. Bahadori, N. Abrams, X. Meng, M. Glick, Y. Liu, M. Hochberg and K. Bergman, "Silicon photonic switch topologies and routing strategies for disaggregated data centers," *IEEE Journal of Selected Topics in Quantum Electronics,* vol. 10.1109/JSTQE.2019.2960950, 2019.

[72] C.-M. Chang, G. d. Valicourt, S. Chandrasekhar and P. Dong, "Differential microring modulators for intensity and phase modulation: theory and experiments," *Journal of Lightwave Technology,* vol. 35, no. 15, 2017.

[73] J. Cardenas, P. A. Morton, J. B. Khurgin, A. Griffith, C. B. Poitras, K. Preston and M. Lipson, "Linearized silicon modulator based on a ring assisted Mach Zehnder interferometer," *Optics Express,* vol. 21, no. 19, 2013.

[74] Z. Wang, S.-J. Chang, C.-Y. Ni and Y. J. Chen, "A high-performance ultra-compact optical interleaver based on double-ring assisted Mach-Zehnder interferometer," *IEEE Photonics Technology Letters,* vol. 19, no. 14, 2007.

[75] Q. Cheng, S. Rumley, M. Bahadori and K. Bergman, "Photonic switching in high performance datacenters [Invited]," *Optics Express,* vol. 26, no. 12, 2018.

[76] M. Bohn, P. Magill, M. Hochberg, D. Scordo, A. Novack and M. Streshinsky, "Next-generation silicon photonic interconnect solutions," in *Optical Fiber Communication Conference*, 2019.

[77] Q. Cheng, L. Y. Dai, M. Bahadori, N. C. Abrams, P. E. Morrissey, M. Glick, P. O'Brien and K. Bergman, "Si/SiN microring-based optical router in switch-and-select topology," in *European Conference on Optical Communication*, 2018.

[78] W. Kabacinski, Nonblocking Electronic and Photonic Switching Fabrics, Springer, 2005.

[79] V. E. Beneš, Mathematical Theory of Connecting Networks and Telephone Traffic, Academic Press, 1965.

[80] D. C. Opferman and N. T. Tsao-wu, "On a class of rearrangeable switching networks, Part I: control algorithm," *The Bell System Technical Journal*, vol. 50, no. 5, 1971.

[81] Q. Cheng, Y. Huang, M. Bahadori, J. Zhou, M. Glick and K. Bergman, "Fabric-wide, penalty-optimized path routing algorithms for optical integrated switches," in *Optical Fiber Communication Conference*, 2019.

[82] Q. Cheng, Y. Huang, M. Bahadori, M. Glick, S. Rumley and K. Bergman, "Advanced routing strategy with highly-efficient fabric-wide characterization for optical integrated switches," in *International Conference on Transparent Optical Networks*, 2018.

[83] C. Clos, "A study of non-blocking switching networks," *The Bell System Technical Journal,* vol. 32, no. 2, 1953.

[84] A. Singh, J. Ong, A. Agarwal, G. Anderson, A. Armistead, R. Bannon, S. Boving, G. Desai, B. Felderman, P. Germano, A. Kanagala, J. Provost, J. Simmons, E. Tanda, J. Wanderer, U. Hölzle, S. Stuart and A. Vahdat, "Jupiter Rising: a decade of Clos topologies and centralized control in Google's datacenter network," in *SIGCOMM*, 2015.

[85] D. Nikolova, S. Rumley, D. Calhoun, Q. Li, R. Hendry, P. Samadi and K. Bergman, "Scaling silicon photonic switch fabrics for data center interconnection networks," *Optics Express,* vol. 23, no. 2, 2015.

[86] A. Bianco, D. Cuda, R. Gaudino, G. Gavilanes, F. Neri and M. Petracca, "Scalability of optical interconnects based on microring resonators," *IEEE Photonics Technology Letters,* vol. 22, no. 15, 2010.

[87] M. Bahadori, M. Nikdast, S. Rumley, L. Y. Dai, N. Janosik, T. V. Vaerenbergh, A. Gazman, Q. Cheng, R. Polster and K. Bergman, "Design space exploration of microring resonators in silicon photonic interconnects: impact of the ring curvature," *Journal of Lightwave Technology,* vol. 36, no. 13, 2018.

