

Optical End-to-End Implementations for Scalable Photonic Computing



Yi Zhang

Wolfson College

University of Oxford

Supervisor: Prof. Harish Bhaskaran

A thesis submitted in partial fulfilment of the requirements for the
degree of

Doctor of Philosophy

Michaelmas 2025

This thesis is dedicated to
my mother Lan Li (李兰) and father Weiguo Zhang (张伟国)
for their continued love and support

Acknowledgements

My DPhil comes to an end in a peaceful winter at twenty-six, after an autumn crowded with thesis writing, journal preparation, an internship, and a constant sense of anxiety about the future. I tried to plan everything ahead—lunch breaks, daily exercise, experiments, manuscripts, and career steps. Looking back, this now feels like the slightly stubborn determination of a teenager who believes she can hold everything in her hands.

The only certainty in the world is uncertainty. People and events rarely follow the schedules we construct. What I can do is accept that anything may happen, planned or unplanned, and allow things and people to move through my life rather than insisting they remain as I imagined. This journey has taught me to appreciate uncertainty, to steady my mind, and to move with change rather than resist it. Spring that follows winter is always more captivating than spring itself.

First, I would like to express my sincere thanks to **Prof. Harish Bhaskaran** for welcoming me into the group and opening the door to a new field. I am grateful for the inspiring conversations we shared over lunch and during coffee walks, and in particular for the supervisor's ice cream. I deeply appreciate the effort he invested in helping me grow into an independent researcher, as well as the patient guidance he offered again and again to strengthen my reasoning, precision, and insight. His own curiosity and lifelong commitment to exploration have always motivated me, and these lessons will stay with me long beyond this thesis.

I would like to thank **Dr. Nikolaos Farmakidis**, from whom I learned most of my experimental skills and understanding. Thank you for guiding me step by step and for setting an example of enthusiasm and dedication to scientific goals. My heartfelt thanks go to **Dr. Mengyun Wang**, who embraced me with kindness from the moment I joined the group. I treasure our dinners at Wolfson and our café adventures across Oxford, during which she always tried to buy me a coffee on our walks. My thanks also go to **Dr. Yu Shu**, who invited me into her beautiful wedding as the after-party host and fulfilled her promise of the bouquet by helping me find a boyfriend. I am sincerely grateful to **Dr. June Sang Lee**, my desk mate, who was always willing to listen, to discuss project details, to stop me when I was heading in the wrong direction, and to help me sketch the bigger picture. Many thanks to **Dr. Samarth Aggarwal**, who was always ready to help and to explore Oxford libraries and Indian restaurants with me. And a very warm thank you to **Dr. Bowei Dong**, who is both talented and humble, and who always listened to my ideas and analysed them thoughtfully with me.

I would also like to say a big thank you to **Dr. Yuhan He** for the time we spent fighting alongside each other as we grew into independent researchers. From the same high school to the same group and the same office in Oxford, I feel truly fortunate to have had you as both a colleague and a friend.

My gratitude extends to all past and present members of the ANE group. I thank **Prof. Zengguang Cheng**, who first introduced me to the group, as well as **Håvard Hem Toftevaag** and **Kairan Huang**, my same-year companions who were always by my side. I would also like to thank **Dr. Mingde Du**, **Dr. Guoce Yang**, **Prof. Wen Zhou**, **Dr. Utku Emre Ali**, **Prof. Angel Ortega**, **Tara Milne**, **Maxine Ong**, **Sijing Zhong**, **Serena Sabnani**, **Sharon Oregel-Chaumont**, **Sirui Liu**, **Zhiyun Xu**, **Zhongyu Tang**, **Barry Hao Yu**, **Jock Dearlove**, **Ruveyda Taşkın**, **Sergei Zotkin**, **Antonios Mikallou**, **Andy Moskalenko**, and **Luci Bywater** for their support, help, and the many memorable moments we shared.

I want to express my special gratitude to **Wolfson College** for four years of support and care. My thanks go to **Paul** and **Martin**, our college minibus drivers, who shuttled me reliably between the college and the department. I also thank the people working behind the scenes who supported my research: **Paul Pattinson** and **Radka Chakalova** for maintaining the cleanroom, and **Paul Warren** for his helpful IT support. My gratitude also goes to my friend **Houke Chen**, who has shared both my joys and my challenges, and whose companionship during our weekly punting sessions has been a constant source of support.

I am grateful to my undergraduate project supervisor, **Prof. Wei Zhang**, for guiding me into the world of research and supporting my application to Oxford. His dedication and rigour have had a lasting impact on me. I also thank **Dr. Yuxing Zhou**, who provided unwavering support during my undergraduate research and set an example of what it means to be a researcher.

Of course, I owe my deepest gratitude to my beloved parents, **Lan Li** and **Weiguo Zhang**, whose unconditional support has sustained me throughout this journey. They have been my strongest foundation, enabling me to explore the world with confidence and without hesitation.

Many thanks to my partner, **Yufeng Zhao**, who has brought new colour into my life, cared for me, and shared every emotion with patience. I look forward to exploring the world together and embracing all the challenges that lie ahead.

Finally, I am grateful to myself—for never giving up and never turning away from who I am.

Yi Zhang
Oxford , 2025

Abstract

Integrated photonics is increasingly explored as a platform for high-speed and energy-efficient computing, yet its scalability remains constrained by the persistent reliance on electronic interfaces to perform the accumulation of multiplexed optical signals and to supply nonlinear activation, thereby limiting cascadability and system bandwidth. This thesis introduces a new class of photonic devices that repurpose heat—traditionally regarded as a parasitic and slow effect—into a functional computational mechanism. Central to this work is the Photonic-Heater-in-Lightpath (PHIL) architecture, in which subwavelength nanoheaters embedded within silicon microring resonators convert dissipated optical energy directly into controllable thermo-optic phase shifts.

The thesis establishes three key capabilities of this architecture. First, it presents the inaugural demonstration of cross-wavelength all-optical encoding and summation, showing that optical power carried at multiple control wavelengths can be selectively transduced onto a spectrally distinct probe wavelength. This establishes PHIL as a mechanism for wavelength-routed accumulation and transformation of optical signals. Second, by exploiting the intrinsic Lorentzian resonance profile of the microring, the device produces continuously reconfigurable nonlinear transfer characteristics. These analogue activation functions emulate neuronal behaviour and provide the nonlinear mapping required for neuromorphic photonic systems. Third, the slow thermal relaxation of the nanoheater is shown to operate as a physical analogue integrator for ultrafast optical signals. The device performs leaky temporal integration of high-repetition-rate pulse trains over a 250 ns thermal window, enabling temporal accumulation with concurrent nonlinear activation entirely within the optical domain.

Together, these results position opto-thermo-optic interactions as a foundational computational resource for integrated photonics. By unifying cross-wavelength encoding and summation, reconfigurable nonlinear activation, and analogue temporal integration within a single CMOS-compatible device, the PHIL platform establishes a coherent framework for all-optical processing across wavelength and time domains. The work demonstrates that slow thermal dynamics can be strategically engineered, not circumvented, to achieve functions difficult or impractical to realise using conventional high-speed modulators. The thesis concludes by outlining future opportunities enabled by this paradigm, including scalable wavelength-multiplexed neuromorphic processors, large-vector optical computing architectures, and fully optical computational pipelines built from densely integrated PHIL-based modules.

Contents

List of Figures	ix
List of Tables	xii
List of Abbreviations	xiii
1 Introduction	1
1.1 Motivation	1
1.1.1 Towards Neuromorphic Computing	1
1.1.2 Photonics for Neuromorphic Computing: Opportunities and Challenges	4
1.2 Objectives and Outline	6
1.2.1 Objectives	6
1.2.2 Outline	7
1.3 Key results	9
1.4 Statement of originality	10
2 Integrated Neuromorphic Photonic Computing: A review	12
2.1 Neuromorphic Photonic Computing	13
2.1.1 Signal Representation and Multiplexing	14
2.1.2 Photonic Implementation of Accumulation and Nonlinearity	16
2.2 Integrated Photonic Accelerator Architectures	17
2.2.1 Space-Division Multiplexing	17
2.2.2 Wavelength-Division Multiplexing	25
2.2.3 Time-Division Multiplexing	30
2.3 Optical Signal Accumulation And Nonlinearity Implementation	34
2.3.1 Through Opto-Electro-Optic Conversions	35
2.3.2 Within the Optical Domain	38
2.4 The Need for All-Optical Nonlinearity Implementation for Scalable Optical Neural Networks	41
2.4.1 Functional Gaps in Photonic Neural Networks	41
2.4.2 Standing-Wave Photonics as an Architectural Precursor	43
2.4.3 From Passive Modulation to All-Optical Computation	45
3 Theory	47
3.1 Dielectric Waveguides	47
3.2 Grating Couplers	51
3.3 Microring Resonators	53
3.3.1 Resonance Condition	55

3.3.2	Transmission and Critical Coupling	56
3.3.3	Spectral Characteristics	58
3.4	Coherent Perfect Absorption	62
4	Methodology: Simulation, Fabrication, and Experimental Setup	70
4.1	Lumerical Simulation	70
4.1.1	MODE Solutions (FDE Solver)	71
4.1.2	FDTD Solutions	71
4.1.3	HEAT Solutions	73
4.1.4	Simulation Workflow	74
4.2	Device Fabrication	75
4.2.1	General Process and Layer Overview	75
4.2.2	Alignment Markers	77
4.2.3	Photonic Layer	77
4.2.4	Antenna Layer	78
4.3	Experiment Configurations	79
4.3.1	Sample Stage	79
4.3.2	Setup for Phase-shift Measurement	80
4.3.3	Setup for Time Integration of Ultrafast Optical Signals . . .	82
5	All-Optical Cross-Wavelength Encoding	84
5.1	Motivation	84
5.2	Photonic-Heater-in-Lightpath (PHIL): Design and Simulations . . .	86
5.2.1	Realisation of Lossy and Lossless Spectral Bands	86
5.2.2	Resonant-Mode Selective Optical Absorption	89
5.2.3	Optically Controlled Nanoheaters	93
5.3	Cross-Wavelength Optical Modulation: Experimental Validation . .	95
5.3.1	PHIL-Loaded Microring Resonator	95
5.3.2	Resonant-Mode Selective Spectral Response	97
5.3.3	Thermo-Optic Cross-Wavelength Signal Encoding	99
5.4	Insertion Loss Analysis	104
5.4.1	Experimental Characterization	105
5.4.2	Nanoheater-Induced Insertion Loss	106
5.4.3	Tunable Absorption Profile	110
5.5	Discussion and Summary	111

6 All-Optical Neuron: Summation and Programmable Nonlinear Activation Across Wavelengths	115
6.1 Motivation	115
6.2 All-Optical Programmable Nonlinearity	117
6.2.1 Engineering of Optical Nonlinearity	117
6.2.2 Modulation Stability and Precision of Nonlinear Response .	121
6.3 Across-Wavelength Optical Accumulation	123
6.3.1 Linear Optical Accumulation Across Two-Wavelength Channels	124
6.3.2 Across-Wavelength Optical Accumulation with Nonlinear Activation	127
6.4 Further Scaling of WDM Channels	130
6.4.1 Optical Accumulation Across Three-Wavelength Channels .	130
6.4.2 Scalability of Multi-Wavelength Channels	135
6.5 Discussion and Summary	138
7 All-Optical Temporal Integration With Nonlinear Activation	140
7.1 Motivation	140
7.2 Temporal Dynamics of Cross-Wavelength Modulation	142
7.2.1 Modulation Speed Characterisation	142
7.2.2 A Leaky Integrator Model	145
7.3 Temporal Integration of Ultrafast Optical Signals with Nonlinear Activation	148
7.3.1 Numerical Simulation of Temporal Integration	148
7.3.2 Experimental Validation of All-Optical Temporal Integration	151
7.3.3 Neural Network Evaluation Using the PHIL Activation Function	154
7.4 Energy Efficiency	157
7.4.1 Comparison of All-Optical Light-Controlled Technologies . .	159
7.5 Discussion and Summary	161
8 Conclusions and Outlook	164
8.1 Conclusions	164
8.2 Outlook	167
List of publications	170
References	172

List of Figures

1.1	Exponential growth of AI computing demands.	2
1.2	Schematic representation of artificial neural networks.	3
2.1	Single neuron framework	13
2.2	Conceptual illustration of multiplexing strategies in integrated photonic neural networks	15
2.3	Integrated photonic architectures based on interferometric mesh . .	19
2.4	Schematics of diffractive optical neural network(DONN)	23
2.5	Schematics of micro-ring resonator (MRR) weight banks	27
2.6	Schematics of photonic crossbar array based on phase-change materials (PCMs)	30
2.7	Photonic accelerators multiplexing space, wavelength and time . . .	32
2.8	Optical/electrical/optical neuron	36
2.9	Wavelength-time neural network with non-linearities applied on-chip	37
2.10	All-optical neuron	39
2.11	Standing-wave engineered CPA in integrated photonics	44
3.1	Dielectric waveguide	48
3.2	Grating coupler	52
3.3	Schematic illustrations of microring resonators (MRR) with coupling parameters.	54
3.4	Coherent perfect absorption via standing wave decomposition . . .	67
4.1	FDTD simulation layout of the PHIL microring device	72
4.2	HEAT simulation layout	73
4.3	Device layout	75
4.4	Fabrication process flow	78
4.5	sample stage	80
4.6	Device characterization setup	81
4.7	Experiment setup for all-optical integration of time-multiplexed signal	83
5.1	Concept of all-optical cross-wavelength encoding	86
5.2	Input lights with equal intensities are sent into silicon waveguide from opposite directions	88
5.3	Zoomed-in view of the nanoheater region	89
5.4	Wavelength-selective optical power absorption	90
5.5	Simulation of PHIL-loaded MRR	91
5.6	Wavelength-selective absorption in PHIL-loaded MRR	93

5.7	Temperature profile of the nanoheater-loaded region	94
5.8	Optical micrograph of the fabricated PHIL-loaded device	96
5.9	PHIL-loaded MRR and transmission spectrum	97
5.10	Spectral response under absorptive-mode excitation	98
5.11	Spectral response under non-absorptive-mode excitation	99
5.12	Control signal characterization setup	100
5.13	Cross-wavelength optical modulation	101
5.14	Phase shift induced by varying CS wavelengths across the absorptive mode	103
5.15	Measured insertion loss of PHIL with referenced to same design without nanoheater	106
5.16	Simulated transmission spectra and absorption profiles in a compact PHIL architecture with varying nanoheater numbers	107
5.17	Simulated insertion loss and attenuation contrast for varying nanoheater numbers	109
5.18	Electrically tunable spectral control in a PHIL-loaded MRR	112
6.1	Concept of all-optical programmable nonlinearity	117
6.2	Setup for nonlinear optical response engineering	118
6.3	Measured modulation traces of probe transmission at different wave- lengths	119
6.4	Extracted nonlinear transfer functions from probe modulation	120
6.5	Modulation stability of the nonlinear optical response	122
6.6	Statistical evaluation of nonlinear modulation stability	123
6.7	Concept of all-optical summation across wavelengths	124
6.8	Setup for incoherent optical addition of two control signals with different wavelengths.	125
6.9	Transmission spectrum of the PHIL device showing two absorptive resonances (<i>a</i> , <i>b</i>) used for multi-wavelength accumulation.	126
6.10	Input power of control signals through the measurement	127
6.11	Measured all-optical summation across two wavelengths	128
6.12	Across-wavelength accumulation with optically programmable trans- fer functions	129
6.13	Setup for incoherent optical addition of three control signals with different wavelengths.	131
6.14	Transmission spectrum of the PHIL device showing three absorptive resonances (<i>a-c</i>) used for multi-wavelength accumulation.	132
6.15	All-optical accumulation across three wavelengths.	133
6.16	Cross-sectional analysis of three-wavelength accumulation under var- ied third-channel power	134

7.1	Square pulse response characterization.	143
7.2	Temporal response of the PHIL device under square-pulse excitation.	144
7.3	Rise and fall time analysis of the temporal response.	145
7.4	Leaky integrator model based on the extracted thermal time constant.	146
7.5	Simulation and experimental demonstration of all-optical temporal integration.	148
7.6	Numerical simulation of thermo-optic leaky integration on ultrafast pulses.	149
7.7	Concept of temporal integration on optical signals	150
7.8	Experiment setup for all-optical integration of time-multiplexed signal	151
7.9	Experimental demonstration of optical temporal integration at 50 GHz	153
7.10	Nonlinear transfer characteristics and precision analysis of optical temporal integration.	154
7.11	PHIL-Based Activation Function Implemented in a Neural Network for Handwritten Digit Recognition.	155
7.12	Performance Comparison Between PHIL- and ReLU-Activated Neural Networks for Handwritten Digit Recognition	156

List of Tables

7.1	Comparison of All-Optical Light-Controlled Technologies	160
-----	---	-----

List of Abbreviations

ADC	Analogue-to-digital converter
AI	Artificial intelligence
ANN	Artificial neural network
ASE	Amplified spontaneous emission
AWG	Arbitrary waveform generator
CMOS	Complementary metal–oxide–semiconductor
CPA	Coherent perfect absorption
CW	Continuous wave
DAC	Digital-to-analogue converter
DEMUX	Wavelength-division demultiplexer
DNN	Deep neural network
DWDM	Dense wavelength-division multiplexing
EAM	Electro-absorption modulator
EBL	Electron-beam lithography
EDFA	Erbium-doped fibre amplifier
EIM	Effective index method
E-O	Electro-optical
EOM	Electro-optical modulator
ENOB	Effective number of bits
FCA	Free-carrier absorption
FCD	Free-carrier dispersion
FDE	Finite-difference eigenmode
FDTD	Finite-difference time-domain
FSR	Free spectral range
FWHM	Full width at half maximum
GST	Ge ₂ Sb ₂ Te ₅ phase-change alloy
I/O	Input–output
IPA	Isopropanol
ITO	Indium tin oxide
LN	Lithium niobate

MAC	Multiply–accumulate
MMI	Multimode interferometer
MRR	Microring resonator
MUX	Wavelength-division multiplexer
MVM	Matrix–vector multiplication
MZI	Mach–Zehnder interferometer
NRMSE	Normalised root mean square error
OBPF	Optical bandpass filter
ONN	Optical neural networks
O-E-O	Optical–electrical–optical
PCMs	Phase-change materials
PD	Photodetector
PHIL	Photonic-heater-in-lightpath
ReLU	Rectified linear unit
RF	Radio frequency
RIE	Reactive ion etching
SDM	Space-division multiplexing
SEM	Scanning electron microscope
Si	Silicon
SiO₂	Silicon dioxide
SNN	Spiking neural networks
SOA	Semiconductor optical amplifier
SOI	Silicon-on-insulator
SNR	Signal-to-noise ratio
TDM	Time-division multiplexing
TE	Transverse electric
TFLN	Thin-film lithium niobate
TIA	Transimpedance amplifier
TIR	Total internal reflection
TM	Transverse magnetic
TOPS	Tera operations per second

VCSEL	Vertical-cavity surface-emitting laser
WDM	Wavelength-division multiplexing
WGM	Whispering gallery mode

1

Introduction

Contents

1.1 Motivation	1
1.1.1 Towards Neuromorphic Computing	1
1.1.2 Photonics for Neuromorphic Computing: Opportunities and Challenges	4
1.2 Objectives and Outline	6
1.2.1 Objectives	6
1.2.2 Outline	7
1.3 Key results	9
1.4 Statement of originality	10

1.1 Motivation

1.1.1 Towards Neuromorphic Computing

The exponential growth of artificial intelligence and data-intensive applications has exposed the inefficiencies of conventional von Neumann architectures[1], in which the physical separation of memory and processing units imposes a fundamental bottleneck on both speed and energy efficiency. As the scale and complexity of deep neural networks increase, the computational cost of training state-of-the-art models now doubles every few months, far outpacing the historical improvement rate described by Moore’s law (Figure 1.1)[2]. This rapid escalation in compute demand underscores a widening disparity between algorithmic progress and the

physical limitations of electronic hardware, driving the urgent search for alternative computing paradigms that can sustain future AI development[3, 4].




Figure 1.1: The evolution of training computational cost for major AI models from 1985 to 2025. Since the advent of GPU computing and deep learning architectures, the computational demand has doubled every few months and requiring petaFLOP-scale resources. This trend highlights the growing mismatch between algorithmic complexity and hardware energy efficiency, motivating the search for new computing paradigms beyond conventional electronics[2]. Reproduced with permission from Mehonic & Kenyon, *Nature* (2022), © Springer Nature.

Neuromorphic computing emerged as a conceptual response to this challenge. Proposed by Carver Mead[5] in the early 1990s, it aims to emulate the brain's co-localized memory and processing units to overcome the von Neumann bottleneck. The human brain exemplifies extreme computational efficiency: it performs on the order of 10^{15} operations per second while consuming less than 20 W of power, far beyond any electronic supercomputer[6, 7]. The efficiency arises from the collective dynamics of billions of neurons and synapses that operate in parallel and exchange information continuously through analogue signals rather than discrete digital states[8].

Building on the early models of McCulloch–Pitts neurons[9] and Hebbian learning[10], artificial neural networks (ANNs) abstract this behaviour into simplified

computational units. Each neuron performs a weighted summation of inputs followed by a nonlinear activation (Figure 1.2(a)), and when assembled into multilayer architectures, these units can approximate arbitrarily complex functions (Figure 1.2(b)), forming the conceptual and algorithmic basis of modern deep learning.

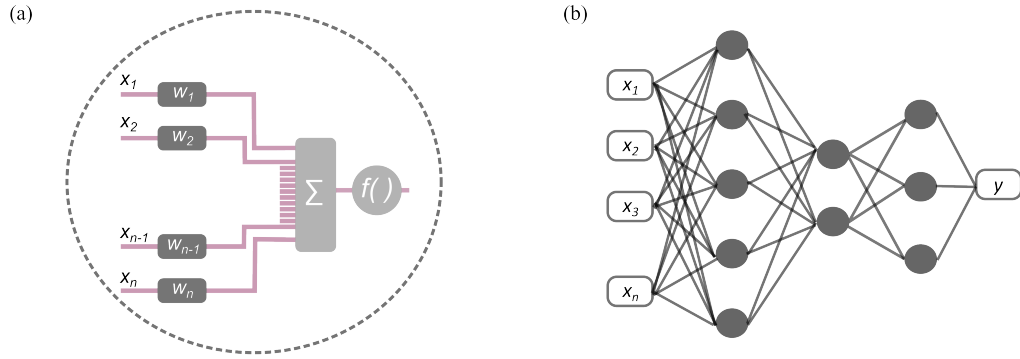


Figure 1.2: Schematic representation of artificial neural networks. (a) Functional abstraction of a single artificial neuron comprising a weighted summation unit and a nonlinear activation function $f(\cdot)$. Each neuron computes $y = f(\sum_i w_i x_i)$, enabling hierarchical feature extraction and nonlinear mapping capabilities characteristic of deep learning architectures. (b) A typical multilayer feed-forward network where signals propagate from input to output through successive hidden layers.

Neuromorphic computing seeks to emulate the parallel, distributed, and event-driven nature of biological neural systems in physical hardware. Electronic implementations of neuromorphic computing, despite substantial advances, remain constrained by the physics of charge transport. The movement of electrons through resistive interconnects inevitably generates Joule heating and restricts communication bandwidth, particularly as devices approach the nanoscale[11, 12]. Even with architectural innovations such as in-memory computing[13–15] and memristive crossbars[16–18], the latency and energy cost associated with sequential read-write operations persist. To achieve brain-like performance, a physical substrate must provide massive parallelism, ultrafast signal propagation, and minimal energy

dissipation per operation, requirements that naturally point toward the domain of photonics.

1.1.2 Photonics for Neuromorphic Computing: Opportunities and Challenges

Photonic systems, by virtue of their passive and linear nature, excel at performing analogue linear operations such as weighting operations and matrix multiplications[19]. Photons travel at the speed of light and do not interact strongly with one another, enabling data to be transmitted and processed across many channels in parallel without mutual interference[20]. Through multiplexing in wavelength[21–24], polarization[25], spatial mode[26, 27], and time[28, 29], photonic circuits can implement high-dimensional vector–matrix linear operations directly in hardware[30, 31]. Moreover, optical interconnects eliminate capacitive charging delays, providing bandwidths that exceed hundreds of gigahertz and interconnect densities unattainable by copper wiring[32, 33]. These attributes make integrated photonics an inherently promising platform for large-scale neuromorphic processors capable of performing trillions of operations per second at sub-picojoule energy cost[31, 34, 35].

Despite this progress, a critical functional gap separates existing photonic accelerators from the full requirements of neuromorphic computing. Modern neural networks require not only linear weighting operations but also the **accumulation of weighted signals** followed by a **nonlinear activation**, which endow learning systems with thresholding and adaptation[36, 37]. Neither of these operations is natively available in the optical domain. Existing wavelength-multiplexed and time-multiplexed photonic architectures almost invariably rely on an opto-electro-optic processing loop: weighted optical signals are detected and accumulated electronically as photoreceiver charge, while nonlinear activation is applied either within

CMOS circuitry or during the subsequent electro-optic conversion where the values are re-modulated into the optical domain[38–44]. This hybrid approach incurs substantial energy overhead and latency, reintroducing the electronic bottleneck that photonics aims to eliminate. Furthermore, on the most widely adopted commercial silicon photonics platform, intrinsic optical nonlinearities, including Kerr[45], free-carrier[46, 47], and excitonic effect[48], are often negligible or limited in dynamic range, hindering scalable integration[32].

Overcoming these challenges requires a **functional optical nonlinear unit** that is simultaneously reconfigurable, cascadable, and compatible with dense integration. Such a mechanism must allow light to control light directly, enabling accumulation, modulation, and nonlinear transformation within a unified optical domain.

A promising conceptual direction emerges from recent advances in standing-wave photonics. In particular, Chen *et al.*[49] demonstrated that counter-propagating waves in a resonant structure give rise to wavelength-dependent standing-wave patterns that permit spatially selective coherent perfect absorption at subwavelength scales. Building on this principle, Lee *et al.* introduced photonic spatio-spectral reconfiguration[50], a mechanism that realises non-volatile, wavelength-selective amplitude tuning by positioning phase-change absorbers at the antinodes of standing-wave patterns. Although Lee’s system was designed for static, non-volatile spectral modulation rather than computation, it revealed a powerful architectural insight: light–matter interaction can be engineered through the spatial field structure, enabling distinct functionalities to be encoded within a single photonic element.

Building on this insight, this thesis develops a new opto-thermo-optic paradigm, implemented through the **Photonic-Heater-in-Lightpath (PHIL)** architecture,

which repurposes standing-wave-mediated absorption into a dynamic computational mechanism. By replacing non-volatile phase-change absorbers with volatile thermo-optic nanoheaters positioned at resonant antinodes, PHIL converts dissipated optical energy into programmable phase shifts and introduces a thermally mediated linkage between wavelength-multiplexed optical signals. This paradigm enables (i) all-optical cross-wavelength encoding, (ii) reconfigurable nonlinear activation, (iii) optical signal accumulation across wavelength channels, and (iv) temporal integration of time-multiplexed optical signals through analogue thermo-optic accumulation. In doing so, this work extends the standing-wave paradigm from static, non-volatile spectral modulation into analogue, cascable computation, addressing a central obstacle in neuromorphic photonics and establishing a unified framework for performing linear and nonlinear operations entirely in the optical domain.

1.2 Objectives and Outline

1.2.1 Objectives

The central objective of this thesis is to develop a unified device-level approach that enables optical signals to be accumulated and nonlinearly transformed within an integrated photonic platform. Although contemporary photonic hardware excels at high-bandwidth linear operations, the lack of a compact and cascable optical element capable of performing both summation and nonlinear activation remains a fundamental constraint on fully optical neuromorphic architectures.

To overcome these limitations, this work investigates opto-thermo-optic interactions as a functional mechanism for cross-wavelength encoding and optical signal accumulation. By shaping this thermal interaction to work together with the nonlinear response of the cavity spectrum, this thesis develops a device architecture

termed the Photonic-Heater-in-Lightpath (PHIL) that enables analogue accumulation and programmable activation to be carried out entirely in the optical domain, without relying on electro-optic conversion.

Through these objectives, the thesis aims to establish a physically grounded and experimentally validated framework that bridges the gap between the ultrafast linearity of photonic hardware and the nonlinear, integrative dynamics required for neuromorphic information processing.

1.2.2 Outline

This thesis is organised to move from foundational principles and physical modelling, through device design and fabrication, to the experimental realisation of all-optical computing functionalities using the PHIL architecture. Each chapter builds upon the preceding one to address the overarching goal of realizing functional, scalable, and energy-efficient nonlinear unit for neuromorphic photonic computation.

- **Chapter 1** introduces the motivation and background of this work, tracing the evolution from von Neumann architectures to neuromorphic computing and identifying photonics as a promising yet incomplete platform for achieving fully optical computation. It concludes with the research objectives and structural overview of the thesis.
- **Chapter 2** provides a comprehensive literature review, summarizing the current state of neuromorphic photonic computing and the underlying physical mechanisms that have been explored to achieve optical nonlinearity. The review identifies the limitations of existing approaches and defines the research gap addressed in this thesis.

- **Chapter 3** develops the theoretical foundation of the PHIL architecture. It introduces the principles of dielectric waveguides, grating couplers, and microring resonators, and provides a detailed formulation of coherent perfect absorption and standing-wave engineering. Together, these concepts establish the physical basis for the wavelength-selective absorption and opto-thermo-optic interactions exploited in PHIL.
- **Chapter 4** describes the device design, fabrication, and experimental setup. It presents the simulation workflow used for concept validation, details the fabrication of the PHIL devices, and outlines the optical characterization techniques employed throughout the experimental studies.
- **Chapter 5** presents the first experimental realisation of the PHIL architecture, establishing its core mechanism for cross-wavelength optical encoding. It demonstrates how optical power injected at a designated control wavelength induces a thermo-optic perturbation that is directly mapped onto a spectrally distinct probe wavelength, thereby validating selective, wavelength-resolved all-optical transduction.
- **Chapter 6** develops this architecture toward neuromorphic functionality. By leveraging the intrinsic Lorentzian response of the microring resonator, the chapter demonstrates reconfigurable nonlinear transfer characteristics that mimic neuronal activation. This section shows how controlled thermo-optic detuning enables analogue, all-optical nonlinear mapping across wavelengths, which is an essential capability for scalable photonic computing
- **Chapter 7** advances the platform into the time domain. Here, the device is shown to perform leaky integration of high-repetition-rate optical pulse trains over the intrinsic thermal relaxation window, enabling temporal accumulation far beyond the thermal bandwidth. The chapter experimentally validates

temporal integration with concurrent nonlinear activation, demonstrating that slow thermo-optic dynamics can serve as a physical analogue integrator for neuromorphic signal processing.

- **Chapter 8** concludes the thesis by summarizing the key findings, highlighting the contributions to thermo-optic device physics and neuromorphic photonics, and outlining future research directions toward scalable, large-vector, and wavelength-multiplexed neuromorphic photonic systems.

1.3 Key results

This thesis has focused on the development and experimental demonstration of a new photonic architecture, termed the PHIL, which enables functional all-optical computation by harnessing coupled opto-thermal and thermo-optic nonlinearities within integrated silicon photonics.

Specifically, the following major results have been achieved:

- **Cross-wavelength all-optical encoding:** A mechanism has been established whereby optical signals at multiple wavelengths are absorbed and accumulated through localized nano-heaters and subsequently re-encoded onto a probe wavelength via the thermo-optic effect. This process achieves fully optical cross-wavelength modulation and integration without electronic intermediaries.
- **Across-wavelength summation:** Linear and multi-channel optical summation has been demonstrated across two and three distinct wavelength channels. The accumulated output exhibits high linearity and multi-bit analogue precision (4–6 bits), confirming wavelength-domain additivity suitable for optical matrix operations.

- **Temporal integration:** The same architecture has been shown to perform leaky integration of ultrafast optical pulse sequences within its intrinsic thermal relaxation time (245 ns). This establishes the first demonstration of *analog temporal accumulation* in a passive photonic device, enabling continuous-time optical memory and signal processing.
- **Optically programmable nonlinear activation:** By tuning the probe wavelength relative to the resonant lineshape, the device exhibits configurable nonlinear transfer functions—including linear, sigmoid-like, and ReLU-like responses—entirely within the optical domain. This functionality emulates neuronal activation behaviour and demonstrates all-optical nonlinearity with analog tunability.

Collectively, these results establish a unified all-optical computing platform capable of performing accumulation, integration, and nonlinear transformation within a single passive photonic unit. The PHIL architecture thus bridges the gap between ultrafast linear photonics and the nonlinear, adaptive operations required for neuromorphic computing, offering a scalable route toward end-to-end all-optical information processing.

1.4 Statement of originality

This thesis presents the original research conducted by the author, Yi Zhang, during her DPhil studies at the University of Oxford. All theoretical analysis, device design, fabrication, and optical experiments described herein were carried out by the author, unless otherwise stated. The 50 GHz measurement experiments were performed at the Aristotle University of Thessaloniki by Dr. Ioannis Roumpos, whose contributions are gratefully acknowledged. Grammar and language proof-reading have been assisted by AI-based language models, ChatGPT. The author

takes full responsibility for the accuracy, interpretation, and presentation of all results contained in this thesis.

2

Integrated Neuromorphic Photonic Computing: A review

Contents

2.1 Neuromorphic Photonic Computing	13
2.1.1 Signal Representation and Multiplexing	14
2.1.2 Photonic Implementation of Accumulation and Nonlinearity	16
2.2 Integrated Photonic Accelerator Architectures	17
2.2.1 Space-Division Multiplexing	17
2.2.2 Wavelength-Division Multiplexing	25
2.2.3 Time-Division Multiplexing	30
2.3 Optical Signal Accumulation And Nonlinearity Implementation	34
2.3.1 Through Opto-Electro-Optic Conversions	35
2.3.2 Within the Optical Domain	38
2.4 The Need for All-Optical Nonlinearity Implementation for Scalable Optical Neural Networks	41
2.4.1 Functional Gaps in Photonic Neural Networks	41
2.4.2 Standing-Wave Photonics as an Architectural Precursor	43
2.4.3 From Passive Modulation to All-Optical Computation	45

This chapter reviews state-of-the-art integrated photonic neural network architectures, focusing on different signal representations (multiplexing methods), associated weighting techniques, and the following weighted signal summation with nonlinearity implementations. Furthermore, it encompasses and discusses the challenges arising from the signal conversion between optical and electronic domains and hardware implementations of non-linearities in the all-optical domain.

2.1 Neuromorphic Photonic Computing

In neural networks, connections between neurons are represented by weighted links, where each input signal is multiplied by a corresponding weight and the weighted inputs are accumulated in neurons of the subsequent layer. Photonic neural networks emulate this neuronal computation in optical hardware, providing a general computing framework in which each neuron performs two fundamental operations [51]: linear accumulation of weighted neural inputs ($x_1w_1 + x_2w_2 + \dots + x_{n-1}w_{n-1} + x_nw_n$), and the following implementation of non-linear activation functions at neural outputs (Figure 2.1(a)).

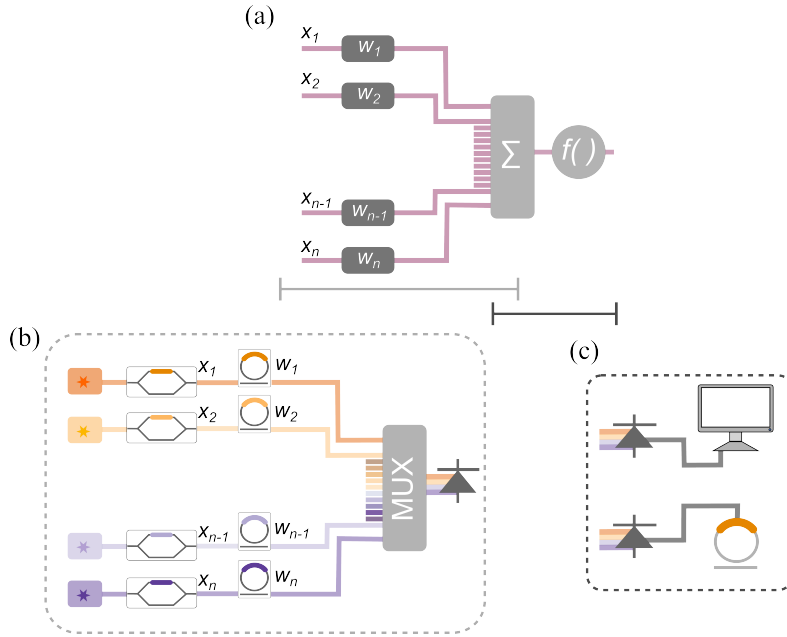


Figure 2.1: Single neuron framework (a) General abstraction of a single neuron, consisting of weighted summation followed by a nonlinear activation function. (b) Example of a wavelength-division-multiplexed (WDM) photonic accelerator implementing parallel weighted multiplication. (c) Nonlinear activation functions implemented through additional electronic circuitry.

Implementing this neuronal operation in photonic hardware requires addressing two complementary tasks: first, representing neural inputs and applying their corresponding weights across optical channels (Figure 2.1(b)); and second, accumulating

these weighted signals and performing the nonlinear transformation that produces the neuron output (Figure 2.1(c)).

2.1.1 Signal Representation and Multiplexing

Multiplexing is a fundamental concept in optical communication and signal processing, referring to the encoding of multiple independent signals into distinct and distinguishable physical channels so that they can be transmitted or processed simultaneously within the same hardware platform.

The key advantage of photonic systems is that optical signals can naturally occupy several orthogonal physical domains, including wavelength, spatial mode, polarization, and time. By assigning different signals to different channels within these domains, multiple data streams can propagate and be processed in parallel within the same photonic circuit. This inherent parallelism is particularly valuable for computing architectures, where neural-network operations such as matrix–vector multiplication directly benefit from the simultaneous manipulation of many signal channels.

In this thesis, multiplexing primarily refers to three physical dimensions that are commonly exploited in integrated photonic computing systems: spatial channels, wavelength channels, and temporal channels. In space-division multiplexing (SDM), independent signals propagate in separate waveguides or spatial modes (Figure 2.2(a)). In wavelength-division multiplexing (WDM), different signals are encoded onto optical carriers with distinct wavelengths that co-propagate within the same waveguide (Figure 2.2(b)). In time-division multiplexing (TDM), signals share the same optical path but occupy different temporal slots, allowing sequential signals to be processed using the same hardware resources (Figure 2.2(c)).

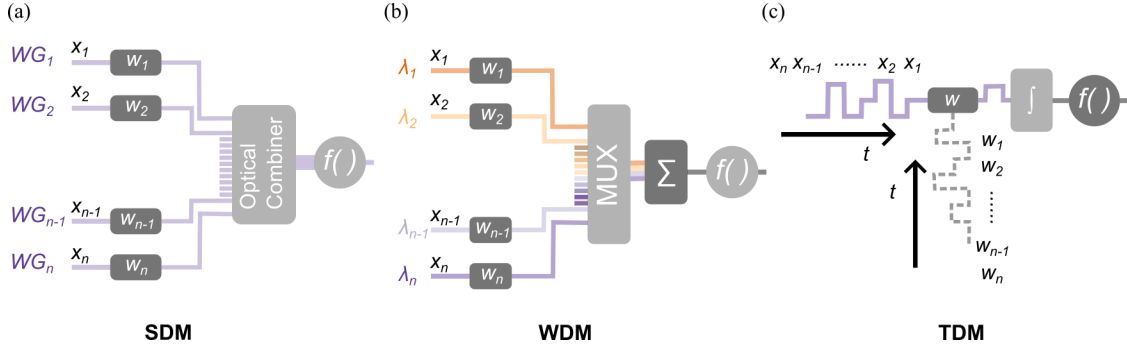


Figure 2.2: Conceptual illustration of multiplexing strategies in integrated photonic neural networks (a) Space-division multiplexing (SDM), where signals propagate in parallel spatial waveguides and are combined by an optical combiner. (b) Wavelength-division multiplexing (WDM), where signals are encoded on different optical wavelengths and merged using a wavelength multiplexer. (c) Time-division multiplexing (TDM), where signals are sequentially processed in time and accumulated through temporal integration.

In such systems, each multiplexed channel can be independently modulated to represent both the input signal and its associated weight, such that each channel carries a weighted optical signal contributing to the overall computation. During optical processing, intermediate computational results may remain distributed across these channels rather than being immediately combined into a single scalar quantity. Parallel optical processing therefore arises naturally within the same photonic hardware platform, enabling the simultaneous manipulation of many weighted signals.

Photonic hardware platforms offer a promising route for accelerating matrix–vector multiplication by exploiting multiplexing across multiple physical domains. Chapter 2.2 reviews representative architectures of integrated photonic accelerators, categorized according to their signal multiplexing strategies and associated weighting techniques. Among these approaches, time-division multiplexing is particularly attractive for increasing the number of accessible computational channels while maintaining a compact system footprint.

2.1.2 Photonic Implementation of Accumulation and Non-linearity

While multiplexing provides an efficient mechanism for representing and weighting neural inputs across optical channels, a complete photonic neuron must additionally perform accumulation of these weighted signals followed by a nonlinear transformation.

Non-linear activation functions, applied to the summation of weighted input signals, are essential for implementing the neuronal thresholding effect. This allows neural networks to extract hierarchical features and approximate complex functions through cascaded layers, preventing them from collapsing into a simple linear model. These will be further discussed in Chapter 2.3 as it explores various hardware implementations of photonic non-linearities, including those achieved through opto-electro-optic signal conversion (Figure 2.1(c)) and purely all-optical processes.

While these developments establish the foundations of neuromorphic photonics, they also expose a deeper architectural limitation: photonic platforms naturally implement weighting operations through linear interference and multiplexing, but they lack a unified mechanism capable of performing both accumulation and non-linear activation in the optical domain. This motivates the development of architectures in which accumulation and nonlinearity are co-localized within the same physical process, enabling more compact and efficient photonic neural implementations.

2.2 Integrated Photonic Accelerator Architectures

2.2.1 Space-Division Multiplexing

Space-division multiplexing (SDM) represents one of the earliest and most direct approaches for implementing parallelism in integrated photonic accelerators. By assigning data to distinct spatial modes or physically separate waveguide paths, SDM allows multiple channels to propagate simultaneously at a common wavelength.

In such architectures, the information carried by each channel is encoded in the optical field associated with its spatial mode. Specifically, the elements of an input vector are mapped onto the amplitudes (or intensities) of optical signals propagating in different spatial channels. Optical weighting is then performed by components that modify the amplitude and/or phase of these fields, including interferometric elements, phase shifters, or spatial transmission profiles. Following propagation through the optical system, the weighted fields are combined either through coherent interference or direct power summation at photodetectors, producing output signals that correspond to weighted sums of the inputs. In this manner, SDM-based photonic circuits naturally implement matrix-vector multiplication through linear optical transformations.

Among the various implementations of SDM-based photonic computing, this section focuses on two representative architectures that illustrate different approaches to realizing spatial optical transformations: interferometric meshes in integrated photonics and diffractive optical neural networks.

Interferometric Mesh

The concept of matrix factorization using programmable linear optics was first introduced in the early work of Reck *et al.*[52], where an $N \times N$ unitary matrix can be decomposed and mapped in the hardware with universal linear optical components[53, 54], such as the Mach-Zehnder interferometer (MZI).

An individual MZI acts as a tunable 2×2 interferometric unit that mixes two optical modes through controlled phase shifts and interference. As shown in Figure 2.3(a), the relation between the input and output optical fields can be expressed as

$$\begin{pmatrix} E_1^{\text{out}} \\ E_2^{\text{out}} \end{pmatrix} = U(\theta, \phi) \begin{pmatrix} E_1^{\text{in}} \\ E_2^{\text{in}} \end{pmatrix}. \quad (2.1)$$

Here, θ controls the effective power splitting ratio between the two output ports, while ϕ introduces an additional tunable phase shift. This transformation belongs to the $SU(2)$ group, representing a lossless rotation in a two-mode optical system.

By cascading multiple MZIs, higher-dimensional unitary transformations can be realized, enabling arbitrary linear optical operations. Each MZI therefore acts as a controllable rotation between two optical modes, forming the fundamental building block for universal linear transformations.

Building on the cascaded MZI mesh architecture (Figure 2.3(b)), Shen *et al.*[35] first implemented an integrated optical neural network (ONN) for deep learning, demonstrating superior computational speed compared to electronic counterparts in conventional inference tasks.

This ONN architecture is depicted in Figure 2.3(c), where each layer comprises an optical interference unit to apply interlayer weights and an off-chip, computer-modelled optical nonlinearity unit. The optical interference unit is experimentally

implemented with 56 programmable MZIs to perform a 4×4 matrix multiplication based on singular value decomposition. By spatially multiplexing guided optical signals and modulating their amplitudes through interference using on-chip thermo-optic phase shifters embedded in each MZI, this approach enables arbitrary unitary matrix multiplication at the speed of light.

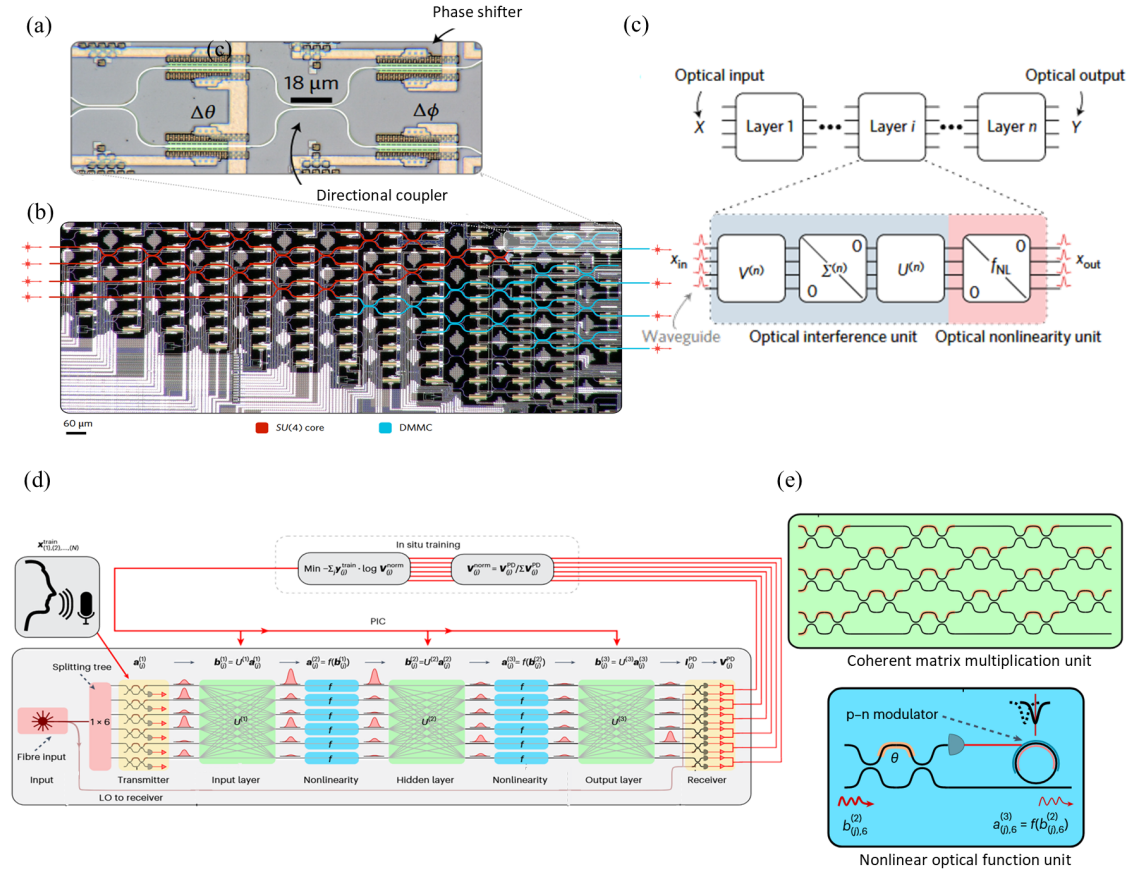


Figure 2.3: Integrated photonic architectures based on interferometric mesh (a) Schematic of phase shifters in the MZI[35]. (b) Experimentally demonstrated optical interference unit[35]. (c) Workflow of the proposed optical neural network (ONN)[35]. Reproduced with permission from Shen et al., *Nature Photonics* (2017), © Springer Nature. (d) Coherent integrated ONN with in situ training[55]. (e) Coherent matrix multiplication unit and nonlinear unit in the ONN[55]. Reproduced with permission from Bandyopadhyay et al., *Nature Photonics* (2024), © Springer Nature.

This MZI-mesh-based architecture can also be extended to implement complex-valued neural networks, in which information is encoded in both optical magnitude and phase, while the interferometric mesh performs the corresponding unitary

transformations through programmable phase shifts and optical interference[38].

In situ training of ONN has also been investigated to achieve accurate inferencing. In 2018, Hughes *et al.*[56] theoretically proposed using adjoint variable methods to derive the photonic analogue of the backpropagation algorithm, and later in 2023, Pai *et al.*[57] validated the in situ backpropagation approach experimentally, employing precise phase control facilitated by an infrared camera. Most recently, a monolithically integrated neural network with six neurons and three layers was demonstrated to perform both inference and in situ training[55] (Figure 2.3(d)). Combining with on-chip programmable nonlinear activation functions (Figure 2.3(e)), this work demonstrates the first single-chip implementation of an optical end-to-end coherent deep neural network.

However, despite the advances to ONN implementation in such an architecture, there are several drawbacks and challenges lying in its design. Interferometric mesh architectures rely on numerous integrated phase shifters for matrix multiplication, requiring precise phase control of each MZI path—a task complicated by fabrication variability and thermal crosstalk[58]. To address this, self-configuration schemes have been explored to enable the ONN to calibrate itself on errors, and recent work has demonstrated that self-configuration algorithms can accurately recover target unitary transformations in large MZI meshes without any additional hardware or prior device characterization, achieving quadratic suppression of fabrication-induced errors and markedly improving ONN inference accuracy on realistic photonic hardware[59].

The relatively large footprint of programmable MZI mesh designs is inevitable, as each MZI typically exceeds $100\ \mu\text{m}$ in length. Moreover, matrix decomposition necessitates additional integrated components, further increasing system complexity[60]. As the matrix size grows, scalability becomes a challenge due to the $\mathcal{O}(N^2)$

scaling of reprogrammable MZI units and the corresponding increase in physical area. Several efforts have been made to improve the scalability of this architecture. In addition to the MZI mesh for complex-valued operation, Zhu *et al.*[61] introduced two compact diffractive cells to incorporate parallel Fourier transformation, thereby reducing the scaling to $\mathcal{O}(N)$. Feng *et al.*[62] further explored pruning techniques, achieving $7\times$ fewer trainable optical components.

Large-scale interferometric meshes are intrinsically spectrally sensitive. Since their operation depends on precise phase control, they generally achieve optimal performance only near a designed wavelength, whereas broadband or multi-wavelength operation requires careful dispersion engineering. This limits their compatibility with additional multiplexing dimensions within the same system[60, 63].

Diffractive Optical Neural Network

Diffractive optical neural networks (DONNs) were first proposed and demonstrated with free-space optics[64–66], as shown in Figure 2.4(a). A DONN consists of cascaded transmissive and/or reflective layers, where each pixel acts as a neuron with a complex-valued transmission coefficient. As light propagates through the network, each layer performs local modulation followed by diffraction-based coupling, governed by the Huygens–Fresnel principle, resulting in a fully connected linear transformation via coherent superposition.

As illustrated in the inset of Figure 2.4(a), the operation of a single layer can be expressed as

$$\tilde{\mathbf{Y}}^{l+1} = \tilde{\mathbf{W}}^{l+1} \left(\tilde{\mathbf{Y}}^l \circ \tilde{\mathbf{B}}^l \right), \quad (2.2)$$

where $\tilde{\mathbf{B}}^l$ denotes the locally trainable complex transmission coefficients, typically implemented as phase-only modulation $\tilde{B}^l(x, y) = e^{j\phi^l(x, y)}$, and $\tilde{\mathbf{W}}^{l+1}$ represents the propagation matrix determined by the Green’s function of free-space diffraction.

In this formulation, diffraction defines the *effective inter-neuron weights*, while the trainable parameters reside in the local modulation coefficients. Consequently, unlike interferometric meshes where individual connection weights are directly programmable, the connectivity in DONNs is fixed by the diffraction kernel and remains invariant during training.

On-chip implementations of DONN have also been explored recently. By leveraging slot-array-based one-dimensional (1D) metasurfaces, coherent inference of parallelly transmitted wavefronts can be precisely controlled[67]. On-chip DONNs based on 1D metasurfaces have been evaluated through theoretical analysis[68, 69] and further validated by multiple experimental implementations[70–72], demonstrating capabilities in matrix–vector multiplication and image classification (Figure 2.4(b)).

As illustrated in Figure 2.4(b), each diffractive unit consists of a subwavelength slot filled with SiO₂ embedded in a silicon waveguide platform [72]. The phase delay imparted to the propagating field is controlled by the slot length, which modifies the effective optical path and thus the local transmission phase. Consequently, each layer implements a trainable diagonal modulation matrix $\tilde{\mathbf{B}}^l = \text{diag}(e^{j\phi_l})$, where ϕ_l is determined by the designed slot lengths.

At the system level, the cascaded operation of multiple layers can be expressed as

$$\mathbf{y} = \mathbf{W}^{(L)} \text{diag}(e^{j\phi_{L-1}}) \cdots \mathbf{W}^{(2)} \text{diag}(e^{j\phi_1}) \mathbf{W}^{(1)} \mathbf{x}, \quad (2.3)$$

where each $\mathbf{W}^{(l)}$ represents the fixed linear transformation arising from on-chip diffraction and waveguide coupling, governed by the system Green’s function. Therefore, while the network appears fully connected, the effective weights are determined by the interplay between fixed propagation-induced mixing and layer-wise trainable phase modulation, rather than being independently programmable.

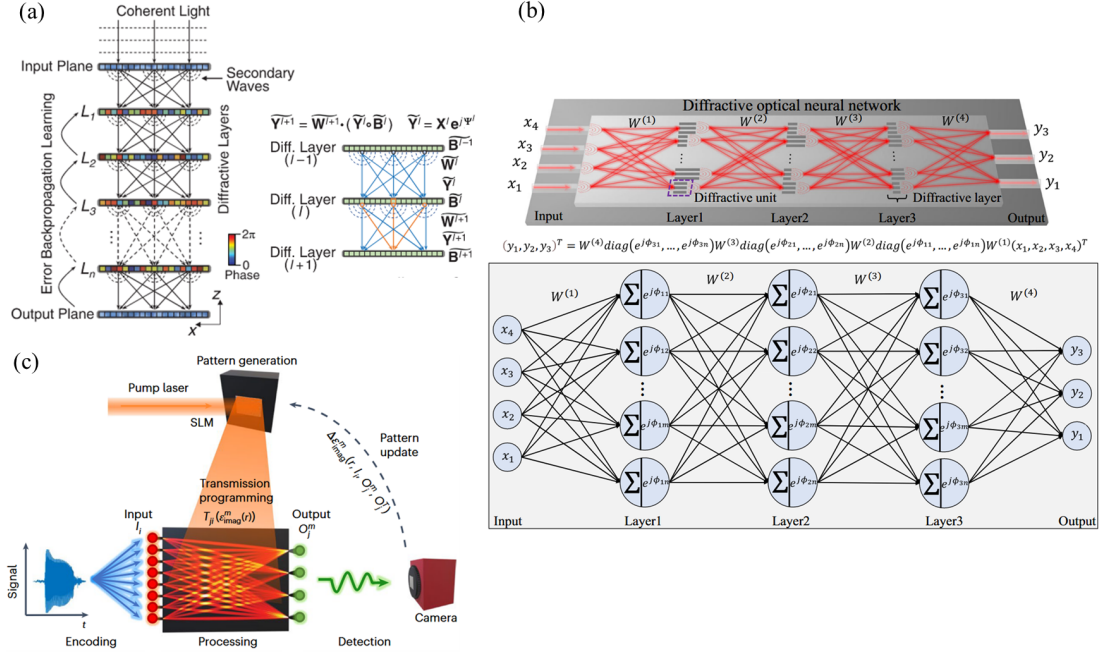


Figure 2.4: Schematics of diffractive optical neural network(DONN) (a) Free-space demonstration of DONN [64]. From Lin et al., *Science* 361, 1004–1008 (2018). Reprinted with permission from AAAS. (b) On-chip demonstration of DONN[72]. Reproduced from Fu et al., *Nature Communications* (2023), licensed under CC BY 4.0. (c) Reconfigurable integrated DONN [73]. Reproduced with permission from Wu et al., *Nature Photonics* (2023), © Springer Nature.

One significant limitation of earlier diffractive neural network designs lies in their reliance on lithographically pre-defined layers, which constrains both flexibility and reconfigurability. Additionally, imperfections introduced during the lithographic process can result in structural defects that degrade the desired optical functionality. To address these limitations, Wu *et al.*[73] proposed a lithography-free approach using a spatial light modulator (SLM) to project a pump laser onto an InGaAsP thin film, thereby creating a reconfigurable photonic network (Figure 2.4(c)). By dynamically modulating the pump pattern projected onto the III–V semiconductor layer, the local refractive index could be tuned, altering the interaction of the input optical signals within the network. This reconfigurability enabled the system to be trained through iterative projected pattern updates, successfully achieving high-accuracy performance in vowel recognition tasks. In addition to direct weight

implementations, the fully connecting property of an on-chip diffractive layer has been investigated in combination with interferometric mesh to significantly reduce spatial redundancy with enhanced energy efficiency [61, 74]. In these architectures, diffractive units function as optical encoders and decoders that compress high-dimensional input signals into lower-dimensional representations while preserving task-relevant information for downstream AI inference.

Architectural Trade-offs

SDM-based architectures present two central limitations that restrict their applicability in large-scale photonic neural networks. The scalability of matrix size is constrained by both physical and architectural factors. Interferometric mesh networks require precise phase alignment and are susceptible to thermal crosstalk and quadratic scaling in the number of phase-tunable components, which in turn impacts layout complexity and footprint. While DONN-based systems offer greater compactness and passive energy efficiency, they often rely on lithographically defined static structures that limit adaptability. Recent progress in reconfigurable DONNs using SLMs shows promise, yet these bulky control interfaces challenge full on-chip deployment.

SDM also exhibits limited flexibility for hybrid parallelism. Both MZI meshes and diffractive networks exhibit strong wavelength sensitivity due to their interference-based weighting mechanisms, making integration with spectral multiplexing architectures technically challenging [63]. The achievable degree of parallelism is further limited by inter-channel crosstalk, phase mismatch, and fabrication-induced asymmetries, which degrade signal isolation and hinder the stable, independent routing of concurrent data streams. These limitations restrict SDM's ability to support high-throughput, multi-kernel operations required for convolutional or attention-based processing.

Together, these constraints have motivated a shift toward spectral and temporal multiplexing schemes, which offer higher channel densities and more compact layouts. The following subsections examine wavelength- and time-multiplexed photonic architectures, highlighting how they address the parallelism bottleneck, while introducing new challenges in implementing optical-domain accumulation and nonlinear activation.

2.2.2 Wavelength-Division Multiplexing

Wavelength-division multiplexing (WDM) is one of the most widely used approaches, offering intuitive many-neuron-to-many-neuron weighted connections across wavelengths with direct matrix mapping onto hardware, enabling highly parallel processing. In this scheme, multiple optical carriers at distinct wavelengths propagate simultaneously within the same waveguide, with each wavelength channel acting as an independent information pathway. Signal operations are typically performed on the amplitude of light, where input data are encoded as intensity modulation on each wavelength channel. The weighted computation is then implemented through wavelength-dependent attenuation or phase control, effectively multiplying the optical signal carried by each wavelength with a corresponding weight. After weighting, the multiplexed optical signals are combined and detected together, producing a physical summation of the weighted contributions and thereby performing parallel matrix-vector multiplication.

Micro-ring Resonator Weight Banks

A common implementation of wavelength-selective weighting in integrated photonic accelerators employs microring resonator (MRR) weight banks. Each microring is designed to resonate at a specific wavelength channel and can therefore selectively control the transmission of that wavelength. By tuning the resonance condition,

typically through thermo-optic or electro-optic phase shifters, the transmission of the corresponding wavelength can be continuously adjusted, allowing the optical amplitude to represent a programmable weight. In this architecture, different wavelength channels are mapped onto different resonant rings, with each ring representing an individual weight element in the matrix.

The first WDM-based signal processing scheme employing such MRR weight banks was demonstrated by Ji *et al.* [75]. As shown in Figure 2.5(a), the input vector is encoded via amplitude modulation of laser outputs at different wavelengths. These wavelength channels are multiplexed and distributed through the photonic circuit, where matrix elements are individually mapped onto an array of wavelength-selective MRRs. Each wavelength-specific MRR incorporates a phase shifter that controls the coupling efficiency at its designated wavelength, thereby implementing the corresponding matrix weight. This configuration simultaneously performs wavelength-division multiplexing and demultiplexing operations while applying the weights. The weighted wavelength channels are then recombined and summed by photodetectors positioned at the output ports, producing the final multiply-accumulate result.

Tait *et al.* then further proposed a "broadcast-and-weight" architecture for optical computing based on MRR weight banks[22, 76–78], and later implemented experimentally [79]. In this architecture (Figure 2.5(b)), the output signals are fed into the input signals and broadcast to all MRR weight banks, making it possible for further implementation of generalized reconfigurable and recurrent neural network models. Both positive and negative weights can be realized in the matrix by implementing both add-drop configuration and balanced photodetector. Later, on-chip integration of MRR weight banks with microcomb laser source was demonstrated by Bai *et al.*[24], achieving 9-bit resolution in weight setting with

dedicated calibration scheme and a record-high computing density of 1.04 TOPS mm^{-2} , while non-linear activation functions were applied off-chip.

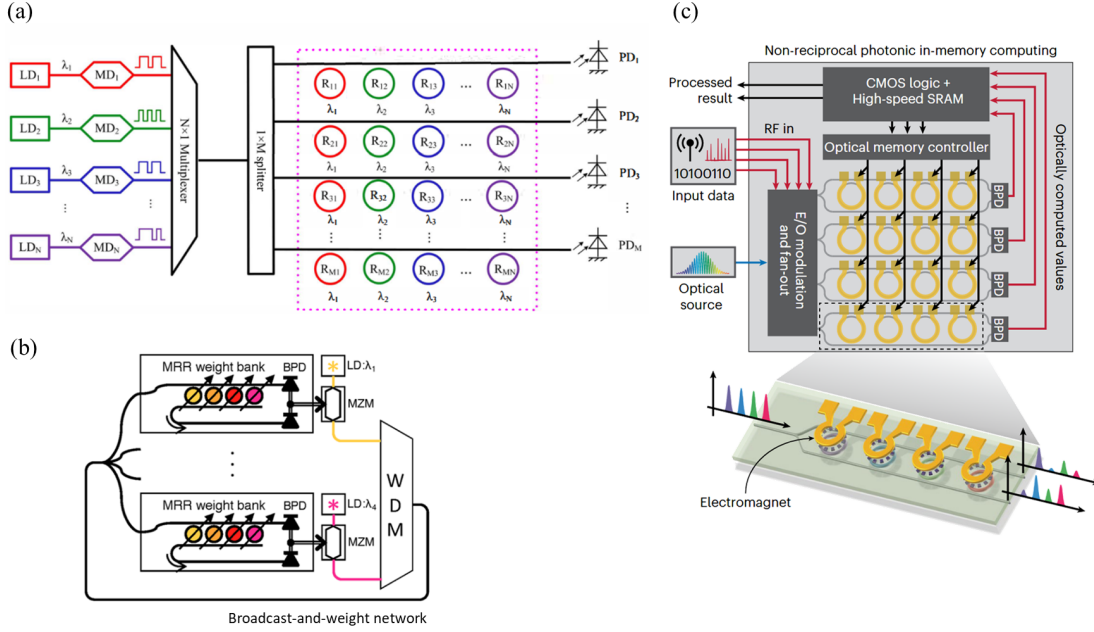


Figure 2.5: Schematics of micro-ring resonator (MRR) weight banks (a) Initial demonstration of MRR weight bank. LD: laser diode. MD: modulator. PD: photodetector[75]. From Ji et al., "On-chip CMOS-compatible optical signal processor", 2012 Asia Communications and Photonics Conference (ACP), IEEE (2012). © 2012 IEEE. Reprinted, with permission. (b) MRR-weight-bank based "broadcast-and-weight" architecture. BPD: balanced photodiode, LD: laser diode, MZM: Mach-Zehnder modulator[79]. Reproduced from Tait et al., *Scientific Reports* (2017), licensed under CC BY 4.0. (c) MRR weight bank integrated with non-volatile material[80]. Reproduced from Pintus et al., *Nature Photonics* (2025), licensed under CC BY 4.0.

A key limitation of scaling up MRR-based systems lies in their sensitivity to environmental variations (e.g., temperature, stress) and fabrication imperfections, which necessitates continuous resonance tracking and control. This requires additional components, such as heater tuning, leading to increased energy consumption and added complexity, while limiting reconfiguration speed to the kHz range[60]. Several techniques have been explored to address these challenges, including post-fabrication trimming using polymers[81] and the integration of transparent phase-change materials[82]. Recently, Pintus *et al.*[80] proposed incorporating a magneto-

optic cladding layer (Ce:YIG) to enable non-volatile, multi-level weight programming in MRR weight banks (Figure 2.5(c)). The non-reciprocal interaction between the cladding and the propagating optical modes enhances both symmetry and extinction ratio for positive and negative weights. Their approach achieves programming speeds around 1 GHz.

Photonic Crossbar Array

Photonic crossbar arrays, as two-dimensional analogues of electronic crossbars, map weight matrices onto a grid of intersecting row and column waveguides. Input signals are encoded as optical intensities injected along the row waveguides. As light propagates, a fraction of the optical power is coupled at each row-column intersection via directional couplers and modulated by integrated attenuators or programmable elements, thereby implementing the corresponding weights. The weighted signals are then accumulated along the column waveguides and detected at the outputs, enabling fully parallel multiply-accumulate operations. When integrated with non-volatile photonic memory as programmable weight, this architecture enables parallel weight multiplication with near-zero static energy consumption.

Phase-change materials (PCMs), most commonly $\text{Ge}_2\text{Sb}_2\text{Te}_5$ (GST), can be integrated atop optical waveguides to function as non-volatile optical weights [83]. As shown in Figure 2.6(a), the waveguide transmission, serving as the weight (a), is modulated by the physical state of the PCM that altered via evanescent optical coupling. Weight encoding is achieved by applying high-energy optical pulses to induce non-volatile phase transitions in the PCM cell, between crystallization (GST_{CRY}) and amorphization (GST_{AM}). The stored weights can subsequently be read out using low-intensity optical pulses (b), with the multiplexed results (c) measured

as the transmitted signal at the waveguide output. These PCM cells can be all-optically addressed with 5-bit precision[84], allowing the all-optical implementation of a neurosynaptic network which is capable of pattern recognition[21].

Feldmann *et al.*[31] proposed the first photonic crossbar array in 2021, with the implementation of optically programmable non-volatile PCM weights, demonstrating accelerated matrix–vector multiplications at 0.256 TOPS (Figure 2.6(b)). This system was integrated with a frequency comb to enable wavelength-multiplexed parallel convolution processing. Each wavelength from the comb source is amplitude-modulated with input data, injected into a certain row, selectively routed into specific columns via custom-designed directional couplers, and then weighted by the corresponding PCM cell. This architecture was further equipped with electrical addressability[85], enabling improved computing accuracy through increased PCM switching contrast achieved by integrated microheaters (Figure 2.6(c)).

Photonic crossbar arrays largely reduce the number of control units down to the matrix size (i.e. $M \times N$). This amplitude-based approach not only enhances fabrication tolerance compared to phase-based encoding, but is also compatible with high-speed components such as electro-absorption modulators[87]. Furthermore, the architecture offers superior computing density, owing to the compact weight cell size and increased parallelism enabled by exploiting orthogonal degrees of freedom. This parallelism can be further extended by integrating both SDM and WDM. As shown in Figure 2.6(d), Dong *et al.*[86] leveraged a reduced-coherence laser source to suppress phase fluctuations caused by interference between different rows, an issue observed with single coherent laser sources. Nevertheless, scaling photonic crossbar arrays to matrices suitable for real-world AI tasks remains challenging. The required input power increases with the number of optical combiners, imposing stringent demands on photodetector performance[88], and factorization

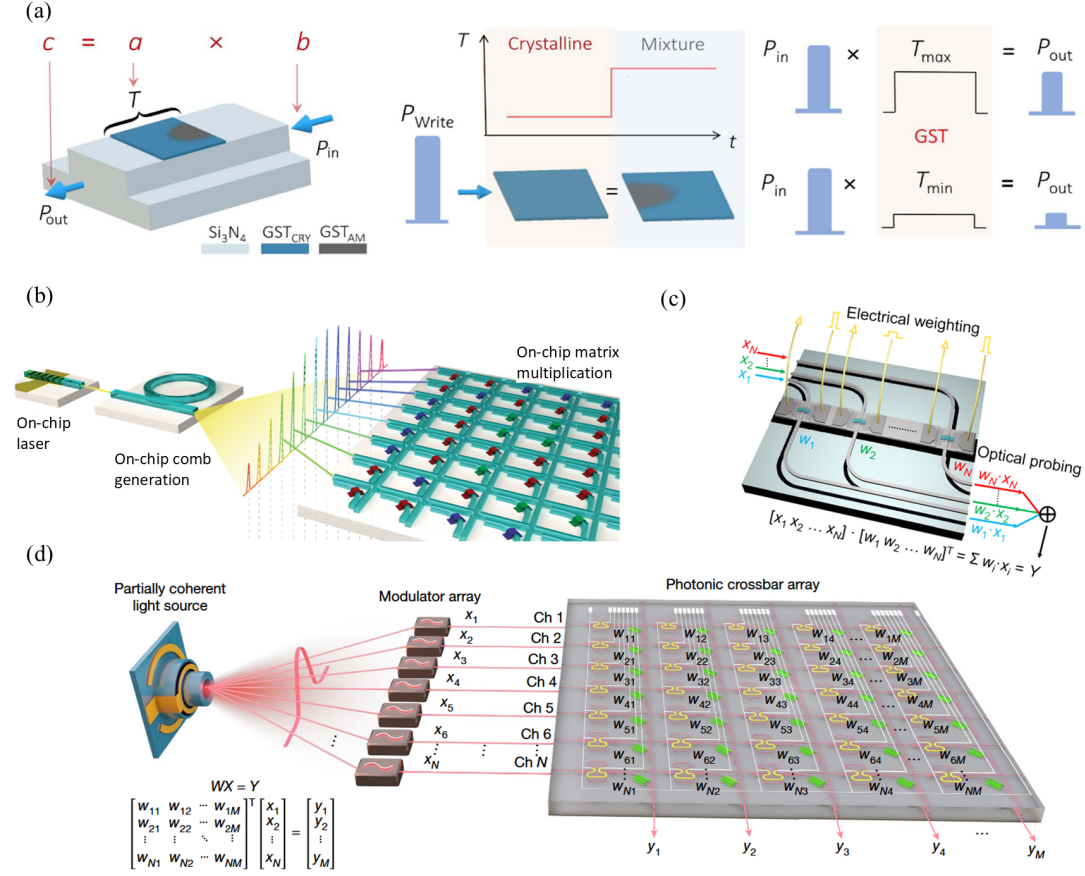


Figure 2.6: Schematics of photonic crossbar array based on phase-change materials (PCMs) (a) PCM cell and working mechanisms[83]. Adapted from Rios et al., *Science Advances* 5, eaat5759 (2019), licensed under CC BY 4.0. (b) Initial demonstration of PCM-based photonic crossbar array with optical programmability[31]. Adapted with permission from Feldmann et al., *Nature* (2021), © Springer Nature. (c) Electrically programmable PCM-based photonic crossbar array[85]. Reproduced from Zhou et al., *Nature Communications* (2023), licensed under CC BY 4.0. (d) Photonic crossbar array with partially coherent light source[86]. Reproduced from Dong et al., *Nature* (2024), licensed under CC BY 4.0.

techniques have been explored to expand the effective size of current crossbar arrays[89].

2.2.3 Time-Division Multiplexing

So far, photonic accelerators map machine learning operations directly in hardware by spatially multiplexing guided optical modes and leveraging orthogonal degrees of freedom of light to perform massively parallel matrix–vector multiplications.

However, the large physical footprints occupied by optical unitary cells and the insertion losses inherent in both active and passive components have made it challenging to further scale weight matrices. Efforts to overcome these limitations by tiling photonic subassemblies, matrix factorization and weight pruning offer partial solutions but remain far from achieving the scale required for the billions of parameters implemented in complex AI tasks.

To expand the computational parameter space, time-division multiplexing (TDM) offloads information processing from the spatial hardware domain into the temporal domain. In this approach, input data are encoded as sequences of optical pulses propagating through the same physical hardware, where each time slot represents a different element of the input vector. This enables a single photonic unit to be reused sequentially. Weighting is achieved through time-synchronised modulation or delayed interference, such that each temporal component experiences a distinct weight. The weighted signals are then accumulated via temporal integration or time-aligned detection, realising matrix–vector multiplication in a time-serial manner. By reusing hardware across time slots, TDM reduces device footprint and relaxes scalability constraints of purely spatial architectures.

Xu *et al.* [90] demonstrated a photonic convolutional accelerator that combines spatial, wavelength, and temporal multiplexing for parallel processing with multiple kernels. As shown in Figure 2.7(a), kernel weights (\mathbf{W}) are encoded in the optical power of different wavelengths, while the input vector (\mathbf{X}) is encoded as a stream of temporal symbols. These are multicast onto the kernel wavelengths via electro-optical modulators with delays matched to the symbol duration. The weighted replicas are summed at a high-speed photodetector, with each time slot yielding a convolution result between \mathbf{X} and \mathbf{W} . Since the size of the input vector \mathbf{X} is decoupled from the hardware topology, this accelerator supports large-scale, arbitrarily sized data processing. In their implementation, image data exceeding

250,000 pixels were processed using 10 convolution kernels mapped across 90 comb lines, achieving a vector computing throughput of 11.3 TOPS.

Dong *et al.*[91] introduced continuous-time data representation into the current PCM-based photonic crossbar array by employing radio-frequency modulation to enhance multiplexing dimensionality, enabling parallel matrix–matrix multiplication (Figure 2.7(b)). Similarly, Meng *et al.*[92] demonstrated a highly integrated MRR-based photonic tensor core that utilizes microwave frequency as an additional degree of freedom, achieving a high computing density of 34.04 TOPS/mm².

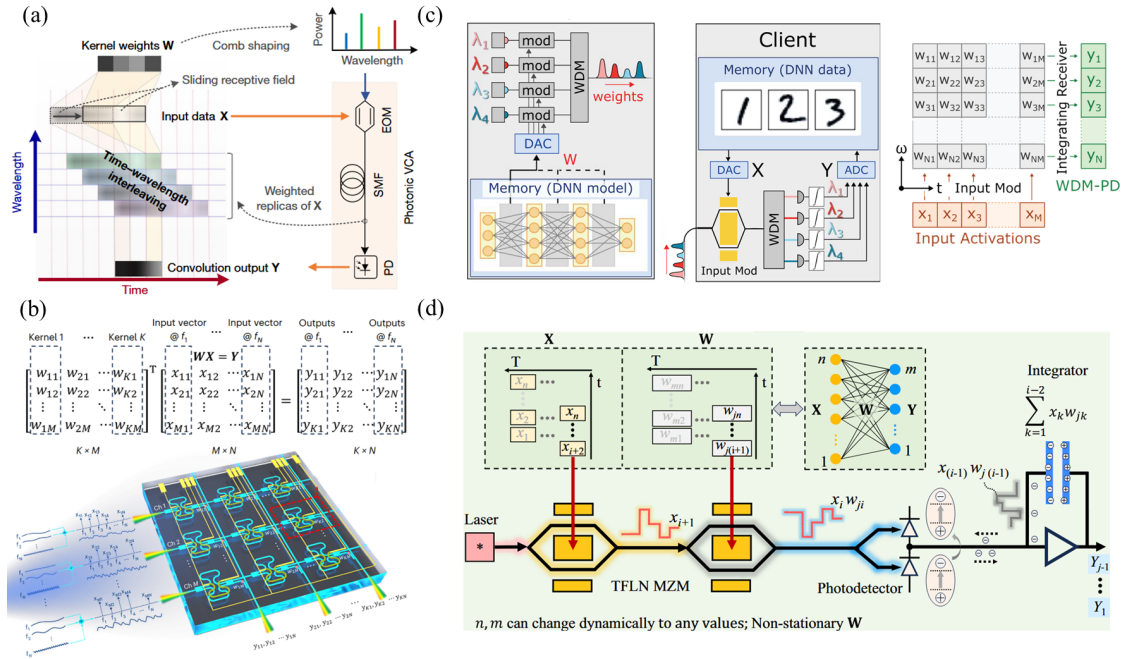


Figure 2.7: Photonic accelerators multiplexing space, wavelength and time. (a) An optical convolutional accelerator interleaving both time and wavelength[90]. Reproduced with permission from Xu et al., *Nature* (2021), © Springer Nature. (b) A photonic tensor core using continuous-time data representation to achieve parallel matrix-matrix multiplication[91]. Reproduced from Dong et al., *Nature Photonics* (2023), licensed under CC BY 4.0. (c) Netcast model, where matrix-vector multiplication is computed on time-integrating receivers at edge devices[44]. From Sludds et al., *Science* 378, 270–276 (2022). Reprinted with permission from AAAS. (d) An integrated photonic tensor core based on only two thin-film lithium niobate modulators[43]. Reproduced from Lin et al., *Nature Communications* (2024), licensed under CC BY 4.0.

By mapping scalar operations in time, TDM enhances system scalability while maintaining a compact footprint. Sludds *et al.*[44] have proposed a photonic edge computing architecture, Netcast (Figure 2.7(c)), in which weights are encoded onto separate wavelengths by smart transceivers integrated along the cloud server and broadcast to edge devices. Input activations are applied in the time domain via a broadband optical modulator, and MAC operations are performed on time-integrating receivers for each wavelength. This architecture offloads the computationally intensive weight training to the cloud, enabling ultra-efficient edge inference with energy usage below one photon per multiply. Lin *et al.*[43] further extended time-domain mapping by implementing both weights and inputs in the temporal domain. They introduced a fully integrated photonic tensor core comprising only two thin-film lithium niobate (TFLN) modulators, a III–V laser, and a charge-integrating photoreceiver, achieving a fan-in size of 131,072 (Figure 2.7(d)).

Beyond mitigating the scaling limitations of spatial multiplexing alone, recent developments unlock higher-dimensional parallelism by increasing the combination with spatial, wavelength, and temporal domains. This multidimensional multiplexing enables flexible and efficient time mapping of both inputs and weights, supporting high-throughput operations such as matrix–vector and matrix–matrix multiplications. By leveraging these orthogonal degrees of freedom, photonic systems can achieve significant gains in computational density, energy efficiency, and scalability—paving the way toward practical photonic platforms for large-scale AI workloads.

2.3 Optical Signal Accumulation And Nonlinearity Implementation

To enable hierarchical feature extraction and providing the mathematical expressiveness required for universal function approximation, neurons in neural networks are expected to respond nonlinearly to the weighted inputs and generate an output to the next-layer neurons as a function of the inputs.

In the current implementation of ONN, nonlinear activation functions are mostly implemented by detecting each optical signal, applying nonlinearities via a conventional digital processor, and then remodulating new optical signals for the next layer. While this method offers the flexibility of digital signal processing, it is constrained by the limited I/O bandwidth of electronic hardware, making it difficult to scale to large matrix dimensions with many optical inputs. Additionally, the reliance on analogue-to-digital conversion introduces latency and limits the processing speed to the GHz-scale clock rates of digital electronics, precisely the bottlenecks ONNs aim to overcome. Therefore, realizing optical nonlinear activation functions in hardware, without repeated opto-electro-optic conversions, is crucial for enabling scalable and high-speed integrated ONNs.

Based on the type of signals within the neurons, photonic neurons can be roughly divided into two categories: One is the optical/electrical/optical neuron, in which the original optical signals will be transformed into electrical ones before sending optical outputs to other neurons. The other one is the all-optical neuron, where no change will happen in the physical form of the optical inputs.

2.3.1 Through Opto-Electro-Optic Conversions

In SDM-based coherent ONNs, signals are weighted via phase modulation and summed all-optically through light interference. However, the implementation of nonlinearities has typically remained off-chip, performed by a digital computer. Williamson *et al.*[93] addressed this limitation by tapping a small portion of the propagating optical signal, converting it into the electrical domain via a photodetector, and electrically amplifying it to drive on-chip phase modulators (Figure 2.8(a)). This approach enables low-threshold nonlinear modulation of the remaining optical signal, demonstrating an integrated solution for nonlinearity implementation.

In WDM-based photonic neural network, the bosonic nature of light, which dictates that photons with different wavelengths interact weakly, grants lossless propagation of densely multiplexed optical signals with ultra-low crosstalk. However, this weak interaction limits the ability to perform algebraic operations directly between incoherent wavelength channels. As a result, weighted optical signals are typically summed at photodetectors and converted into electrical signals via opto-electric conversion. And then the non-linearities are either applied off-line with computer, with a high performance cost to digitize and remodulate signals, or through on-chip electro-optic conversions.

In 2016, Nahmias *et al.*[40] experimentally demonstrated a semiconductor laser neuron that leveraged WDM for multichannel optical fan-in (Figure 2.8(b)). Weighted optical signals were summed via a photodetector and used to drive a laser cavity, thereby enabling nonlinear activation and re-emission into the optical domain without requiring additional electronic circuitry.

Integrated nonlinearities have also been explored by using the intrinsic nonlinear electro-optic response of MRR modulators. For instance, Tait *et al.*[94] proposed

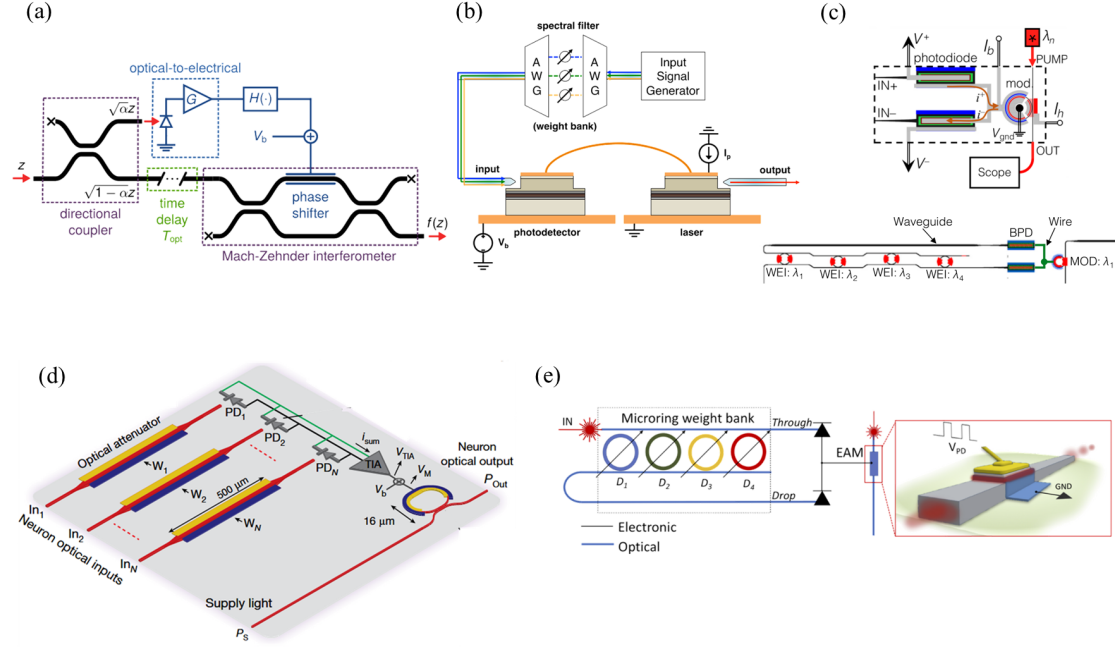


Figure 2.8: Optical/electrical/optical neuron (a) A photodetector-driven MZI-based nonlinear activation function for interferometric meshes[93]. From Williamson et al., "Reprogrammable Electro-Optic Nonlinear Activation Functions for Optical Neural Networks", IEEE. © 2019 IEEE. Reprinted, with permission. (b) An integrated analogue O/E/O link for wavelength multiplexed optical data[40]. Reproduced with permission from Nahmias et al., *Applied Physics Letters* (2016), © AIP Publishing. (c) A WDM silicon photonic modulator neuron[94]. Reproduced with permission from Tait et al., *Physical Review Applied* (2019), © American Physical Society. (d) An N-input photonic neuron, in which the photocurrent generated from the photodiodes is converted to voltage and amplified with a transimpedance amplifier (TIA) to drive the MRR modulator[42]. Reproduced with permission from Ashtiani et al., *Nature* (2022), © Springer Nature. (e) An electro-absorption modulator (EAM)-based photonic neuron for wavelength multiplexed signals[39]. Reproduced from Amin et al., *APL Materials* (2019), licensed under CC BY 4.0.

a silicon photonic modulator neuron based on MRRs (Figure 2.8(c)), employing a balanced photodiode to sum multiple wavelength channels, implement excitatory (positive) and inhibitory (negative) weights, and drive a MRR modulator to impose nonlinear transformation on a newly generated optical carrier. Building on this approach, Ashtiani *et al.*[42] adopted the MRR-modulator-based neuron in their on-chip demonstration of a photonic neural network. As shown in Figure 2.8(d), incoming optical signals are weighted using optical PIN attenuators, and then summed through an array of parallel photodetectors, with the resulting photocur-

rent converted and amplified by a transimpedance amplifier (TIA). The amplified voltage drives an MRR modulator to realize a ReLU-like nonlinear activation function, where the neuron output is generated by modulating an auxiliary continuous-wave optical carrier.

Similarly, an indium-tin-oxide (ITO)-based electro-absorption modulator (EAM) was employed to realize nonlinear activation functions during the electro-optical conversion[39]. As depicted in Figure 2.8(e), weighted WDM inputs are accumulated in the balanced photodiode, generating a photovoltage that modulates free-carrier absorption in an ITO layer monolithically integrated with a silicon photonic waveguide, thereby applying non-linear activation functions to the transmitted optical carrier.

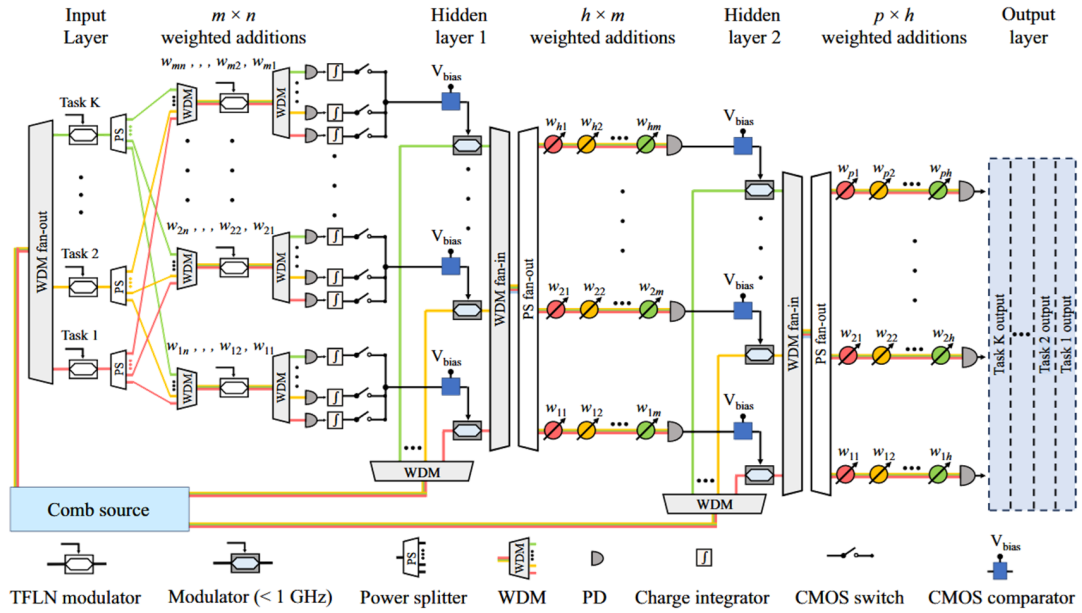


Figure 2.9: Wavelength-time neural network with non-linearities applied on-chip [43]. Reproduced from Lin et al., *Nature Communications* (2024), licensed under CC BY 4.0.

In current TDM-based ONNs, where dot-product operations are mapped in the temporal domain, time-multiplexed signals are converted by photoreceivers and accumulated as photogenerated charge in electrical integrators. In these cases,

nonlinearities are typically applied off-chip using digital computation. Lin *et al.*[43] theoretically proposed an on-chip approach to implement nonlinearities by employing photodetector-driven modulators in a hybrid ONN architecture that interleaves both TDM and WDM, as illustrated in Figure 2.9. In this design, charge-integrating photoreceivers compute vector–vector dot products at each wavelength channel. The output timing of each photoreceiver is controlled by commercial complementary metal–oxide–semiconductor (CMOS) switches, while a CMOS comparator selects the maximum between the input and reference voltages, thereby implementing a rectified linear unit (ReLU) activation function for each vector–vector dot product and encoding the result onto a new optical carrier.

This proposed architecture offers flexible scalability, enabled by interleaving both time and wavelength dimensions from the input layer to the first hidden layer. It also provides, in theory, unlimited cascadability, because information is re-encoded onto a newly generated carrier at each nonlinear layer, mitigating signal attenuation and preserving the signal-to-noise ratio (SNR) across successive stages.

2.3.2 Within the Optical Domain

The feasibility of all-optical neural networks remains uncertain due to the lack of an ideal all-optical nonlinear element. Photons, while nearly perfect carriers for long-distance, low-loss information transmission owing to their weak interactions, are inherently difficult to manipulate. In all-optical neurons, where signals remain entirely in the optical domain, nonlinear activation functions must be implemented through the interaction between the optical field and the material, which poses significant challenges due to the fundamentally weak nature of optical nonlinearities.

Mourgias-Alexandris *et al.*[95] demonstrated an all-optical neuron (Figure 2.10(a)) based on a semiconductor optical amplifier–loaded Mach–Zehnder interferometer (SOA-MZI), which is deeply saturated by two high-power continuous-wave signals (λ_0, λ_1) to achieve power equalization. A subsequent SOA acts as a cross-gain modulation–based wavelength converter (XGM-WC). This architecture enables the all-optical implementation of a sigmoid activation function on WDM-weighted signals. However, its energy consumption becomes significant when extended to multiple layers, primarily due to the continuous electrical biasing of all SOAs and the additional optical power required to maintain deep saturation in the SOA-MZI.