[88] C. L. Manganelli, P. Pintus, F. Gambini, D. Fowler, M. Fournier, S. Faralli, C. Kopp and C. J. Oton, "Large-FSR thermally tunable double-ring filters for WDM applications in silicon photonics," *IEEE Photonics Journal,* vol. 9, no. 1, 2017.

[89] "AIM Photonics," [Online]. Available: http://www.aimphotonics.com/.

[90]  M. Bahadori, A. Gazman, N. Janosik, S. Rumley, Z. Zhu, R. Polster, Q. Cheng and K. Bergman, "Thermal rectification of integrated microheaters for microring resonators in silicon photonics platform," *Journal of Lightwave Technology,* vol. 36, no. 3, 2018.

[91]  Y. Ma, Y. Zhang, S. Yang, A. Novack, R. Ding, A. E.-J. Lim, G. Lo, T. Baehr-Jones and M. Hochberg, "Ultralow loss single layer submicron silicon waveguide crossing for SOI optical interconnect," *Optics Express,* vol. 21, no. 24, 2013.

[92]  Y. Zhang, A. Hosseini, X. Xu, D. Kwong and R. T. Chen, "Ultralow-loss silicon waveguide crossing using Bloch modes in index-engineered cascaded multimode-interference couplers," *Optics Letters,* vol. 38, no. 18, 2013.

[93]  Y. Liu, J. M. Shainline, X. Zeng and M. A. Popović, "Ultra-low-loss CMOS-compatible waveguide crossing arrays based on multimode Bloch waves and imaginary coupling," *Optics Letters,* vol. 39, no. 2, 2014.

[94]  M. Bahadori, S. Rumley, H. Jayatilleka, K. Murray, N. Jaeger, L. Chrostowski, S. Shekhar and K. Bergman, "Crosstalk penalty in microring-based silicon photonic interconnect systems," *Journal of Lightwave Technology,* vol. 34, no. 17, 2016.

[95]  W. Sacher, J. Mikkelsen, Y. Huang, J. Mak, Z. Yong, X. Luo, Y. Li, P. Dumais, J. Jiang, D. Goodwill, E. Bernier, P. Lo and J. Poon, "Monolithically integrated multilayer silicon nitride-on-silicon waveguide platforms for 3-D photonic circuits and devices," *Proceedings of the IEEE,* vol. 106, no. 12, 2018.

[96]  A. Annoni, E. Guglielmi, M. Carminati, S. Grillanda, P. Ciccarella, G. Ferrari, M. Sorel, M. J. Strain, M. Sampietro, A. Melloni and F. Morichetti, "Automated routing and control of silicon photonic switch fabrics," *IEEE Journal of Selected Topics in Quantum Electronic,* vol. 22, no. 6, 2016.

[97]  Y. Huang, Q. Cheng, N. C. Abrams, J. Zhou, S. Rumley and K. Bergman, "Automated calibration and characterization for scalable integrated optical switch fabrics without built-in power monitors," in *European Conference on Optical Communication*, 2017.

[98]  Y. Huang, Q. Cheng and K. Bergman, "Automated calibration of balanced control to optimize performance of silicon photonic switch fabrics," in *Optical Fiber Communication Conference*, 2018.

[99]  Y. Huang, Q. Cheng and K. Bergman, "Crosstalk-aware calibration for fast and automated functionalization of photonic integrated switch fabrics," in *Conference on Lasers and Electro-Optics*, 2018.

[100]  Y. Huang, Q. Cheng and K. Bergman, "Advanced control for crosstalk minimization in MZI-based silicon photonic switches," in *IEEE Optical Interconnects Conference*, 2018.

[101] J. C. Mikkelsen, W. D. Sacher and J. K. S. Poon, "Dimensional variation tolerant silicon-on-insulator directional couplers," *Optics Express,* vol. 22, no. 3, 2014.

[102] N. Dupuis, B. G. Lee, A. V. Rylyakov, D. M. Kuchta, C. W. Baks, J. S. Orcutt, D. M. Gill, W. M. J. Green and C. L. Schow, "Design and fabrication of low-insertion-loss and low-crosstalk broadband 2×2 Mach-Zehnder silicon photonic switches," *Journal of Lightwave Technology,* vol. 33, no. 17, 2015.

[103] P. Dumais, D. J. Goodwill, D. Celo, J. Jiang, C. Zhang, F. Zhao, X. Tu, C. Zhang, S. Yan, J. He, M. Li, W. Liu, Y. Wei, D. Geng, H. Mehrvar and E. Bernier, "Silicon photonic switch subsystem with 900 monolithically integrated calibration photodiodes and 64-fiber package," *Journal of Lightwave Technology,* vol. 36, no. 2, 2018.

[104] M. S. Hai, M. Moayedi Pour Fard, D. An, F. Gambini, S. Faralli, G. Preve and O. Liboiron-Ladouceur, "Automated characterization of SiP MZI-based switches," in *IEEE Optical Interconnects Conference*, 2015.

[105] D. A. B. Miller, "Setting up meshes of interferometers - reversed local light interference method," *Optics Express,* vol. 25, no. 23, 2017.

[106] S. Suda, H. Matsuura, K. Tanizawa, K. Suzuki, K. Ikeda, H. Kawashima and S. Namiki, "Fast and accurate automatic calibration of a 32 × 32 silicon photonic strictly-non-blocking switch," in *Photonics in Switching*, 2017.

[107] A. Gazman, E. Manzhosov, C. Browning, M. Bahadori, Y. London, L. Barry and K. Bergman, "Tapless and topology agnostic calibration solution for silicon photonic switches," *Optics Express,* vol. 26, no. 25, 2018.

[108] H. Çam and J. A. B. Fortes, "Work-efficient routing algorithms for rearrangeable symmetrical networks," *IEEE Transactions on Parallel and Distributed Systems,* vol. 10, no. 7, 1999.

[109] K. Y. Lee, "A new Beneš network control algorithm," *IEEE Transactions on Computers,* vol. 36, no. 6, 1987.

[110] M. Ding, Q. Cheng, A. Wonfor, R. V. Penty and I. H. White, "Routing algorithm to optimize loss and IPDR for rearrangeably non-blocking integrated optical switches," in *Conference on Lasers and Electro-Optics*, 2015.

[111] Y. Qian, H. Mehrvar, H. Ma, X. Yang, K. Zhu, H. Fu, D. Geng, D. Goodwill, P. Dumais and E. Bernier, "Crosstalk optimization in low extinction-ratio switch fabrics," in *Optical Fiber Communication Conference*, 2014.

[112] Q. Cheng, Y. Huang, M. Bahadori, J. Zhou, M. Glick and K. Bergman, "Fabric-wide, penalty-optimized path routing algorithms for integrated optical switches," in *Optical Fiber Communication Conference*, 2019.

[113] D. Zheng, J. D. Doménech, W. Pan, X. Zou, L. Yan and D. Pérez, "Low-loss broadband 5 × 5 non-blocking Si3N4 optical switch matrix," *Optics Letters,* vol. 44, no. 11, 2019.

[114] J. Xing, Z. Li, P. Zhou, X. Xiao, J. Yu and Y. Yu, "Nonblocking 4x4 silicon electro-optic switch matrix with push-pull drive," *Optics Letters,* vol. 38, no. 19, 2013.

[115] R. A. Spanke and V. E. Benes, "N-stage planar optical permutation network," *Applied Optics,* vol. 26, no. 7, 1987.

[116] T. Shimoe, K. Hajikano and K. Murakami, "Path-independent insertion loss optical space switch," in *Optical Fiber Communication Conference*, 1987.

[117] I. H. White, K. A. Williams, R. V. Penty, T. Lin, A. Wonfor, E. T. Aw, M. Glick, M. Dales and D. McAuley, "Control architecture for high capacity multistage photonic switch circuits," *Journal of Optical Networking,* vol. 6, no. 2, 2007.

[118] Q. Cheng, M. Bahadori and K. Bergman, "Advanced path mapping for silicon photonic switch fabrics," in *Conference on Lasers and Electro-Optics*, 2017.

[119] N. Dupuis and B. G. Lee, "Impact of topology on the scalability of Mach–Zehnder-based multistage silicon photonic switch networks," *Journal of Lightwave Technology,* vol. 36, no. 3, 2018.

[120] Q. Chen, F. Zhang, R. Ji, L. Zhang and L. Yang, "Universal method for constructing N-port non-blocking optical router based on 2 × 2 optical switch for photonic networks-on-chip," *Optics Express,* vol. 22, no. 10, 2014.

[121] S. Rumley, M. Bahadori, K. Wen, D. Nikolova and K. Bergman, "PhoenixSim: crosslayer design and modeling of silicon photonic interconnects," in *International Workshop on Advanced Interconnect Solutions and Technologies for Emerging Computing Systems*, 2016.