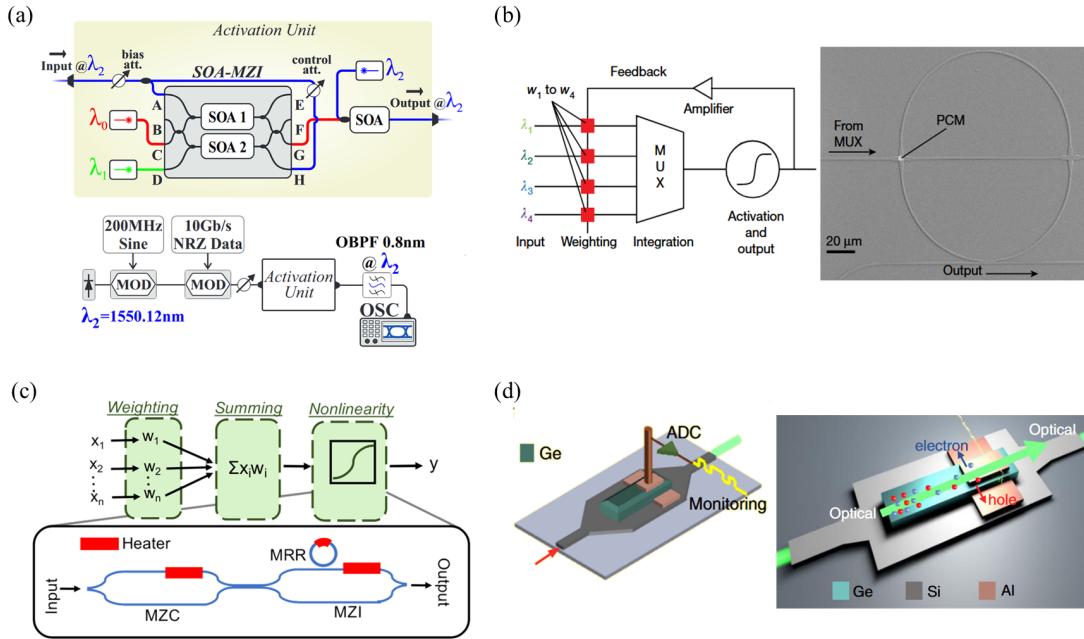


Figure 2.10: All-optical neuron (a) A semiconductor-optical-amplifier (SOA)-based all-optical neuron with sigmoid activation function[95]. Reproduced from Mourgias-Alexandris *et al.*, *Optics Express* (2019), licensed under CC BY 4.0. (b) A PCM-based spiking neuron where weighted inputs are thresholded with a PCM cell on ring resonator[21]. Reproduced with permission from Feldmann *et al.*, *Nature* (2019), © Springer Nature. (c) An all-optical neuron with reconfigurable nonlinear activation functions[96]. Reproduced from Jha *et al.*, *Optics Letters* (2020), licensed under CC BY 4.0. (d) An all-optical nonlinear unit based on a Ge-Si photodiode[97]. Reproduced from Shi *et al.*, *Nature Communications* (2022), licensed under CC BY 4.0.

PCMs, known for their non-volatility and energy efficiency in optical linear oper-

ations, have also been demonstrated to enable all-optical nonlinear activation by leveraging the inherent nonlinear spectral response of a ring resonator[21]. The nonlinear spiking behaviour arises from sudden changes in the material's optical properties. As illustrated in Figure 2.10(b), an output pulse is triggered only when the sum of the weighted input powers exceeds the PCM switching threshold, inducing a phase transition from the crystalline to the amorphous state.

Reconfigurable nonlinear activation functions in all-optical neurons were demonstrated by Jha *et al.* in 2020, using a cavity-loaded MZI implemented on a silicon photonic platform. By exploiting the free-carrier dispersion (FCD) effect in silicon, the device was able to emulate various activation functions, including sigmoid, radial basis, clamped rectified linear unit, and softplus, with tunable thresholds. As illustrated in Figure 2.10(c), these nonlinear responses can be flexibly programmed via embedded thermal heaters, enabling task-specific optimization in neuromorphic computing. However, the reliance on external electrical circuits to control the heaters introduces additional system complexity and contributes to increased energy consumption due to the continuous electrical biasing required.

The free-carrier absorption (FCA) effect in germanium has also been explored to implement all-optical nonlinear activation functions, as illustrated in Figure 2.10(d). Shi *et al.*[97] engineered the Ge-Si photodiode by shortening the electrode length to enhance the FCA response, thereby enabling a saturable absorber-like transfer function on the transmitted optical signal while simultaneously preserving electrical readout for in-situ monitoring. This design enables both all-optical nonlinearity and non-intrusive monitoring within a compact footprint of $4.3\,\mu\text{m} \times 8\,\mu\text{m}$. The functionality was experimentally validated in a three-layer photonic neural network. However, since the nonlinear transfer function is applied directly to the same optical carrier, it introduces power attenuation at each layer. This cumulative loss becomes

significant in cascaded architectures, necessitating repeated optical amplification, as implemented in their experimental demonstration.

2.4 The Need for All-Optical Nonlinearity Implementation for Scalable Optical Neural Networks

2.4.1 Functional Gaps in Photonic Neural Networks

The pursuit of scalable ONNs is motivated by the intrinsic advantages of photonic hardware, including massive parallelism, low latency, and ultrahigh bandwidth. These properties make integrated photonic platforms attractive for accelerating artificial intelligence inference and signal processing workloads. Multiplexing strategies such as wavelength-division and time-division schemes exploit orthogonal spectral or temporal degrees of freedom to encode independent data streams, enabling high-throughput linear computation within compact circuit footprints. As reviewed in Section 2.2, such architectures have demonstrated performance at the scale of trillions of operations per second while maintaining low static power consumption. These results reflect the maturity of photonic circuits for implementing large-scale linear transformations.

Despite these advances, existing photonic accelerators still lack the essential functional primitives required for implementing complete neural computation in the optical domain, namely activation and accumulation, which must be applied at every neuron in every layer[36, 37]. Although photonic devices can implement weighted linear transformations efficiently, the absence of an effective and scalable nonlinear activation mechanism within the optical layer has remained a fundamental limitation. As discussed in Section 2.3, most ONNs currently adopt a hybrid strategy in which multiplication is performed in the optical domain, but

activation and signal accumulation are delegated to electronic subsystems. Each conversion between optical and electrical domains introduces latency, dissipates energy, and adds hardware complexity, offsetting the benefits of optical throughput and parallelism. Furthermore, the performance becomes constrained by the electronic bandwidth of transimpedance amplifiers and modulators. As systems based on wavelength or time multiplexing expand in scale, these conversion steps become a growing obstacle to full optical scalability.

To overcome this bottleneck, an ideal nonlinear element should operate entirely in the optical domain while being compact, low-power, and compatible with dense WDM/TDM integration. However, achieving such functionality is inherently challenging because photons do not naturally interact. Various material and structural routes have been investigated to emulate nonlinear activation, including instantaneous Kerr effects and free-carrier dispersion in silicon, electro-absorption in III–V semiconductors[95], two-dimensional materials[98], and PCMs[84, 99, 100]. Kerr and carrier-based mechanisms offer femtosecond response but require high optical intensities and provide limited modulation depth. PCMs, on the other hand, deliver large refractive-index contrast and multilevel tunability with non-volatility, but constrained by endurance. These trade-offs have motivated the exploration of nonlinearities that combine analogue reconfigurability with optical cascadability and modest energy consumption.

This gap motivates the search for an alternative class of nonlinear mechanisms. What is required is a process that allows light to modulate light in a way that is programmable while supporting cascaded operation and integration with wavelength- and time-multiplexed schemes.

2.4.2 Standing-Wave Photonics as an Architectural Precursor

Rather than searching exclusively for new material platforms, it is therefore instructive to return to the underlying wave-physics mechanisms capable of mediating such interactions. In this regard, coherent perfect absorption (CPA) provides a particularly useful conceptual foundation.

CPA arises when coherent input fields are arranged such that they interfere destructively in all radiative channels while constructively enhancing absorption within the medium. This effect first established as the time-reversed analogue of lasing[101, 102] and later generalized across plasmonic, dielectric, and metamaterial platforms[103, 104]. Because the absorbed power depends sensitively on the relative phase and amplitude of the incident beams, CPA enables one optical channel to modulate the dissipation experienced by another without invoking intrinsic material nonlinearity. A detailed theoretical treatment of the scattering formalism underlying CPA, including its eigenmode structure, zero-determinant condition, and relation to guided-wave standing modes, is provided in Chapter 3. This principle has directly inspired standing-wave engineered photonic structures, in which spatial interference patterns concentrate optical energy at subwavelength absorbers and thereby provide a route to programmable light–matter interaction.

Building on this foundation, a series of recent works have demonstrated how standing-wave interference can be harnessed to engineer CPA-like behaviour in integrated photonic platforms through spatially localized absorption. As shown in Figure 2.11(a), Bruck *et al.* [105] theoretically evaluated implementing CPA in a silicon-on-insulator (SOI) waveguide using plasmonic nanoantennas positioned at field maxima of specific propagating modes. Their numerical simulations showed extinction ratios exceeding 25 dB with transmission losses below 1% in the transparent state,

establishing that broadband absorbers strategically positioned relative to standing-wave nodes could achieve mode-selective attenuation without active tuning.

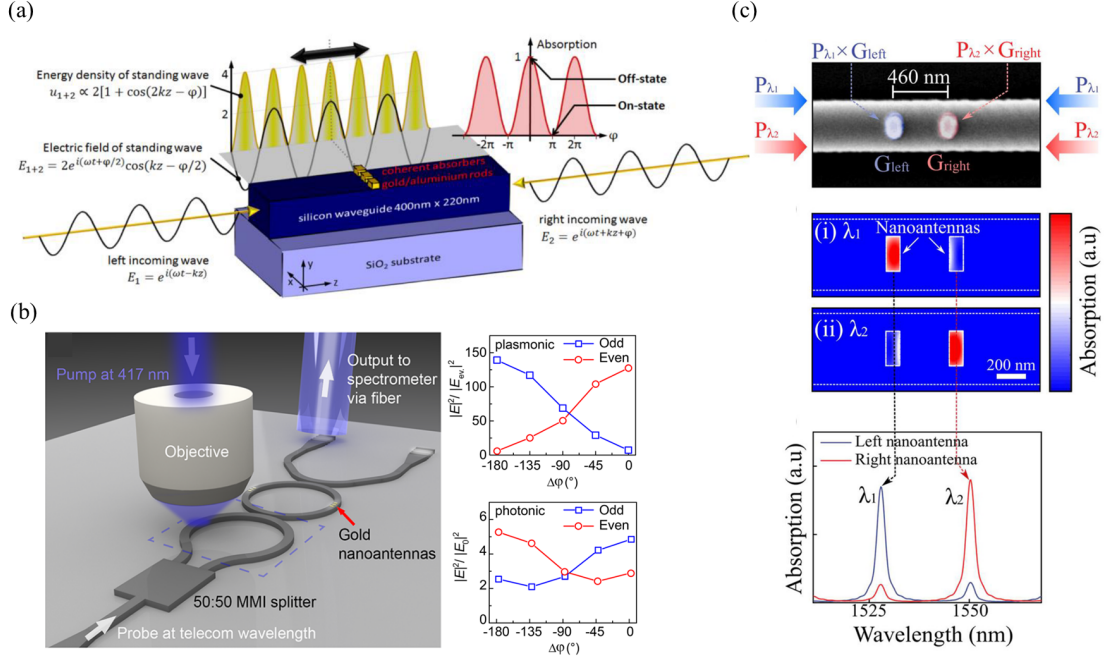


Figure 2.11: Standing-wave engineered CPA in integrated photonics. (a) Selective attenuation of even-order standing-wave modes achieved by embedding gold nanoantennas at antinodes within a silicon waveguide [105]. Reproduced from Bruck and Muskens, *Optics Express* (2013), licensed under CC BY 4.0. (b) Phase-sensitive hybrid photon-plasmon coupling in a ring resonator, realised by placing gold nanoantennas at controlled field positions [49]. Reproduced with permission from Chen et al., *Nano Letters* (2018), © American Chemical Society. (c) Spatio-spectral reconfiguration achieved by placing PCM nanoantennas at wavelength-dependent antinodes, allowing selective amplitude tuning at different wavelengths[50]. Reproduced from Lee et al., *Nanophotonics* (2024), licensed under CC BY 4.0.

Chen *et al.*[49] extended this concept to MRR, exploiting phase-sensitive hybrid photon-plasmon coupling in standing-wave configurations formed by counter-propagating clockwise and counter-clockwise modes. As illustrated in Figure 2.11(b)), nanoantennas placed at selected resonance antinodes coupled selectively to odd or even whispering-gallery modes, allowing photo-plasmon coupling to be dynamically modulated through phase control at the input.

Most recently, Lee *et al.*[50] combined standing-wave photonics with non-volatile PCMs to achieve spatio-spectral reconfiguration. As depicted in Figure 2.11(c),

their architecture exploited the wavelength dependence of standing-wave antinode positions: by placing GST nanoantennas at the spatial maximum of one resonance but the minimum of another, they achieved selective absorption at one wavelength while maintaining transparency at adjacent channels. This approach enabled multi-level, wavelength-addressable modulation with non-volatile state retention, demonstrating independent control over multiple spectral channels within a single microring element.

These standing-wave photonic architectures present an alternative to conventional wavelength-division systems by embedding wavelength-dependent functionality directly into the spatial field distribution. Rather than requiring separate resonators or spatial paths for each channel, they harness the intrinsic interference patterns of counter-propagating waves to spatially encode spectral information within a shared structure. This capability positions standing-wave photonics as a promising foundation for more advanced optical signal processing, yet its demonstrated functionality to date has remained fundamentally passive.

2.4.3 From Passive Modulation to All-Optical Computation

While standing-wave photonics has demonstrated elegant methods for wavelength-selective modulation, its application has largely remained confined to static, passive spectral filtering. In the current demonstrations, absorbers interact with the optical field at isolated antinodes, enabling binary or multilevel extinction at selected resonances, but lacking investigation on dynamic high-precision tunability, cross-channel interaction, or cumulative signal processing. Furthermore, these systems rely on spatial field selectivity alone, without exploiting temporal dynamics or energy accumulation mechanisms, which are essential for implementing neural operations such as summation and nonlinear activation.

This thesis revisits the thermo-optic effect, traditionally regarded as a slow and parasitic mechanism in photonic systems, and reinterprets it as a viable route to achieving functional optical nonlinearity. Through the deliberate engineering of spatially localized absorption and thermal diffusion, optical signals can be transferred across wavelengths via thermally mediated interactions, enabling direct coupling between distinct spectral channels within a single photonic circuit.

Building on this principle, a new architecture termed the Photonic-Heater-in-Lightpath (PHIL) is developed, in which subwavelength nanoheaters embedded within MRRs at precise standing-wave antinode locations. These nanoheaters convert dissipated optical energy into localized, controllable thermo-optic phase shifts, allowing the resonator's spectral response to be dynamically reconfigured by optical inputs alone.

Crucially, the temporal dynamics of thermal relaxation, often viewed as a limitation, are repurposed here as a computational feature. The finite thermal decay time defines an effective integration window, during which successive high-speed optical pulses can be accumulated into a single analogue state. This state is then mapped onto the spectral transmission of the MRR, governed by its Lorentzian resonance profile, yielding a programmable nonlinear response. Together, these mechanisms establish PHIL as a compact, cascable, and wavelength-multiplexed unit for all-optical computation, extending the standing-wave paradigm from passive filtering into dynamic signal processing and neuromorphic functionality.

3

Theory

Contents

3.1 Dielectric Waveguides	47
3.2 Grating Couplers	51
3.3 Microring Resonators	53
3.3.1 Resonance Condition	55
3.3.2 Transmission and Critical Coupling	56
3.3.3 Spectral Characteristics	58
3.4 Coherent Perfect Absorption	62

The development of an integrated photonic device relies on several fundamental components, including grating couplers for coupling light from optical fibres into on-chip circuits, dielectric waveguides for guiding light, and microring resonators for selectively coupling resonant light into closed-loop paths. This chapter introduces the theoretical principles underlying these elements, as well as the concept of coherent perfect absorption, which forms the basis of the device design.

3.1 Dielectric Waveguides

The propagation of light in dielectric waveguides relies on the principle of total internal reflection (TIR) occurring at the interface between a high-index core and lower-index cladding layers. These waveguides serve as the fundamental building blocks in integrated photonic circuits, guiding optical signals between functional

units with minimal loss.

A typical slab waveguide consists of a core layer with refractive index n_1 , bounded above and below by cladding layers with refractive indices n_2 and n_3 , where $n_1 > n_2, n_3$ (Figure 3.1(a)).

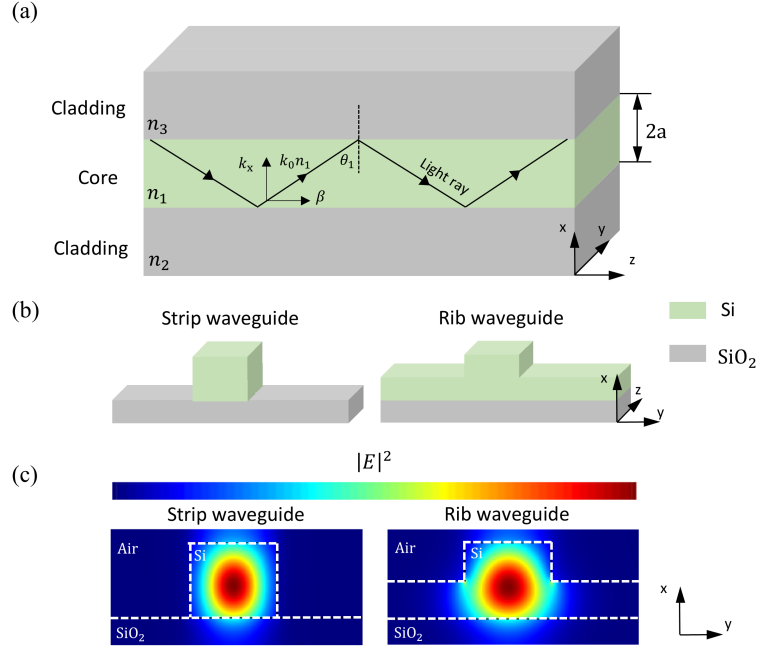


Figure 3.1: Dielectric waveguide (a) Slab waveguide and light confinement via total internal reflection. (b) Strip and rib waveguide geometries on SOI platforms. (c) Simulated TE-mode field intensity $|E|^2$ for both geometries, showing differences in confinement and symmetry.

Light propagates along the \hat{z} -axis and is confined in the transverse \hat{x} -direction by TIR. The electric and magnetic fields in such a waveguide are governed by Maxwell's equations. For a linear, isotropic, and source-free dielectric medium, the curl equations are given by[106]:

$$\begin{aligned}\nabla \times \mathbf{E} &= -\mu_0 \frac{\partial \mathbf{H}}{\partial t} \\ \nabla \times \mathbf{H} &= \varepsilon_0 n^2(x, y) \frac{\partial \mathbf{E}}{\partial t}\end{aligned}\tag{3.1}$$

Assuming a time-harmonic field dependence $e^{j(\omega t - \beta z)}$, the electric and magnetic

fields are expressed as:

$$\mathbf{E}(x, y, z, t) = \mathbf{E}(x, y)e^{j(\omega t - \beta z)}, \quad \mathbf{H}(x, y, z, t) = \mathbf{H}(x, y)e^{j(\omega t - \beta z)} \quad (3.2)$$

Substituting into Maxwell's equations yields the spatial field equations for $\mathbf{E}(x, y)$ and $\mathbf{H}(x, y)$, with β representing the propagation constant along the z -axis. In a slab waveguide with no variation in the y -direction ($\partial/\partial y = 0$), the field equations simplify and decouple into two modal families:

For Transverse Electric (TE) modes, the electric field is polarized in the y -direction ($E_y \neq 0$, $E_x = E_z = 0$), and the non-zero magnetic components are H_x and H_z . The wave equation becomes:

$$\frac{d^2 E_y}{dx^2} + (k_0^2 n^2(x) - \beta^2) E_y = 0 \quad (3.3)$$

For Transverse Magnetic (TM) modes, the magnetic field is polarized in the y -direction ($H_y \neq 0$, $H_x = H_z = 0$), and the non-zero electric field components are E_x and E_z . The TM wave equation is:

$$\frac{d}{dx} \left(\frac{1}{n^2(x)} \frac{dH_y}{dx} \right) + \left(k_0^2 - \frac{\beta^2}{n^2(x)} \right) H_y = 0 \quad (3.4)$$

These equations describe field confinement in the core region and exponential decay in the cladding when the following condition is satisfied:

$$k_0^2 n_1^2 > \beta^2 > \max(k_0^2 n_2^2, k_0^2 n_3^2) \quad (3.5)$$

The effective index is defined as $n_{\text{eff}} = \beta/k_0$. To obtain physically meaningful solutions, boundary conditions at the dielectric interfaces must be satisfied. Specifically,

the continuity of the transverse electric field and its derivative at the core-cladding boundaries ($x = \pm a$) requires that:

$$E_y(x) \quad \text{and} \quad \frac{dE_y(x)}{dx} \quad \text{are continuous at} \quad x = \pm a \quad (3.6)$$

These continuity conditions arise from Maxwell's equations and ensure that there are no surface charges or discontinuities at the dielectric interfaces. Applying them leads to a set of transcendental equations whose discrete solutions correspond to the allowed guided modes in the waveguide. Each solution gives a specific propagation constant β , or equivalently, a discrete effective index n_{eff} , characterizing a supported mode. Under these conditions, the field distributions can be computed analytically for symmetric slab waveguides or numerically for more complex index profiles.

In practice, integrated photonic platforms employ waveguides with confinement in both x - and y -directions, such as strip and rib geometries (Figure 3.1(b)). The electromagnetic fields in these 2D waveguides are formally governed by the vectorial Helmholtz equation. However, in dielectric waveguides supporting quasi-TE or quasi-TM modes, one field component typically dominates. Under this approximation, the dominant field component can be described by the scalar Helmholtz equation:

$$\nabla^2 E_z(x, y) + (\beta^2 - k_0^2 n^2(x, y)) E_z(x, y) = 0 \quad (3.7)$$

$$\nabla^2 H_z(x, y) + (\beta^2 - k_0^2 n^2(x, y)) H_z(x, y) = 0 \quad (3.8)$$

Due to the lack of symmetry in 2D structures, analytical solutions are generally not available. To approximate 2D waveguide modes analytically, the Effective Index

Method (EIM) can be used[107]. This method decomposes the 2D waveguide into orthogonal 1D slabs:

1. Solve the vertical (y -axis) slab waveguide to obtain $n_{\text{eff,vertical}}(x)$.
2. Use this as the refractive index profile in the x -direction to solve the horizontal slab.

The EIM provides accurate approximations when refractive index variations along the two transverse axes are weakly coupled, but becomes less effective for complex geometries. Therefore, numerical methods are employed in this thesis to calculate mode profiles. Figure 3.1(c) shows simulated TE-mode intensity distributions for strip and rib waveguides, obtained using the Lumerical finite-difference eigenmode (FDE) solver, which directly solves Maxwell's equations in a 2D cross-sectional geometry using a finite-difference approach. These mode profiles affect key photonic design metrics such as coupling efficiency, bending loss, and mode overlap. In this work, rib waveguides (half-etched) on silicon-on-insulator (SOI) substrates are used due to their strong confinement and fabrication compatibility

3.2 Grating Couplers

Edge and grating couplers [108] are two mostly applied techniques to couple light from optical fibres into integrated waveguides. While edge couplers offer advantages such as high coupling efficiency and broad operational bandwidth, their use is limited by the requirement for physical access to the chip edge, which can hinder experimental flexibility. To support vertical light coupling and simplify alignment during measurements, this thesis adopts grating couplers.

As illustrated in Figure 3.2(a), the grating coupler comprises a periodic array of partially etched silicon structures on a SOI substrate. These structures redirect incident light from free space into the in-plane guided mode, or vice versa, by satisfying a specific phase-matching condition derived from the Bragg diffraction principle [108]:

$$n_i k_0 \sin \theta_i + \frac{2\pi m}{p} = \beta \quad (3.9)$$

In this expression, n_i is the refractive index of the incident medium (typically air), k_0 is the free-space wave vector, θ_i is the incident angle, m is the diffraction order, p is the grating period, and $\beta = \frac{2\pi n_{\text{eff}}}{\lambda}$ represents the propagation constant of the guided mode. Assuming first-order diffraction ($m = 1$), the grating period can be




Figure 3.2: Grating coupler. (a) Cross-sectional schematic of a uniform grating coupler. θ_i : incident angle; $n_{\text{tooth}}, n_{\text{etch}}$: effective index of the unetched tooth and etched. (b) Top view of an apodised grating coupler used in this work. The fill factor increases gradually along the z -axis, from a lower value at the input to a higher value at the output, to improve mode matching and coupling efficiency.

written as:

$$p = \frac{\lambda}{n_{\text{eff}} - \sin \theta_i} \quad (3.10)$$

The effective refractive index n_{eff} of the grating region is approximated by a weighted average of the refractive indices of the etched and unetched regions:

$$n_{\text{eff}} = ff \cdot n_{\text{tooth}} + (1 - ff) \cdot n_{\text{etch}} \quad (3.11)$$

Here, n_{tooth} and n_{etch} denote the refractive indices of the silicon tooth and the etched trench (typically air), respectively. The fill factor ff is defined as the ratio of the tooth width to the grating period, given by $ff = (p - l_e)/p$, where l_e is the etched length in one grating period. This formulation corresponds to the geometry shown in Figure 3.2(a), where light is incident at angle θ_i and guided along the z -direction within the waveguide layer.

To further improve the coupling efficiency, apodised grating coupler[109] with gradual variation in the ff are applied in this thesis (Figure 3.2(b)), and incident angle θ_i is fixed at 8° during all subsequent optical measurements[110]. For broader wavelength coupling, the use of 3D grating couplers[111] is recommended for future work.

3.3 Microring Resonators

Integrated microring resonators (MRRs) are composed of two-dimensional waveguides that confine and guide light in a closed-loop geometry, enabling selective coupling to input and output waveguides. The theoretical framework discussed in this section is based on the treatment presented in [112]. A typical MRR structure consists of a circular waveguide evanescently coupled to one or more straight bus waveguides[113], as illustrated in Figure 3.3. The most basic configuration involves

unidirectional coupling between a single bus waveguide and the ring, as depicted in Figure 3.3(a).

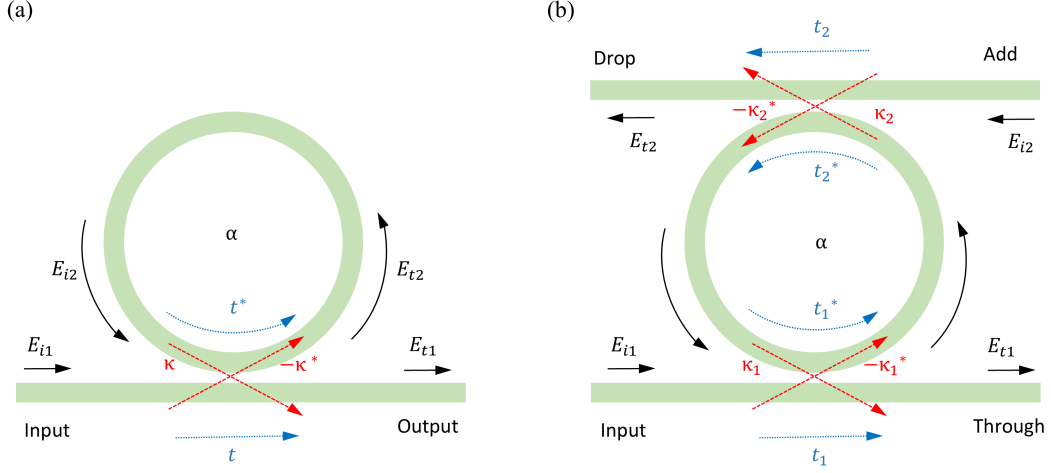


Figure 3.3: Schematic illustrations of microring resonators (MRR) with coupling parameters. (a) All-pass configuration with a single bus waveguide. The input field E_{i1} is partially coupled into the ring via evanescent cross-coupling coefficients κ with transmission coefficient t . The intracavity field accumulates a round-trip loss α and re-couples back to produce the output field E_{t1} . (b) Add-drop configuration with symmetric input and output couplers. Fields are coupled into and out of the ring through two couplers characterized by (t_1, κ_1) and (t_2, κ_2) , respectively. The resonator enables spectral filtering and routing between the input, through, and drop ports.

The electromagnetic field interaction at the coupling region can be described by the following matrix equation:

$$\begin{pmatrix} E_{t1} \\ E_{t2} \end{pmatrix} = \begin{pmatrix} t & \kappa \\ -\kappa^* & t^* \end{pmatrix} \begin{pmatrix} E_{i1} \\ E_{i2} \end{pmatrix} \quad (3.12)$$

where t and κ denote the transmission and evanescent cross-coupling coefficients, respectively, and t^* and κ^* represent their complex conjugates. Under the assumption of a lossless coupling mechanism, the transformation matrix satisfies the unitarity condition, ensuring energy conservation:

$$|t|^2 + |\kappa|^2 = 1 \quad (3.13)$$

The field in the ring accumulates attenuation and phase delay with each round trip. Assuming $E_{i1} = 1$ for model simplification, the feedback relation becomes:

$$E_{i2} = \alpha e^{j\phi} E_{t2} \quad (3.14)$$

where α is the loss coefficient of the ring ($\alpha = 1$ for lossless system) and ϕ is the phase delay per round trip, given by:

$$\phi = \frac{\omega L}{c} \quad (3.15)$$

Here, $L = 2\pi r$ is the ring circumference with r being the radius measured from the centre of the ring to the centre of the waveguide. The phase velocity of the ring mode is $c = \frac{c_0}{n_{\text{eff}}}$, where c_0 is the speed of light in vacuum and n_{eff} is the effective refractive index of the ring mode.

The angular frequency is defined as $\omega = kc_0$ and the vacuum wavenumber k relates to the wavelength λ by $k = \frac{2\pi}{\lambda}$. Substituting these definitions, the total phase delay becomes:

$$\phi = \frac{\omega L}{c} = \frac{kc_0 L}{c} = k \cdot n_{\text{eff}} \cdot 2\pi r = \frac{2\pi \cdot n_{\text{eff}} \cdot 2\pi r}{\lambda} = n_{\text{eff}} \frac{4\pi^2 r}{\lambda}. \quad (3.16)$$

3.3.1 Resonance Condition

To build up field inside the ring (on resonance), light completing one round trip should be in phase with newly coupled light to interfere constructively. This requires:

$$\phi + \phi_t = 2\pi m, \quad m \in \mathbb{Z} \quad (3.17)$$

where ϕ_t represents the additional phase shift introduced by the coupler, which is typically $\frac{\pi}{2}$ for directional couplers[114, 115], and m is the resonance index that uniquely labels each resonant mode supported by the ring. This condition can also be written in the form as below:

$$\beta L = 2\pi m, \quad m \in \mathbb{Z} \quad (3.18)$$

where $\beta = \frac{2\pi n_{\text{eff}}}{\lambda}$ is the propagation constant.

3.3.2 Transmission and Critical Coupling

By combining 3.12, 3.13 and 3.14, we can obtain the following field equations:

$$E_{t1} = \frac{-\alpha + t \cdot e^{-j\phi}}{-\alpha t^* + e^{-j\phi}}, \quad (3.19)$$

$$E_{i2} = \frac{-\alpha \kappa^*}{-\alpha t^* + e^{-j\phi}}, \quad (3.20)$$

$$E_{t2} = \frac{-\kappa^*}{1 - \alpha t^* e^{j\phi}}. \quad (3.21)$$

Solving the system yields the through-port power transmission:

$$P_{t1} = |E_{t1}|^2 = \frac{\alpha^2 + |t|^2 - 2\alpha|t| \cos(\phi + \phi_t)}{1 + \alpha^2|t|^2 - 2\alpha|t| \cos(\phi + \phi_t)} \quad (3.22)$$

and the circulating power in the ring can be written as:

$$P_{i2} = |E_{i2}|^2 = \frac{\alpha^2(1 - |t|^2)}{1 + \alpha^2|t|^2 - 2\alpha|t| \cos(\phi + \phi_t)} \quad (3.23)$$

$$t = |t|e^{j\phi_t} \quad (3.24)$$

where $|t|$ denotes the coupling loss. Under the resonance condition given in 3.17, the through-port power transmission and circulating power can be obtained as:

$$P_{t1} = \left(\frac{\alpha - |t|}{1 - \alpha|t|} \right)^2 \quad (3.25)$$

$$P_{i2} = \left(\frac{\alpha - \alpha|t|}{1 - \alpha|t|} \right)^2 \quad (3.26)$$

At critical coupling, when $\alpha = |t|$, the internal loss equals to the coupling losses, resulting in:

$$P_{t1} = 0 \quad (3.27)$$

In an add-drop resonator with two couplers (Figure 3.3(b)), characterized by (t_1, κ_1) and (t_2, κ_2) , the half round trip loss and phase can be written as $\alpha_{0.5}$ and $\phi_{0.5}$, respectively, where $\alpha = \alpha_{0.5}^2$ and $\phi = 2\phi_{0.5}$. As a result, the corresponding electric field in the through port E_{t1} and drop port E_{t2} become the following:

$$E_{t1} = \frac{t_1 - t_2^* \alpha e^{-j\phi}}{1 - t_1^* t_2^* \cdot \alpha e^{-j\phi}} \quad (3.28)$$

$$E_{t2} = \frac{-\kappa_1^* \kappa_2 \cdot \alpha_{0.5} e^{-j\phi_{0.5}}}{1 - t_1^* t_2^* \cdot \alpha e^{-j\phi_{0.5}}} \quad (3.29)$$

and the drop-port power transmission at resonance is:

$$P_{t2, \text{Resonance}} = |E_{t2}|^2 = \frac{(1 - |t_1|^2)(1 - |t_2|^2) \cdot \alpha}{(1 - \alpha|t_1 t_2|)^2} \quad (3.30)$$

For an ideal lossless system ($\alpha = 1$) with symmetric coupling ($t_1 = t_2$), the field amplitude at the through port (E_{t1}) becomes 0 due to complete destructive interference. For a passive resonator system with a fixed loss constant α , the

maximum power transfer to the drop port ($P_{t2,\text{Resonance}}$) occurs when the through-port transmission is minimized ($P_{t1} = 0$). This critical coupling condition is achieved when:

$$\alpha = \frac{|t_1|}{|t_2|} \quad (3.31)$$

3.3.3 Spectral Characteristics

The spectral characteristics of MRRs enable their implementation in various applications, particularly as optical filters that selectively transmit specific wavelength bands at resonance. The filtering performance of these devices is characterized by several key figures of merit that quantify their spectral behaviour.

The Free Spectral Range (FSR) represents one of the most fundamental parameters, defined as the spectral spacing between adjacent resonance peaks in either frequency or wavelength domain. This parameter establishes the periodicity of the resonator's transmission spectrum and determines the repeat interval of the device's spectral response.

According to the resonance condition in 3.18, for two adjacent resonances (m and $m + 1$):

$$\beta_1 L = 2\pi m \quad (3.32)$$

$$\beta_2 L = 2\pi(m + 1) \quad (3.33)$$

Taking the difference:

$$(\beta_2 - \beta_1)L = 2\pi \quad (3.34)$$

$$\Delta\beta \cdot L = 2\pi \quad (3.35)$$

Using the chain rule, we get:

$$\Delta\beta = \left(\frac{d\beta}{d\lambda} \right) \cdot \Delta\lambda \quad (3.36)$$

Substituting $\Delta\beta$ in :

$$\left(\frac{d\beta}{d\lambda} \right) \cdot \Delta\lambda \cdot L = 2\pi \quad (3.37)$$

Solving for $\Delta\lambda$ yields the FSR:

$$\text{FSR} = |\Delta\lambda| = \left| \frac{2\pi}{L} \cdot \left(\frac{d\beta}{d\lambda} \right)^{-1} \right| \quad (3.38)$$

When neglecting the wavelength dependence of n_{eff} , we have:

$$\frac{d\beta}{d\lambda} \approx -\frac{\beta}{\lambda} = -\frac{2\pi n_{\text{eff}}}{\lambda^2} \quad (3.39)$$

Therefore:

$$\text{FSR} \approx \left| \frac{2\pi}{L} \cdot \left(-\frac{\lambda^2}{2\pi n_{\text{eff}}} \right) \right| \quad (3.40)$$

$$= \frac{\lambda^2}{n_{\text{eff}}L} \quad (3.41)$$

When wavelength dependence cannot be neglected, the group index $n_g = n_{\text{eff}} - \lambda \frac{dn_{\text{eff}}}{d\lambda}$ must be used. Since $\frac{d\beta}{d\lambda} = -\frac{2\pi n_g}{\lambda^2}$, the FSR becomes

$$\text{FSR} = |\Delta\lambda| = \frac{\lambda^2}{n_g L} \quad (3.42)$$

FSR determines the maximum channel spacing in wavelength division multiplexing applications and sets the upper limit for the number of resolvable spectral features within the device's operating bandwidth.

The spectral linewidth of a resonance, characterized by the full width at half maximum (FWHM), represents another critical performance metric for ring resonators. This parameter quantifies the 3 dB bandwidth of the resonance profile and directly relates to the quality factor of the device.

Beginning with the normalized drop port transmission coefficient from the coupled-mode analysis in 3.29 and 3.30:

$$\left| \frac{-\kappa_1^* \kappa_2 \alpha_{1/2} e^{i\phi_{1/2}}}{1 - t_1^* t_2 \alpha e^{i\phi}} \right|^2 = \frac{1}{2} \cdot \frac{|\kappa_1|^2 |\kappa_2|^2 \alpha}{(1 - \alpha |t_1 t_2|)^2} \quad (3.43)$$

Assuming purely real coupling coefficients and negligible losses, the transmission intensity becomes:

$$\frac{(\kappa_1 \kappa_2 \alpha_{1/2})^2}{1 - 2t_1 t_2 \alpha \cos(\phi) + (t_1 t_2 \alpha)^2} = \frac{1}{2} \cdot \frac{(\kappa_1 \kappa_2 \alpha_{1/2})^2}{(1 - t_1 t_2 \alpha)^2} \quad (3.44)$$

Cross-multiplying and rearranging the above equation gives:

$$2(1 - t_1 t_2 \alpha)^2 = 1 - 2t_1 t_2 \alpha \cos(\phi) + (t_1 t_2 \alpha)^2 \quad (3.45)$$

Using the Taylor expansion for small phase deviation from resonance, we get:

$$\cos(\phi) \approx 1 - \frac{\phi^2}{2} \quad (3.46)$$

Substituting the $\cos(\phi)$ in equation 3.45 and solving for ϕ^2 :

$$\phi^2 = \frac{(1 - t_1 t_2 \alpha)^2}{t_1 t_2 \alpha} \quad (3.47)$$

Assuming symmetric coupling $t_1 = t_2 = t$ and negligible round-trip loss:

$$\phi = \sqrt{\frac{(1 - t^2)^2}{t^2}} = \frac{1 - t^2}{t} \quad (3.48)$$

Converting from phase to wavelength using the relationship between phase deviation and wavelength shift:

$$2\delta\lambda = \frac{\lambda^2}{\pi L n_{\text{eff}}} \cdot \frac{1 - t^2}{t} \quad (3.49)$$

For weak coupling $\kappa^2 \ll 1$, and using the approximation $1 - t^2 \approx \kappa^2$:

$$\text{FWHM} = 2\delta\lambda \approx \frac{\kappa^2 \lambda^2}{\pi L n_{\text{eff}}} \quad (3.50)$$

The finesse F of the ring indicates the "contrast" between resonance and non-resonance, quantifying the spectral selectivity of the device. It is mathematically expressed as the ratio between the FSR and the resonance linewidth:

$$F = \frac{\text{FSR}}{\text{FWHM}} = \pi \cdot \frac{t}{1 - t^2} \approx \frac{\pi}{\kappa^2} \quad (3.51)$$

where the approximation holds for weak coupling conditions ($\kappa^2 \ll 1$).

The quality factor Q provides another measure of resonator performance, defined as the ratio of the operating wavelength to the spectral width of the resonance:

$$Q = \frac{\lambda}{\text{FWHM}} = \frac{n_{\text{eff}} L}{\lambda} \cdot \frac{t}{1 - t^2} = \frac{n_{\text{eff}} L}{\lambda} \cdot F \quad (3.52)$$

Q quantifies the number of oscillations before energy decay, and a higher Q indicates lower loss in the resonator system. This provides physical insight into the resonator's ability to confine and sustain optical energy, making it a crucial parameter for applications requiring high spectral purity or extended photon lifetimes.

One important contributor to resonator loss in microring devices is bending radiation loss, which depends strongly on the ring radius. When light propagates in a curved waveguide, the optical mode shifts toward the outer edge of the bend, and if the radius becomes too small, part of the optical field radiates into the surrounding cladding. Consequently, smaller radii generally lead to increased bending loss and reduced resonator quality factor[116, 117].

In this work, different microring radii were adopted for simulation and experiment. In the simulations, a compact radius of 4 μm was used to verify the physical mechanism while reducing computational cost. In the fabricated devices, a larger radius of 30 μm was chosen to suppress bending-induced radiation loss and improve resonator performance. In addition, the larger ring circumference leads to a smaller free spectral range, allowing more resonant wavelengths within the same spectral window and thus supporting a greater number of wavelength channels.

3.4 Coherent Perfect Absorption

Coherent perfect absorption (CPA) is a phenomenon where incoming electromagnetic radiation is completely extinguished through coherent illumination of an absorber, effectively representing a time-reversed process of lasing[101, 102]. This is achieved through constructive standing-wave interference at the absorber, generated by two counter-propagating waves of equal amplitude.

For a simplified two-port CPA system, the counter-propagating input waves along z -axis are expressed as [103, 118]:

$$E_k^{(\text{in})} = \mathcal{E}_k^{(\text{in})} e^{j(kz - \omega t)} \quad (3.53)$$

$$E_{-k}^{(\text{in})} = \mathcal{E}_{-k}^{(\text{in})} e^{-j(kz + \omega t)} \quad (3.54)$$

where $\mathcal{E}_k^{(\text{in})}$ and $\mathcal{E}_{-k}^{(\text{in})}$ denote the complex amplitudes, k is the wave vector, t is the time variable, and ω is the angular frequency. The corresponding output fields are:

$$E_k^{(\text{out})} = \mathcal{E}_k^{(\text{out})} e^{j(kz - \omega t)} \quad (3.55)$$

$$E_{-k}^{(\text{out})} = \mathcal{E}_{-k}^{(\text{out})} e^{-j(kz + \omega t)} \quad (3.56)$$

Assuming a reciprocal and symmetric two-port system with reflection coefficient r and transmission coefficient t , the input and output fields are related by the scattering matrix:

$$\begin{pmatrix} \mathcal{E}_k^{(\text{out})} \\ \mathcal{E}_{-k}^{(\text{out})} \end{pmatrix} = \begin{pmatrix} r & t \\ t & r \end{pmatrix} \begin{pmatrix} \mathcal{E}_k^{(\text{in})} \\ \mathcal{E}_{-k}^{(\text{in})} \end{pmatrix}, \quad (3.57)$$

Perfect absorption occurs when the outgoing fields vanish for a non-zero input vector,

$$\mathcal{E}_k^{(\text{out})} = \mathcal{E}_{-k}^{(\text{out})} = 0. \quad (3.58)$$

This condition is satisfied when the scattering matrix possesses a zero eigenvalue, indicating that one eigenmode of the input field produces no outgoing radiation and is therefore completely absorbed by the system.

The eigenvalues of the symmetric matrix are

$$\lambda_{\pm} = r \pm t. \quad (3.59)$$

Therefore, CPA occurs when either

$$r + t = 0 \quad \text{or} \quad r - t = 0. \quad (3.60)$$

For an ultrathin absorbing layer whose thickness is much smaller than the optical wavelength, the transmitted field can be expressed as the sum of the incident field and the forward-scattered field from the absorber. This leads to the relation

$$t = 1 + r. \quad (3.61)$$

Substituting this relation into the CPA condition $r + t = 0$ yields

$$t = -r = \frac{1}{2}. \quad (3.62)$$

Under single-side illumination, the intensity absorption coefficient is

$$A = 1 - |r|^2 - |t|^2 = \frac{1}{2}. \quad (3.63)$$

Therefore, the CPA condition corresponds to a single-side absorption of 50% [104]. When two coherent beams with equal amplitude illuminate the absorber from opposite directions with the appropriate phase relation, the outgoing waves cancel at both ports, resulting in zero reflected and transmitted fields and complete absorption of the incoming energy within the structure.

The scattering matrix under the CPA condition becomes:

$$\begin{pmatrix} \mathcal{E}_k^{(\text{out})} \\ \mathcal{E}_{-k}^{(\text{out})} \end{pmatrix} = S_{\text{CPA}} \begin{pmatrix} \mathcal{E}_k^{(\text{in})} \\ \mathcal{E}_{-k}^{(\text{in})} \end{pmatrix}, \quad S_{\text{CPA}} = 0.5 \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} \quad (3.64)$$

where the non-unitary nature of the matrix reflects energy dissipation in the absorber.

If the input waves have equal amplitude \mathcal{E}_0 , the complex amplitudes in equation 3.53 and 3.54 are written as:

$$\mathcal{E}_k^{(\text{in})} = \mathcal{E}_0 e^{j\theta_k} \quad (3.65)$$

$$\mathcal{E}_{-k}^{(\text{in})} = \mathcal{E}_0 e^{j\theta_{-k}} \quad (3.66)$$

The total input field can then be described as :

$$E^{(\text{in})}(z, t) = E_k^{(\text{in})} + E_{-k}^{(\text{in})} \quad (3.67)$$

$$= \mathcal{E}_0 e^{-j\omega t} \left(e^{j(kz+\theta_k)} + e^{-j(kz-\theta_{-k})} \right) \quad (3.68)$$

Define the phase difference between two input waves as:

$$\Delta\theta = \theta_k - \theta_{-k}$$

Then the total input field can be written as a superposition of even and odd standing-wave modes after applying Euler's formula:

$$E^{(\text{in})}(z, t) = \mathcal{E}_0 e^{-j\omega t} \left(e^{j(kz+\theta_k)} + e^{-j(kz-\theta_{-k})} \right) \quad (3.69)$$

$$= 2\mathcal{E}_0 e^{-j\omega t} e^{j\left(\frac{\theta_k+\theta_{-k}}{2}\right)} \left(\cos kz \cos \frac{\Delta\theta}{2} + j \sin kz \sin \frac{\Delta\theta}{2} \right) \quad (3.70)$$

$$= \sqrt{2} \left(\Lambda_{\text{even}}^{(\text{in})} \cos kz + j \Lambda_{\text{odd}}^{(\text{in})} \sin kz \right) e^{-j\omega t} \quad (3.71)$$

$$\Lambda_{\text{odd}}^{(\text{in})} = \sqrt{2}\mathcal{E}_0 \sin\left(\frac{\Delta\theta}{2}\right) e^{j(\frac{\theta_k + \theta_{-k}}{2})} \quad (3.72)$$

$$\Lambda_{\text{even}}^{(\text{in})} = \sqrt{2}\mathcal{E}_0 \cos\left(\frac{\Delta\theta}{2}\right) e^{j(\frac{\theta_k + \theta_{-k}}{2})} \quad (3.73)$$

where $\Lambda_{\text{odd}}^{(\text{in})}$ is the amplitude of the odd (antisymmetric) standing wave mode, and $\Lambda_{\text{even}}^{(\text{in})}$ is the amplitude of the even (symmetric) standing wave mode.

As depicted in Figure 3.4 (a), two counter-propagating input waves $E_k^{(\text{in})}$ and $E_{-k}^{(\text{in})}$, with a phase difference of $\Delta\theta = (2n + 1)\pi$ ($n \in \mathbb{Z}$), interfere to form a purely odd standing wave Λ_{odd} , which is characterized by a node at the absorber location. This results in a minimized electric field at the absorber, leading to negligible interaction and complete transmission into the output ports. In contrast, when the input waves are in phase ($\Delta\theta = 2n\pi$ ($n \in \mathbb{Z}$)), as shown in Figure 3.4 (b), an even standing wave Λ_{even} with an antinode at the absorber is formed. This configuration maximizes the electric field at the absorber, enabling coherent perfect absorption and yielding zero output fields ($E_k^{(\text{out})} = E_{-k}^{(\text{out})} = 0$). Therefore, by strategically controlling the phase of one of the travelling waves, the intensity of the other can be modulated via the subwavelength absorber.

The transformation in equation 3.71 is described by the beam-splitter-like matrix:

$$\begin{pmatrix} \Lambda_{\text{even}}^{(\text{in})} \\ \Lambda_{\text{odd}}^{(\text{in})} \end{pmatrix} = S_{\text{BS}} \begin{pmatrix} \mathcal{E}_k^{(\text{in})} \\ \mathcal{E}_{-k}^{(\text{in})} \end{pmatrix}, \quad S_{\text{BS}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad (3.74)$$

In the same way, the relation between the output field and standing-to-travelling wave modes can be written as following:

$$\begin{pmatrix} \mathcal{E}_k^{(\text{out})} \\ \mathcal{E}_{-k}^{(\text{out})} \end{pmatrix} = S_{\text{BS}} \begin{pmatrix} \Lambda_{\text{even}}^{(\text{out})} \\ \Lambda_{\text{odd}}^{(\text{out})} \end{pmatrix}$$

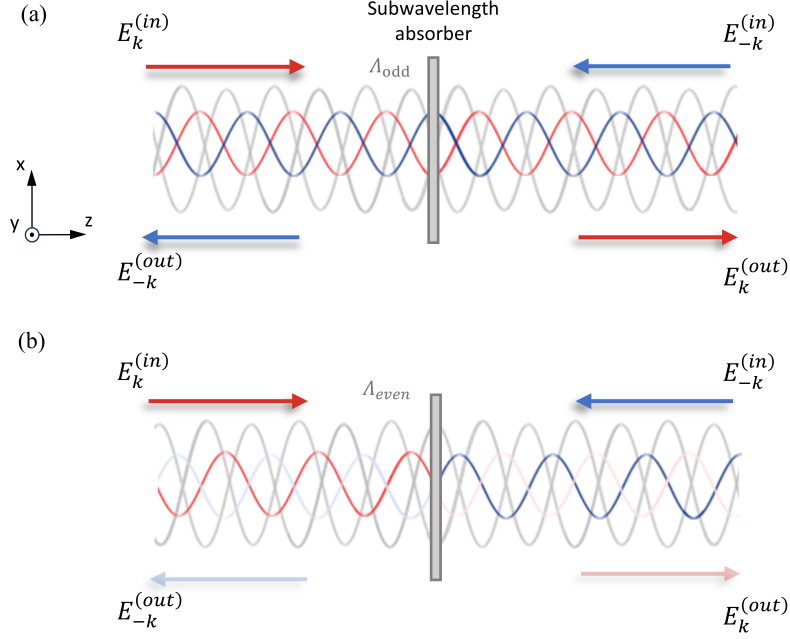


Figure 3.4: Coherent perfect absorption via standing wave decomposition. (a) Two counter-propagating input waves $E_k^{(in)}$ and $E_{-k}^{(in)}$ with a phase difference of $\Delta\theta = \pi$ form a pure odd standing wave Λ_{odd} with a node at the absorber. The electric field is minimized at the absorber location, resulting in no interaction and full transmission into the output ports. (b) When the input waves are in phase, an even standing wave Λ_{even} with an antinode at the absorber is formed. This leads to complete absorption by the subwavelength absorber and zero output fields ($E_k^{(out)} = E_{-k}^{(out)} = 0$). The phase-controlled decomposition into standing wave modes enables selective absorption in integrated photonic systems.

which indicates that while one of the standing-waves is fully extinguished, the survived standing-wave mode is split into two output travelling waves.

This principle has then been investigated in integrated photonic waveguides with subwavelength absorbers loaded in the light propagation direction[105]. The standing wave formed by the photonic mode in a dielectric waveguide with two equal-amplitude counter-propagating waves can be described as:

$$E_{\text{tot}}(z, t) = E_0 e^{-j(\beta z + \omega t)} + E_0 e^{j(\beta z - \omega t + \Delta\theta)} \quad (3.75)$$

$$= 2E_0 \cos\left(\omega t - \frac{\Delta\theta}{2}\right) \cos\left(\beta z - \frac{\Delta\theta}{2}\right) \quad (3.76)$$

where E_0 is the amplitude of the input waves, propagation constant (guided-mode wave number) $\beta = \frac{2\pi n_{\text{eff}}}{\lambda}$, and n_{eff} is the effective refractive index of the waveguide. Assuming $\Delta\theta = 0$, the resulting field exhibits a spatially periodic distribution with fixed nodes and antinodes at any given time, characterized by a spatial period of $\frac{\lambda}{2n_{\text{eff}}}$.

Similarly, by precisely placing subwavelength absorbers onto MRRs, selective coupling to specific resonant optical modes can be achieved under a standing-wave incident field, which allows the implementation of CPA in integrated transmission devices of greater relevance for real-world applications [49]. In an MRR with counter-propagating inputs with the same amplitude, the standing wave intensity follows:

$$I(l) = |E_{\text{tot}}(l)|^2 = \left| E_0 e^{j(\beta l + \omega t + \Delta\theta)} + E_0 e^{j(\omega t - \beta l)} \right|^2 \quad (3.77)$$

$$= 2E_0^2 [1 + \cos(2\beta l + \Delta\theta)] \quad (3.78)$$

According to the resonance condition in 3.18, we get:

$$I(l) = 2E_0^2 \left[1 + \cos \left(\frac{4\pi m l}{L} + \Delta\theta \right) \right] \quad (3.79)$$

where l denotes the position along the ring in the clockwise direction from the out-coupling region, L is the ring perimeter, and m is the mode number (resonance index), corresponding to the number of resonance wavelengths supported within the MRR cavity. These resonant modes fall into two categories of whispering-gallery-mode (WGM) resonances, odd and even modes, which appear alternately at successive resonant wavelengths. When the counter-propagating input waves are fully in phase ($\Delta\theta = 0$), these two groups of resonance modes form out-of-phase

standing waves at midpoint of the MRR ($l = \frac{L}{4}$): odd modes exhibit a node (zero electric field), while even modes exhibit an antinode with a field intensity of $4E_0^2$, oscillating spatially at $l = \frac{L}{4} + \frac{\lambda}{2n_{\text{eff}}}$.

By manipulating the phase difference of two input waves $\Delta\theta$ or repositioning the absorber along the propagation direction, selective absorption of individual resonant modes can be achieved[49, 105]. The correlation between the spatial location of standing wave and the wavelength of counter-propagating input waves ($l = \frac{L}{4} + \frac{\lambda}{2n_{\text{eff}}}$) enables strategic control over the interaction between the optical field and integrated absorber through wavelength selection.

Previous approach exploited on-chip CPA with embedded, reconfigurable nanostrips to achieve individual manipulation of different wavelengths[50] within a single device. However, each nanostrip interacts only with a single antinode of a selected standing-wave mode, resulting in limited modulation contrast and posing challenges for advanced computing applications. In this thesis, we investigate a photon-driven, non-localized thermo-optic effect among an array of wavelength-selective absorptive nanostrips to enable cross-wavelength optical end-to-end encoding. The underlying mechanisms and implementation strategies are detailed in the following section.

Methodology: Simulation, Fabrication, and Experimental Setup

Contents

4.1 Lumerical Simulation	70
4.1.1 MODE Solutions (FDE Solver)	71
4.1.2 FDTD Solutions	71
4.1.3 HEAT Solutions	73
4.1.4 Simulation Workflow	74
4.2 Device Fabrication	75
4.2.1 General Process and Layer Overview	75
4.2.2 Alignment Markers	77
4.2.3 Photonic Layer	77
4.2.4 Antenna Layer	78
4.3 Experiment Configurations	79
4.3.1 Sample Stage	79
4.3.2 Setup for Phase-shift Measurement	80
4.3.3 Setup for Time Integration of Ultrafast Optical Signals	82

This chapter presents the simulation tools, device fabrication processes and experimental setups used for device characterization .

4.1 Lumerical Simulation

This section summarizes the numerical simulation tools employed to design and analyse the Photonic-Heater-in-Lightpath (PHIL) devices. All simulations were performed using the *Lumerical Photonic Design Suite* (Ansys, Inc.), which provides

complementary solvers for optical and thermal modelling. The optical domain was studied using the MODE and FDTD solvers, while the thermal response was analysed using the HEAT solver. Together, these modules provide a self-consistent understanding of light-matter interaction, thermal dissipation, and device reconfigurability.

4.1.1 MODE Solutions (FDE Solver)

The *Finite-Difference Eigenmode* (FDE) solver in *Lumerical* MODE was used to calculate the effective refractive index, mode field distribution, and confinement factors of the guided modes in silicon-on-insulator (SOI) waveguides, as indicated in Figure 3.1(c).

The simulations defined a cross-sectional geometry corresponding to a half-etched SOI waveguide, consisting of a 220 nm silicon device layer with a 110 nm etch depth on top of a 3 μm buried oxide (BOX) layer. This rib configuration ensures single-mode operation at 1550 nm while maintaining sufficient mode confinement for efficient coupling to the nanoheaters. Material dispersions were taken from the built-in Lumerical database. The eigenmode analysis identified the fundamental TE_0 mode, which was used to determine the effective index (n_{eff}), group index, and mode overlap with the nanoheaters. These parameters were subsequently imported into FDTD simulations to construct realistic three-dimensional models of the optical fields.

4.1.2 FDTD Solutions

The *Finite-Difference Time-Domain* (FDTD) solver was employed to simulate optical interference and absorption within the PHIL-loaded waveguide and microring resonator (Figure 4.1).




Figure 4.1: FDTD simulation layout of the PHIL microring device. Optical excitation is launched from the input waveguide (blue arrows), and the transmitted field is monitored at the output. Perfectly matched layer (PML) boundaries (yellow box) surround the simulation region to absorb outgoing radiation and prevent artificial reflections

The model incorporated titanium nanoheaters positioned at the field antinodes to achieve wavelength-selective absorption. Mesh refinement was applied adaptively, with manual overrides in regions around the nanoheaters, particularly within the evanescent field region of the microring, to ensure higher spatial resolution and more realistic modelling of light-matter interaction. Perfectly Matched Layer boundary conditions were applied to all outer simulation boundaries to prevent artificial reflections.

Excitation was provided by an eigenmode source launched in the bus waveguide to selectively excite the fundamental TE mode. Frequency-domain power monitors were placed both at drop waveguide and nanoheater region to record the transmission spectra and local field distributions. An embedded analysis group was used to quantify and export the spatially resolved absorbed optical power within the nanoheater region. The resulting absorbed power density maps were subsequently

imported into the *HEAT* solver as localized heat sources.

4.1.3 HEAT Solutions

The *Lumerical* HEAT solver was employed to simulate both the steady-state and transient thermal responses of the PHIL device under optical excitation. The optical absorption data obtained from the FDTD simulations were subsequently imported as localized heat sources, directly coupling between optical absorption and thermal dissipation ensuring consistency between optical absorption and subsequent heat generation. The geometry of the simulation domain and the corresponding material stack used in the HEAT solver are illustrated in Figure 4.2

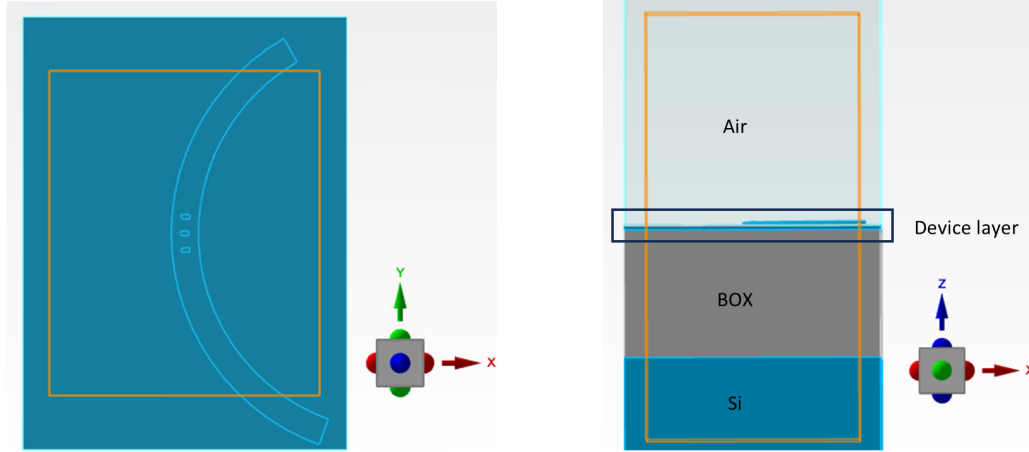


Figure 4.2: Layout of the thermal (HEAT) simulation. The left panel shows the top view of the simulated device region, while the right panel illustrates the vertical stack consisting of the silicon handle layer, buried oxide (BOX), device layer, and air cladding.

The fixed-temperature thermal boundary condition was applied at the bottom surface of the substrate to maintain a constant temperature of 300 K. The top

of the simulation region was filled with air, where convective boundary conditions were defined at the material interfaces to model natural heat exchange with the environment. The thermal properties of the materials were defined according to standard values in the Lumerical database, with silicon thermal conductivity set to $148 \text{ W}\cdot\text{m}^{-1}\cdot\text{K}^{-1}$ and silica to $1.4 \text{ W}\cdot\text{m}^{-1}\cdot\text{K}^{-1}$.

The solver then computed the three-dimensional temperature distribution within the waveguide and surrounding cladding. Both steady-state temperature maps and time-dependent thermal decay curves were extracted.

4.1.4 Simulation Workflow

In this work, the three solvers were used in a sequentially coupled workflow:

1. **MODE (FDE):** Established the guided optical modes and effective refractive indices.
2. **FDTD:** Modelled the full-field optical interactions and generated absorbed power distributions within the nanoheaters.
3. **HEAT:** Translated these absorbed power maps into temperature rise and decay profiles, revealing the temporal response of the device.

This hierarchical simulation framework provided a comprehensive, self-consistent picture of optical-thermal coupling in the PHIL system, bridging the electromagnetic, thermal, and functional domains that underpin the device operation. These simulations lay the theoretical foundation for the subsequent design, fabrication, and experimental validation of the Photonic-Heater-in-Lightpath architecture.

4.2 Device Fabrication

Device fabrication begins with diced $10 \times 10 \text{ mm}^2$ SOI chips. Each chip consists of a 220-nm-thick silicon device layer on top of a $3\text{-}\mu\text{m}$ buried oxide (BOX) layer.

4.2.1 General Process and Layer Overview

Three key layers are structured in each device (Figure 4.3): alignment markers, photonic elements (waveguides and grating couplers), and metallic antennas. Layouts are generated using Python with the `gdsshelplers` package and exported in GDS-II format for electron beam lithography (EBL).

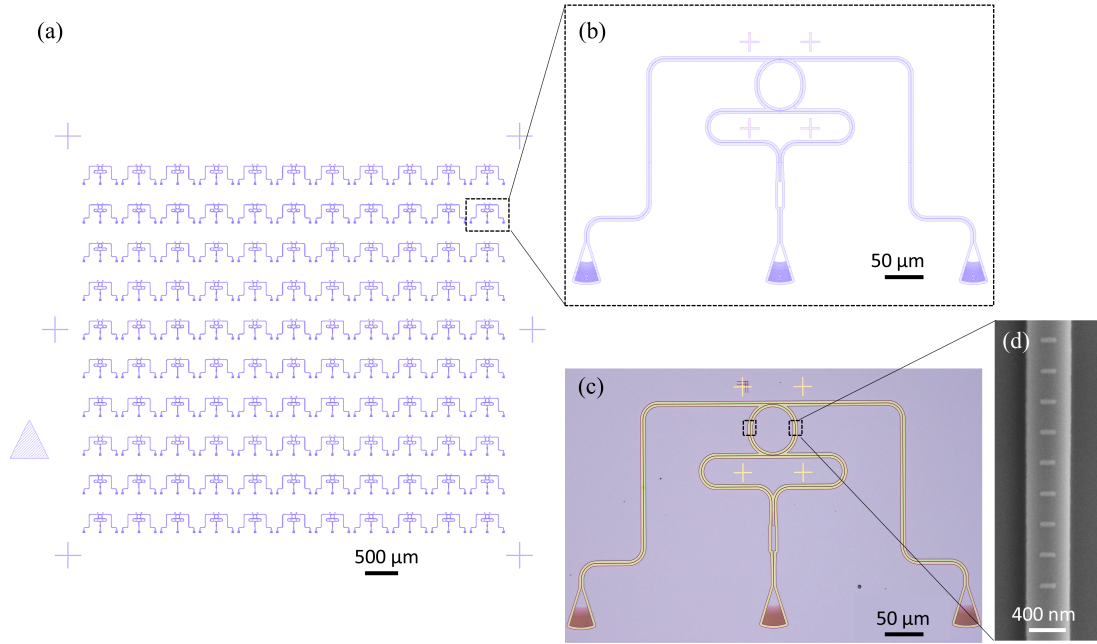


Figure 4.3: Device layout (a) Overview of the full chip layout designed for fabrication. (b) Layout for a single device. (c) Optical microscope image of a single device. (d) Scanning electron microscope (SEM) image of the antenna-loaded region.

The general fabrication process consists of the following steps:

1. **Surface Cleaning and Preparation:** Chips are sequentially cleaned in acetone and isopropanol (IPA) for 5 minutes each. Ultrasonic agitation is

applied during the initial cleaning step to remove residues from the wafer dicing process but is omitted in subsequent layers to avoid damaging previously fabricated structures. After drying under a nitrogen gas flow, the chips are treated with low-power oxygen plasma to enhance photoresist adhesion.

2. **Photoresist Spin-Coating:** A CSAR 65 positive electron-beam resist is used in this thesis for most patterning layers. The resist is spin-coated at 4000 rpm for 45 seconds, followed by a soft bake at 150 °C for 3 minutes to remove residual moisture and solvents.
3. **Electron-Beam Lithography (EBL):** Patterning is performed using a JEOL JBX5500 system operated at 50 kV. A beam current of 100 pA is used, and the exposure dose for CSAR 65 is set to 180 $\mu\text{C}/\text{cm}^2$.
4. **Development:** After exposure, the resist is developed in AR 600-546 for 30 seconds, followed by sequential rinses in methyl isobutyl ketone (MIBK) and IPA for 15 seconds each. The chip is then dried using a nitrogen gas flow.
5. **Etching or Material Deposition:** Photonic structures are defined using reactive ion etching (RIE) with a gas mixture of CHF_3/SF_6 . structures are etched for 75 seconds at a depth of 110 nm. Metal layers are deposited by either thermal evaporation (for gold) or electron-beam evaporation (for titanium), depending on the application.
6. **Lift-Off:** For metal patterning, lift-off is carried out by immersing the sample in heated MicropositTM Remover 1165 until the undesired metal is fully dissolved. The process is followed by rinsing in acetone and IPA, and completed with nitrogen gas drying.

7. **Inspection and Quality Control:** Patterned features are inspected under an optical microscope or scanning electron microscope (SEM) to verify critical dimensions, alignment accuracy, and structural integrity before proceeding to the next fabrication step.

The design and fabrication considerations for each layer are discussed in detail in the following subsections.

4.2.2 Alignment Markers

Two types of markers are used in this thesis to enable precise matching of layers: global markers (Figure 4.3 (a)) for chip-level alignment and local markers (Figure 4.3 (b)) for device-level precision. During chip-level alignment, two global markers positioned on opposite sides of the chip are used to calibrate its location within the EBL coordinate system and to correct for any rotational misalignment. Subsequently, device-level alignment is performed by detecting a local marker specific to each device. These local markers are strategically placed within the same writing field ($100 \times 100 \mu\text{m}^2$) as the critical device features (waveguides and antennas) to eliminate stage movement between marker detection and pattern exposure, thereby ensuring precise layer-to-layer alignment.

The fabrication process for alignment markers is illustrated in Figure 4.4(a). The markers are defined using EBL, followed by the thermal evaporation of a 5 nm Cr / 75 nm Au metal stack, and completed with an overnight lift-off in heated MicropositTM Remover 1165.

4.2.3 Photonic Layer

The fabrication flow of the photonic layer is shown in Figure 4.4(b). Photonic structures are defined using EBL and transferred to the silicon layer via RIE.

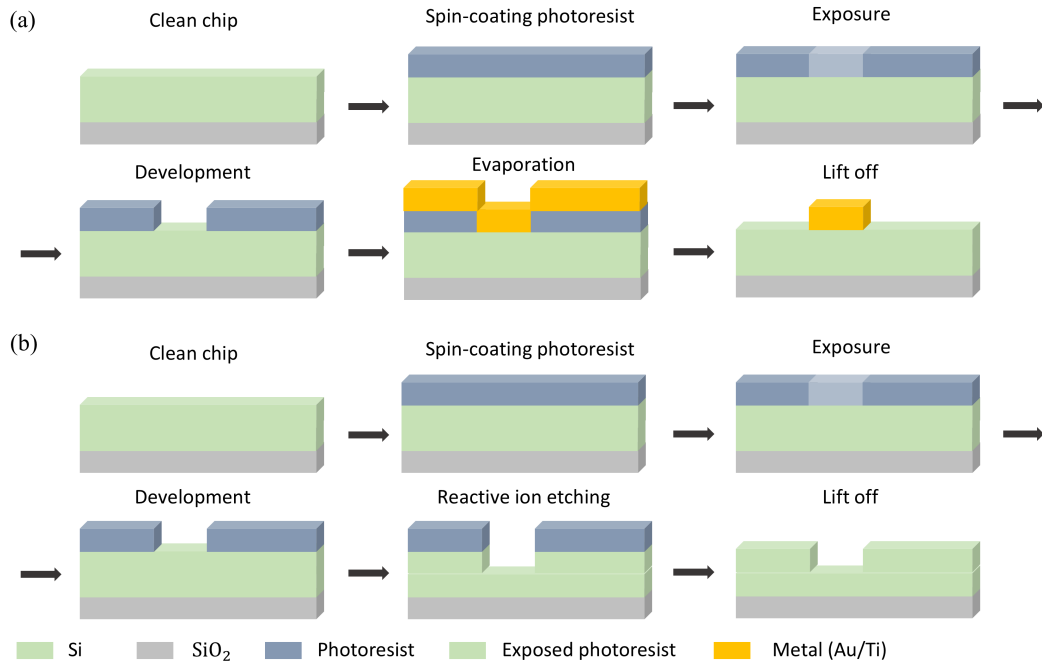


Figure 4.4: Fabrication process flow (a) Fabrication process for gold alignment markers or titanium nanoantenna structures using metal evaporation and lift-off. (b) Fabrication process for silicon waveguides using RIE following e-beam lithography patterning.

An etch depth of 110 nm is achieved using a CHF₃/SF₆ gas mixture, forming waveguides, grating couplers, and ring resonators, as shown in Figure 4.3 (c). Following etching, the remaining photoresist is removed by immersing the sample in heated MicropositTM Remover 1165 for approximately 30 minutes. The chip is then rinsed with acetone and IPA, and dried using a nitrogen gas flow.

4.2.4 Antenna Layer

In this thesis, arrays of subwavelength antennas are integrated on both sides of the ring resonator to enable wavelength-selective absorption of optical signals and function as nanoheaters that induce non-local thermo-optic effects. Titanium is selected as the antenna material due to its strong optical absorption and high melting point.

A third EBL step is performed to define the antenna pattern, following the same

process as the previous run. After development, a 20-nm titanium (Ti) layer is deposited via electron-beam evaporation to form the antennas. The unwanted Ti is then removed by immersing the sample in heated MicropositTM Remover 1165 overnight. The chip is subsequently rinsed with acetone and IPA, and dried with a nitrogen gas flow. The resulting antenna pattern is shown in Figure 4.3 (d).

It is worth noting that the EBL auto-focus system was malfunctioning for the majority of this DPhil project, necessitating a manual focusing procedure. Following the standard calibration routine, the electron beam's focus and astigmatism were manually adjusted on the peripheral region of the sample. Final optimization was achieved using contamination dots: by directing the electron beam slightly away from the device area for approximately 10 seconds, a bright, circular contamination dot was formed, indicating optimal calibration conditions.

4.3 Experiment Configurations

This section describes the optical measurement setup, including the sample stage and customized experimental configurations developed specifically for this thesis.

4.3.1 Sample Stage

Figure 4.5 presents photos of the measurement stage for optical experiments in this thesis. Optical signals are delivered to the device through grating couplers using multi-channel fibre arrays. Two positioning stages enable precise alignment of the sample and fibre arrays along the X, Y, and Z (vertical) axes. To correct for rotational misalignment, rotation stages are integrated with the translation stages. The alignment process is assisted by visual feedback from a lateral camera and an overhead microscope system, consisting of a Thorlabs CCD camera and a




Figure 4.5: Photos of the sample stage (a) with a zoomed photo (b) of the sample and fibre array

Navitar zoom lens. This platform ensures stable fibre–chip coupling and provides a consistent baseline for all device characterisation.

4.3.2 Setup for Phase-shift Measurement

Having established the physical measurement environment, the next step is to introduce the optical configuration used throughout this thesis. All experiments, regardless of whether they involve cross-wavelength encoding, multi-channel addition, or time-domain integration, are built upon a common two-line measurement architecture. At its core is the phase-shift measurement, which serves as the primary diagnostic used to quantify the thermo-optic response of the device.

The phase-shift measurement is performed using the setup shown in Figure 4.6. A tunable laser (TSL-550, Santec) serves as the probe source to measure the device spectrum within the telecommunications band (1500–1600 nm). The probe signal is injected into the device through the 10% port of a 90/10 fibre coupler,

with its polarization adjusted using a polarization controller. A circulator is included to isolate and protect instrumentation by blocking back-reflections from the chip.




Figure 4.6: Device characterization setup.

A second laser source (N7711A, Keysight) is used as the control signal to induce an all-optical spectral shift. This pump signal, operated in continuous-wave (CW) mode, is injected into the device through the 90% port of another 90/10 coupler, with polarization likewise controlled. The laser output power is varied from 0 to 34 mW, corresponding to 0–1.4 mW after coupling into the chip, to gradually increase the local temperature near the nanoantennas and induce a thermo-optic spectral shift.

Device transmission spectra under different control powers are collected at the output following an optical tunable filter (OTF-350, Santec) with a 2-nm bandwidth to suppress the control signal.

The phase-shift configuration establishes the fundamental two-line measurement

architecture used throughout this thesis. All other experiments are implemented by extending this baseline arrangement with additional optical components tailored to the specific measurement requirement.

4.3.3 Setup for Time Integration of Ultrafast Optical Signals

This high-speed measurement build directly on the same probe-control framework but incorporate additional components capable of generating and shaping picosecond pulse streams. This experiment was performed using the experimental setup shown in Figure 7.8. A tunable laser source (TSL-550, Santec) generated the probe signal with an output power of 25 mW. The probe was injected into the device through the 10% port of a 90/10 coupler, with its polarization adjusted using a polarization controller.

A second laser source (AQ4321A, ANDO-TSL) provided the 4 mW control signal. The control signal was first polarization-adjusted and then modulated by a Mach-Zehnder modulator (MZM, 40 GHz bandwidth), driven by an arbitrary waveform generator (AWG, Keysight 8194A; 120 GSa/s, 45 GHz bandwidth) to produce picosecond optical pulses. The electrical drive from the AWG was amplified by an RF amplifier (SHF L806A, 40 GHz bandwidth), delivering a 5 V_{pp} signal to the MZM.

The modulated optical output was amplified by an erbium-doped fiber amplifier (EDFA), and amplified spontaneous emission (ASE) noise was suppressed using a 1-nm optical bandpass filter (OBPF). This pulse train was subsequently passed through a second MZM (10 GHz bandwidth) to generate 20 ns temporal envelopes, driven by a low-speed AWG (ARB RIDER AWG-4081; 1.2 GSa/s, 300 MHz bandwidth) providing a 5.5 V_{pp} electrical signal. After modulation, the optical

All-Optical Cross-Wavelength Encoding

Contents

5.1 Motivation	84
5.2 Photonic-Heater-in-Lightpath (PHIL): Design and Simulations	86
5.2.1 Realisation of Lossy and Lossless Spectral Bands	86
5.2.2 Resonant-Mode Selective Optical Absorption	89
5.2.3 Optically Controlled Nanoheaters	93
5.3 Cross-Wavelength Optical Modulation: Experimental Validation	95
5.3.1 PHIL-Loaded Microring Resonator	95
5.3.2 Resonant-Mode Selective Spectral Response	97
5.3.3 Thermo-Optic Cross-Wavelength Signal Encoding	99
5.4 Insertion Loss Analysis	104
5.4.1 Experimental Characterization	105
5.4.2 Nanoheater-Induced Insertion Loss	106
5.4.3 Tunable Absorption Profile	110
5.5 Discussion and Summary	111

5.1 Motivation

As discussed in Section 2.2, WDM increases information-processing capacity by transmitting data in parallel on multiple wavelengths. However, data encoding still depends on electronic peripherals, and inter-wavelength data transfer is typically accomplished via repeated electro-optical (E-O) conversions.

In reconfigurable photonics, achieving both low optical loss and reduced electronic

overhead remains challenging. Carrier-based electro-optic modulators provide large bandwidths for ultrafast weight updates and transfer-function implementation, but they introduce non-negligible insertion loss through the free-carrier plasma-dispersion effect and require interfacing with costly high-speed digital-to-analogue and analogue-to-digital converters (DACs/ADCs). Thermo-optic phase shifters with the merit of low insertion loss and straightforward fabrication[119, 120], are widely used in silicon photonics to actively reconfigure the phase of propagating light[22, 35, 79], while their reconfigurability is still governed by electronics. The reliance on repeated electronic implementations contributes to increased device footprint, greater system complexity, and higher energy consumption. PCMs, such as $\text{Ge}_2\text{Sb}_2\text{Te}_5$, enable all-optical reconfiguration with nanosecond-scale laser pulses[21, 31], but suffer substantial optical loss due to the large extinction coefficient κ in the crystalline phase, leading to indiscriminate attenuation of all propagating optical modes and hindering cascadability.

These limitations highlight the need for alternative approaches that maintain low optical loss while enabling direct optical-domain encoding and inter-wavelength data transfer. In the following section, I will introduce a strategy that addresses these challenges through an all-optically controlled thermo-optic phase shifter based on nanoscale metallic heaters. In this scheme (Figure 5.1), the optical power of incoming signals (control signal) is absorbed and the encoded information is transferred onto a spectrally distinct probe signal via the thermo-optic effect with minimum optical loss.




Figure 5.1: Concept of all-optical cross-wavelength encoding. The control optical power (control signal, CS) is converted to heat and shifts the ring spectrum via the thermo-optic effect. The CS power is then encoded in the amplitude of a probe wavelength (probe signal, PS).

5.2 Photonic-Heater-in-Lightpath (PHIL): Design and Simulations

To implement an all-optical reconfigurable unit with minimal insertion loss, photonic integrated circuits are engineered to exhibit distinct lossy and near-lossless spectral bands. The former is employed for modulation (control signals, CS) and the latter supports low-loss readout and facilitates cascability (probe signals, PS).

5.2.1 Realisation of Lossy and Lossless Spectral Bands

The underlying mechanism hinges on the creation of a stationary interference pattern produced by two counter-propagating coherent optical waves (i.e. CPA). The field distribution of such a system is described by a periodically varying electric field in space with spatially fixed nodes (locations of minimum electric field) and antinodes (locations of maximum electric field) at any arbitrary point in time (t). The spatial periodicity of these is derived by the superposition of two counter-

propagating waves in space (z) and time (t) as:

$$\begin{aligned} E(z, t) &= E_0 \cos\left(\frac{2\pi n_{\text{eff}} z}{\lambda} - \omega t + \Delta\varphi\right) + E_0 \cos\left(\frac{2\pi n_{\text{eff}} z}{\lambda} + \omega t\right) \\ &= 2E_0 \cos\left(\frac{2\pi n_{\text{eff}} z}{\lambda} + \frac{\Delta\varphi}{2}\right) \cos\left(\omega t - \frac{\Delta\varphi}{2}\right) \end{aligned} \quad (5.1)$$

where E_0 is the amplitude of the input waves and $\Delta\varphi$ is the phase difference between two coherent inputs.

Assuming $\Delta\varphi = 0$, positions along the propagating direction that satisfy even multiples of a quarter wavelength

$$z = 2n \cdot \frac{\lambda}{4n_{\text{eff}}}, \quad n = (\dots, -3, -2, -1, 0, 1, 2, 3, \dots) \quad (5.2)$$

form *antinodes* (field maxima), while odd multiples of a quarter wavelength

$$z = (2n + 1) \cdot \frac{\lambda}{4n_{\text{eff}}}, \quad n = (\dots, -3, -2, -1, 0, 1, 2, 3, \dots) \quad (5.3)$$

form *nodes* (field minima).

For different wavelengths, the spatial period of the antinodes $P = \frac{\lambda}{2n_{\text{eff}}}$ varies accordingly. This effect is illustrated using finite-difference time-domain simulations (FDTD, *Lumerical*) in Figure 5.2. This simulation considers counter-propagating light in a silicon rib waveguide at distinct wavelengths. Absorptive titanium nanoheaters, with dimensions of 20 nm in thickness, 50 nm in width (along the propagation direction), and 150 nm in length (perpendicular to the propagation direction), were placed on the waveguide at a position approximately $6\mu\text{m}$ from the waveguide centre along the propagation axis. The nanoheaters were arranged with a spacing at 297 nm, corresponding to the antinode period at a wavelength of 1550 nm. The resulting spatial E-field distribution reveals the formation of periodic nodes and antinodes under coherent counter-propagating excitation.

Crucially for this work, wavelengths λ_1 (1525.6 nm) and λ_2 (1550.3 nm) can be selected here such that their standing-wave patterns are nearly out of phase at a




Figure 5.2: Input lights with equal intensities are sent into silicon waveguide from opposite directions. Two absorptive nanoheaters are designed to spatially match the anti-nodes of light λ_1 (1525.6 nm) (top) while overlapping with the nodes of light λ_2 (1550.3 nm) (bottom).

specific location, with the antinodes of one wavelength coinciding with the nodes of the other and vice versa. As shown in Figure 5.3, at the nanoheater position the field intensity exhibits a maximum for λ_1 while simultaneously reaching a minimum for λ_2 .

This wavelength selectivity is reflected in the absorption spectrum, (Figure 5.4(a)), where I swept the injected wavelengths from 1500 nm to 1600 nm, and calculated the total power absorbed by the nanoheaters by integrating over the height of the structure. The total absorption experienced by λ_1 is 17 times higher than that of λ_2 , where the structure is nearly transparent. The combined effects of the standing-wave field distribution and the absorptive nanoheaters thus establish alternating lossy and lossless regions in the wavelength domain.

Figure 5.4(b) further illustrates the spatial absorption profile under a 1-mW optical input. At λ_1 , nearly all absorption occurs within the nanoheaters, leading to strong attenuation, whereas at λ_2 the waveguide remains largely unperturbed, with residual absorption limited by the finite nanoheater dimensions and their spatial placement.

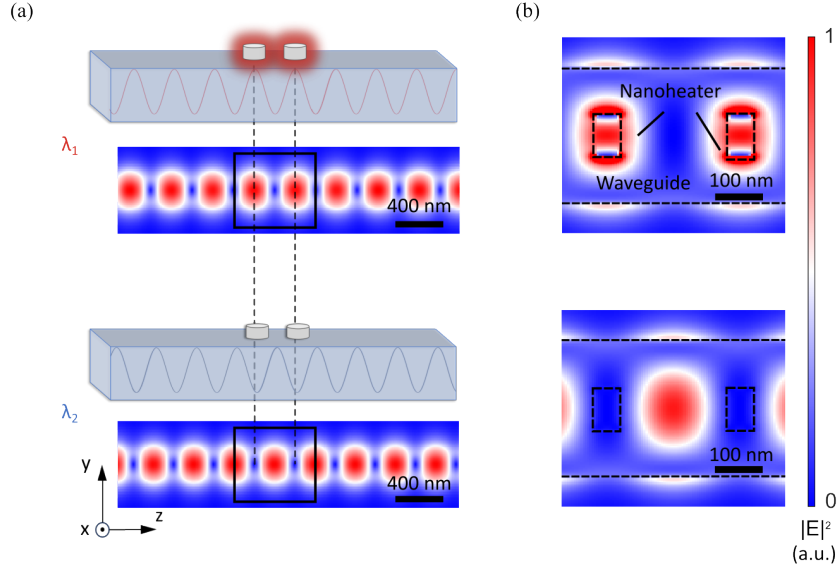


Figure 5.3: Zoomed-in view of the nanoheater region. (a) Two absorptive nanoheaters are designed to spatially match the anti-nodes of light λ_1 (1525.6 nm) while overlapping with the nodes of light λ_2 (1550.3 nm). (b) Normalized electric field intensity ($|E|^2$) distribution at the nanoheater locations for the two wavelengths. The nanoheaters have dimensions of 20 nm (thickness), 50 nm (width), and 150 nm (length), and are placed on a 500 nm-wide silicon rib waveguide.

5.2.2 Resonant-Mode Selective Optical Absorption

This concept can then be integrated on ring resonators in order to discretize the coupled wavelengths and exploit the oscillatory nature of standing waves inside the MRR to enhance wavelength selectivity, enabling further experimental realization.

To couple light into the ring, the perimeter (L) of the ring must be an integer multiple of the input light wavelength:

$$L = \frac{m\lambda}{n_{\text{eff}}}, \quad (5.4)$$

where m denotes the mode number, corresponding to the number of resonance




Figure 5.4: Wavelength-selective optical power absorption. (a) Total optical power absorbed by the nanoheaters under excitation with different wavelengths. The absorption reaches a maximum at $\lambda_1 = 1525.6$ nm (red star) and a minimum at $\lambda_2 = 1550.3$ nm (blue star), corresponding to the alignment of the nanoheaters with an antinode and a node of the standing wave, respectively. (b) Optical power absorbed per unit area by the nanoheaters under a 1 mW optical input at λ_1 and λ_2 .

wavelengths supported within the MRR cavity. These modes primarily fall into two categories of whispering-gallery-mode (WGM) resonances: odd and even modes, which appear alternately at successive resonant wavelengths. And two adjacent resonant wavelengths of the MRR can then be written as

$$\lambda_1 = \frac{L \cdot n_{\text{eff}}}{2n}, \quad \lambda_2 = \frac{L \cdot n_{\text{eff}}}{2n + 1}, \quad (5.5)$$

where λ_1 is an even-mode resonant wavelength, and λ_2 is the following odd-mode resonant wavelength.

Numerical simulations of the electric intensity profiles at the resonant modes of an MRR under counter-propagating excitations were performed using the FDTD method (*Lumerical*). As shown in Figure 5.5(a), the results present standing-wave

patterns with enhanced optical confinement, exhibiting periodic oscillations along the ring. The spatial periodicity can be expressed as

$$x = l \cdot \frac{\lambda}{2n_{\text{eff}}}, \quad (5.6)$$

where l denotes the position index along the ring, ranging from 0 to m , with m the mode number. The phase difference $\Delta\varphi$ between the counter-propagating wave is set at 0 in the simulation.

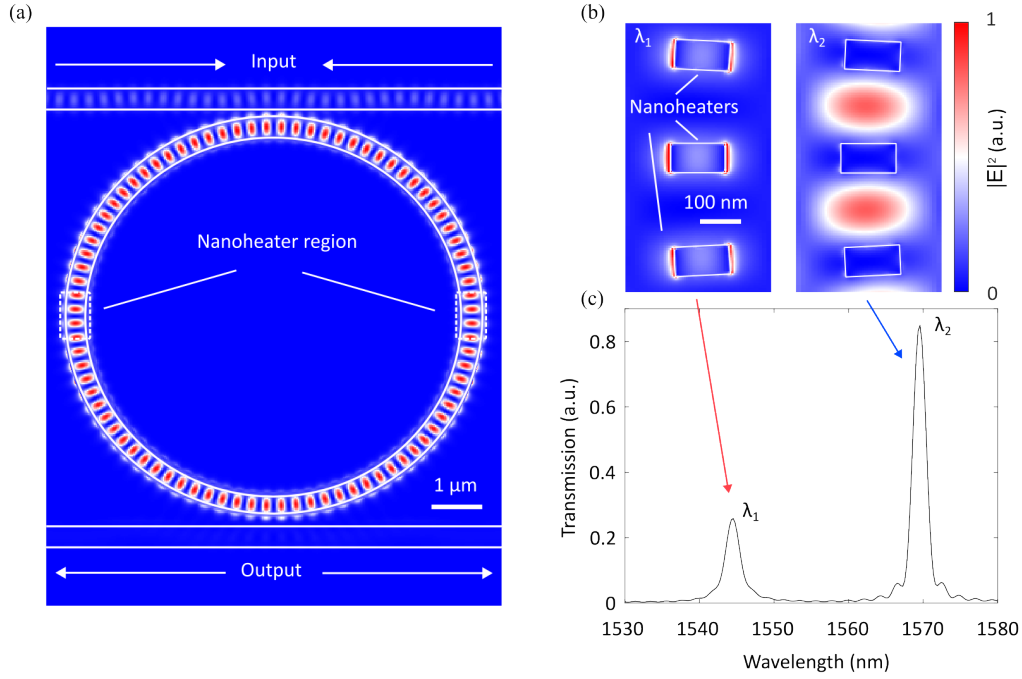


Figure 5.5: Simulation of PHIL-loaded MRR. (a) Simulated field intensity distribution in an MRR showing the nanoheater region where absorption occurs. (b) Zoomed-in electric field intensity ($|E|^2$) profiles at the nanoheater positions for two adjacent resonant wavelengths, λ_1 and λ_2 . At λ_1 , the nanoheater overlap with field maxima, whereas at λ_2 , the nanoheaters coincide with field minima. (c) Transmission spectrum of the MRR drop port, showing strong attenuation at λ_1 and a nearly transparent response at λ_2 .

For a specific location, $x_{\text{out-of-phase}}$, where two adjacent resonances are nearly out of phase, the spatial shift between their antinodes and nodes should be approximately half of the spatial period. This condition gives

$$l_1 \cdot \frac{\lambda_1}{2n_{\text{eff}}} = (l_2 + \frac{1}{2}) \cdot \frac{\lambda_2}{2n_{\text{eff}}}. \quad (5.7)$$

Combining with equation 5.5 and 5.6 yields the corresponding index and position:

$$l_1 = l_2 = l_{\text{out-of-phase}} = n, \quad x_{\text{out-of-phase}} = \frac{L}{4}. \quad (5.8)$$

Thus, these two groups of resonance modes will form out-of-phase standing-wave patterns at the midpoint of the MRR ($L/4$), where the antinodes of one group overlap with the nodes of the other. Consequently, placing absorptive nanoheaters periodically around $L/4$, as depicted in Figure 5.5(b), enables selective coupling of one resonant mode to the nanoheaters, while the adjacent mode is spatially bypassed. This effect is evident in the transmission spectrum at the drop port (Figure 5.5(c)), where the even mode (λ_1) exhibits significantly reduced transmission and a broadened resonant peak, indicating strong optical loss inside the ring, whereas the odd mode (λ_2) propagates with minimal perturbation.

Figure 5.6(a) visualizes the spatial absorption profile, where the optical power of λ_1 is predominantly absorbed by the nanoheaters through coupling at the field antinodes.

In contrast, the adjacent resonant wavelength remains nearly unaffected, aside from minor overlap at the nanoheater edges. This behaviour results in the wavelength-selective absorption spectrum shown in the Figure 5.6(b), where the alternating absorption-transparency response between odd and even modes highlights the intrinsic wavelength selectivity of counter-propagating excited MRR.