[122] I. d. Miguel, R. J. Durán, T. Jiménez, N. Fernández, J. C. Aguado, R. M. Lorenzo, A. Caballero, I. T. Monroy, Y. Ye, A. Tymecki, I. Tomkos, M. Angelou, D. Klonidis, A. Francescon, D. Siracusa and E. Salvadori, "Cognitive dynamic optical networks," *Journal of Optical Communications and Networking,* vol. 5, no. 10, 2013.

[123] D. Rafique and L. Velasco, "Machine learning for network automation: overview, architecture,and applications [Invited Tutorial]," *Journal of Optical Communications and Networking,* vol. 10, no. 10, 2018.

[124] D. Côté, "Using machine learning in communication networks [Invited]," *Journal of Optical Communications and Networking,* vol. 10, no. 10, 2018.

[125] D. Zibar, L. H. H. d. Carvalho, M. Piels, A. Doberstein, J. Diniz, B. Nebendahl, C. Franciscangelis, J. Estaran, H. Haisch, N. G. Gonzalez, J. C. R. F. d. Oliveira and I. T. Monroy, "Application of machine learning techniques for amplitude and phase noise characterization," *Journal of Lightwave Technology,* vol. 33, no. 7, 2015.

[126] D. C. Kilper, M. Bhopalwala, H. Rastegarfar and W. Mo, "Optical power dynamics in wavelength layer software defined networking," in *Advanced Photonics*, 2015.

[127] A. S. Ahsan, C. Browning, M. S. Wang, K. Bergman, D. C. Kilper and L. P. Barry, "Excursion-free dynamic wavelength switching in amplified optical networks," *Journal of Optical Communications and Networking,* vol. 7, no. 9, 2015.

[128] A. K. Srivastava, Y. Sun, J. L. Zyskind and J. W. Sulhoff, "EDFA transient response to channel loss in WDM transmission system," *IEEE Photonics Technology Letters,* vol. 9, no. 3, 1997.

[129] C. Tian and S. Kinoshita, "Analysis and control of transient dynamics of EDFA pumped by 1480- and 980-nm lasers," *Journal of Lightwave Technology,* vol. 21, no. 8, 2003.

[130] D. A. Mongardien, S. Borne, C. Martinelli, C. Simonneau and D. Bayart, "Managing channels add/drop in flexible networks based on hybrid raman / Erbium amplified spans," in *European Conference on Optical Communication*, 2006.

[131] E. A. Barboza, C. J. A. Bastos-Filho, J. F. Martins-Filho, U. C. d. Moura and J. R. F. d. Oliveira, "Self-adaptive Erbium-doped fiber amplifiers using machine learning," in *International Microwave & Optoelectronics Conference*, 2013.

[132] P. J. Lin, "Reducing optical power variation in amplified optical network," in *International Conference on Communication Technology*, 2003.

[133] N. Sambo, F. Cugini, G. Bottari, P. Iovanna and P. Castoldi, "Routing and spectrum assignment for super-channels in flex-grid optical networks," in *European Conference on Optical Communication*, 2012.

[134] J. Junio, D. C. Kilper and V. W. S. Chan, "Channel power excursions from single-step channel provisioning," *Journal of Optical Communications and Networking,* vol. 4, no. 9, 2012.

[135] K. Ishii, J. Kurumida and S. Namiki, "Wavelength assignment dependency of AGC EDFA gain offset under dynamic optical circuit switching," in *Optical Fiber Communication Conference*, 2014.

[136] U. Moura, M. Garrich, H. Carvalho, M. Svolenski, A. Andrade, F. Margarido, A. C. Cesar, E. Conforti and J. Oliveira, "SDN-enabled EDFA gain adjustment cognitive methodology for dynamic optical networks," in *European Conference on Optical Communication*, 2015.

[137] C. M. Bishop, Pattern Recognition and Machine Learning, Springer, 2007.

[138] J. D. Reis, M. Garrich, D. M. Pataca, J. C. M. Diniz, V. N. Rozental, L. H. H. Carvalho, E. C. Magalhães, U. Moura, N. G. Gonzalez, J. R. F. Oliveira and J. C. R. F. Oliveira, "Flexible optical transmission systems for future networking," in *International Telecommunications Network Strategy and Planning Symposium*, 2014.