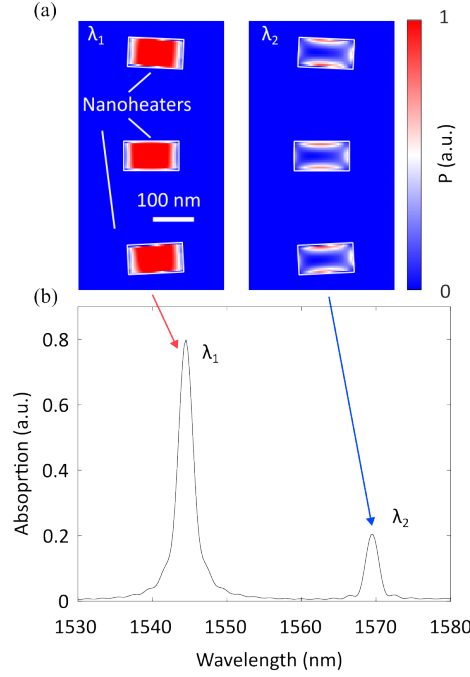


Figure 5.6: Wavelength-selective absorption in PHIL-loaded MRR. (a) Normalized optical power absorbed per unit area at the nanoheater region under a 1-mW optical input at λ_1 and λ_2 . (b) Absorption spectrum highlights the strong absorption of λ_1 and the minimum optical loss at λ_2 .

5.2.3 Optically Controlled Nanoheaters

The temperature profile (Figure 5.7) resulting from optical absorption was simulated using the HEAT solver (*Lumerical*), with the spatial absorption profile at the nanoheater region that obtained from the FDTD solver loaded as the heat source. In both simulations and following experiments, the buried oxide (BOX) layer thickness was set to $2\ \mu\text{m}$, while the cladding material was modelled as air. The temperature at the bottom of the thermal simulation region (in the Si substrate) was fixed at the ambient value of 300 K, and a convection boundary condition of $10\ \text{W}/\text{m}^2 \cdot \text{K}$ was applied between the device layer and the surrounding air. The HEAT solver discretises and solves the heat transport equation in three dimensions on a finite-element mesh extending beyond the heat source region.

For a 1 mW optical excitation with a temporal width of $5\ \mu\text{s}$ launched into the

PHIL-loaded MRR at λ_1 , heat generated by optical absorption leaks to the entirety of the waveguide (Figure 5.7), raising local temperature between nanoheaters to approximately 307 K. In contrast, excitation at the adjacent resonance λ_2 induces negligible heating, with the temperature remaining close to room level.

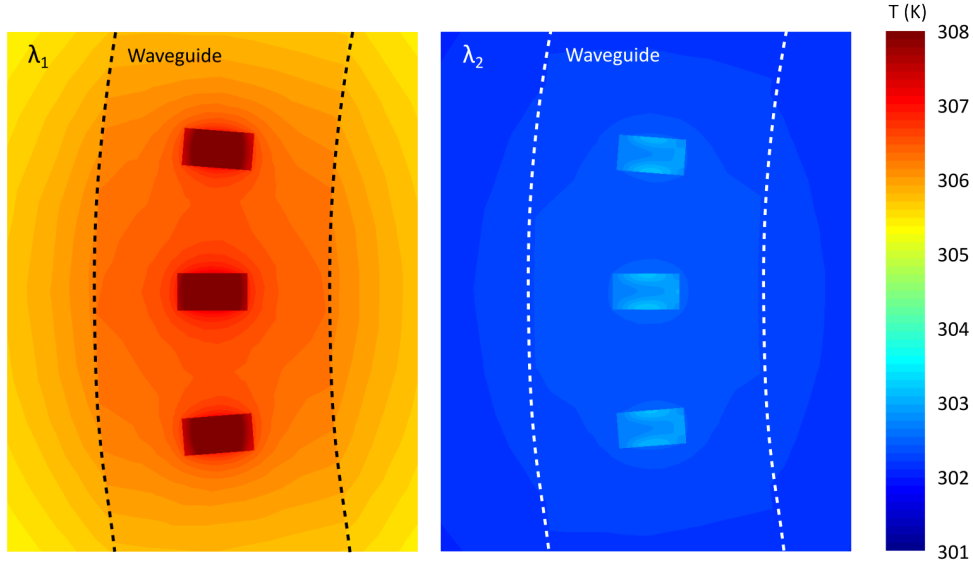


Figure 5.7: Temperature profile of the nanoheater-loaded region. Heat generated from optical absorption raises the local temperature (T) of the surrounding waveguide, simulated using *Lumerical* HEAT.

The observed temperature distribution highlights the spatial non-locality of the heating process, arising from the high thermal conductivity of silicon ($148 \text{ W/m} \cdot \text{K}$). This not only demonstrates that strong absorption is confined to resonant modes with proper spatial alignment, but also suggests a potential route for transferring information from λ_1 to λ_2 through temperature-dependent phase shifts in the heated waveguide, enabled by the relatively large thermo-optic coefficient of silicon ($1.8 \times 10^{-4} \text{ K}^{-1}$ at 1550 nm). Moreover, the negligible heating observed for λ_2 reconfirms that wavelengths not overlapping with the nanoheaters experience minimal optical loss, validating both the low-loss transmission pathway and the wavelength-selective nature of the proposed mechanism.

5.3 Cross-Wavelength Optical Modulation: Experimental Validation

5.3.1 PHIL-Loaded Microring Resonator

In the previous section, two operating wavelengths, λ_1 and λ_2 were introduced, which experience different absorptive behaviours due to their spatial overlap with the nanoheaters. For clarity in the following discussion, λ_1 is defined as the *control signal* and λ_2 as the *probe signal*. The *control signal* is preferentially absorbed by the nanoheaters to generate local heating, while the *probe signal* is used to probe the resulting thermo-optic modulation with minimal loss.

The proposed PHIL design was experimentally implemented on a racetrack MRR with a radius of 30 μm . As shown in Figure 5.8, input light was coupled on-chip through the central grating coupler from a fibre array and directed into a 1:2 multimode interferometer (MMI) splitter. The MMI divides the signal into two equal-intensity paths, which are coupled into the MRR from opposite directions. The clockwise and counter-clockwise travelling waves produce a standing-wave pattern within the ring with varied spatial periodicity of $\frac{\lambda}{2n_{\text{eff}}}$. *Control signals* (CS, red arrows) are expected to be selectively absorbed at the antinodes by the embedded nanoheaters, while *probe signals* (PS, blue arrows) simultaneously experience thermo-optic modulation induced by the absorbed CS. After circulating within the MRR, optical signals were outcoupled through the drop waveguides, where the transmission response was collected to characterize the device performance.

Figure 5.9(a) presents a false-coloured scanning electron microscope (SEM) image of the PHIL-loaded MRR, where nanoheaters are positioned at $x = \pm \frac{L}{4}$ (corresponding to the middle left and right sections of the ring). Each section contains an array of nine nanoheaters, as shown in Figure 5.9(b). These titanium nanoheaters, chosen for their strong optical absorption and high melting point, have dimensions




Figure 5.8: Optical micrograph of the fabricated PHIL-loaded device. A 1:2 multimode interferometer (MMI) splitter divides the input into two paths with nearly equal intensity, launching counter-propagating waves into the MRR. The superposition of clockwise and counter-clockwise modes forms a standing-wave pattern inside the ring. Control signals (CS, red arrows) are selectively absorbed by the nanoheaters embedded along the MRR, while probe signals (PS, blue arrows) are simultaneously coupled into the resonator to experience the thermo-optic modulation induced by the absorbed control signals.

of 70 nm in width, 250 nm in length, and 20 nm in height. They are placed with a centroidal pitch of 308 nm, corresponding to $\frac{\lambda}{2n_{\text{eff}}}$, ensuring spatial alignment with the optical antinodes of the even-mode resonant wavelengths while coinciding with the nodes of the odd-mode wavelengths.

The experimentally measured transmission spectrum of the device (Figure 5.9(c)) shows typical periodic transmission peaks of an add-drop ring resonator, yet with the creation of alternating low- and high-quality factor resonances corresponding to the engineered absorption of the even-mode transmission. Specifically, the even-mode resonance exhibits an extra attenuation of around 3dB compared with the odd-mode. These observations are in good qualitative agreement with the simulated absorption spectrum (Figure 5.5).

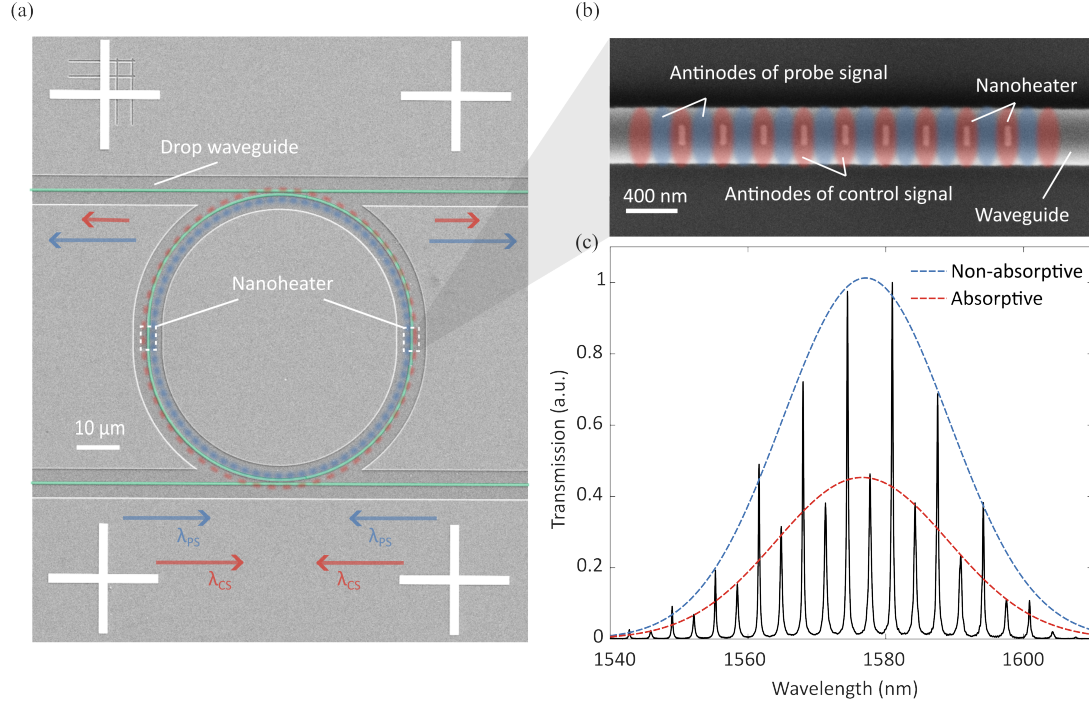


Figure 5.9: PHIL-loaded MRR and transmission spectrum. (a) False-coloured SEM image of a PHIL-loaded MRR, with arrays of nine nanoheaters placed on each side of the ring. (b) Zoom-in false-coloured SEM of the fabricated PHIL region, showing titanium nanoheaters patterned to overlap with the antinodes of the control signal (CS) while coinciding with the nodes of the probe signal (PS). The heaters are deposited with a pitch of 308 nm, enabling selective optical absorption of CS without significantly perturbing PS. (c) Experimentally measured transmission spectrum of the device. Periodic resonances exhibit distinct absorptive and non-absorptive behaviour. The blue curve corresponds to a Gaussian fit of the reference spectrum with minimal absorption, while the red curve includes the additional loss induced by the PHIL nanoheaters.

5.3.2 Resonant-Mode Selective Spectral Response

The ability of the absorptive modes to induce spectrum shifts within the MRR and thereby modulate other co-propagating optical signals was then experimentally evaluated. With the experiment depicted in the section 4.3.2, optical power was injected at the peak wavelength of an absorptive resonance (CS), and the corresponding device spectra under varying CS power were recorded, as shown in Figure 5.10(a). A progressive redshift in the resonance spectrum was observed as the CS power increased from 0 to 1.4 mW.

Figure 5.10 (b) quantifies the spectral shift of the MRR as a function of CS input power, revealing a linear dependence with a sensitivity of $0.04 \pi/\text{mW}$ when the CS is applied at an absorptive resonance. This behaviour confirms that heat generated through CS absorption diffuses into the guiding structure, producing a broadband thermo-optic shift. Consequently, this result indicates that the amplitude information of the CS can be mapped onto a spectrally distinct probe signal (PS) through the thermo-optic effect.

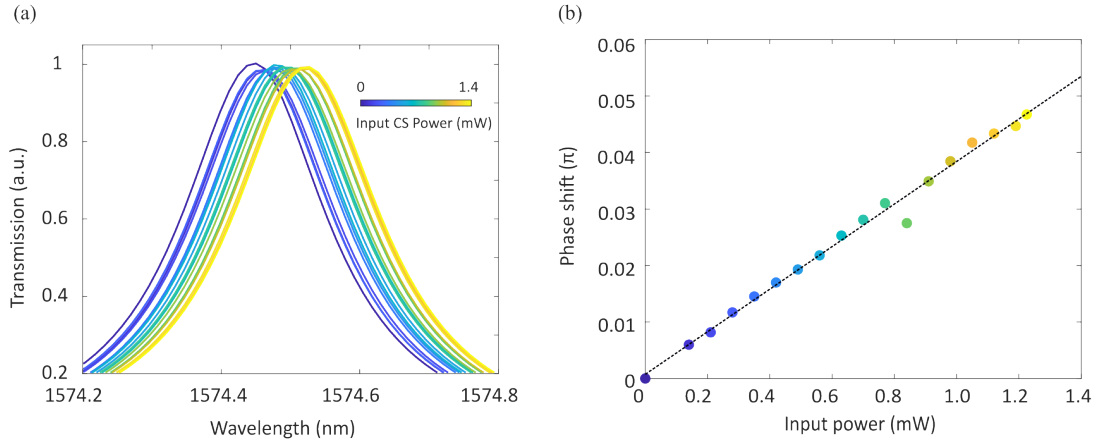


Figure 5.10: Spectral response under absorptive-mode excitation. (a) Transmission spectrum measured at a non-absorptive resonance while increasing the input power coupled into the absorptive mode. (b) Extracted resonance shift as a function of control signal power, showing a nearly linear phase shift with increasing input power, with an efficiency of approximately $0.04 \pi/\text{mW}$.

When the probe and control wavelengths are swapped (placing the CS at a non-absorptive resonance), the spectral shift drops significantly to $0.0087 \pi/\text{mW}$ (Figure 5.11). These results support the non-reciprocal operational principle of the device: the spectral response is only modulated by the absorptive modes, while the non-absorptive modes of the device retain low optical loss. Collectively, the resonant-mode selectivity of the mechanism was confirmed experimentally, where specific resonances are strongly attenuated while adjacent modes propagate with little perturbation.

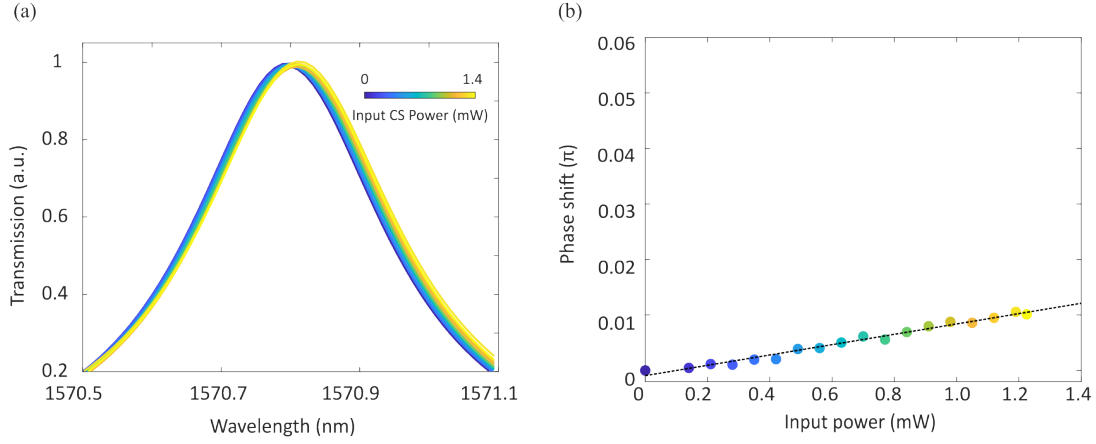


Figure 5.11: Spectral response under non-absorptive-mode excitation. (a) Transmission spectrum measured at a absorptive resonance while increasing the input power coupled into the non-absorptive mode. (b) Extracted resonance shift as a function of control signal power, showing only a negligible phase response compared to the absorptive-mode case, with a reduced slope of $0.0087 \pi/\text{mW}$.

5.3.3 Thermo-Optic Cross-Wavelength Signal Encoding

Then I proceeded to demonstrate all-optical cross-wavelength modulation with the experiment setup depicted in Figure 5.12, to show that the amplitude information of the CS can be mapped onto a spectrally distinct PS through the thermo-optic effect.

A tunable laser (TSL-550, Santec) was used as the probe source, and the probe signal was injected into the device in continuous-wave mode via the 10% port of a 90/10 combiner. The probe wavelength was selected at the right half-maximum of a non-absorptive resonance peak to monitor the device response through transmission changes induced by the total optical power during phase modulation.

A second tunable laser (TSL-570, Santec) was used as the control source to induce all-optical spectral shifts. Its optical power was first amplified using an erbium-doped fiber amplifier (EDFA), and amplified spontaneous emission (ASE) noise was suppressed using an optical tunable filter (OTF-320, Santec). The amplified

control signal was injected into the device in continuous-wave mode via the 99% port of a 99/1 splitter, followed by the 90% port of a 90/10 combiner. Throughout the experiment, the control signal power was monitored through the 1% port of the 99/1 splitter.

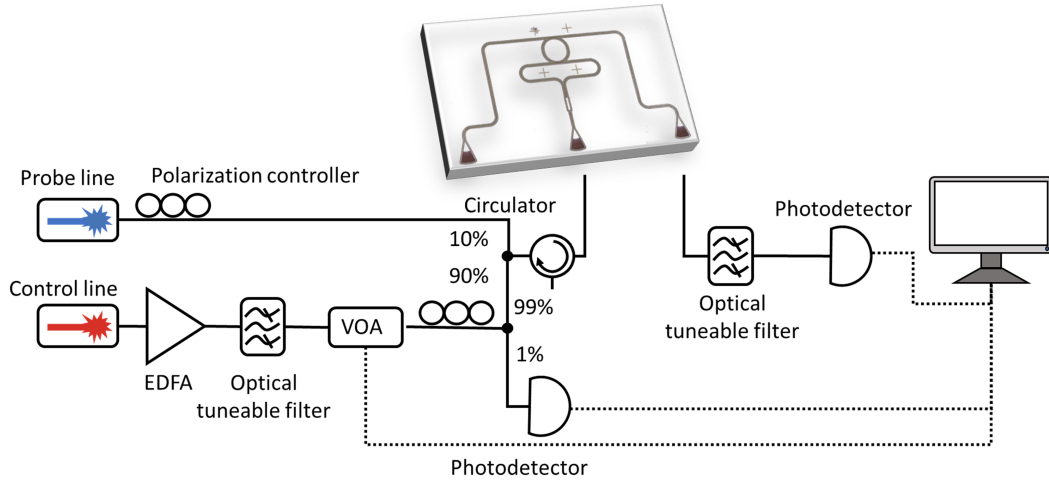


Figure 5.12: Control signal characterization setup

The device output was filtered by another optical tuneable filter (OTF-320, Santec) to isolate the probe signal. Only the probe-line transmission after device was measured at the photodetector.

As shown in Figure 5.13(a), the probe laser was biased on the slope of the PS resonance while a separate control laser located at the peak of a nearby CS resonance. Absorption of CS power induced thermo-optic tuning of the cavity, shifting the resonance and modulating the detected PS transmission. For each CS power sweep, the raw transmission T was recorded and the fractional change was computed as:

$$\Delta T (\%) = 100 \times \frac{T - T_0}{T_0}, \quad (5.9)$$

where T_0 is the transmission at the minimum CS power in that sweep (baseline). To mitigate nonuniform sampling and quantify uncertainty, the data are grouped into 0.1 mW power bins, and only bins with $n \geq 15$ samples are retained. For each bin, the mean $\overline{\Delta T}$ and standard deviation $\sigma_{\Delta T}$ are recorded and plotted as a solid curve with a $\pm 1\sigma$ shaded band in Figure 5.13 (b).




Figure 5.13: Cross-wavelength optical modulation. (a) the PS and the CS resonances. The PS readout wavelength is fixed on the PS resonance shoulder, while the CS wavelength is chosen on the peak of resonance. (b) Static transfer characteristic: PS fractional transmission change $\Delta T = (T - T_0)/T_0$ (in %) versus input CS power. For each dataset, I baseline at the minimum CS power T_0 , compute ΔT , and bin in power with width 0.1 mW. Only bins with $n \geq 15$ samples are kept. Then the per-bin mean (solid line) and the $\pm 1\sigma$ band (shaded) are plotted, with five representative error bars shown for clarity and labelled by their corresponding standard deviations.

To avoid visual clutter, only five representative error bars are displayed, evenly spaced across the CS power range, with each labelled by its corresponding $\pm\sigma$ value. The resulting transfer curve is monotonic, showing a stable one-to-one mapping from CS input to PS output. A quasi-linear trend is observed, arising from the combination of the linear thermo-optic response and the PS bias on the resonance shoulder, where the Lorentzian shape is approximately linear. This drives the nearly linear resonance shift across the near-linear portion of the Lorentzian line-shape. The uncertainty band remains narrow compared to the overall excursion, indicating good repeatability of the static modulation.

Since the whole spectrum of the MRR is shifted during the optical modulation, the absorption profile changes accordingly at CS. To further examine the wavelength dependence of this absorption-induced response, in the following demonstration, the probe laser was still biased on the slope of the non-absorptive resonance, while the control laser was no longer fixed at the resonance peak but instead tuned across the absorptive mode, as marked in Figure 5.14 (a). A zoomed-in view of the absorptive mode in Figure 5.14 (b) shows the set of selected CS wavelengths across the Lorentzian resonance, colour-coded from blue (near-resonance peak) to magenta (off-resonance detuning).

By varying the CS wavelength along the resonance while extending the input CS power range to 3.8 mW, I quantified how the detuning from the absorption maximum influences the strength and efficiency of the induced modulation. Figure 5.14 (c) plots the fractional PS transmission change versus CS input power, with plot colour indicating the corresponding CS wavelengths. Each transmission response at different CS wavelengths was processed by binning the power axis (bin width 0.1 mW) and averaging the PS transmission across samples in each bin. The shaded band shows ± 1 standard deviation to indicate modulation repeatability. Notably, the closer CS detuning to the resonance peak, a steeper slope was obtained at PS transmission change with low input CS power (inset figure), reflecting higher modulation efficiency.

To extract phase shifts induced by different CS wavelengths from the transmission change and decouple the impact of PS selection, a wavelength sweep of the device was firstly recorded, where the FSR was calculated, and then a Lorentzian profile was fitted to the target PS resonance. The normalized Lorentzian is then inverted separately on its rising and falling edges using shape-preserving cubic interpolation, enabling robust mapping from transmission to wavelength across monotonic branches. For each modulation trace, the measured PS transmission is converted

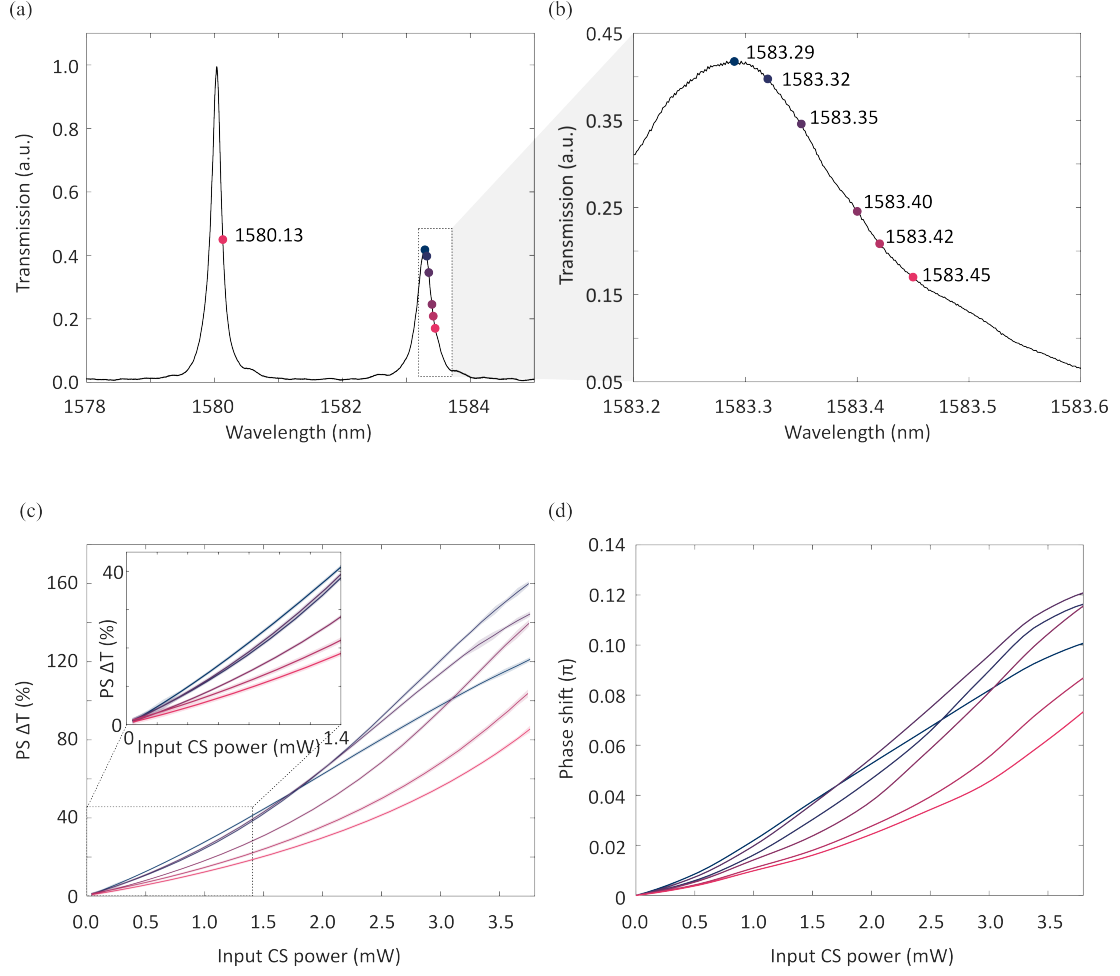


Figure 5.14: Phase shift induced by varying CS wavelengths across the absorptive mode. (a) Transmission spectrum of the MRR highlighting the absorptive resonance, with coloured markers indicate the selected CS wavelengths. (b) Zoomed-in view of the absorptive resonance showing the different CS positions. (c) Measured PS transmission change as a function of CS input power for different CS wavelengths across the absorptive mode. At low input CS power (inset), fixing at the peak of absorptive resonance provides the largest modulation depth. (d) Retrieved phase shift as a function of CS input power for different CS wavelengths. The results show wavelength-dependent sensitivity, confirming that the thermo-optic modulation strength varies with the detuning of the CS from the resonance peak. Colours are consistent across panels.

to its corresponding wavelength shift, and the relative phase shift is computed as:

$$\phi_k = 2\pi \times \frac{\lambda_{\text{ref}} - \lambda_k}{\text{FSR}}.$$

Repeats at the same CS power are averaged. The results, summarized in Figure

5.14(d), closely track the amplitude trends: Detunings closer to the absorptive peak produce larger phase shifts at low power, while further-detuned CS inputs show reduced sensitivity.

Moreover, as the input power increases, cumulative heating induces a progressive redshift of the entire resonance. This causes the near-peak CS wavelengths to reach maximum absorption earlier, after which the absorption decreases, evidenced by the gradual flattening of the modulation plots. Specifically, beyond approximately 1.5 mW, the modulation depth of the on-resonance CS is overtaken by that of slightly detuned CS wavelengths. Conversely, CS wavelengths initially positioned further from the peak (e.g., magenta curves) demonstrate increasing modulation efficiency at higher powers, as the redshift brings them closer to peak absorption.

To maximize modulation efficiency in the low-power regime, subsequent experiments in this thesis will fix the CS wavelength at the centre of the absorptive resonance.

5.4 Insertion Loss Analysis

Unlike conventional add-drop microrings, where nearly all the out-coupled signal is directed into a single drop port, the PHIL design symmetrically couples counter-propagating waves into two output ports. In the current experimental setup, only one port is collected, which limits the maximum detectable output to approximately 50% of the total signal, corresponding to a baseline ~ 3 dB coupling penalty unrelated to nanoheater absorption.

It should be noted that this apparent penalty arises from the measurement configuration rather than a fundamental limitation of the device. In a neuromorphic

photonic architecture, both output ports can be routed to subsequent neurons, allowing the optical power to be fully utilized through network interconnections. In this context, the symmetric power splitting can in fact be advantageous, as it naturally enables signal fan-out to multiple downstream nodes without requiring additional optical splitting components. Therefore, while the present characterization exhibits a ~ 3 dB detection loss, the symmetric output configuration does not represent an intrinsic efficiency loss at the system level and may even facilitate more efficient network-level connectivity.

5.4.1 Experimental Characterization

To experimentally quantify the insertion loss associated with the nanoheater array, the transmission characteristics of the device are measured and compared with a reference structure.

In order to decouple geometry-induced losses, which may be mitigated through optimized output collection schemes, from absorption arising specifically from the nanoheater array, a reference device with identical geometry but no nanoheaters is employed in the subsequent analysis.

As discussed above, the insertion loss (IL) of PHIL devices varies with the operating wavelength (non-absorptive vs. absorptive). To evaluate the IL of the reported PHIL configuration (9×2 nanoheaters), I measured several PHIL devices with identical design parameters and compared their performance against a co-fabricated nanoheater-free reference. This reference shares the same MMI coupler, waveguide routing, and symmetric out-coupling architecture, enabling isolation of the nanoheater-induced absorption. As shown in the Figure 5.15, the average IL at non-absorptive wavelengths is ~ 1.7 dB, while absorptive resonances exhibit stronger attenuation around 6 dB.




Figure 5.15: Measured insertion loss of PHIL with referenced to same design without nanoheater.

5.4.2 Nanoheater-Induced Insertion Loss

The IL of the PHIL design also depends on the number of integrated nanoheaters. To further explore this dependence and to quantify the contribution of each heater element, a series of 3D FDTD simulations were performed in *Lumerical* on a compact PHIL structure with a ring radius of $R = 4 \mu\text{m}$. The number of nanoheaters N was swept from 1×2 to 11×2 , and the absorption profile was evaluated at both the non-absorptive (λ_{PS}) and absorptive (λ_{CS}) wavelengths (Figure 5.16).

The transmission spectra (top panels) are normalized to that of a nanoheater-free reference device to isolate loss induced by the nanoheater array. As expected, the resonance peak at the absorptive wavelength (λ_{CS}) decreases and widens progressively with more heaters, indicating increased attenuation. In contrast, the transmission at the non-absorptive wavelength (λ_{PS}) remains largely unchanged, indicating minimal modal perturbation at those positions.

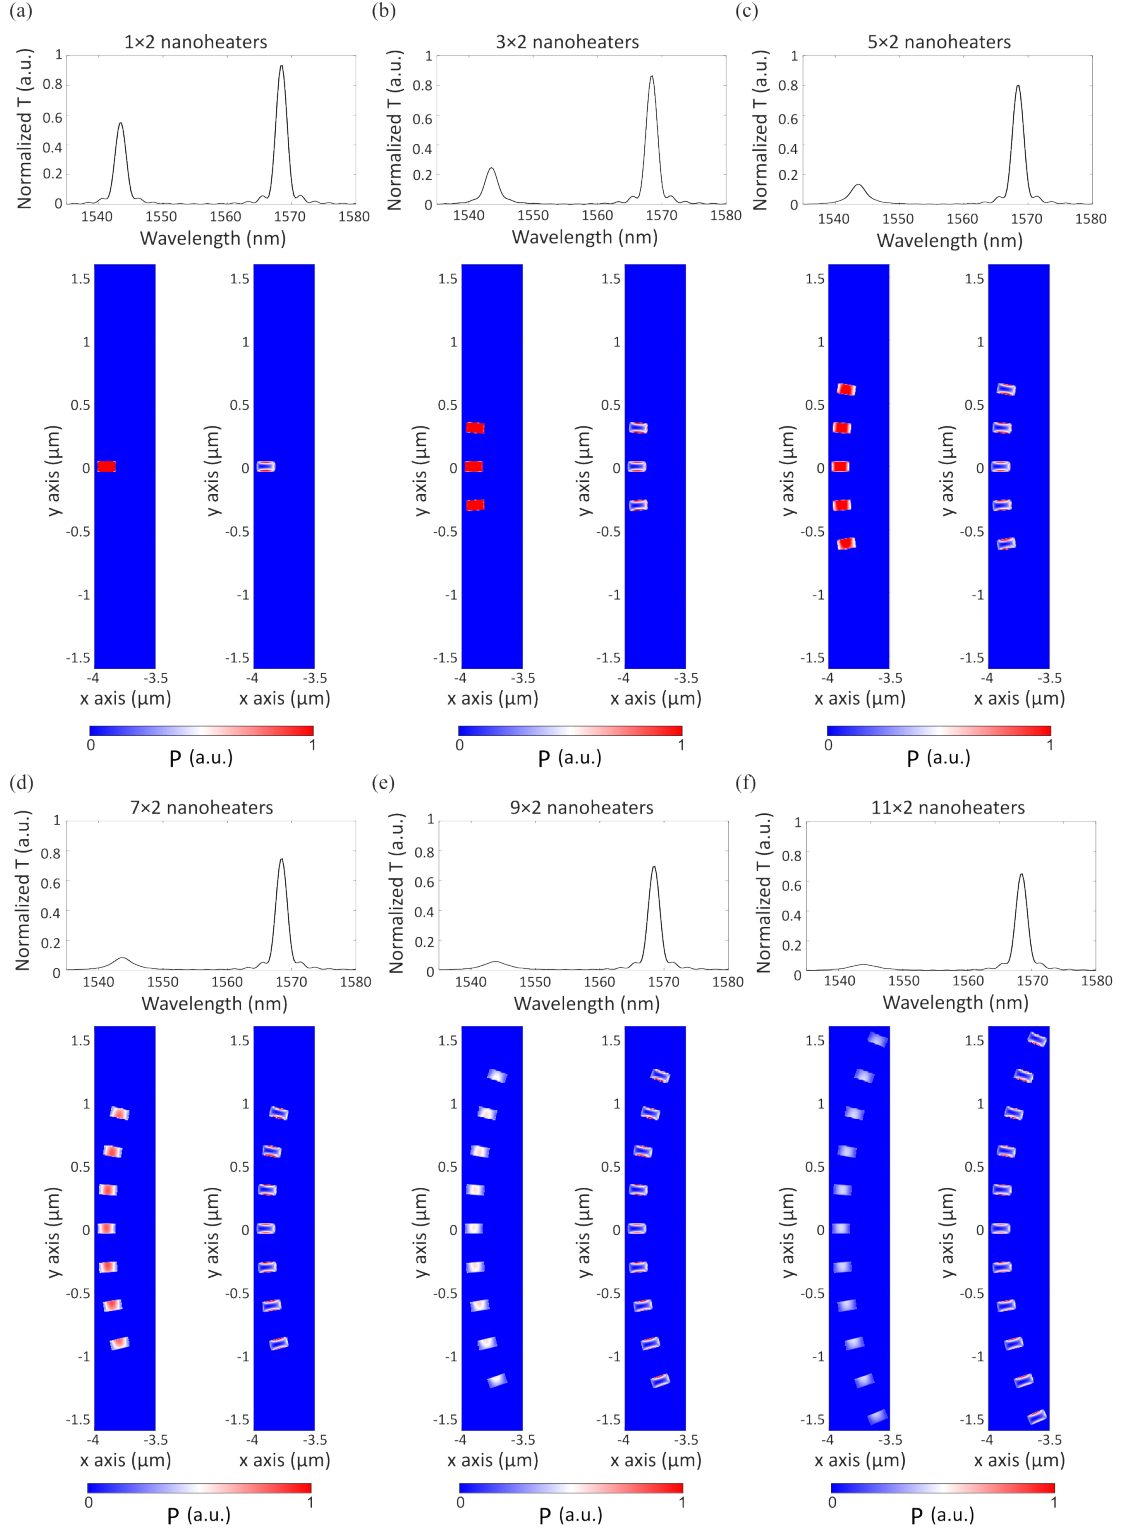


Figure 5.16: Simulated transmission spectra and absorption profiles in a compact PHIL architecture with varying nanoheater numbers. (a)–(f) correspond to PHIL structures with 1×2 , 3×2 , 5×2 , 7×2 , 9×2 and 11×2 nanoheaters, respectively. The top subpanel shows the normalized transmission spectrum, scaled to the transmission of a nanoheater-free reference design. The bottom subpanels display the normalized optical power absorbed per unit area in the nanoheater regions at the absorptive wavelength (λ_{CS} , left) and the non-absorptive wavelength (λ_{PS} , right), extracted from FDTD simulations under a 1 mW optical input.

The bottom subpanels show the normalized optical power absorbed per unit area in the nanoheater regions, extracted at both λ_{CS} and λ_{PS} . A key observation is that although the total absorption increases with heater number, the *absorbed power per heater* decreases, suggesting a global renormalization of the standing-wave field.

This effect can be explained by considering the power balance in a lossy resonator[121]. In the steady-state regime, the circulating field amplitude adjusts in response to increased dissipative loss, while maintaining the resonance condition. The stored energy U in the cavity is related to the input power P_{in} and total loss rate γ by

$$P_{\text{in}} = \gamma U = \gamma \cdot \frac{|A|^2}{\omega_0} \quad \Rightarrow \quad |A|^2 \propto \frac{P_{\text{in}} \cdot \omega_0}{\gamma} \quad (5.10)$$

where $|A|^2$ denotes the intracavity field intensity and ω_0 is the resonance frequency. As more nanoheaters are introduced, γ increases due to additional absorption, leading to a reduction in the field amplitude $|A|^2$. Because all heaters are positioned at antinodes of the standing wave, this amplitude scaling leads to uniformly lower absorption across the entire array.

To quantify the impact of nanoheater number on device performance, the IL at both λ_{CS} and λ_{PS} was extracted from the simulated transmission spectra and plotted as a function of heater count, as shown in Figure 5.17(a). At λ_{CS} , the IL increases rapidly with the number of nanoheaters, reaching approximately 12.3 dB at $N = 11 \times 2$. However, the scaling is sub-linear, with a visibly reduced slope beyond $N = 7 \times 2$. This reflects the saturation mechanism discussed earlier: as more absorbers are introduced, the standing-wave amplitude decreases due to increased loss, reducing the absorption efficiency per heater despite continued total attenuation growth.

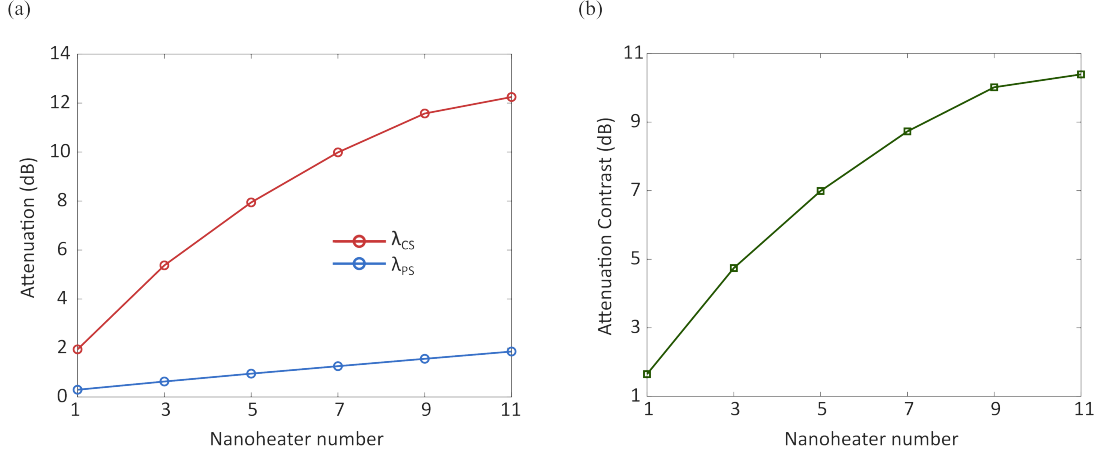


Figure 5.17: Simulated insertion loss and attenuation contrast for varying nanoheater numbers. (a) Insertion loss at the absorptive (λ_{CS}) and non-absorptive (λ_{PS}) wavelengths extracted from FDTD simulations of compact PHIL structures with increasing heater count. (b) Corresponding attenuation contrast $\Delta IL = IL_{\lambda_{CS}} - IL_{\lambda_{PS}}$, indicating saturation of modulation depth beyond $N = 9 \times 2$.

In contrast, the IL at λ_{PS} increases approximately linearly with heater number, reaching ~ 2.0 dB at $N = 11 \times 2$. This corresponds to an average parasitic loss of ~ 0.09 dB per nanoheater. Since the field intensity is minimized at the heater sites at λ_{PS} , the loss arises primarily from passive scattering and finite-volume perturbations, and remains relatively insensitive to the number or alignment of heaters.

The resulting *attenuation contrast*, defined as $\Delta IL = IL_{\lambda_{CS}} - IL_{\lambda_{PS}}$, is plotted in Figure 5.17(b). The contrast increases with nanoheater number but begins to plateau beyond $N = 9 \times 2$, saturating around 10.5 dB. This reflects the diminishing marginal absorption gain per additional heater, while parasitic loss continues to accumulate linearly at λ_{PS} .

Specifically, for the reported nanoheater configuration (9×2), the simulated absolute IL at λ_{CS} (~ 11.5 dB) exceeds the experimentally measured value (~ 6 dB).

This discrepancy can be attributed to three main factors: (i) Field localization with ring radius, where a smaller radius confines stronger optical intensities, resulting in greater absorption by identical nanoheaters; (ii) fabrication-induced variations, including refractive-index nonuniformities across the wafer and deviations in nanoheater periodicity, which modify the effective standing-wave period and lead to imperfect matching between the designed nanoheater spacing and the standing-wave antinode distribution; and (iii) residual phase differences between counter-propagating modes, introduced after the MMI splitting, shift the alignment between antinodes and nanoheaters, thereby reducing the effective absorption efficiency and also increasing insertion loss at λ_{PS} .

Taken together, this comparative analysis between experiment and simulation not only confirms the cascadability of the PHIL architecture, but also defines a practical design envelope for PHIL-based absorptive modulators. Heater arrays in the range of 7×2 to 9×2 offer an effective trade-off between achieving high attenuation contrast and maintaining low baseline loss. Beyond this range, performance improvements become marginal, while IL at the non-absorptive wavelength continues to accumulate.

5.4.3 Tunable Absorption Profile

To mitigate the phase differences between counter-propagating mode and compensate for fabrication-induced variations, the spectral response of the PHIL device can be reconfigured by actively tuning the relative optical phase difference between the two counter-propagating inputs. As illustrated in Figure 5.18(a), incorporating a thermo- or electro-optic phase shifter in one arm allows direct modulation of the interference condition within the MRR, thereby shifting the spatial distribution of the standing-wave pattern. As derived in equation 3.79, this modulation is

governed by the phase delay $\Delta\Phi$ between the input arms, which dynamically reconfigures the spatial overlap between the standing-wave antinodes and the absorptive nanoheaters embedded in the ring.

In the experimental demonstration, a thermo-optic phase shifter was integrated after the 1:2 MMI splitter and before coupling into the MRR. Nichrome (NiCr), a commercially established alloy composed of 80% nickel and 20% chromium, was chosen as the heater material for its high electrical resistivity ($\sim 10^{-6} \Omega \cdot \text{m}$), which ensures efficient Joule heating, and its chemical and thermal stability, attributed to the self-passivating chromium oxide layer that prevents further oxidation under elevated temperatures. The phase shifter is driven by an external electrical bias applied via gold contact pads and accessed using RF probes.

As $\Delta\Phi$ is varied by tuning the input voltage, the transmission spectrum of the MRR evolves correspondingly (Figure 5.18(b)). The odd-order resonances transition from transparent to fully absorptive, while the even-order resonances exhibit the inverse behaviour. Concurrently, the attenuation contrast between absorptive and non-absorptive wavelengths can be optimized (Figure 5.18(c)). This phase-controlled tuning mechanism provides fine-grained control over the device's spectral characteristics, adding reconfigurability in the phase domain alongside the wavelength domain. Moreover, it can be exploited to enhance tolerance to fabrication-induced asymmetries and introduces an additional degree of functional flexibility essential for scalable photonic computing architectures.

5.5 Discussion and Summary

This chapter presented the design, implementation, and characterization of an all-optical cross-wavelength modulation framework enabled by the PHIL architecture.

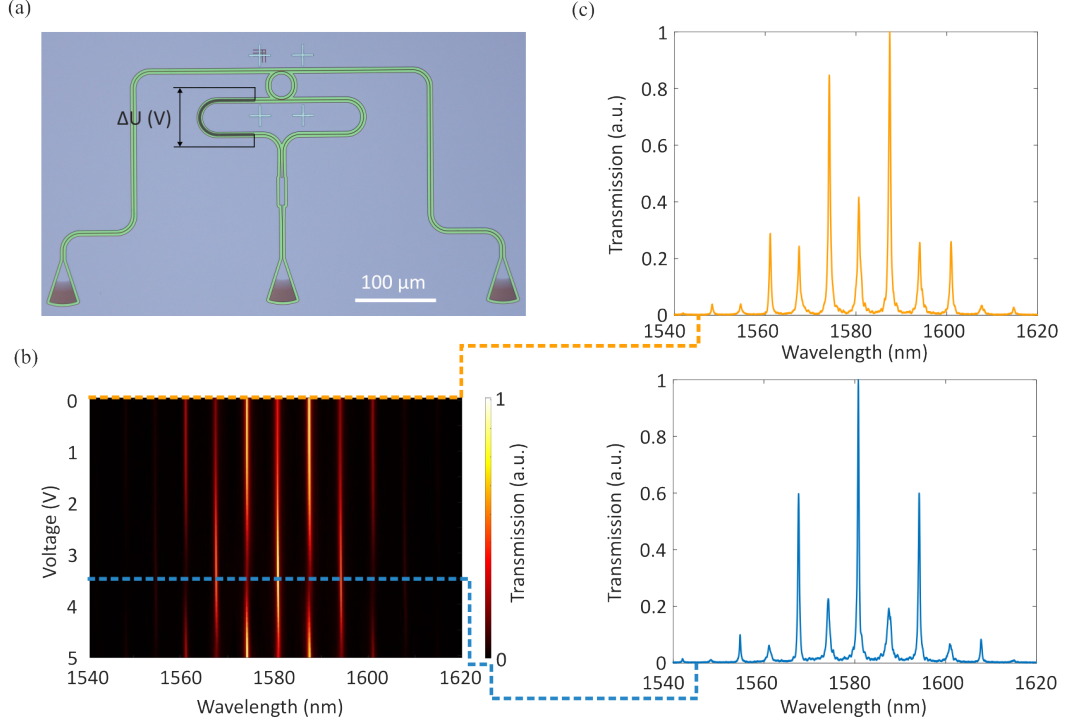


Figure 5.18: Electrically tunable spectral control in a PHIL-loaded MRR. (a) Optical micrograph of the test device incorporating a thermo-optic phase shifter on one arm after a 1:2 MMI splitter. The heater, driven by an external bias ΔU , enables dynamic tuning of the relative optical phase between the two input paths. (b) Measured transmission spectra as a function of heater voltage and wavelength. The standing-wave spatial distribution within the ring is modulated via the phase shift $\Delta\Phi$, leading to mode-dependent changes in resonance absorption. (c) Transmission spectra at two representative voltages: 0 V (top, orange) and 3.2 V (bottom, blue). A clear reversal in resonance depth is observed between even- and odd-order modes, demonstrating controllable spectral selectivity via phase-domain modulation.

Building on wavelength-selective absorption engineering in MRRs, this work introduced a reconfigurable mechanism whereby one optical wavelength modulates another purely via the thermo-optic effect, without requiring electrical interfaces or external modulators.

The underlying principle relies on spatially engineered absorption through precise alignment of nanoheaters with the standing-wave field distribution. At the control wavelength, the optical field overlaps strongly with the nanoheater array, enabling efficient photothermal conversion. The generated heat diffuses through the waveguide core and modifies the local refractive index, thereby shifting the resonance

condition at the spectrally distinct probe wavelength, which experiences minimal modal overlap with the heater regions. As a result, amplitude information from the control signal is transduced onto the probe channel via a shared thermal reservoir, enabling fully optical, cross-spectral modulation.

Experimental demonstrations validated this selective absorption and cross-wavelength signal encoding behaviour, including programmable resonance shaping and thermo-optic modulation of a probe signal by a spectrally distinct control signal. Systematic measurements confirmed PHIL’s ability to apply controllable attenuation to the probe channel, with modulation strength governed by the control signal’s resonance alignment and optical power.

Insertion loss analysis further revealed both the scalability and practical limits of the PHIL design. Measurements and simulations jointly confirmed that absorption is strongly wavelength-selective and tunable through the number of integrated nanoheaters. Importantly, simulations identified a saturation trend in total attenuation due to standing-wave amplitude renormalization, which limits the effective contribution of additional heaters. Reported configuration (9×2 nanoheaters) achieve high attenuation contrast while maintaining low parasitic loss (< 2 dB) at the non-absorptive wavelength.

Potential thermal crosstalk may arise when such thermo-optic interactions are implemented in densely integrated photonic circuits. In contrast to conventional thermo-optic phase shifters, which typically employ extended metallic heaters with lateral dimensions of several tens of micrometres and require electrical driving circuitry, the nanoheater elements used in the PHIL architecture operate at nanometre scale and are optically actuated. This localized heating mechanism substantially reduces both the total thermal load and the spatial extent of heat diffusion. Moreover, the standing-wave-engineered absorption profile concentrates optical energy near

the nanoheater locations, further promoting spatially confined photothermal conversion with limited lateral heat spreading. In large-scale implementations, residual thermal interactions between neighbouring elements may be further mitigated using established foundry-compatible approaches, including thermal isolation trenches, undercut structures, or optimized heater spacing. These considerations indicate that the PHIL architecture remains compatible with dense photonic integration while maintaining controllable thermal interactions.

Overall, this chapter provides the first comprehensive validation of PHIL as a passive yet programmable photonic element for reconfigurable signal routing and all-optical modulation. These results establish the experimental and conceptual foundation for subsequent chapters, where the same principle is extended to across-wavelength summation and dynamic time-domain computations enabled by thermal integration.

6

All-Optical Neuron: Summation and Programmable Nonlinear Activation Across Wavelengths

Contents

6.1	Motivation	115
6.2	All-Optical Programmable Nonlinearity	117
6.2.1	Engineering of Optical Nonlinearity	117
6.2.2	Modulation Stability and Precision of Nonlinear Response	121
6.3	Across-Wavelength Optical Accumulation	123
6.3.1	Linear Optical Accumulation Across Two-Wavelength Channels	124
6.3.2	Across-Wavelength Optical Accumulation with Nonlinear Activation	127
6.4	Further Scaling of WDM Channels	130
6.4.1	Optical Accumulation Across Three-Wavelength Channels	130
6.4.2	Scalability of Multi-Wavelength Channels	135
6.5	Discussion and Summary	138

6.1 Motivation

Wavelength-multiplexed optical signals enable simultaneous and parallel information processing across multiple wavelength channels, providing a natural foundation for high-throughput and low-latency computing. However, the bosonic and weakly interacting nature of photons imposes a fundamental constraint: optical signals at different wavelengths cannot be directly summed or weighted within

the optical domain. In most existing photonic neural network implementations, weighted additions are performed at photodetectors[39–42, 122], where optical fields are first converted into electrical currents. Nonlinear activation functions are then applied either electronically[27, 31, 77, 91] or during subsequent E–O conversions[39, 42, 94]. These repeated conversions introduce additional latency and energy overhead, increasing both the device footprint and system complexity, while fundamentally limiting the scalability of recurrent and multilayer photonic neural networks.

These limitations highlight the need for an intrinsically nonlinear, cascadable, and fully optical processing element capable of performing both summation and nonlinear transformation entirely within the optical domain. In this chapter, I leverage the PHIL architecture as an all-optical neuron that performs both linear summation and nonlinear activation directly within the optical domain. By harnessing the intrinsic nonlinearity embedded in the microring spectrum, PHIL naturally produces a programmable transfer function under optical excitation, eliminating the need for electrical biasing or active modulation. The associated heat-dissipation dynamics inherently couple optical power across wavelengths, giving rise to linear summation and nonlinear activation that are realized self-consistently within a single photonic device.

6.2 All-Optical Programmable Nonlinearity

While the previous chapter demonstrated that the resonance shift of the MRR scales linearly with CS power, I proceed to show that the shape of the resonance itself can be exploited to implement programmable nonlinear activation functions entirely in the optical domain (Figure 6.1), eliminating the need for external circuitry, a known bottleneck in photonic accelerators.




Figure 6.1: Concept of all-optical programmable nonlinearity. By tuning the probe wavelength along the Lorentzian resonance of the microring, different nonlinear transfer functions can be achieved entirely in the optical domain.

6.2.1 Engineering of Optical Nonlinearity

The all-optical nonlinearity engineering was performed using the setup shown in Figure 6.2.

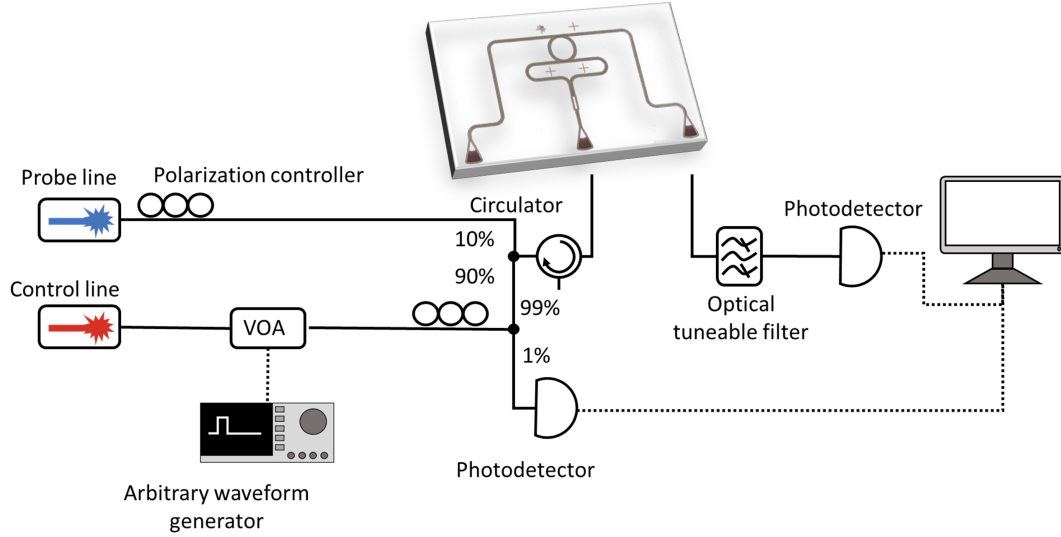


Figure 6.2: Setup for nonlinear optical response engineering.

A tunable laser (TSL-550, Santec) was used as the probe source, and the probe signal was injected into the device as a continuous wave through the 10% port of a 90/10 coupler, with its polarization adjusted using a polarization controller. The probe wavelength was chosen around the non-absorptive resonance to examine different transfer functions (1587.26 nm, 1587.28 nm, and 1587.30 nm, respectively).

A second laser source (N7711A, Keysight) was used as the control signal to induce all-optical spectral shifts. Its incident power was adjusted using a variable optical attenuator (V1550, Thorlabs), which was driven by an electrical pulse generator (AFG3151C, Tektronix). The control signal was delivered to the device in CW mode via the 99% port of a 99/1 coupler, followed by the 90% port of a 90/10 coupler. Polarization was adjusted using a second polarization controller. Throughout the experiment, the control signal power was monitored via the 1% port of the 99/1 coupler.

With the selected probe wavelength, the input power of the control signal was varied from 0 to 1.6 mW (after the grating coupler). As shown in Figure 6.3, the probe responses exhibit distinct modulation traces depending on the spectral position of the probe wavelength relative to the Lorentzian resonance. At the resonance centre, the probe transmission decreases non-linearly with increasing control power, following the Lorentzian profile of the cavity (Figure 6.3(a)).




Figure 6.3: Measured modulation traces of probe transmission at different wavelengths. Output probe responses (ΔT) recorded under periodic modulation of the control signal (λ_{CS}) show distinct nonlinear behaviours when the probe wavelength (λ_{PS}) is tuned across the resonance. The transmission decreases, inverts, or saturates depending on the probe detuning, indicating wavelength-dependent nonlinear response.

When the probe is slightly detuned to the longer wavelength side, the resonance shift driven by the thermo-optic effect causes the probe transmission to first increase and then decrease, indicating that the resonance has crossed the probe position (Fig-

ure 6.3(b)). Further detuning away from the centre leads to a monotonic increase in transmission that gradually approaches a flat region near the resonance peak, where the response saturates due to the limited spectral slope (Figure 6.3(c)).

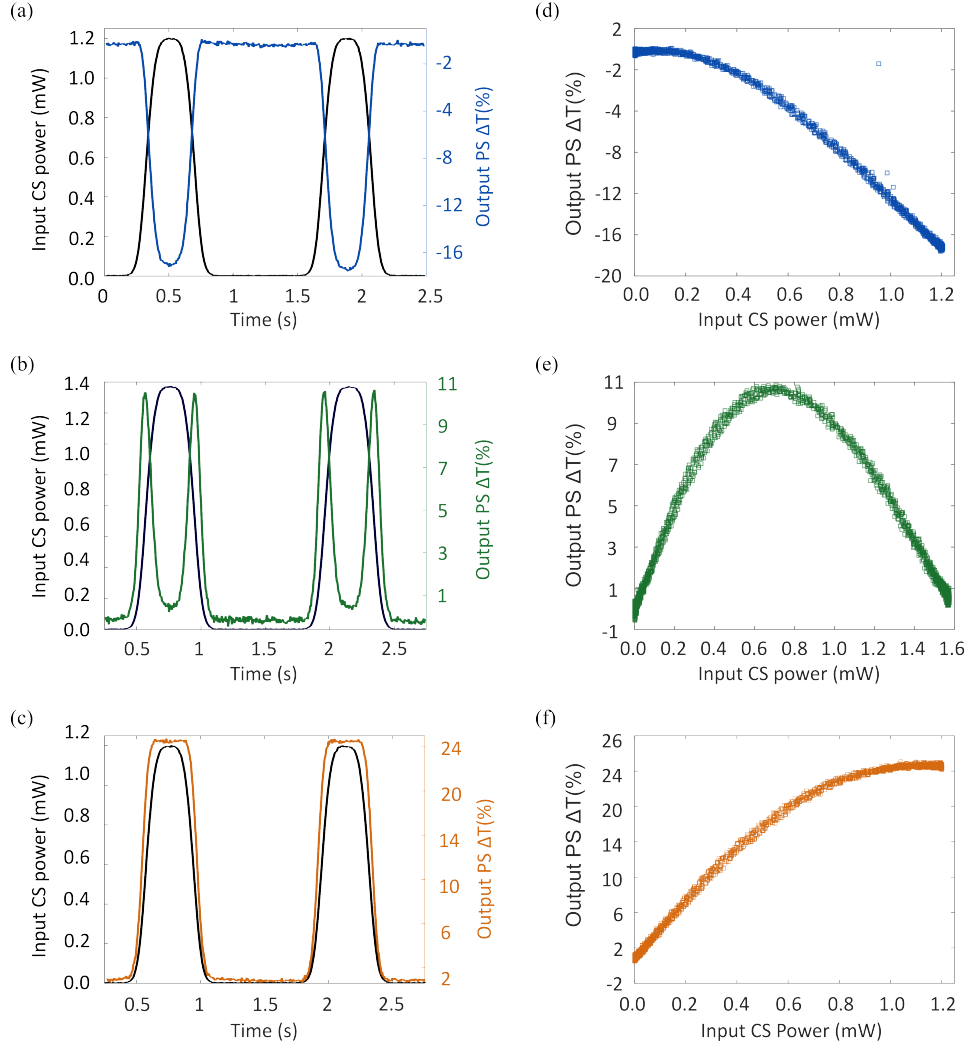


Figure 6.4: Extracted nonlinear transfer functions from probe modulation. Zoomed-in modulation traces (a–c) and corresponding transfer curves (d–f) of the probe signal at three wavelengths reveal Lorentzian-shaped responses, confirming the programmable optical nonlinearities realized all-optically within a single MRR

The corresponding steady-state responses extracted from these modulation traces are shown in Figure 6.4. The zoomed-in traces (a–c) and their extracted transfer curves (d–f) clearly demonstrate that the resonance shape can be directly mapped

into different nonlinear transfer functions. Depending on the probe selection, PHIL produces a Lorentzian-type attenuation, symmetric, or saturating nonlinear activations as responses to the input power of control signal.

These results confirm that programmable optical nonlinearities can be realized purely by selecting the probe wavelength, without modifying the device structure or applying any electrical bias, thereby establishing a fully optical and reconfigurable mechanism for nonlinear activation within a single MRR.

6.2.2 Modulation Stability and Precision of Nonlinear Response

The modulation stability and precision of the all-optical nonlinear response were further evaluated to assess the reproducibility of the thermo-optic activation under repeated excitation cycles. As shown in Figure 6.5, the probe signal was continuously monitored over multiple modulation sweeps spanning one hour. Across all three probe wavelengths, the transmission traces maintain highly consistent amplitude and waveform throughout successive cycles, with no observable baseline drift or distortion. These observations indicate that the localized heating and cooling processes within the PHIL-loaded microring are fully reversible and remain unaffected by cumulative thermal loading or material relaxation effects during quasi-static modulation under open laboratory conditions.

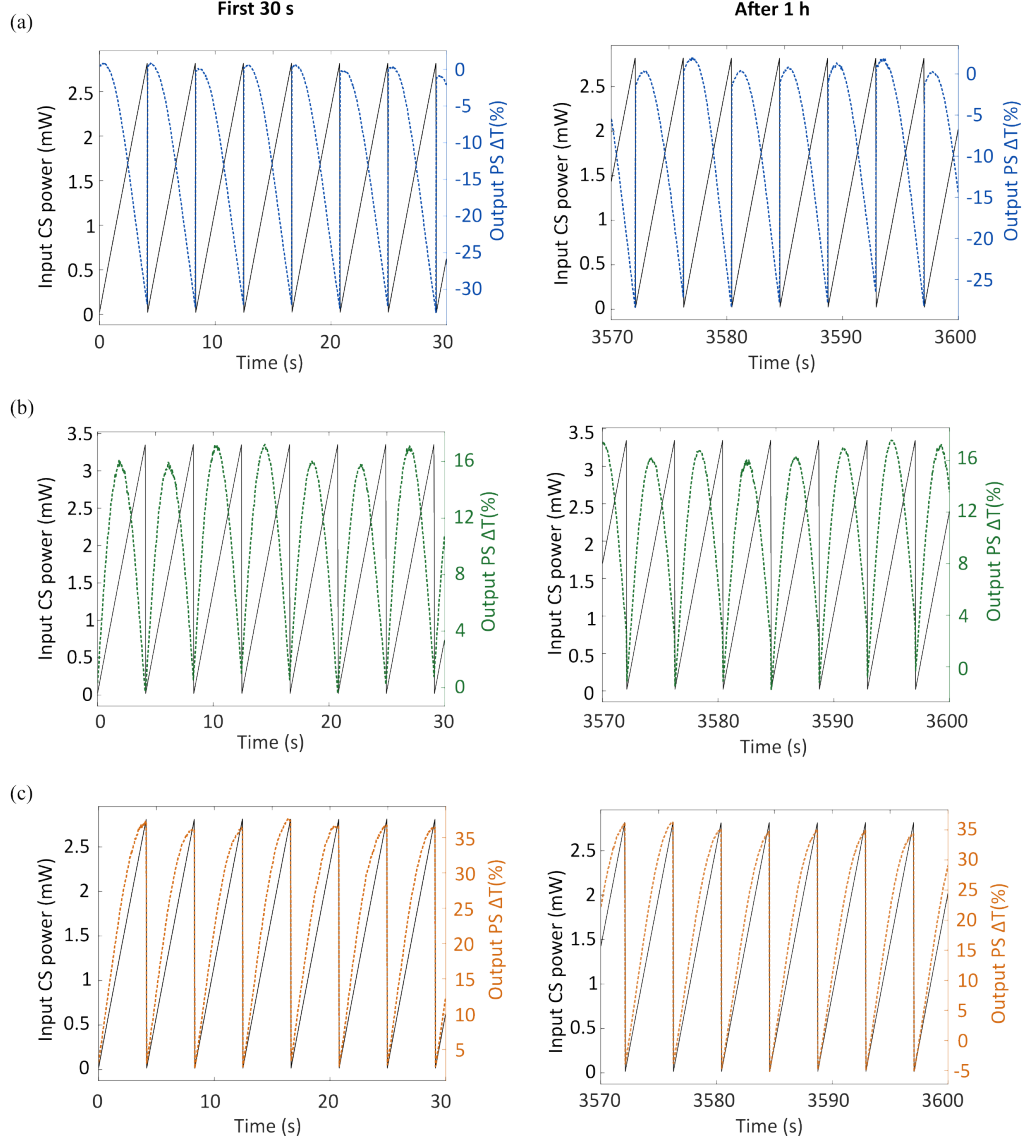


Figure 6.5: Modulation stability of the nonlinear optical response. Continuous modulation traces of the probe signal (ΔT) recorded over 1 hour for three representative probe wavelengths demonstrate stable and repeatable thermo-optic nonlinear modulation. Each probe wavelength maintains its characteristic response shape throughout multiple modulation cycles without noticeable baseline drift or amplitude degradation, confirming the intrinsic stability of the all-optical nonlinearity.

To quantify the precision and repeatability of the nonlinear modulation, the mean and standard deviation of the probe transmission change were extracted from the repeated modulation cycles, as summarized in Figure 6.6. The fluctuation of the modulation amplitude remains below $\pm 1.7\%$ across all probe wavelengths, demon-

strating excellent cycle-to-cycle consistency of the nonlinear response.




Figure 6.6: Statistical evaluation of nonlinear modulation stability. Mean and standard deviation of the probe transmission change (ΔT) extracted from repeated modulation cycles over 1 hour show minimal fluctuation ($< \pm 1.7\%$) across all measured probe wavelengths. The small variance indicates excellent long-term reproducibility and thermal robustness of the PHIL design under continuous optical excitation.

This small statistical spread corresponds to an effective modulation precision of approximately 6-bit, calculated from $BR = \log_2(1/\sigma)$ with $\sigma = 0.017$. These results confirm the stable and highly repeatable implementation of the all-optical nonlinear activation, establishing a reliable foundation for the realization of the all-optical neuron discussed in the following sections.

6.3 Across-Wavelength Optical Accumulation

Having established that the resonance shape of the cavity spectrum can be harnessed to realize programmable nonlinear activation at a single wavelength, I now apply this functionality across wavelength channels. In this section, I explore how the PHIL architecture enables all-optical accumulation of optical power across dis-

tinct wavelengths, where incoherent control signals jointly contribute to a common thermal field (Figure 6.7).




Figure 6.7: Concept of all-optical summation across wavelengths. Two independent control signals (λ_{CS-a} , λ_{CS-b}) are simultaneously injected into the PHIL-loaded microring, where their absorbed optical powers are accumulated in the form of heat. The resulting temperature rise induces a combined thermo-optic redshift of the resonance, which is read out by a probe signal (λ_{PS}). The probe transmission encodes the total absorbed optical energy, realizing linear optical summation across distinct wavelength channels.

6.3.1 Linear Optical Accumulation Across Two-Wavelength Channels

The summation of two control signals with different wavelengths was performed using the setup shown in Figure 6.8. A tunable laser (TSL-550, Santec) served as the probe source, operating in CW mode. The probe signal was injected into the device through the 10% port of a 90/10 coupler, with its polarization adjusted using a polarization controller. The probe wavelength was selected at the half-maximum of a non-absorptive resonance peak to monitor the device response through transmission changes induced by total optical power during phase modulation.

Two continuous-wave signals from independent laser sources (N7711A, Keysight) were used as control inputs carrying the information to be accumulated. These signals, labelled a and b , were tuned to distinct absorptive resonances (Figure 6.9),

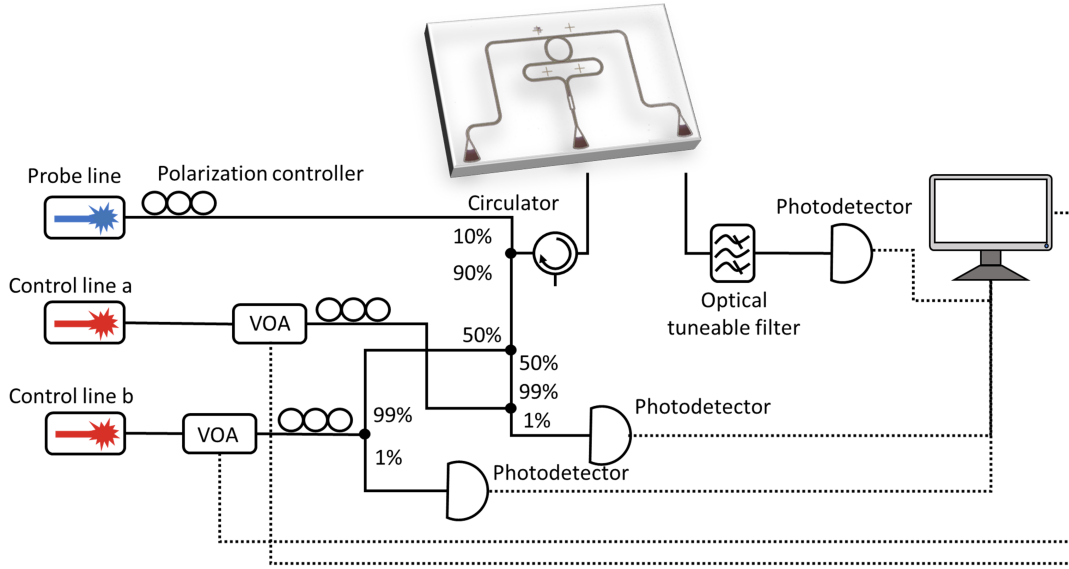


Figure 6.8: Setup for incoherent optical addition of two control signals with different wavelengths.

and combined using a 50/50 optical coupler. The combined control signal was then injected into the PHIL device through the 90% port of a 90/10 coupler. Two variable optical attenuators (V1550, Thorlabs) encoded the data by modulating the transmission of each control line, with the modulation sequence generated and synchronized by a computer-controlled data acquisition (DAQ) system. After calibrating the attenuators, 101 evenly spaced power levels were assigned to each control signal by varying the applied voltage from 0 to 5 V.

As illustrated in Figure 6.10, during this measurement, Control signal *a* (red) increased by one power step every 101 time steps, while Control signal *b* (blue) increased by one power step at each time step. This configuration enabled the measurement of all pairwise combinations between the two input channels, yielding a total of 10,201 distinct addition events for analysis.

The output of the device was filtered using an optical tunable filter (OTF-320, Santec) so that only the probe-line transmission was measured. Throughout the




Figure 6.9: Transmission spectrum of the PHIL device showing two absorptive resonances (*a*, *b*) used for multi-wavelength accumulation. Each control signal is assigned to one of these absorptive wavelengths, while a low-loss probe wavelength (blue box, inset) is positioned around a transparent resonance for readout.

experiment, the two input control signals were continuously monitored via the 1% port of a 99/1 coupler and simultaneously recorded to calculate the theoretical summation results.

The measured transmission change of the probe signal (ΔT) as a function of the two control-signal powers is shown in Figure 6.11(a). As both control signals were tuned to absorptive resonances, their optical powers were jointly converted into heat within the nanoheater region and subsequently encoded onto the probe signal through the accumulated thermo-optic shift. The resulting two-dimensional surface exhibits a monotonic increase in ΔT with total input power, indicating that the probe transmission faithfully represents the summed optical energy from both channels.

In Figure 6.11(b), the comparison between the measured and theoretical additive responses shows a nearly unity correlation, verifying that the PHIL device performs optical-domain summation with minimal cross-channel interference.

The histogram of the deviation between the measured and ideal outputs yields

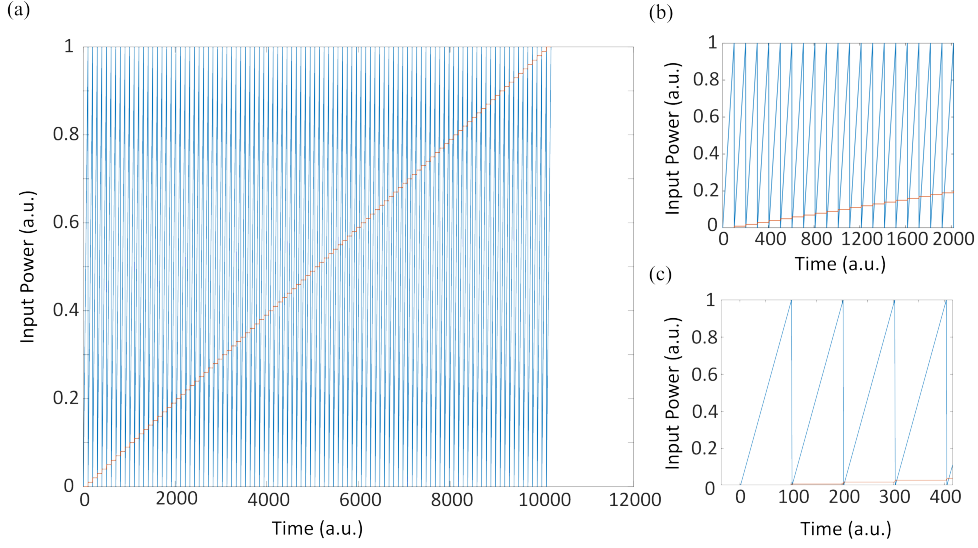


Figure 6.10: Input power of control signals through the measurement. (a) Each control signal is assigned 101 power steps, control signal a (red line) moves one power step at every 101 time-steps while control signal b (blue line) will move one power steps at each time step. 10,201 addition events were measured through the experiment. (b, c) Zoomed-in view at selected time steps

a mean error of -0.0041 with a minimal deviation of 0.0181 (Figure 6.11(c)). This highlights the linearity and reproducibility of the thermo-optic accumulation process under simultaneous multi-wavelength excitation.

6.3.2 Across-Wavelength Optical Accumulation with Non-linear Activation

In addition to the demonstration of linear summation of incoherent control signals, the probe wavelength was varied to apply different activation functions to the across-wavelength summation results, as shown in Figure 6.12(a). When the probe wavelength is positioned at different detunings along the Lorentzian resonance, the accumulated thermo-optic shift transforms the summed optical input into distinct nonlinear mappings. Specifically, probe positions at the resonance centre (Fig. 6.12(b), at a slight right detuning (Figure 6.12(c)), and further along the

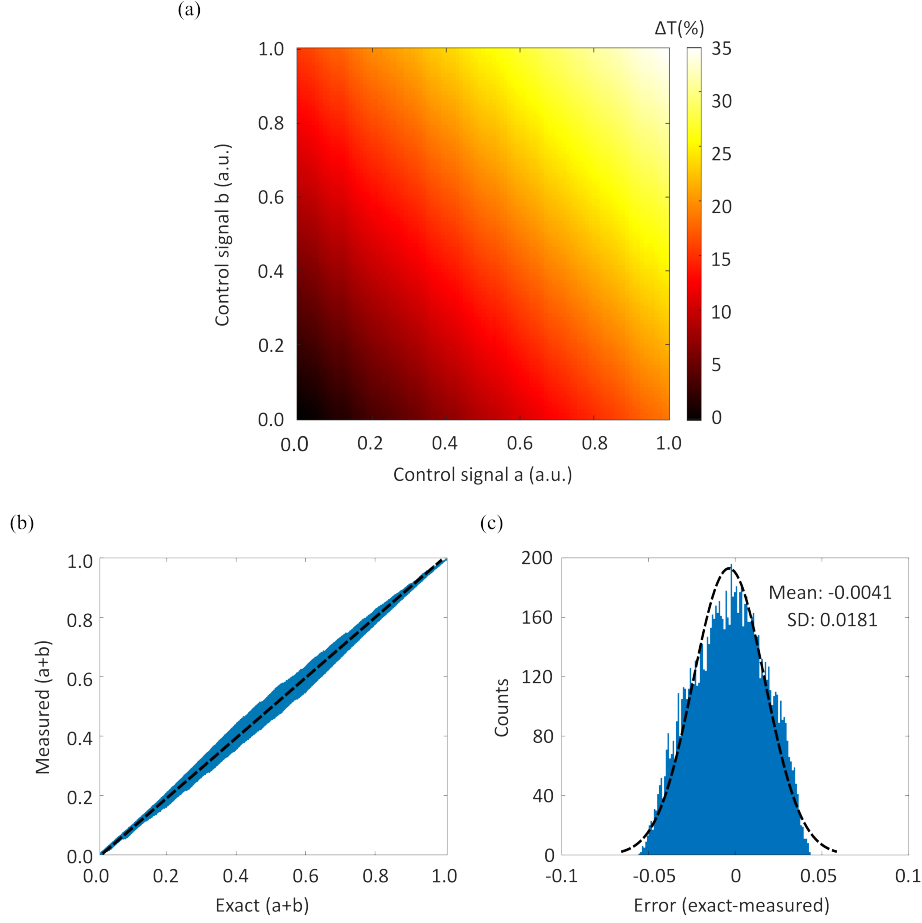


Figure 6.11: Measured all-optical summation across two wavelengths. (a) Measured transmission change of the probe signal (ΔT) as a function of the two independent control-signal powers (a , b), showing a linear increase with total optical input power. (b) Comparison between the measured and ideal additive responses, exhibiting a near-unity correlation and confirming linear summation behaviour. (c) Histogram of the deviation between measured and ideal outputs, yielding a mean error of -0.0041 and a standard deviation of 0.0181 .

right slope (Figure 6.12(d)) yield attenuating, symmetric, and saturating transfer functions, respectively. This behaviour reproduces the programmable nonlinear activation previously demonstrated in the single-wavelength case, now applied directly to the summed optical inputs across multiple wavelength channels.

The larger variation observed in the middle range of the input power arises from the discrete power-sweep configuration used during measurement. In the experimental sequence, Control signal b was swept through 101 discrete power levels for each

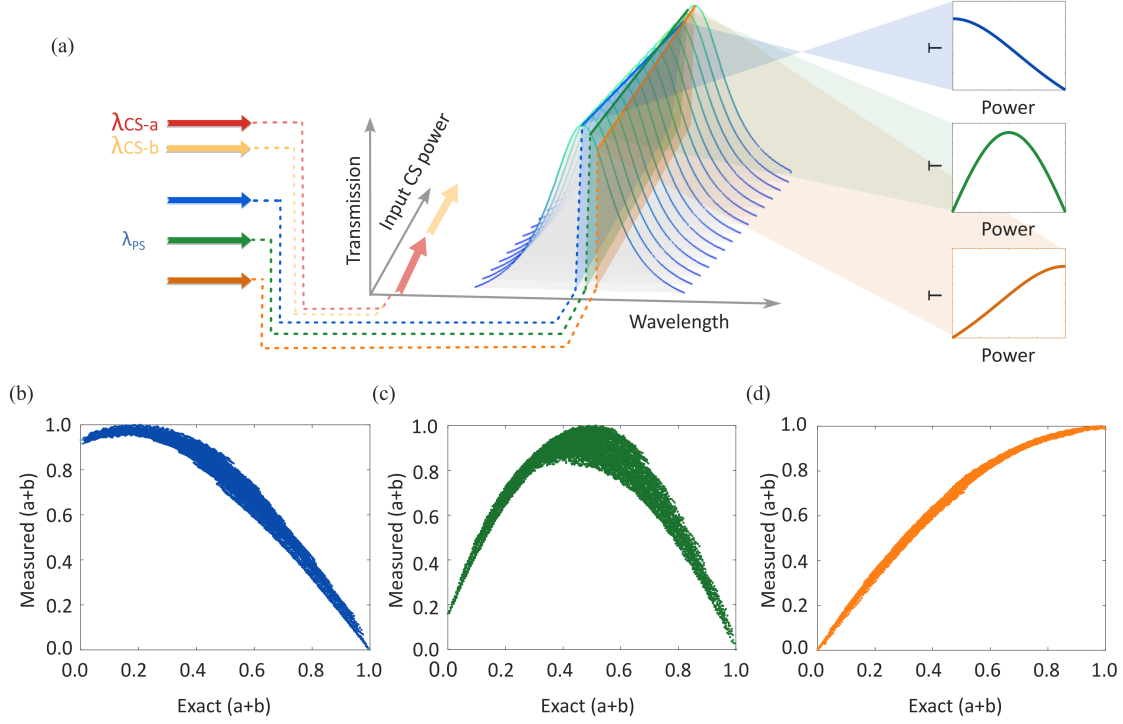


Figure 6.12: Across-wavelength accumulation with optically programmable transfer functions. (a) Concept of across-wavelength accumulation with programmable nonlinearity, where optical powers from multiple incoherent control signals (λ_{CS-a} , λ_{CS-b}) are converted into heat and collectively shift the microring resonance through the thermo-optic effect. By selecting different probe wavelengths (λ_{PS}) along the resonance slope, distinct nonlinear transfer characteristics can be applied to the accumulated optical energy. (b–d) Measured nonlinear transfer functions applied to the across-wavelength summation results, obtained at different probe wavelengths along the resonance profile. Each probe position produces a distinct nonlinear mapping.

fixed level of Control signal a , generating a total of 10,201 addition events. When plotted against the total input power $P_{total} = P_a + P_b$, the data points follow a triangular distribution because a larger number of (P_a, P_b) pairs yield intermediate total powers than the extreme low or high values. As a result, a denser population of summation events occurs near the midrange of the total input power, resulting in slightly larger variation in that region. Despite this, the overall transfer characteristics remain smooth and reproducible, verifying that the PHIL device consistently applies nonlinear transformations to the optically accumulated input.

This demonstrated across-wavelength optical summation and nonlinear activation establish an all-optical reconfigurable framework in which the thermo-optic phase shifter is fully driven by incident optical signals. Optical powers carried on distinct wavelengths are accumulated in the form of heat and subsequently encoded onto a probe wavelength through the thermo-optic effect, entirely eliminating the need for electronic control. This mechanism enables direct optical-domain addition and nonlinear transformation within a single MRR, forming the foundational operation for scalable WDM-based photonic computing.

6.4 Further Scaling of WDM Channels

To further evaluate the scalability of the PHIL system in the wavelength domain, I experimentally extended the optical accumulation from two to three control wavelengths. Each control signal was assigned a distinct absorptive resonance within the same microring cavity, while the probe signal was monitored at a low-loss wavelength.

6.4.1 Optical Accumulation Across Three-Wavelength Channels

The addition measurement of three different control signals was performed using the setup shown in Figure 6.13. A tunable laser (N7711A, Keysight) served as the probe source, operating in CW mode. The probe signal was injected into the device through one port of a 25/25/25/25 fiber coupler and its wavelength was selected at a low-loss resonance peak to monitor the device response through transmission changes induced by total optical power during phase modulation.

Three CW signals from independent laser sources, TSL-550 (Santec), TSL-570 (Santec), and N7711A (Keysight), were used as control inputs carrying the information

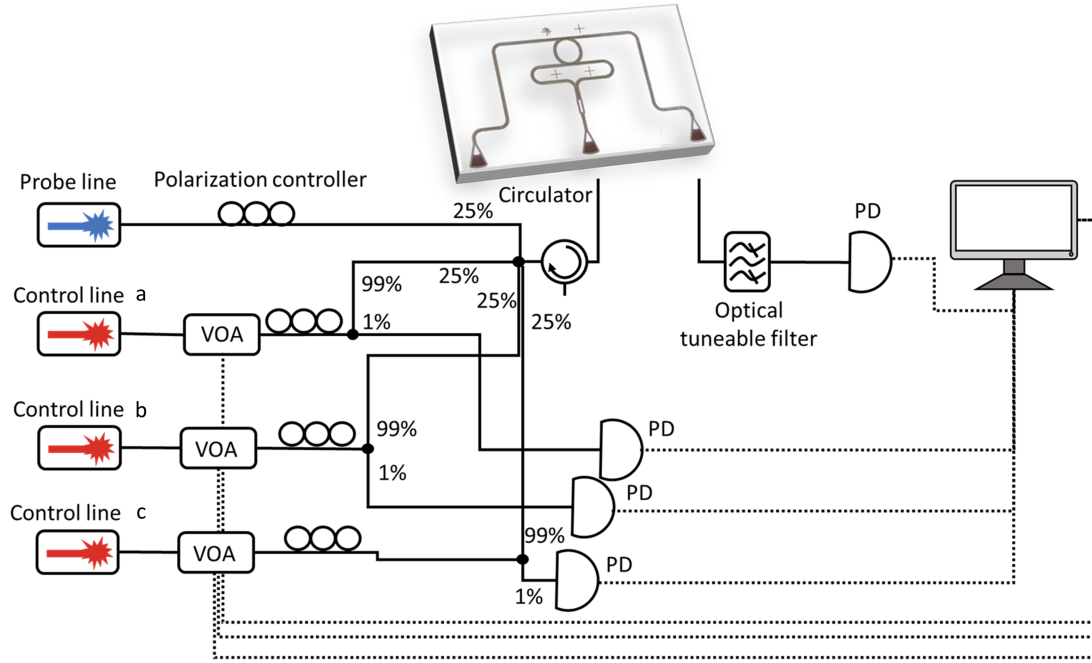


Figure 6.13: Setup for incoherent optical addition of three control signals with different wavelengths.

to be accumulated. These signals, labeled a , b , and c , were set to distinct absorptive wavelengths (Figure 6.14), combined, and injected into the device through the remaining three ports of the same 25/25/25/25 coupler. Each control line was modulated by a variable optical attenuator (VOA, V1550, Thorlabs), with transmission levels controlled by a data acquisition (DAQ) system. After calibration, 26 evenly spaced power levels were assigned to each control signal by varying the applied voltage from 0 to 5 V, resulting in a total of $26^3 = 17,576$ input power combinations.

As shown in Figure 6.15(a), the probe-signal modulation ΔT increases smoothly with the combined optical power of the three control inputs, indicating the cumulative heating effect produced by simultaneous multi-wavelength excitation. The response surface exhibits a monotonic and nearly planar dependence on the total absorbed power, confirming that the PHIL device integrates optical energy from distinct spectral channels in an incoherent yet highly linear manner. The comparison




Figure 6.14: Transmission spectrum of the PHIL device showing three absorptive resonances (a – c) used for multi-wavelength accumulation. Each control signal is assigned to one of these absorptive wavelengths, while a low-loss probe wavelength (blue box, inset) is positioned around a transparent resonance for readout.

between measured and ideal additive responses in Figure 6.15(b) reveals excellent linear correlation and the corresponding error distribution in Figure 6.15(c) follows a Gaussian profile centred near zero ($mean = 0.0162$, $\sigma = 0.0437$), corresponding to an effective analog precision of approximately 4.5 bits.

To further examine the consistency of multi-wavelength accumulation and quantify the contribution of each spectral channel, I analysed two-dimensional response slices under fixed control-signal c powers, as shown in Figure 6.16. Each row corresponds to a constant optical power level of signal c , ranging from 0.004mW to 0.454mW while the powers of signals a and b were systematically varied.

The resulting probe modulation maps (left panels, Figure 6.16(a-c)) exhibit uniformly increasing intensity gradients with the total absorbed optical power. The measured outputs (middle panels, Figure 6.16(d-f)) maintain excellent agreement with the ideal additive response across all operating conditions, confirming that the thermal accumulation remains linear and independent of the specific wavelength combination.




Figure 6.15: All-optical accumulation across three wavelengths. (a) Measured probe-signal modulation ΔT as a function of three control signals (a , b , and c). (b) Comparison between measured and ideal additive responses, indicating excellent linearity. (c) Error histogram of the measured results, yielding a Gaussian distribution with mean 0.0162 and standard deviation 0.0437, corresponding to ~ 4.5 -bit effective precision.

The error histograms (right panels, Figure 6.16(g-i)) follow narrow Gaussian distributions with standard deviations of approximately 0.05, corresponding to an effective analogue precision between 4 and 5 bits. Importantly, no observable bias or broadening is detected at higher input powers. These results validate that the PHIL device performs stable and repeatable accumulation over a wide dynamic range, maintaining high fidelity across multiple wavelength channels.

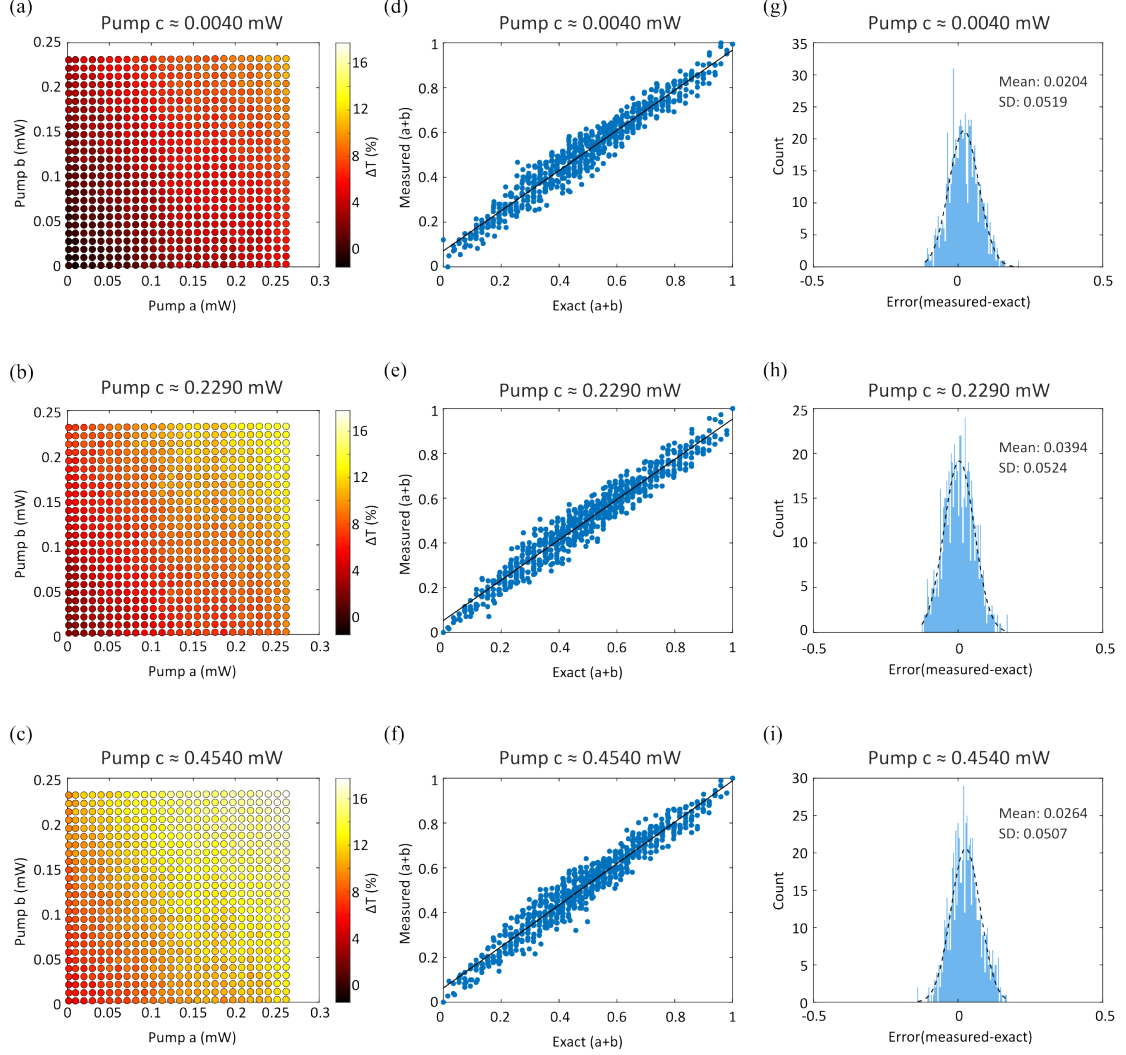


Figure 6.16: Cross-sectional analysis of three-wavelength accumulation under varied third-channel power. Each row corresponds to a fixed control signal c power of 0.0040, 0.2290, and 0.4540 mW, respectively. Left (a-c): measured probe-signal modulation ΔT as a function of control signals a and b . Middle(d-f): comparison between measured and ideal additive responses, confirming high linearity. Right(g-i): error histograms showing Gaussian distributions with mean deviations below 0.04 and standard deviations of ~ 0.05 , indicating ~ 4 -5 bit analogue precision across operating conditions.