[139] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone and S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies," *IEEE Communications Magazine,* vol. 47, no. 11, 2009.

[140] H. Zang, J. P. Jue and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Optical Networks Magazine,* vol. 1, no. 1, 2000.

[141] R. Wang and B. Mukherjee, "Provisioning in elastic optical networks with non-disruptive defragmentation," *Journal of Lightwave Technology,* vol. 31, no. 15, 2013.

[142] M. Zhang, C. You, H. Jiang and Z. Zhu, "Dynamic and Adaptive Bandwidth Defragmentation in Spectrum-Sliced Elastic Optical Networks with Time-Varying Traffic," *Journal of Lightwave Technology,* vol. 32, no. 5, 2014.

[143] M. Zhang, Y. Yin, R. Proietti, Z. Zhu and S. J. B. Yoo, "Spectrum defragmentation algorithms for elastic optical networks using hitless spectrum retuning techniques," in *Optical Fiber Communication Conference*, 2013.

[144] T. Takagi, H. Hasegawa, K. Sato, Y. Sone, A. Hirano and M. Jinno, "Disruption minimized spectrum defragmentation in elastic optical path networks that adopt distance adaptive modulation," in *European Conference on Optical Communication*, 2011.

[145] F. Cugini, F. Paolucci, G. Meloni, G. Berrettini, M. Secondini, F. Fresi, N. Sambo, L. Potì and P. Castoldi, "Push-pull defragmentation without traffic disruption in flexible grid optical networks," *Journal of Lightwave Technology,* vol. 31, no. 1, 2013.

[146] T. Horikawa, D. Shimura and T. Mogami, "Low-loss silicon wire waveguides for optical integrated circuits," *MRS Communications,* vol. 6, no. 1, 2016.

[147] Z. Sheng, Z. Wang, C. Qiu, L. Li, A. Pang, A. Wu, X. Wang, S. Zou and F. Gan, "A compact and low-loss MMI coupler fabricated with CMOS technology," *IEEE Photonics Journal,* vol. 4, no. 6, 2012.

[148] X. Wang, X. Quan, M. Liu and X. Cheng, "Silicon-nitride-assisted edge coupler interfacing with high numerical aperture fiber," *IEEE Photonics Technology Letters,* vol. 31, no. 5, 2019.

[149] X. Ji, F. A. S. Barbosa, S. P. Roberts, A. Dutt, J. Cardenas, Y. Okawachi, A. Bryant, A. L. Gaeta and M. Lipson, "Ultra-low-loss on-chip resonators with sub-milliwatt parametric oscillation threshold," *Optica,* vol. 4, no. 6, 2017.

[150] H. Yang, Q. Cheng, R. Chen and K. Bergman, "Polarization-diversity microring-based optical switch fabric in a switch-and-select architecture," in *Optical Fiber Communication Conference*, 2020.

[151] O. A. J. Gordillo, S. Chaitanya, Y.-C. Chang, U. D. Dave, A. Mohanty and M. Lipson, "Plug-and-play fiber to waveguide connector," *Optics Express,* vol. 27, no. 15, 2019.

[152] Y. Shen, P. Samadi, Z. Zhu, A. Gazman, E. Anderson, D. Calhoun, M. Hattink and K. Bergman, "Software-defined networking control plane for seamless integration of silicon photonics in Datacom networks," in *European Conference on Optical Communication*, 2017.

[153] X. Qiang, X. Zhou, J. Wang, C. M. Wilkes, T. Loke, S. O'Gara, L. Kling, G. D. Marshall, R. Santagati, T. C. Ralph, J. B. Wang, J. L. O'Brien, M. G. Thompson and J. C. F. Matthews, "Large-scale silicon quantum photonics implementing arbitrary two-qubit processing," *Nature Photonics,* vol. 12, 2018.

[154] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund and M. Soljačić, "Deep learning with coherent nanophotonic circuits," *Nature Photonics,* vol. 11, 2017.

[155] A. Mohanty, Q. Li, M. A. Tadayon, S. P. Roberts, G. R. Bhatt, E. Shim, X. Ji, J. Cardenas, S. A. Miller, A. Kepecs and M. Lipson, "Reconfigurable nanophotonic silicon probes for sub-millisecond deep-brain optical stimulation," *Nature Biomedical Engineering,* vol. 4, no. 2, 2020.