6.4.2 Scalability of Multi-Wavelength Channels

To assess the scalability of the PHIL device under increasing WDM operation, I evaluated its analogue precision across one, two, and three simultaneous optical channels. The effective bit resolution (BR) was determined from the standard deviation of the normalized output error relative to the ideal additive or integrated response, defined as

$$BR = \log_2 \left(\frac{1}{\sigma} \right),$$

where σ represents the standard deviation of the residual error. This metric, equivalent to the effective number of bits (ENOB) in analogue electronic systems, serves as a measure of computational precision and linearity for optical accumulation processes.

(i) Single-channel operation. For single-wavelength excitation, the device demonstrated a nearly ideal thermo-optic response with minimal nonlinearity or noise-induced variation. From the statistical spread of the nonlinear transfer measurements (Figure 6.6), the standard deviation of the response was approximately $\sigma = 0.0117$, corresponding to an analogue precision of $BR \approx 6.4$ bits. This high intrinsic precision defines the upper limit of the system performance when only a single absorptive resonance is excited.

(ii) Two-channel operation. When two absorptive control wavelengths were simultaneously modulated, the measured probe response exhibited excellent agreement with the expected additive behaviour. The corresponding error histogram yielded $\sigma = 0.0181$ (Figure 6.11), corresponding to an effective resolution of ~ 5 -6 bits. The negligible mean bias confirms that the thermal summation of two incoherent wavelengths remains highly linear and independent, with minimal inter-channel crosstalk.

(iii) Three-channel operation. To investigate precision scaling with additional multiplexed inputs, a third wavelength channel was introduced. The deviation increased moderately to $\sigma \approx 0.04$ (Figure 6.15), resulting in a resolution of ~ 4.5 bits, corresponding to about $2^{4.5} \approx 22$ distinguishable accumulation levels. This precision remains well within the range typically considered sufficient for analogue neural computing and inference tasks, where 4–6 bit resolution is generally adequate[123].

This reduction primarily arises from shared thermal loading and stochastic variations among multiple wavelength paths. Nevertheless, the retention of >4 -bit precision under concurrent three-wavelength excitation confirms robust multi-channel accumulation fidelity.

(iv) Physical origins of reduced precision. The observed sub-linear degradation in BR with increasing N_λ can be attributed to three main factors. First, in the current experiment, the total CS power was kept constant. As the number of wavelength channels increases, the available optical power per channel decreases, leading to a smaller modulation depth at each wavelength and, therefore, a reduced overall signal-to-noise ratio. Second, each additional laser path introduces independent noise sources, including intensity fluctuations, wavelength jitter, and insertion loss variations, which accumulate as uncorrelated additive noise in the measured output. Together, these two effects produce a combined error variance that scales approximately as

$$\sigma^2(N_\lambda) \approx \sigma_0^2 + N_\lambda (\sigma_{\text{laser}}^2 + \sigma_{\text{path}}^2) + \sigma_{\text{NL}}^2(P_{\text{tot}}, N_\lambda),$$

where σ_0 accounts for baseline readout noise, and σ_{NL} represents the additional curvature arising from nonlinear thermal loading. Although the total control power P_{tot} was kept nearly constant, σ_{NL} can still increase with N_λ because each wavelength excites a slightly different absorption profile and resonance position. The

resulting non-uniform heating and partial thermal overlap alter the local thermo-optic efficiency, introducing small cross-channel correlations that contribute to the nonlinear variance term.

A further contribution arises from the wavelength-dependent transmission of the grating couplers. Because the coupler response follows a Gaussian envelope, different control wavelengths experience slightly different coupling efficiencies, resulting in unequal absorbed powers even when identical power is applied before the grating coupler. This spectral imbalance alters the thermo-optic transduction slopes among channels, introducing an additional variance term σ_{wg} . Pre-calibrating the per-channel transduction coefficients α_i and applying a power pre-emphasis $P_{\text{set},i} \rightarrow P_{\text{set},i}/k_i$ can restore equalized absorption and suppress this wavelength-dependent error source. Alternatively, employing broadband edge couplers[124, 125] or 3D couplers[126] with flatter spectral response can further minimize wavelength-dependent coupling variations, reducing σ_{wg} without requiring per-channel calibration.

(v) Scaling projection. Assuming uncorrelated noise, the total deviation can be approximated as $\sigma_N \propto \sqrt{N_\lambda} \sigma_1$, yielding the empirical precision scaling

$$\text{BR}_N = \log_2 \left(\frac{1}{\sqrt{N_\lambda} \sigma_1} \right).$$

Under this scaling, the PHIL device is expected to maintain $\text{BR} > 4$ bits even for $N_\lambda > 10$, provided that inter-channel correlations and thermal cross-coupling remain negligible. Improvements in heater geometry, thermal isolation, and per-channel power normalization are expected to further mitigate these effects. Importantly, even the experimentally demonstrated ~ 4.5 -bit resolution at three wavelengths exceeds the precision required for neuromorphic inference tasks[127, 128]. These results confirm that the PHIL architecture provides a scalable foundation

for multi-wavelength analogue optical computing with minimal degradation in computational fidelity.

6.5 Discussion and Summary

This chapter demonstrated a fully optical neuron based on the PHIL architecture, which performs both summation and nonlinear activation entirely within the optical domain. Through localized optical absorption and thermo-optic modulation in a microring resonator, the PHIL device enables distinct wavelength channels to interact via a shared thermal field, forming a native optical summation node without electro-optical conversion. The resonance shape of the MRR defines programmable nonlinear transfer functions, ranging from Lorentzian attenuation to saturating and symmetric responses, achieved solely by tuning the probe wavelength. This establishes a reconfigurable, bias-free mechanism for neuron-like activation in integrated photonics.

The experimental results confirm stable and precise operation with an effective 5–6-bit analogue resolution in single- and two-wavelength operation, and approximately 4.5-bit precision for three-wavelength accumulation. These measurements verify that optical powers carried at different wavelengths can be linearly summed and nonlinearly transformed through thermo-optic coupling within one device. The modest precision degradation with increased wavelength count arises mainly from power imbalance and wavelength-dependent coupling, both of which can be mitigated through optimized heater geometry, broadband couplers, and per-channel calibration.

Overall, the PHIL platform introduces a scalable route toward wavelength-multiplexed, energy-efficient, and reconfigurable photonic computing. By mediating optical interactions through heat rather than electrical charge, it unifies accumulation,

activation, and reconfiguration in a single passive component. The resulting all-optical neuron serves as a fundamental building block for large-scale photonic neural and analogue processors capable of end-to-end computation without electrical intermediaries.

All-Optical Temporal Integration With Nonlinear Activation

Contents

7.1 Motivation	140
7.2 Temporal Dynamics of Cross-Wavelength Modulation	142
7.2.1 Modulation Speed Characterisation	142
7.2.2 A Leaky Integrator Model	145
7.3 Temporal Integration of Ultrafast Optical Signals with Nonlinear Activation	148
7.3.1 Numerical Simulation of Temporal Integration	148
7.3.2 Experimental Validation of All-Optical Temporal Inte- gration	151
7.3.3 Neural Network Evaluation Using the PHIL Activation Function	154
7.4 Energy Efficiency	157
7.4.1 Comparison of All-Optical Light-Controlled Technologies	159
7.5 Discussion and Summary	161

7.1 Motivation

While photonic neural networks have demonstrated exceptional promise for parallel matrix–vector operations, their scalability remains fundamentally limited. The large physical footprints of optical unitary cells, together with insertion losses inherent to both active and passive components, have made scaling up weight matrices highly challenging[20]. Various architectural strategies, such as tiling of photonic sub-assemblies[129], matrix factorization[89], and weight pruning[130],

provide only partial relief, falling short of the scalability required for the billion-parameter workloads characteristic of modern AI.

To expand the processing capacity within a compact footprint, the temporal degree of freedom of photons has also been increasingly exploited as an additional multiplexing dimension[28, 43]. Temporal multiplexing, often in combination with wavelength or spatial multiplexing, has enabled advanced photonic accelerator architectures capable of real-time image and video recognition[43, 90]. Despite these advances, direct all-optical processing of large input vectors with nonlinear transformation remains an unmet challenge. Temporal accumulation of weighted optical signals has thus far only been demonstrated electronically, through photoreceiver charge integration[43, 44], while nonlinear activation between network layers is typically implemented either offline[38] or through additional optoelectronic conversions[39–42].

In this chapter, I exploit the intrinsic thermal relaxation dynamics of PHIL as a means to realize all-optical cross-time integration. By harnessing the slow heat dissipation process as a controllable temporal memory, successive optical inputs are accumulated as thermal states within the waveguide and continuously transduced into phase shifts of a probe signal. The inherently leaky nature of the thermal response transforms what is traditionally viewed as a limitation, the slow thermo-optic dynamics, into a powerful mechanism for analogue, all-optical temporal integration. PHIL effectively converts the dissipative dynamics of heat into a functional optical leaky integrator, capable of performing temporal integration, decay, and nonlinear transformation entirely within a single passive photonic element, thereby bridging the spectral and temporal dimensions of optical computing into a unified framework that supports accumulation across both wavelength and time domains without any electro-optic conversion.

7.2 Temporal Dynamics of Cross-Wavelength Modulation

Building upon the across-wavelength summation and nonlinear activation demonstrated in the previous chapters, this section focuses on the temporal characteristics of the PHIL device. Because the thermo-optic effect operates over a finite heat dissipation timescale, the device possesses an intrinsic temporal constant that defines its memory and integration behaviour. To quantify this effect, I experimentally characterized the transient response of the probe signal under square control pulses with varying durations.

7.2.1 Modulation Speed Characterisation

The dynamic response of the device was characterized using the setup in Figure 7.1. The probe wavelength was aligned to the half-maximum of a non-absorptive resonance, as this operating point maximizes sensitivity and ensures a linear response to variations in the overall input power during phase shifting. A separate laser served as the control line to induce all-optical spectral shifts. The control signal power was modulated by an electro-optic modulator, amplified by a following erbium-doped fibre amplifier (EDFA), and launched into the device as rectangular pulses with optimized polarization.

Pulse widths were varied from 10 ns to 500 ns. At the output, an optical tunable filter with a 1 nm bandwidth suppressed the residual control signal, isolating the probe for detection. The probe was then amplified by a pre-amplifier to compensate for insertion loss and measured with a 125 MHz photodiode.

The modulation results on the probe signal transmission with different control signal pulse width are shown in Figure 7.2. The modulation traces clearly exhibit

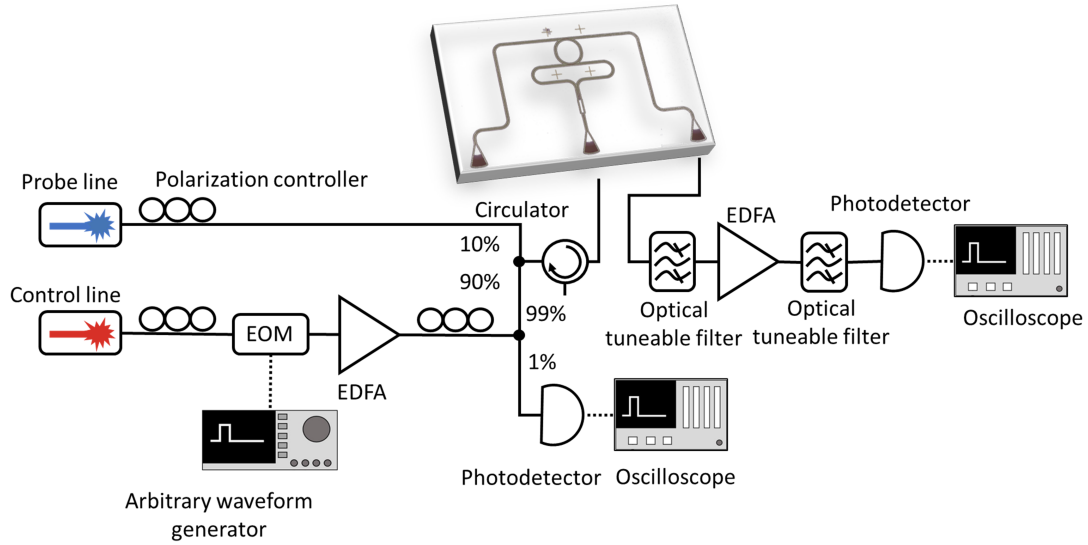


Figure 7.1: Square pulse response characterization. Transmission change at probe signal in response to input control signal with pulse width from 5 ns to 500 ns.

an exponential rise and decay characteristic of the thermo-optic process, confirming that the PHIL response is governed by the intrinsic heat accumulation and dissipation dynamics of the subwavelength heaters. As the control-pulse width increases, the probe transmission gradually approaches a higher steady-state level, reflecting the progressive build-up of thermal energy within the cavity. Short pulses (e.g., 5 ns to 20 ns) produce only partial modulation due to insufficient heating time, whereas longer pulses allow the system to approach the thermal equilibrium within this duration.

To further quantify the characteristic timescales of the thermo-optic process, the transient responses corresponding to different control-pulse widths were analysed. Figure 7.3(a) overlays the temporal responses of the probe transmission under control pulses of different widths, directly comparing the heating and cooling transients. All responses exhibit an almost identical initial slope, confirming that the heating dynamics are governed by a common thermal time constant independent




Figure 7.2: Temporal response of the PHIL device under square-pulse excitation. Normalized transmission change of the PS in response to CS pulses with widths ranging from 5 ns to 500 ns.

of the input duration. As the control pulse remains on, the temperature continues to increase until the absorbed power and thermal dissipation reach equilibrium, giving rise to the observed plateau for longer pulse widths.

From the temporal response measured under a 400 ns control pulse, the 10%–90% rise and fall times were extracted through exponential fitting to quantify the heating and cooling dynamics of the device (Figure 7.3(b)). The transient segments were independently fitted using a bi-exponential function of the form

$$y(t) = a_1 e^{b_1 t} + a_2 e^{b_2 t},$$

implemented in MATLAB. Separate fits were employed for the heating and cooling processes to account for their distinct thermal boundary conditions. This bi-exponential model accurately reproduces both the rapid local temperature change near the nanoheaters and the slower global thermal relaxation of the waveguide. The extracted time constants are approximately $\tau_{\text{rise}} = 89.6$ ns and $\tau_{\text{fall}} = 841.2$ ns.

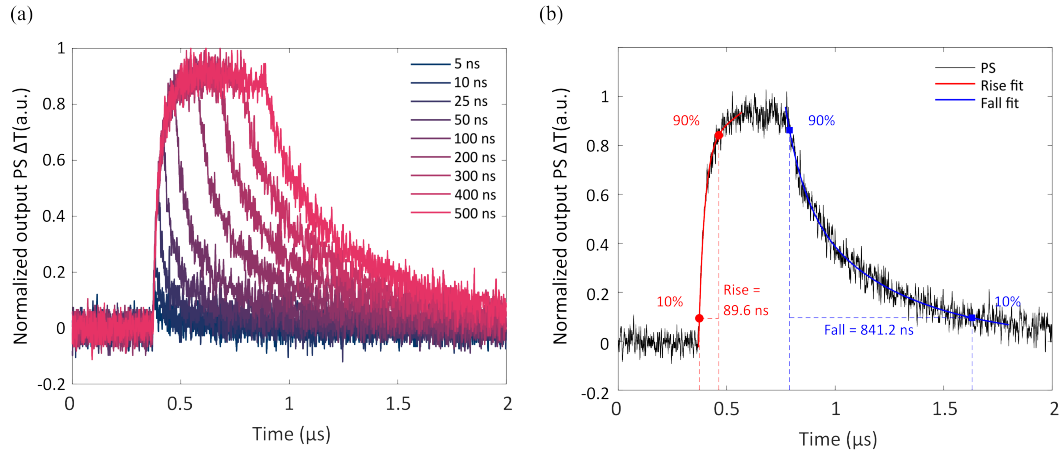


Figure 7.3: Rise and fall time analysis of the temporal response. (a) Overlaid probe-signal responses for control pulses of different widths, showing the transition from transient to steady-state behaviour as the pulse width exceeds 100 ns. A plateau appears beyond this duration, marking the onset of thermal saturation. (b) Extracted 10%–90% rise and fall times obtained from exponential fitting of the heating and cooling edges, yielding a rise time of approximately 89.6 ns and a fall (cooling) time of 841.2 ns.

This pronounced asymmetry between the fast heating and slow cooling responses aligns with the physical nature of PHIL: thermal energy is accumulated during optical excitation and quickly transferred to the nearby detecting region (~ 150 nm), where it is sensed by the probe signal, while the subsequent dissipation of heat occurs gradually over a much longer timescale. This behaviour establishes PHIL as a promising platform for all-optical across-time integration.

7.2.2 A Leaky Integrator Model

To quantitatively describe this thermally mediated integration process, it is essential to determine the characteristic cooling timescale that governs the rate of energy dissipation in the PHIL device. The exponential decay of the probe transmission during the cooling phase provides a direct measure of this relaxation behaviour, from which the thermal time constant τ_{cool} is extracted at the $1/e$ point of the fitted curve (Figure 7.4(a)). This follows the same convention as in first-order RC circuits,

where the time constant τ corresponds to the time required for the response to decay to $1/e$ (approximately 37%) of its initial value [131]. Here, $\tau_{\text{cool}} \approx 245$ ns represents the effective thermal memory of the device, the duration over which residual heat from preceding optical pulses continues to influence subsequent ones.

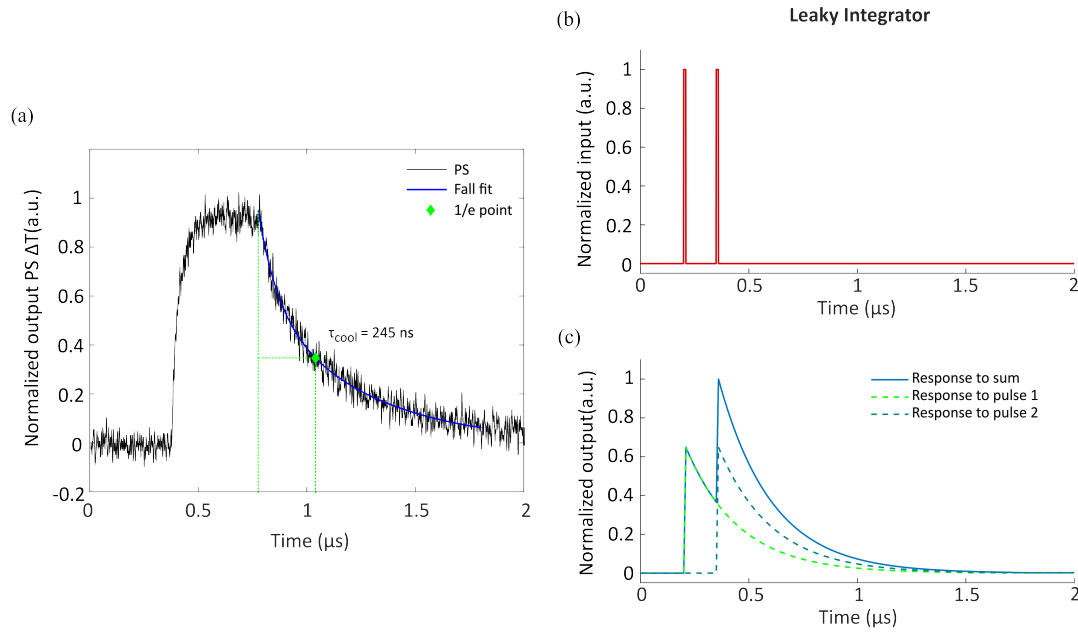


Figure 7.4: Leaky integrator model based on the extracted thermal time constant. (a) Exponential fit of the probe-signal decay (PS) following optical excitation, used to extract the thermal time constant $\tau = 245$ ns corresponding to the $1/e$ cooling point. (b)(c) Conceptual representation of the leaky integrator behaviour, where two optical input pulses are temporally integrated within the thermal window. The output response (blue) follows an exponential decay governed by τ .

Building upon this experimentally determined cooling constant, the temporal evolution of the PHIL device can be described analytically using a first-order leaky-integrator model. The system dynamics are governed by

$$\frac{dy(t)}{dt} = -\frac{y(t)}{\tau_{\text{cool}}} + x(t),$$

where $x(t)$ represents the optical input pulse train and $y(t)$ denotes the normalized thermal response of the device. The first term captures the continuous heat

dissipation governed by τ_{cool} , while the second term accounts for the instantaneous temperature rise induced by the incident optical power. This equation is solved numerically using a forward Euler method in MATLAB to reconstruct the transient thermal dynamics under arbitrary optical excitations.

Figure 7.4(b–c) conceptually illustrates the constructed leaky-integrator model by simulating two consecutive optical inputs with 10 ns pulse width and 150 ns interval. The resulting output exhibits a cumulative response that integrates the two excitations within the thermal window defined by τ_{cool} , demonstrating the characteristic leaky-integration behaviour of the model.

The integration ability of PHIL is further examined by applying pulse-train excitations with temporal durations comparable to the measured thermal constant τ_{cool} . Figure 7.5 presents the system response under a sequence of five 35 ns pulses separated by 15 ns (250 ns in total). The simulated output in Figure 7.5(a) exhibits cumulative heating during the pulse sequence, followed by a gradual exponential decay after the excitation ends.

The same optical excitation was applied experimentally (Figure 7.5(b)). The PS response amplitude increases progressively with the number of input pulses, indicating cumulative thermal accumulation within the integration window defined by τ_{cool} . After the final pulse, the signal decays exponentially as the stored heat dissipates, reproducing the leaky integration behaviour predicted by the model. The excellent agreement between the measured and simulated responses confirms that PHIL operates as a first-order thermal integrator, capable of performing real-time accumulation of optical signals in the time domain.

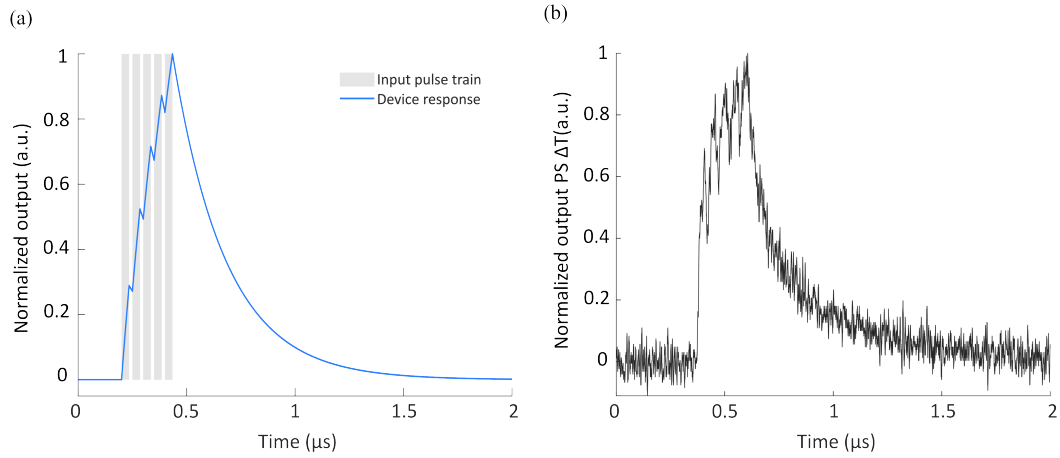


Figure 7.5: Simulation and experimental demonstration of all-optical temporal integration. (a) Simulated output of the leaky integrator model under a pulse-train excitation consisting of 5×35 ns pulses separated by 15 ns. The accumulated thermal response (blue curve) rises stepwise with each pulse before decaying exponentially after the final excitation. (b) Experimental measurement of the PHIL device under identical excitation conditions, showing excellent agreement with the model and confirming that the heat-dissipation dynamics accurately reproduce the leaky-integration process in the time domain.

7.3 Temporal Integration of Ultrafast Optical Signals with Nonlinear Activation

Building on the leaky-integrator dynamics established earlier, I next demonstrate temporal integration of GHz-rate optical inputs by leveraging the inherently slower, MHz-scale heat dissipation process.

7.3.1 Numerical Simulation of Temporal Integration

To verify this behaviour, I firstly applied the built numerical model based on the first-order leaky-integrator equation with a measured cooling time constant of $\tau_{\text{cool}} = 245$ ns to simulate the device response under a train of 500 ultrafast optical pulses (20 ps width, 40 ps period) at various input power, as shown in Figure 7.6(a). The calculated output shows an exponential rise and decay profile

consistent with thermal accumulation and dissipation dynamics.

Increasing the input power results in proportionally higher integration peaks, reflecting larger temperature offsets due to cumulative optical absorption. The corresponding transfer curve (Figure 7.6(b)) shows a linear dependence of the normalized peak output on input power, mimicking the linear accumulation of heating effect.

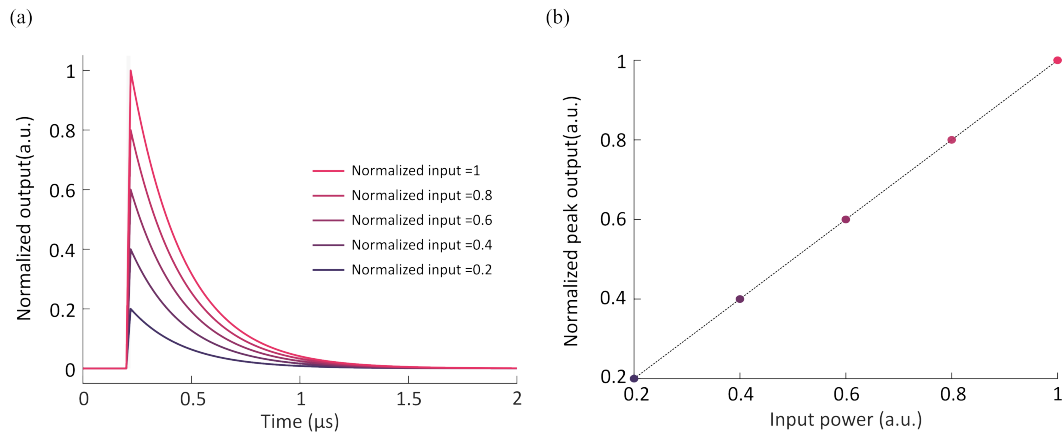


Figure 7.6: Numerical simulation of thermo-optic leaky integration on ultrafast pulses. (a) Simulated responses under a train of 500 ultrafast pulses (20 ps width, 40 ps period) for different input amplitudes, showing cumulative heating followed by exponential decay with $\tau_{\text{cool}} = 245$ ns. (b) Corresponding transfer curve of normalized peak output as a function of input power.

The slow thermal relaxation of the PHIL structure provides a natural integration window, within which transient optical excitations are accumulated as a continuous thermal response. The thermal time constant therefore determines the effective temporal integration window of the PHIL device: optical pulses arriving within this interval accumulate through the thermo-optic response, whereas signals separated by longer delays decay as the stored heat dissipates.

The idea of using PHIL as an all-optical temporal integrator is sketched in Figure 7.7. A sequence of short optical control pulses (CS) is launched into the absorptive

channel, while a low-power PS monitors the induced phase shift through the thermo-optic effect. Each control pulse deposits heat that slightly increases the local temperature of the waveguide, and when successive pulses arrive within the thermal time constant τ_{cool} , their effects add up to produce an integrated thermo-optic shift on the probe signal. After the excitation ceases, the temperature decays exponentially, yielding the characteristic output of a leaky integrator. In this way, the instantaneous optical power of the pulse train is converted into a time-integrated optical response.




Figure 7.7: Concept of temporal integration on optical signals. High-speed control pulses at wavelength λ_{cs} are injected into the PHIL unit and undergo wavelength-selective absorption at the nanoheater sites. The absorbed optical energy accumulates as heat over the device's thermo-optic time constant, producing a leaky integration of the pulse train. A weak probe at λ_{ps} reads the thermo-optic resonance shift, with the Lorentzian lineshape imparting the nonlinear transfer. The probe output thus reflects both the accumulated energy and the cavity-defined nonlinear activation.

Experimental verification of this concept at GHz repetition rates requires dedicated high-speed optical instrumentation. As such facilities were not available in Oxford, the measurements were performed at the Aristotle University of Thessaloniki by Dr. Ioannis Roumpos. The following section presents these results, confirming the predicted leaky-integration behaviour under realistic experimental conditions.

7.3.2 Experimental Validation of All-Optical Temporal Integration

The experimental validation of the temporal integration capability of PHIL was conducted using the setup shown in Figure 7.8, where high-speed optical signals (CS) were applied at a modulation rate of 50 GHz, corresponding to 20 ps optical pulses. These signals were delivered in bursts lasting 20 ns and separated by 1 μ s intervals to ensure complete thermal recovery between integration events. The CS was injected into the device at an absorptive wavelength, while the PS was monitored at a fixed spectral position on the resonance slope to detect the thermo-optic modulation. The experimental setup is detailed in the section 4.3.3.

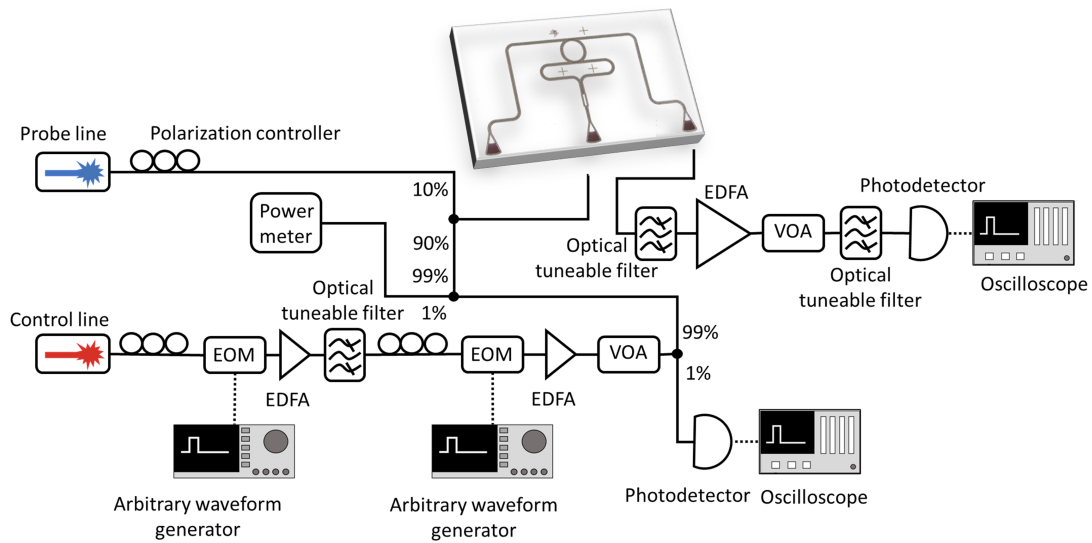


Figure 7.8: Experiment setup for all-optical integration of time-multiplexed signal

Figure 7.9 displays the measurement results. The left panels (a-d) show different CS pulse trains, including a baseline sequence of alternating ones and zeros and three randomized patterns (Random 1–3), each containing an equal number of logical ones and zeros (500 each). This configuration ensures identical total optical energy across all patterns while varying the temporal order of excitation.

The right panels (Figure 7.9(e-h)) present the measured PS responses under different average CS powers. As the CS bursts are absorbed and converted into heat, the induced thermo-optic index change accumulates and modulates the PS transmission. The peak amplitude of the PS response corresponds to the integrated energy of each burst, whereas the subsequent exponential decay reflects the gradual thermal dissipation.

The dependence of modulation amplitude on average CS power confirms that PHIL accumulates optical energy within its thermal relaxation window, performing analogue integration of high-speed optical signals. This experimentally observed behaviour closely aligns with the simulated results discussed above (Figure 7.6), validating the leaky-integrator model and confirming the thermal accumulation mechanism underlying the all-optical temporal integration process.

To further quantify this behaviour, Figure 7.10(a) analyses the integration peak amplitude as a function of CS average power. The spectral shift of the PS resonance follows a nonlinear mapping due to the Lorentzian lineshape of the MRR, giving rise to a ReLU-like transfer function that is inherently useful for neuromorphic processing[132, 133].

The reproducibility of the integration response was evaluated across multiple random pulse sequences, with the standard deviation of integration peaks plotted in Figure 7.10(b). The corresponding dynamic bit resolution, defined as

$$\text{BR} = \log_2\left(\frac{1}{\sigma}\right),$$

where σ denotes the standard deviation of the integration peaks relative to their mean value, is shown in Figure 7.10(c). PHIL achieves a mean resolution of approximately 5 bits across the tested power range, making it sufficient for analogue neural computing and inference applications[123].

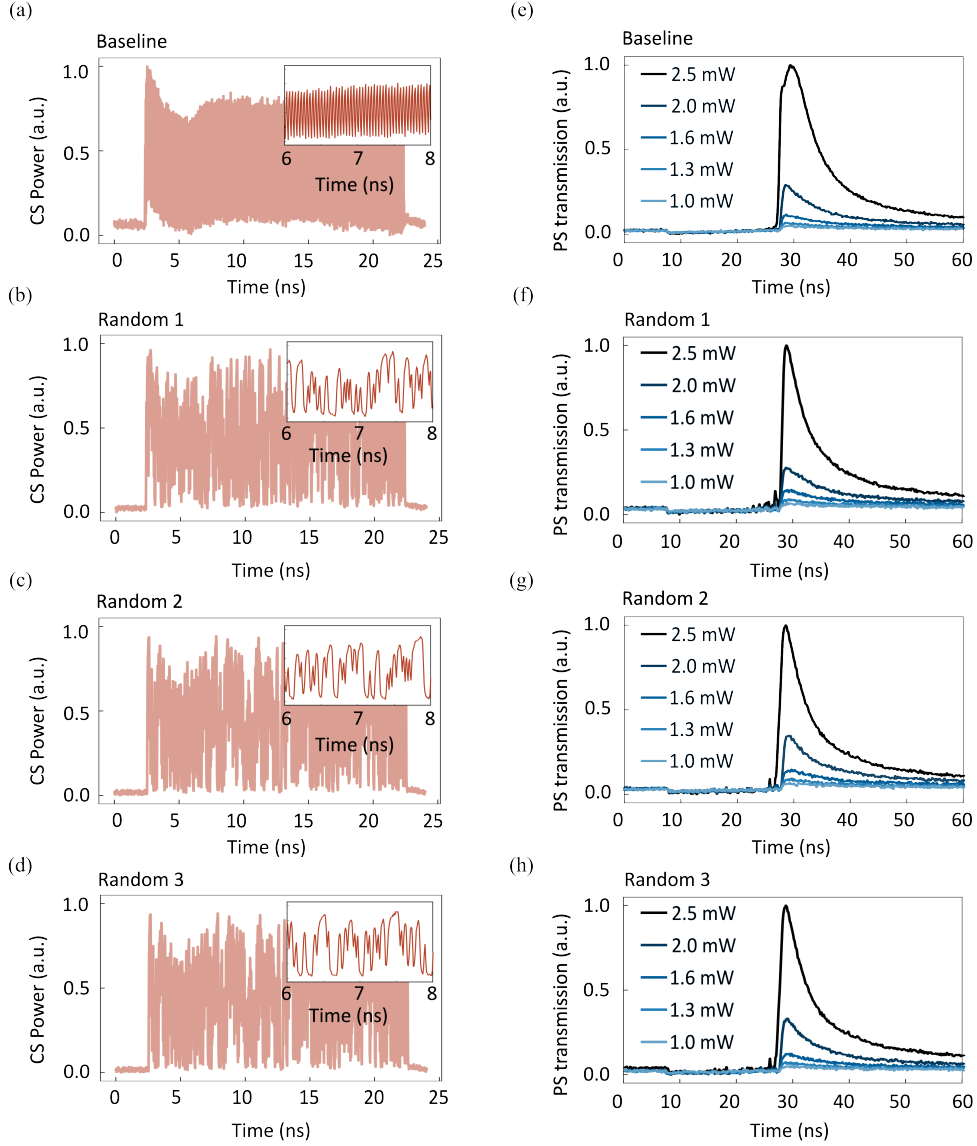


Figure 7.9: Experimental demonstration of optical temporal integration at 50 GHz. (a–d) Different CS pulse trains with a repetition rate of 50 GHz, each maintaining the same total number of logical ones and zeros (500 ones and 500 zeros) as the baseline signal. (e–h) The corresponding PS responses show the optical integration behaviour under varying average CS powers.

These results confirm that the PHIL device operates as a first-order thermal integrator, capable of performing precise, all-optical temporal integration of vectors up to size 1000 within a 20 ns window and nonlinear transformation in the time domain. Even for picosecond-scale optical excitations, the PHIL structure effectively

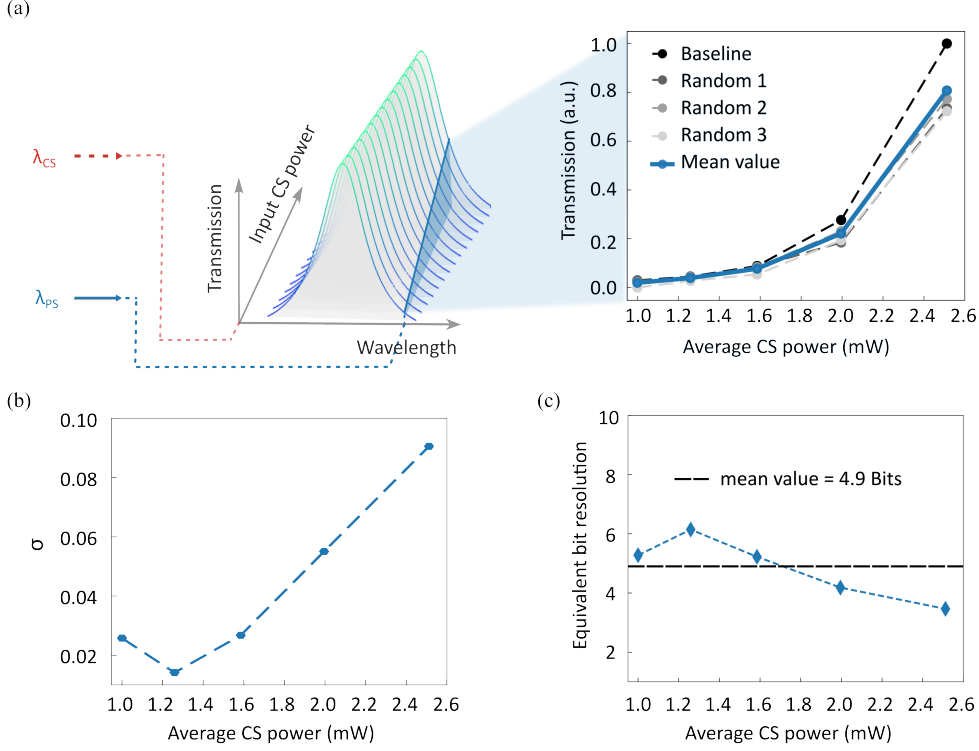


Figure 7.10: Nonlinear transfer characteristics and precision analysis of optical temporal integration. (a) Accumulated optical power of the CS is integrated over time through heat, resulting in a resonance shift of the PS spectrum due to the thermo-optic effect. The extracted integration peak values at the PS wavelength exhibit a ReLU-like nonlinear transfer function arising from the intrinsic Lorentzian response of the microring resonator. (b) Standard deviation of the integration-peak variations from the mean value at different CS input powers. (c) Corresponding dynamic bit resolution, calculated as $BR = \log_2(1/\sigma)$, showing an average precision of approximately 5 bits, which is sufficient for neuromorphic and AI computing applications.

integrates their energy within a sub-microsecond thermal window, demonstrating its potential for ultrafast all-optical temporal integration.

7.3.3 Neural Network Evaluation Using the PHIL Activation Function

To further demonstrate the computational relevance of the experimentally measured nonlinear response, I evaluated its applicability as an activation function

for neuromorphic inference tasks. Specifically, I implemented the PHIL-derived transfer function as a custom activation function within a neural network framework. The measured PHIL curve was first normalized to the range $[-0.2, 1.0]$ and fitted using a piecewise-linear differentiable function compatible with PyTorch (Figure 7.11(a)). This activation was then integrated into a fully connected neural network (784–256–128–10), corresponding to the MNIST handwritten-digit recognition task, replacing the conventional ReLU activation (Figure 7.11(b)).

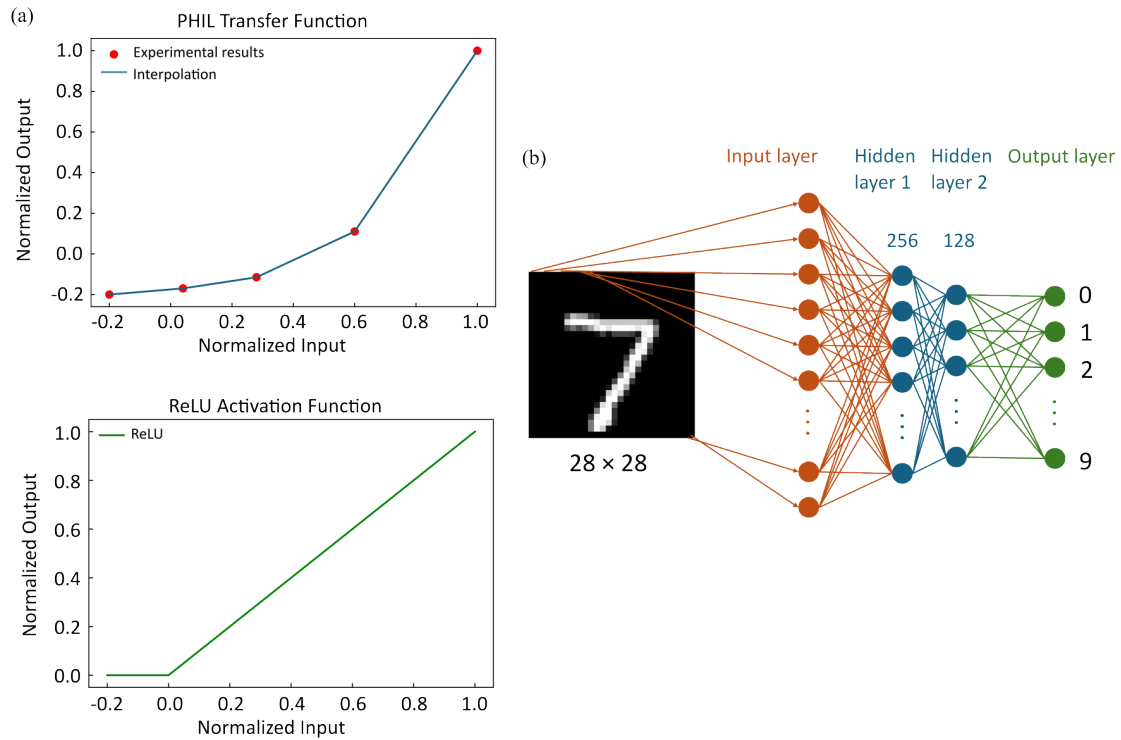


Figure 7.11: PHIL-Based Activation Function Implemented in a Neural Network for Handwritten Digit Recognition. (a) Top: Experimentally measured PHIL transfer curve (red circles) with interpolated nonlinear activation function (blue line) used in the network model. Bottom: Rectified linear unit (ReLU) activation function. (b) Architecture of the fully connected neural network (784–256–128–10) for MNIST classification, where the PHIL-based activation replaces the conventional ReLU function.

During training, the network was optimized using the Adam optimizer with a learning rate of 3×10^{-4} for 20 epochs and a batch size of 64. Input data were pre-normalized to match the activation range, ensuring a fair comparison between the PHIL-based and ReLU-based networks. The convergence behaviour and classifica-

tion performance were analysed under identical training conditions in Figure 7.12.

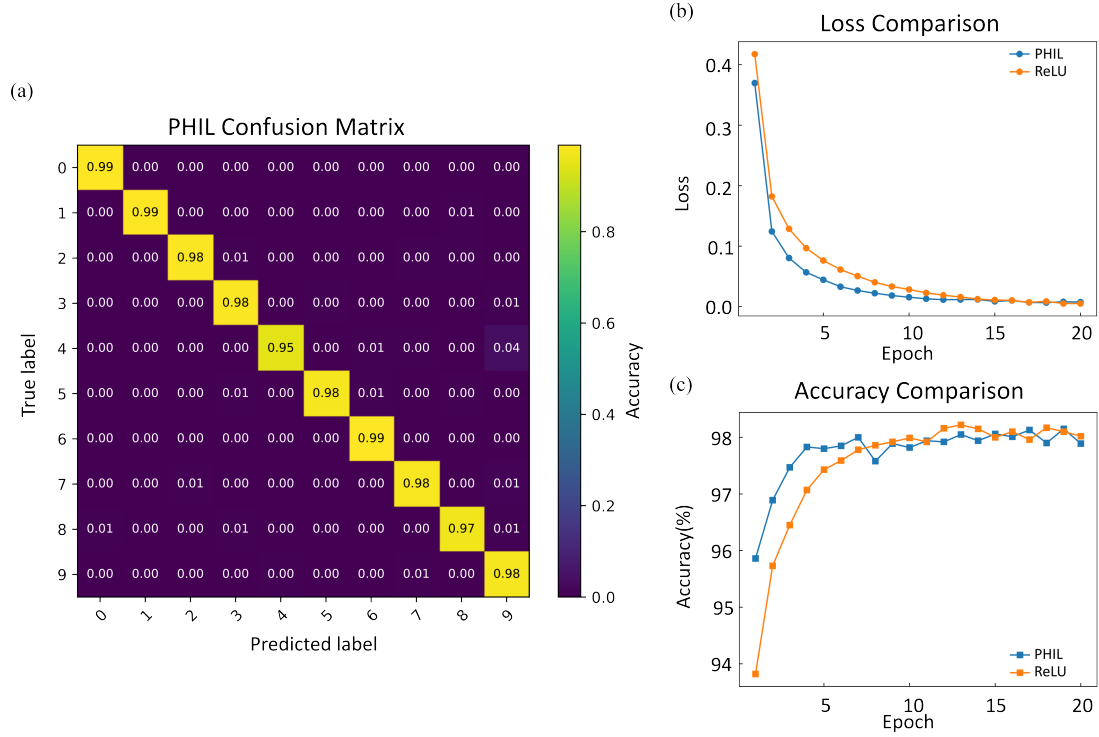


Figure 7.12: Performance Comparison Between PHIL- and ReLU-Activated Neural Networks for Handwritten Digit Recognition. (a) Confusion matrix of the PHIL-activated network showing accurate digit classification across all categories with a test accuracy of 97.2%, compared to 97.8% for the ReLU baseline. (b, c) Training loss and validation accuracy versus epoch for PHIL and ReLU activations, demonstrating nearly identical convergence behaviour.

The confusion matrix (Figure 7.12) demonstrates high classification fidelity across all ten digits, with uniformly strong diagonal elements and minimal inter-class confusion, indicating consistent nonlinear separability across input categories.

The training and validation curves (Figure 7.12(b, c)) reveal nearly identical convergence dynamics between the PHIL and ReLU activations. Both exhibit rapid loss minimization within the first few epochs and smooth accuracy saturation thereafter. The PHIL-activated network achieved a final test accuracy of 97.2%, comparable to 97.8% obtained with the standard ReLU activation, implying that

the PHIL-induced nonlinearity preserves stable gradient propagation and does not introduce vanishing or exploding gradients during back propagation. This close correspondence confirms that the optical transfer function implemented by the PHIL device is computationally equivalent to conventional electronic nonlinearities in practical inference tasks.

7.4 Energy Efficiency

While the preceding sections focused on demonstrating accurate temporal integration and nonlinear transformation entirely within the optical domain, a quantitative assessment of the optical energy consumption provides a more complete understanding of PHIL’s practical viability. In this section, I estimate the total energy required per integration window and normalize it to the equivalent multiply-accumulate (MAC) operations, enabling a direct comparison with other state-of-the-art all-optical computing platforms.

Energy Estimation

At a 50 GHz repetition rate, the pulse duration is

$$T_{\text{pulse}} = \frac{1}{50 \text{ GHz}} = 20 \text{ ps.} \quad (7.1)$$

Within a 20 ns integration window:

$$N_{\text{pulses}} = \frac{20 \text{ ns}}{20 \text{ ps}} = 1000 \text{ pulses.} \quad (7.2)$$

Assuming half of the pulses correspond to logical “1” and half to logical “0”, approximately 500 active pulses contribute to the energy accumulation. The nonlinear

activation is triggered at an optical power of ≤ 2.6 mW.

Energy per pulse

$$E_{\text{CS,pulse}} = 52 \text{ fJ.} \quad (7.3)$$

Hence, each “1” pulse costs approximately 52 fJ. The total control-signal energy within a 20 ns window is:

$$E_{\text{CS,total}} = 500 \times 52 \text{ fJ} = 26,000 \text{ fJ} = 26 \text{ pJ.} \quad (7.4)$$

The probe signal contributes a steady power of 0.1 mW across the integration window:

$$E_{\text{probe}} = 2 \text{ pJ.} \quad (7.5)$$

Therefore, the total energy per integration window is:

$$E_{\text{total}} = E_{\text{CS,total}} + E_{\text{probe}} = 28 \text{ pJ.} \quad (7.6)$$

Normalized Energy per MAC

Each 20 ns integration window accumulates approximately 500 operations (MAC equivalents). The normalized energy per MAC operation is therefore:

$$E_{\text{per-operation}} = \frac{28 \text{ pJ}}{500} = 0.056 \text{ pJ} = 56 \text{ fJ.} \quad (7.7)$$

7.4.1 Comparison of All-Optical Light-Controlled Technologies

As summarized in Table 7.1, the PHIL system uniquely performs temporal integration entirely in the optical domain while enabling programmable nonlinearity, as demonstrated in Chapter 6. This is a functionality not achievable by other all-optical technologies. Its opto-thermal design allows summation of thousands of ultrafast optical inputs over a 254 ns thermal window, making it inherently suitable for neuromorphic computing tasks requiring accumulation and thresholding over time.

Compared to PCMs, PHIL is volatile but enables real-time, analogue tunability. PCMs are non-volatile and ideal for static reconfiguration[84, 99, 100], yet their nanosecond-scale switching hinders high-speed integration. Relative to ultrafast Kerr or 2D material devices, PHIL trades off switching speed for the ability to perform analogue accumulation and nonlinear activation in a single passive platform, leveraging its slow-heat-based integration dynamics.

In summary, PHIL occupies a distinct niche among all-optical integrative computing schemes: it offers repeatable, energy-efficient accumulation with multi-bit analogue precision, balancing performance and tunability where other mechanisms prioritize either speed or non-volatility. While its thermal time constant (245 ns) is slower than sub-picosecond nonlinearities, integration occurs over this window independent of the incoming pulse rate, enabling compatibility with ultrafast (50 GHz) pulse streams and WDM operation, a regime difficult to access using carrier-dynamics-based photonic neuromorphic hardware.

Table 7.1: Comparison of All-Optical Light-Controlled Technologies

Technology	Energy per Operation	Speed (Response/BW)	Precision	Optical Temporal Integration + Nonlinear Activation	Ref
Photonic Heater-in-Lightpath (PHIL) Integrator	~ 56 fJ	~ 130 ns time constant (≈ 1.2 MHz bandwidth). Can integrate ultrafast (~ 50 GHz) optical pulses over this window	~ 5 -bit resolution demonstrated	Yes	This work
Phase-Change Material (PCM)	\sim tens of pJ	~ 10 – 100 ns optical pulse required for phase change	Multi-level analog phase possible (~ 5 -bit)	No	Li <i>et al.</i> , 2019[84]; Farmakidis <i>et al.</i> , 2019[99]; He <i>et al.</i> , 2024[100]
Free-Carrier Optical Injection	\sim hundreds of fJ	\sim hundreds of ps (carrier lifetime limited)	*	No	Preston <i>et al.</i> , 2008[134]; Shi <i>et al.</i> , 2022[97]
2D Material All-Optical Modulator	\sim tens of fJ	\sim hundreds of fs	*	No	Ono <i>et al.</i> , 2020[98]
Semiconductor Optical Amplifier (SOA)	~ 3 pJ (~ 100 pJ if electrical bias included)	\sim hundreds of ps	Sigmoid activation fitted to logistic function with high agreement (NRMSE < 0.08 , dynamic range ~ 27 dB for cross-connect networks)	Yes—but with electrical bias	Mourgias-Alexandris <i>et al.</i> , 2019[95]; Kravtsov <i>et al.</i> , 2011[135]; Shi <i>et al.</i> , 2020[136]

7.5 Discussion and Summary

Building upon the cross-wavelength accumulation and activation demonstrated in the previous chapter, this work extends the PHIL architecture into the time domain, establishing a unified platform for all-optical computation across both wavelength and temporal dimensions. By leveraging the intrinsic thermal relaxation dynamics of the cavity, the device performs analogue temporal integration of high-speed optical pulse streams without any electronic assistance. The experimentally extracted time constant of approximately 245 ns defines an effective integration window over which successive optical events are accumulated, thereby mapping ultrafast optical inputs (at repetition rates up to 50 GHz) onto a slow, thermally mediated response.

Although the thermal relaxation time is on the order of hundreds of nanoseconds, it does not limit the optical processing bandwidth of the system. The optical inputs operate at repetition rates of up to 50 GHz, corresponding to pulse intervals of approximately 20 ps, meaning that thousands of optical events can be accumulated within a single thermal integration window. From a system perspective, this time scale remains substantially shorter than the latency requirements of most real-time AI applications, which typically operate at microsecond-to-millisecond decision times[137–139]. Therefore, the thermal dynamics primarily define the effective memory depth of the photonic leaky integrator rather than the intrinsic computation speed of the optical signals.

The dynamic experiments confirm that the device output follows an exponential decay consistent with the analytical leaky-integrator model, and the integration peak scales monotonically with pulse number and average power. This behaviour verifies that the thermo-optic nonlinearity, traditionally regarded as a slow parasitic effect, can be harnessed as a functional mechanism for analogue temporal computing. The

experimentally derived activation curve further exhibits a ReLU-like nonlinearity, which was validated through neural-network simulations achieving a 97.2 % recognition accuracy on the MNIST dataset—comparable to standard electronic ReLU activations. These results demonstrate that the PHIL-based integrator can perform energy accumulation and nonlinear transformation entirely in the optical domain, offering a passive yet reconfigurable analogue computing primitive.

The energy and speed characteristics of the system were benchmarked against other all-optical computing platforms (Table 7.1). The PHIL device achieves an energy cost of approximately 56 fJ per multiply–accumulate (MAC) operation while maintaining compatibility with 50 GHz optical pulse trains, representing a favourable balance between efficiency and precision. This combination of tunable analogue response, multi-bit precision, and passive operation distinguishes PHIL as a versatile building block for photonic neuromorphic architectures.

Although the thermal response could limit the maximum neuron state update rate to the megahertz range, this constraint can be mitigated if faster dynamics are required for specific applications. In principle, the thermal time constant can be engineered through device geometry, heater dimensions, or improved thermal isolation, providing a pathway to tune the trade-off between integration memory depth and response speed.

In summary, this chapter demonstrates all-optical temporal integration as a natural extension of the PHIL framework, unifying wavelength- and time-domain computation within a single device. The results highlight the feasibility of performing sequential analogue processing through purely optical means, bridging the gap between ultrafast signal dynamics and slow thermal memory. Future developments will focus on cascading temporally integrative units to implement recurrent or reservoir architectures, optimising the thermal design for faster response while

preserving analogue fidelity, and integrating wavelength multiplexing for concurrent spatial-temporal processing. Together, these capabilities advance the PHIL platform toward a scalable and energy-efficient foundation for end-to-end optical neuromorphic computing.

8

Conclusions and Outlook

Contents

8.1 Conclusions	164
8.2 Outlook	167

8.1 Conclusions

This thesis has established and experimentally validated a unified framework for all-optical information processing based on the Photonic-Heater-in-Lightpath architecture. By embedding subwavelength absorbers directly within standing-wave resonant fields, this design converts heat dissipation into programmable thermo-optic modulation, enabling incoherent addition, nonlinear activation, and temporal integration within an entirely optical signal pathway. The work reinterprets the thermo-optic effect, from a parasitic delay mechanism to a functional nonlinearity governing optical computation.

All-Optical Cross-Wavelength Encoding

Chapter 5 established the physical foundation of PHIL by demonstrating wavelength-selective absorption within a standing-wave resonator. Counter-propagating optical modes in a silicon MRR generated spatially fixed nodes and antinodes. Titanium nanoheaters were aligned to antinodes of selected resonances, producing controlled

optical absorption that induced a broadband thermo-optic phase shift. A weak probe signal travelling through field minima experienced minimal loss while sensing the induced refractive-index change. This effect provided end-to-end optical modulation across spectrally separated wavelengths, forming the basis of optically driven signal control and cross-channel encoding without E-O conversion.

All-Optical Summation and Programmable Activation Across Wavelengths

Chapter 6 extended the PHIL concept to multi-wavelength accumulation and spectrally engineered nonlinearities. Multiple control wavelengths were injected simultaneously at absorptive resonances, their total optical powers accumulated incoherently in the form of heat. The shared probe channel captured this cumulative modulation, realising all-optical summation across WDM channels. By positioning the probe wavelength along the Lorentzian resonance, different nonlinear transfer functions were generated within the same passive unit. These programmable nonlinearities were applied directly to the summed optical input, demonstrating optical accumulation and nonlinear activation within a single element.

All-Optical Temporal Integration.

Chapter 7 demonstrated cross-time accumulation by exploiting the intrinsic thermal relaxation of the device as an integration window. Ultrafast control pulses at a 50 GHz repetition rate were passively integrated over a leaky time constant of approximately ~ 250 ns, producing a leaky-integrator response that accumulates optical energy rather than instantaneous intensity. The integrated probe output exhibited a nonlinear transfer curve determined by the resonator lineshape, thus performing both temporal summation and nonlinear activation simultaneously. Because integration depends on total absorbed energy rather than instantaneous thermal switching speed, the device naturally bridges picosecond optical dynamics with MHz-scale thermal evolution while maintaining ~ 5 -bit analogue precision

and energy efficiency near 56 fJ per equivalent MAC operation. Network-level simulations further verified that the experimentally measured PHIL activation reproduces electronic ReLU behaviour, achieving comparable inference accuracy on benchmark tasks while remaining optically driven.

The temporal response of PHIL is therefore governed by the thermal relaxation time of the device, measured here to be approximately ~ 250 ns. This timescale is slower than the intrinsic switching speeds of electronic logic gates or electro-optic modulators, which typically operate in the sub-nanosecond regime. However, the PHIL architecture compensates for this slower physical nonlinearity through the inherent parallelism of photonics. WDM and TDM enable many optical channels to be processed simultaneously within a single resonator, allowing high aggregate throughput even when the local nonlinear response evolves on thermal timescales.

Consequently, the architecture trades per-operation speed for parallel optical bandwidth and passive hardware simplicity, offering a practical route toward energy-efficient optical neural processing without electronic activation stages.

Device-Level Metrics and Scalability.

Comprehensive characterization confirmed that the PHIL architecture meets the requirements for scalable, low-loss, and thermally stable photonic computing. The measured probe-band insertion loss was approximately 1.7 dB, while absorption at the control wavelengths introduced a loss of around 6 dB, providing strong attenuation contrast without compromising transparency for non-absorptive channels. Long-term open-loop operation demonstrated exceptional spectral stability, with resonance drift below $\pm 1.7\%$ over an hour of continuous modulation, confirming the robustness of the opto-thermal mechanism and its immunity to cumulative thermal drift.

Furthermore, scalability across multiple WDM channels was validated both experimentally and numerically: two- and three-wavelength summation showed linear additive behaviour with near-Gaussian error distribution, corresponding to an effective precision of 4–5 bits. Finite-difference thermal simulations revealed that heat diffusion remains confined within $\sim 2\,\mu\text{m}$ around the heater, ensuring minimal thermal crosstalk between channels and compatibility with high-density photonic integration. These combined results establish that PHIL can be densely tiled within foundry-scale silicon photonic platforms while retaining its analogue accuracy and low-loss performance.

8.2 Outlook

The results of this work collectively redefine the role of thermo-optic processes in integrated photonics. Traditionally considered a speed-limiting mechanism, thermal dynamics are here transformed into a computational resource, enabling accumulation, nonlinear activation, and analogue memory directly in the optical domain. This approach opens a path toward photonic devices that unite high bandwidth with controllable nonlinear responses within one compact, passive structure.

A natural next step is the integration of the PHIL element with programmable photonic weight matrices, such as PCM-based crossbar arrays. In existing photonic accelerators, matrix multiplication can already be performed optically, but nonlinear activation typically requires electrical or electro-optic conversion. Embedding PHIL units at the outputs of photonic weighting networks would provide passive, reconfigurable nonlinearities without disrupting the optical signal flow. Such integration would complete a fully optical processing pipeline, enabling matrix multiplication, accumulation, and activation, to be performed within a single

photonic layer.

Because the PHIL mechanism offers analogue tunability, low insertion loss, and compatibility with WDM and TDM, it can be seamlessly embedded in hybrid photonic tensor-core architectures. In this framework, multiplexed optical carriers perform matrix operations in parallel across wavelengths, while sequential temporal channels handle high-throughput data streams. PHIL units positioned after each weighted summation node can simultaneously provide nonlinear activation and analogue accumulation across both multiplexing domains, enabling large-scale, energy-efficient, and fully optical neural inference.

Beyond computing, the passive and wavelength-selective nature of PHIL opens opportunities in photonic interconnects and cross-domain signal modulation. Its ability to selectively respond, accumulate and remap information from multiple optical carriers enables local processing functions within larger optical networks, such as cross-wavelength modulation, spectral conditioning[140], or power-aware routing[141, 142]. The low insertion loss and strong thermal confinement demonstrated in this work ensure that PHIL elements can be densely deployed as wavelength-aware interconnects or adaptive routing nodes, supporting large-scale, reconfigurable optical networks.

Beyond the silicon platform, the underlying principle of spatially localized absorption and thermo-optic coupling can be extended to diverse material systems. Implementations on lithium niobate could combine fast electro-optic modulation with PHIL’s analogue accumulation, enabling hybrid EO–TO computing cores. Germanium–silicon or silicon carbide platforms could further extend operation to mid-infrared and visible spectral regions. The structural simplicity and foundry compatibility of the PHIL design make it adaptable across these materials, posi-

tioning it as a universal nonlinear building block for integrated photonic technologies.

In summary, the PHIL architecture offers a practical and versatile route to optical neural processors that merge the high parallelism of WDM and the high bandwidth of TDM with optically programmable nonlinearities. Its integration into photonic weight matrices and interconnect fabrics could enable truly end-to-end optical computation, spanning weighting, accumulation, and nonlinear activation on a single, scalable photonic platform.

List of Publications

Journal Articles

1. **Zhang, Y.** *et al.* All-optical temporal integration mediated by subwavelength heat antennas. *Nat Commun* **17**, 989 (2026).
2. He, Y.†, Farmakidis, N.†, Aggarwal, S., Dong, B., Lee, J. S., Wang, M., **Zhang, Y.**, Parmigiani, F. & Bhaskaran, H. Energy-efficient integrated electro-optic memristors. *Nano Lett.* **24**, 16325–16332 (2024).

Conference Presentations

1. **Zhang, Y.** *et al.* Bosonic control in integrated photonics. In *Proc. SPIE* PC12889, Integrated Optics: Devices, Materials, and Technologies XXVIII, PC128890R (2024).
2. **Zhang, Y.** *et al.* Coupling multiwavelength signals using light on an integrated photonics platform. In *Proc. SPIE* PC13110, Active Photonic Platforms (APP) 2024, PC131101B (2024) (**invited**).
3. He, Y., Farmakidis, N., Aggarwal, S., Lee, J. S., Wang, M., Dong, B., Shu, Y., **Zhang, Y.**, Parmigiani, F. & Bhaskaran, H. Electro-optical phase-change devices approaching single pico-joule switching. In *Proc. SPIE* PC12889, Integrated Optics: Devices, Materials, and Technologies XXVIII, PC128890A (2024).
4. Zhong, S., Ortega-Gomez, A., Dong, B., Farmakidis, N., Lee, J. S., Sabnani, S., Wang, M., Du, M., **Zhang, Y.**, Aggarwal, S., Yang, G. & Bhaskaran, H. Photonic probabilistic computing with arbitrary distributions. In *CLEO 2025*, paper AA128-4 (2025).

Patents

1. Bhaskaran, H., **Zhang, Y.**, Farmakidis, N. & Lee, J. S. All-optical functional unit. US patent (US20250208481A1, 2025).

References

1. Von Neumann, J. First draft of a report on the EDVAC. *IEEE Annals of the History of Computing* **15**, 27–75 (1993).
2. Mehonic, A. & Kenyon, A. J. Brain-inspired computing needs a master plan. en. *Nature* **604**, 255–260 (Apr. 2022).
3. Big data needs a hardware revolution. en. *Nature* **554**. Bandiera__abtest: a Cg_type: Editorial Number: 7691 Subject_term: Computer science, Mathematics and computing, 145–146 (Feb. 2018).
4. Does AI have a hardware problem? en. *Nature Electronics* **1**. Number: 4, 205–205 (Apr. 2018).
5. Mead, C. Neuromorphic electronic systems. *Proceedings of the IEEE* **78**. Conference Name: Proceedings of the IEEE, 1629–1636 (Oct. 1990).
6. Madhavan, A. Brain-Inspired Computing Can Help Us Create Faster, More Energy-Efficient Devices —If We Win the Race. en. *NIST*. Last Modified: 2025-02-11T12:57:05:00 (Mar. 2023).
7. *When will computer hardware match the human brain? by Hans Moravec*
8. Roy, K., Jaiswal, A., *et al.* Towards spike-based machine intelligence with neuromorphic computing. en. *Nature* **575**. Number: 7784, 607–617 (Nov. 2019).
9. McCulloch, W. S. & Pitts, W. A logical calculus of the ideas immanent in nervous activity. en. *The bulletin of mathematical biophysics* **5**, 115–133 (Dec. 1943).
10. Hebb, D. O. *The Organization of Behavior: A Neuropsychological Theory* en. Google-Books-ID: dZ0eDiLTwuEC (Wiley, 1949).
11. Torres, F., Basaran, A. C., *et al.* Thermal Management in Neuromorphic Materials, Devices, and Networks. en. *Advanced Materials* **35**. _eprint: <https://advanced.onlinelibrary.wiley.com/doi/pdf/10.1002/adma.202205098>, 2205098 (2023).
12. Hoffmann, A., Ramanathan, S., *et al.* Quantum materials for energy-efficient neuromorphic computing: Opportunities and challenges. *APL Materials* **10**, 070904 (July 2022).
13. Ovshinsky, S. R. Reversible Electrical Switching Phenomena in Disordered Structures. *Physical Review Letters* **21**, 1450–1453 (Nov. 1968).
14. Zhang, W., Deringer, V. L., *et al.* Density-functional theory guided advances in phase-change materials and memories. en. *MRS Bulletin* **40**, 856–869 (Oct. 2015).
15. Zhang, W. & Ma, E. Single-element glass to record data. en. *Nature Materials* **17**. Number: 8, 654–655 (Aug. 2018).
16. Zhang, W., Mazzarello, R., *et al.* Designing crystallization in phase-change materials for universal memory and neuro-inspired computing. en. *Nature Reviews Materials* **4**. Number: 3, 150–168 (Mar. 2019).
17. published, N. M. *Intel Announces Optane DC Persistent Memory Is Sampling Now, With Broad Availability In 2019* en. May 2018.

18. Sebastian, A., Le Gallo, M., *et al.* Tutorial: Brain-inspired computing using phase-change memory devices. *Journal of Applied Physics* **124**, 111101 (Sept. 2018).
19. Ambs, P. Optical Computing: A 60Year Adventure. *Advances in Optical Technologies* **2010**, 1–16 (Jan. 2010).
20. Prucnal, P. R., Shastri, B. J., *et al.* *Neuromorphic Photonics* (CRC Press, Boca Raton, May 2017).
21. Feldmann, J., Youngblood, N., *et al.* All-optical spiking neurosynaptic networks with self-learning capabilities. en. *Nature* **569**, 208–214 (May 2019).
22. Tait, A. N., Wu, A. X., *et al.* Microring Weight Banks. *IEEE Journal of Selected Topics in Quantum Electronics* **22**. Conference Name: IEEE Journal of Selected Topics in Quantum Electronics, 312–325 (Nov. 2016).
23. Marin-Palomo, P., Kemal, J. N., *et al.* Microresonator-based solitons for massively parallel coherent optical communications. en. *Nature* **546**, 274–279 (June 2017).
24. Bai, B., Yang, Q., *et al.* Microcomb-based integrated photonic processing unit. en. *Nature Communications* **14**, 66 (Jan. 2023).
25. Lee, J. S., Farmakidis, N., *et al.* Polarization-selective reconfigurability in hybridized-active-dielectric nanowires. *Science Advances* **8**, eabn9459 (June 2022).
26. Zhang, Y., Zhang, R., *et al.* Architecture and Devices for Silicon Photonic Switching in Wavelength, Polarization and Mode. *Journal of Lightwave Technology* **38**. Conference Name: Journal of Lightwave Technology, 215–225 (Jan. 2020).
27. Meng, X., Zhang, G., *et al.* Compact optical convolution processing unit based on multimode interference. en. *Nature Communications* **14**, 3000 (May 2023).
28. Yamaguchi, T., Arai, K., *et al.* Time-domain photonic image processor based on speckle projection and reservoir computing. en. *Communications Physics* **6**, 1–10 (Sept. 2023).
29. Giamougiannis, G., Tsakyridis, A., *et al.* Neuromorphic silicon photonics with 50 GHz tiled matrix multiplication for deep-learning applications. *Advanced Photonics* **5**, 016004 (Feb. 2023).
30. Nahmias, M. A., de Lima, T. F., *et al.* Photonic Multiply-Accumulate Operations for Neural Networks. en. *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1–18 (Jan. 2020).
31. Feldmann, J., Youngblood, N., *et al.* Parallel convolutional processing using an integrated photonic tensor core. en. *Nature* **589**. Number: 7840, 52–58 (Jan. 2021).
32. Miller, D. A. B. Device Requirements for Optical Interconnects to Silicon Chips. *Proceedings of the IEEE* **97**, 1166–1185 (July 2009).
33. Sun, C., Wade, M. T., *et al.* Single-chip microprocessor that communicates directly using light. en. *Nature* **528**, 534–538 (Dec. 2015).

34. Miller, D. A. B. Attojoule Optoelectronics for Low-Energy Information Processing and Communications. *Journal of Lightwave Technology* **35**, 346–396 (Feb. 2017).
35. Shen, Y., Harris, N. C., *et al.* Deep learning with coherent nanophotonic circuits. en. *Nature Photonics* **11**, 441–446 (July 2017).
36. LeCun, Y., Bengio, Y., *et al.* Deep learning. en. *Nature* **521**, 436–444 (May 2015).
37. Silver, D., Huang, A., *et al.* Mastering the game of Go with deep neural networks and tree search. en. *Nature* **529**, 484–489 (Jan. 2016).
38. Zhang, H., Gu, M., *et al.* An optical neural chip for implementing complex-valued neural network. en. *Nature Communications* **12**, 457 (Jan. 2021).
39. Amin, R., George, J. K., *et al.* ITO-based electro-absorption modulator for photonic neural activation function. *APL Materials* **7**, 081112 (Aug. 2019).
40. Nahmias, M. A., Tait, A. N., *et al.* An integrated analog O/E/O link for multi-channel laser neurons. en. *Applied Physics Letters* **108**, 151106 (Apr. 2016).
41. George, J. K., Mehrabian, A., *et al.* Neuromorphic photonics with electro-absorption modulators. EN. *Optics Express* **27**, 5181–5191 (Feb. 2019).
42. Ashtiani, F., Geers, A. J., *et al.* An on-chip photonic deep neural network for image classification. en. *Nature* **606**. Number: 7914, 501–506 (June 2022).
43. Lin, Z., Shastri, B. J., *et al.* 120 GOPS Photonic tensor core in thin-film lithium niobate for inference and in situ training. en. *Nature Communications* **15**, 9081 (Oct. 2024).
44. Sludds, A., Bandyopadhyay, S., *et al.* Delocalized photonic deep learning on the internet’ s edge. *Science* **378**, 270–276 (Oct. 2022).
45. Leuthold, J., Koos, C., *et al.* Nonlinear silicon photonics. en. *Nature Photonics* **4**. Number: 8, 535–544 (Aug. 2010).
46. Almeida, V. R., Barrios, C. A., *et al.* All-optical control of light on a silicon chip. en. *Nature* **431**, 1081–1084 (Oct. 2004).
47. Xu, Q. & Lipson, M. Carrier-induced optical bistability in silicon ring resonators. en. *Optics Letters* **31**, 341 (2006).
48. Boyd, R. W. *Nonlinear optics* 3rd ed. eng. Open Library ID: OL16843330M (Academic Press, Amsterdam, Boston, 2008).
49. Chen, B., Bruck, R., *et al.* Hybrid Photon–Plasmon Coupling and Ultrafast Control of Nanoantennas on a Silicon Photonic Chip. *Nano Letters* **18**, 610–617 (Jan. 2018).
50. Lee, J. S., Farmakidis, N., *et al.* Spatio-spectral control of coherent nanophotonics. en. *Nanophotonics* (Jan. 2024).
51. Shastri, B. J., Tait, A. N., *et al.* Photonics for artificial intelligence and neuromorphic computing. en. *Nature Photonics* **15**. Number: 2, 102–114 (Feb. 2021).
52. Reck, M., Zeilinger, A., *et al.* Experimental realization of any discrete unitary operator. *Physical Review Letters* **73**, 58–61 (July 1994).

53. Miller, D. A. B. Self-configuring universal linear optical component [Invited]. EN. *Photonics Research* **1**, 1–15 (June 2013).
54. Carolan, J., Harrold, C., *et al.* Universal linear optics. *Science* **349**, 711–716 (Aug. 2015).
55. Bandyopadhyay, S., Sludds, A., *et al.* Single-chip photonic deep neural network with forward-only training. en. *Nature Photonics* **18**, 1335–1343 (Dec. 2024).
56. Hughes, T. W., Minkov, M., *et al.* Training of photonic neural networks through in situ backpropagation and gradient measurement. EN. *Optica* **5**, 864–871 (July 2018).
57. Pai, S., Sun, Z., *et al.* Experimentally realized in situ backpropagation for deep learning in photonic neural networks. *Science* **380**, 398–404 (Apr. 2023).
58. Shokraneh, F., Geoffroy-Gagnon, S., *et al.* Towards Phase-Error- and Loss-Tolerant Programmable MZI-Based Optical Processors for Optical Neural Networks in 2020 IEEE Photonics Conference (IPC) (Sept. 2020), 1–2.
59. Hamerly, R., Bandyopadhyay, S., *et al.* Robust Zero-Change Self-Configuration of the Rectangular Mesh in 2021 Optical Fiber Communications Conference and Exhibition (OFC) (June 2021), 1–3.
60. Peserico, N., Shastri, B. J., *et al.* Integrated Photonic Tensor Processing Unit for a Matrix Multiply: A Review. *Journal of Lightwave Technology* **41**, 3704–3716 (June 2023).
61. Zhu, H. H., Zou, J., *et al.* Space-efficient optical computing with an integrated chip diffractive neural network. en. *Nature Communications* **13**, 1044 (Feb. 2022).
62. Feng, C., Gu, J., *et al.* A Compact Butterfly-Style Silicon Photonic–Electronic Neural Chip for Hardware-Efficient Deep Learning. *ACS Photonics* **9**, 3906–3916 (Dec. 2022).
63. Farmakidis, N., Dong, B., *et al.* Integrated photonic neuromorphic computing: opportunities and challenges. en. *Nature Reviews Electrical Engineering* (June 2024).
64. Lin, X., Rivenson, Y., *et al.* All-optical machine learning using diffractive deep neural networks. *Science* **361**, 1004–1008 (Sept. 2018).
65. Qian, C., Lin, X., *et al.* Performing optical logic operations by a diffractive neural network. en. *Light: Science & Applications* **9**, 59 (Apr. 2020).
66. Zhou, T., Lin, X., *et al.* Large-scale neuromorphic optoelectronic computing with a reconfigurable diffractive processing unit. en. *Nature Photonics* **15**, 367–373 (May 2021).
67. Wang, Z., Li, T., *et al.* On-chip wavefront shaping with dielectric metasurface. en. *Nature Communications* **10**, 3547 (Aug. 2019).
68. Zarei, S., Marzban, M.-r., *et al.* Integrated photonic neural network based on silicon metalines. EN. *Optics Express* **28**, 36668–36684 (Nov. 2020).
69. Fu, T., Zang, Y., *et al.* On-chip photonic diffractive optical neural network based on a spatial domain electromagnetic propagation model. EN. *Optics Express* **29**, 31924–31940 (Sept. 2021).

70. Wang, Z., Chang, L., *et al.* Integrated photonic metasystem for image classifications at telecommunication wavelength. en. *Nature Communications* **13**, 2131 (Apr. 2022).
71. Yan, T., Yang, R., *et al.* All-optical graph representation learning using integrated diffractive photonic computing units. *Science Advances* **8**, eabn7630 (June 2022).
72. Fu, T., Zang, Y., *et al.* Photonic machine learning with on-chip diffractive optics. en. *Nature Communications* **14**, 70 (Jan. 2023).
73. Wu, T., Menarini, M., *et al.* Lithography-free reconfigurable integrated photonic processor. en. *Nature Photonics* **17**, 710–716 (Aug. 2023).
74. Xu, Z., Zhou, T., *et al.* Large-scale photonic chiplet Taichi empowers 160-TOPS/W artificial general intelligence. *Science* **384**, 202–209 (Apr. 2024).
75. Ji, R., Zhang, L., *et al.* On-chip CMOS-compatible optical signal processor in 2012 Asia Communications and Photonics Conference (ACP) (Nov. 2012), 1–3.
76. Tait, A. N., Nahmias, M. A., *et al.* Broadcast and Weight: An Integrated Network For Scalable Photonic Spike Processing. *Journal of Lightwave Technology* **32**. Conference Name: Journal of Lightwave Technology, 4029–4041 (Nov. 2014).
77. Tait, A. N., Ferreira de Lima, T., *et al.* Continuous Calibration of Microring Weights for Analog Optical Networks. *IEEE Photonics Technology Letters* **28**, 887–890 (Apr. 2016).
78. Tait, A. N., Lima, T. F. d., *et al.* Multi-channel control for microring weight banks. EN. *Optics Express* **24**, 8895–8906 (Apr. 2016).
79. Tait, A. N., de Lima, T. F., *et al.* Neuromorphic photonic networks using silicon photonic weight banks. en. *Scientific Reports* **7**, 7430 (Aug. 2017).
80. Pintus, P., Dumont, M., *et al.* Integrated non-reciprocal magneto-optics with ultra-high endurance for photonic in-memory computing. en. *Nature Photonics* (Oct. 2024).
81. Farmakidis, N., Yu, H., *et al.* Scalable High-Precision Trimming of Photonic Resonances by Polymer Exposure to Energetic Beams. *Nano Letters* **23**, 4800–4806 (June 2023).
82. Chen, R., Fang, Z., *et al.* Non-volatile electrically programmable integrated photonics with a 5-bit operation. en. *Nature Communications* **14**, 3465 (June 2023).
83. Ríos, C., Youngblood, N., *et al.* In-memory computing on a photonic platform. *Science Advances* **5**, eaau5759 (Feb. 2019).
84. Li, X., Youngblood, N., *et al.* Fast and reliable storage using a 5 bit, nonvolatile photonic memory cell. EN. *Optica* **6**, 1–6 (Jan. 2019).
85. Zhou, W., Dong, B., *et al.* In-memory photonic dot-product engine with electrically programmable weight banks. en. *Nature Communications* **14**. Number: 1, 2887 (May 2023).
86. Dong, B., Brückerohoff-Plückelmann, F., *et al.* Partial coherence enhances parallelized photonic computing. en. *Nature* **632**, 55–62 (Aug. 2024).

87. Brücknerhoff-Plückelmann, F., Bente, I., *et al.* Hybrid Electro-Optic Crossbar Array for Matrix-Vector Multiplications in 2023 Conference on Lasers and Electro-Optics (CLEO) (May 2023), 1–2.
88. Li, X., Youngblood, N., *et al.* On-chip Phase Change Optical Matrix Multiplication Core en. in 2020 IEEE International Electron Devices Meeting (IEDM) (IEEE, San Francisco, CA, USA, Dec. 2020), 7.5.1–7.5.4.
89. Aggarwal, S., Dong, B., *et al.* Reduced rank photonic computing accelerator. *Optica* **10**. ADS Bibcode: 2023Optic..10.1074A, 1074 (Aug. 2023).
90. Xu, X., Tan, M., *et al.* 11 TOPS photonic convolutional accelerator for optical neural networks. en. *Nature* **589**, 44–51 (Jan. 2021).
91. Dong, B., Aggarwal, S., *et al.* Higher-dimensional processing using a photonic tensor core with continuous-time data. en. *Nature Photonics* **17**, 1080–1088 (Dec. 2023).
92. Meng, X., Shi, N., *et al.* High-integrated photonic tensor core utilizing high-dimensional lightwave and microwave multidomain multiplexing. en. *Light: Science & Applications* **14**, 27 (Jan. 2025).
93. Williamson, I. A. D., Hughes, T. W., *et al.* Reprogrammable Electro-Optic Nonlinear Activation Functions for Optical Neural Networks. en. *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1–12 (Jan. 2020).
94. Tait, A. N., Ferreira de Lima, T., *et al.* Silicon Photonic Modulator Neuron. *Physical Review Applied* **11**, 064043 (June 2019).
95. Mourgias-Alexandris, G., Tsakyridis, A., *et al.* An all-optical neuron with sigmoid activation function. EN. *Optics Express* **27**, 9620–9630 (Apr. 2019).
96. Jha, A., Huang, C., *et al.* Reconfigurable all-optical nonlinear activation functions for neuromorphic photonics. en. *Optics Letters* **45**, 4819 (Sept. 2020).
97. Shi, Y., Ren, J., *et al.* Nonlinear germanium-silicon photodiode for activation and monitoring in photonic neuromorphic networks. en. *Nature Communications* **13**, 6048 (Oct. 2022).
98. Ono, M., Hata, M., *et al.* Ultrafast and energy-efficient all-optical switching with graphene-loaded deep-subwavelength plasmonic waveguides. en. *Nature Photonics* **14**, 37–43 (Jan. 2020).
99. Farmakidis, N., Youngblood, N., *et al.* Plasmonic nanogap enhanced phase-change devices with dual electrical-optical functionality. *Science Advances* **5**, eaaw2687 (Nov. 2019).
100. He, Y., Farmakidis, N., *et al.* Energy-Efficient Integrated Electro-Optic Memristors. *Nano Letters* **24**, 16325–16332 (Dec. 2024).
101. Chong, Y. D., Ge, L., *et al.* Coherent Perfect Absorbers: Time-Reversed Lasers. *Physical Review Letters* **105**, 053901 (July 2010).
102. Wan, W., Chong, Y., *et al.* Time-Reversed Lasing and Interferometric Control of Absorption. *Science* **331**, 889–892 (Feb. 2011).
103. Baranov, D. G., Krasnok, A., *et al.* Coherent perfect absorbers: linear control of light with light. en. *Nature Reviews Materials* **2**, 17064 (Oct. 2017).

104. Zhang, J., MacDonald, K. F., *et al.* Controlling light-with-light without nonlinearity. en. *Light: Science & Applications* **1**, e18–e18 (July 2012).
105. Bruck, R. & Muskens, O. L. Plasmonic nanoantennas as integrated coherent perfect absorbers on SOI waveguides for modulators and all-optical switches. EN. *Optics Express* **21**, 27652–27661 (Nov. 2013).
106. Okamoto, K. in *Fundamentals of Optical Waveguides (Second Edition)* (ed Okamoto, K.) Second Edition, 13–55 (Academic Press, Burlington, 2006).
107. Mashanovich, G. Z., Reed, G. T., *et al.* en. in *Silicon Photonics Section: 2* _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9780470994535.ch2>, 15–46 (John Wiley & Sons, Ltd, 2008).
108. Cheng, L., Mao, S., *et al.* Grating Couplers on Silicon Photonics: Design Principles, Emerging Trends and Practical Issues. en. *Micromachines* **11**. Number: 7, 666 (July 2020).
109. Bozzola, A., Carroll, L., *et al.* Optimising apodized grating couplers in a pure SOI platform to -0.5 dB coupling efficiency. eng. *Optics Express* **23**, 16289–16304 (June 2015).
110. Ríos Ocampo, C. A. *Phase-change materials for photonic memories and optoelectronic applications* English. <http://purl.org/dc/dcmitype/Text> (University of Oxford, 2016).
111. Gehring, H., Eich, A., *et al.* Broadband out-of-plane coupling at visible wavelengths. EN. *Optics Letters* **44**, 5089–5092 (Oct. 2019).
112. en. in *Integrated Ring Resonators: The Compendium* (ed Rabus, D. G.) 3–40 (Springer, Berlin, Heidelberg, 2007).
113. Schweikert, C., Klenk, R. H., *et al.* Thermally Robust Silicon Integrated Ring Resonators With Postprocessing Sensitivity Enhancement for Biosensing. *Journal of Lightwave Technology* **43**, 4819–4825 (May 2025).
114. Yariv, A. & Yeh, P. *Photonics: Optical Electronics in Modern Communications* en. Google-Books-ID: B2xwQgAACAAJ (Oxford University Press, 2007).
115. en. in *Fundamentals of Photonics Section: 9* _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/0471213748.ch9>, 310–341 (John Wiley & Sons, Ltd, 1991).
116. Bogaerts, W., De Heyn, P., *et al.* Silicon microring resonators. en. *Laser & Photonics Reviews* **6**. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/lpor.201100017>, 47–73 (2012).
117. Chrostowski, L. & Hochberg, M. *Silicon Photonics Design: From Devices to Systems* (Cambridge University Press, Cambridge, 2015).
118. Vetlugin, A. N. Coherent perfect absorption of quantum light. *Physical Review A* **104**, 013716 (July 2021).
119. Jacques, M., Samani, A., *et al.* Optimization of thermo-optic phase-shifter design and mitigation of thermal crosstalk on the SOI platform. EN. *Optics Express* **27**, 10456–10471 (Apr. 2019).
120. Qiang, X., Zhou, X., *et al.* Large-scale silicon quantum photonics implementing arbitrary two-qubit processing. en. *Nature Photonics* **12**, 534–539 (Sept. 2018).

121. Haus, H. A. *Waves and Fields in Optoelectronics* en. Google-Books-ID: AQUpAQAAMAAJ (Prentice-Hall, 1984).
122. Bangari, V., Marquez, B. A., *et al.* Digital Electronics and Analog Photonics for Convolutional Neural Networks (DEAP-CNNs). *IEEE Journal of Selected Topics in Quantum Electronics* **26**. Conference Name: IEEE Journal of Selected Topics in Quantum Electronics, 1–13 (Jan. 2020).
123. Gokmen, T., Rasch, M. J., *et al.* The marriage of training and inference for scaled deep learning analog hardware in *2019 IEEE International Electron Devices Meeting (IEDM)* (Dec. 2019), 22.3.1–22.3.4.
124. Marchetti, R., Lacava, C., *et al.* Coupling strategies for silicon photonics integrated chips [Invited]. EN. *Photonics Research* **7**, 201–239 (Feb. 2019).
125. Mu, X., Wu, S., *et al.* Edge Couplers in Silicon Photonic Integrated Circuits: A Review. en. *Applied Sciences* **10**, 1538 (Jan. 2020).
126. Luo, H., Xie, F., *et al.* Low-loss and broadband fiber-to-chip coupler by 3D fabrication on a silicon photonic platform. EN. *Optics Letters* **45**, 1236–1239 (Mar. 2020).
127. Banner, R., Nahshan, Y., *et al.* Post training 4-bit quantization of convolutional networks for rapid-deployment. en.
128. Park, E., Yoo, S., *et al.* Value-aware Quantization for Training and Inference of Neural Networks arXiv:1804.07802 [cs]. Apr. 2018.
129. Shi, J., Pollard, M. E., *et al.* Photonic crystal and quasi-crystals providing simultaneous light coupling and beam splitting within a low refractive-index slab waveguide. en. *Scientific Reports* **7**, 1812 (May 2017).
130. Fujisawa, S., Yaman, F., *et al.* Weight Pruning Techniques Towards Photonic Implementation of Nonlinear Impairment Compensation Using Neural Networks. EN. *Journal of Lightwave Technology* **40**, 1273–1282 (Mar. 2022).
131. Alexander, C. K. & Sadiku, M. N. O. Chapter 7 First-Order Circuits. en.
132. Nair, V. & Hinton, G. E. Rectified Linear Units Improve Restricted Boltzmann Machines. en.
133. Agarap, A. F. *Deep Learning using Rectified Linear Units (ReLU)* arXiv:1803.08375 [cs]. Feb. 2019.
134. Preston, K., Dong, P., *et al.* High-speed all-optical modulation using polycrystalline silicon microring resonators. *Applied Physics Letters* **92**, 151104 (2008).
135. Kravtsov, K., Fok, M. P., *et al.* Ultrafast all-optical implementation of a leaky integrate-and-fire neuron. EN. *Optics Express* **19**, 2133–2147 (Jan. 2011).
136. Shi, B., Calabretta, N., *et al.* Deep Neural Network Through an InP SOA-Based Photonic Integrated Cross-Connect. *IEEE Journal of Selected Topics in Quantum Electronics* **26**. Conference Name: IEEE Journal of Selected Topics in Quantum Electronics, 1–11 (Jan. 2020).
137. Chen, J., Le, D. V., *et al.* TimelyNet: Adaptive Neural Architecture for Autonomous Driving with Dynamic Deadline. *ACM Trans. Embed. Comput. Syst.* **24**, 140:1–140:23 (Sept. 2025).

- 138. Gehrig, D. & Scaramuzza, D. Low-latency automotive vision with event cameras. *Nature* **629**, 1034–1040 (2024).
- 139. Lacava, A., Bonati, L., *et al.* dApps: Enabling real-time AI-based Open RAN control. *Computer Networks* **269**, 111342 (Sept. 2025).
- 140. Ujjwal, U., Mahala, N., *et al.* Dynamic adaptive spectrum allocation in flexible grid optical network with multi-path routing. en. *IET Communications* **15**.
_eprint: <https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/cmu2.12046>, 211–223 (2021).
- 141. Zhao, T., Fan, X., *et al.* A Resource-Adaptive Routing Scheme with Wavelength Conflicts in Quantum Key Distribution Optical Networks. en. *Entropy* **25**, 732 (May 2023).
- 142. Walkowiak, K., Klinkowski, M., *et al.* Dynamic Routing in Spectrally Spatially Flexible Optical Networks with Back-to-Back Regeneration. en. *Journal of Optical Communications and Networking* **10**, 523 (May 2018).