



UNIVERSITY OF TRENTO

DEPARTMENT OF PHYSICS

DOCTORAL THESIS

**Microring Based
Neuromorphic Photonics**

Davide Bazzanella

Supervisor: Prof. Paolo Bettotti
Co-Supervisor: Mattia Mancinelli, Ph.D.

May 2022

Ai miei genitori, Susam e Paolo

Abstract

This manuscript investigates the use of microring resonators to create all-optical reservoir-computing networks implemented in silicon photonics. Artificial neural networks and reservoir-computing are promising applications for integrated photonics, as they could make use of the bandwidth and the intrinsic parallelism of optical signals. This work mainly illustrates two aspects: the modelling of photonic integrated circuits and the experimental results obtained with all-optical devices. The modelling of photonic integrated circuits is examined in detail, both concerning fundamental theory and from the point of view of numerical simulations. In particular, the simulations focus on the nonlinear effects present in integrated optical cavities, which increase the inherent complexity of their optical response. Toward this objective, I developed a new numerical tool, PRECISE, which can simulate arbitrary circuits, taking into account both linear propagation and nonlinear effects. The experimental results concentrate on the use of SCISSORs and a single microring resonator as reservoirs and the complex perceptron scheme. The devices have been extensively tested with logical operations, achieving bit error rates of less than 10^{-5} at 16 Gbps in the case of the complex perceptron. Additionally, an in-depth explanation of the experimental setup and the description of the manufactured designs are provided. The achievements reported in this work mark an encouraging first step in the direction of the development of novel networks that employ the full potential of all-optical devices.

Acknowledgements

This work would not have been possible without the help of several people.

Technical Acknowledgements

Concerning the development of PRECISE, I must first and foremost thank Mattia Mancinelli, with whom I developed the initial core of the software. Furthermore, I am in debt to Prof. Paolo Bettotti and Massimo Borghi for their help in reorganising all the bits and pieces of theory and simulation.

Regarding the MMI scripts, I thank Mattia Mancinelli for the theoretical description and Nicolò Leone for the code review. Moreover, I thank Clara Zaccaria for the experimental measures.

The experimental setup in the RF laboratory has been built together with Mattia Mancinelli, and I have to thank him and Prof. Paolo Bettotti for sharing their (infinite) wisdom on practical matters. Likewise, I would like to thank Stefano Biasi and Riccardo Franchi for the numerous occasions on which we collaborated. Last but not least, I express my gratitude to the technical staff, especially to Nicola Furlan, Claudio Salomon, and Enrico Moser, whose help certainly simplified my life in the laboratory.

Finally, I must acknowledge the support of the European Research Council through the advanced grant BACKUP (788793).

Personal Acknowledgements

I want to thank Prof. Paolo Bettotti for his supervision and for sharing precious knowledge and insight into the world of academic research. Likewise, I must thank Prof. Lorenzo Pavesi for creating the collaborative environment at the NanoLab and always being available for us students, fostering our professional growth.

I have to thank Mattia, whose explosive personality has been, for me, an unrivalled source of strength. Together, in the lab, we faced many problems, but we also shared many achievements, often chanting “Error-free! Error-free!”. I will always look up to you and how you tackle life inside and outside the lab.

I thank my colleagues and friends Clara and Nicolò, with which I had the pleasure to share this journey. Similarly, I thank all the team at the NanoLab for the time spent together. In addition, I thank all the people who passed through the RF laboratory and enjoyed our humour, terrific and terrible at the same time, for playing along.

Finally, I must thank all my dear friends outside the university and my family that always supported me. In particular, I thank my parents, Susam and Paolo, whose tireless support allowed me to pursue the path that brought me here.

To all of you, as Douglas Adams once wrote,

So long, and thanks for all the fish.

Contents

List of Figures	xvi
List of Tables	xviii
List of Abbreviations	xix
1 Introduction	1
1.1 The BACKUP Project	1
1.1.1 Ethical aspects	2
1.2 Thesis outline	3
2 Neuromorphic Photonics	5
2.1 Artificial neural networks	5
2.1.1 Types of networks	6
2.1.2 Network optimisation	8
2.1.3 Reservoir computing	9
2.2 Neuromorphic computing	9
2.3 Neuromorphic photonics	10
2.3.1 RCN in photonics	12
3 Silicon Photonics Fundamentals	15
3.1 Theoretical description	16
3.1.1 Light propagation in waveguides	16
3.1.2 Mode coupling	20
3.1.3 Fundamental structures	23
3.2 Nonlinear Optics in Silicon	31
3.2.1 Optical Effects	32
3.2.2 Free Carrier Effects	33
3.2.3 Thermal Effects	34
3.3 Nonlinear optics in waveguides	34
3.3.1 Free carrier generation	37
3.3.2 Heat generation	37
3.3.3 Nonlinear effects in microring resonators	38

4	Integrated Photonics Modelling	41
4.1	PRECISE	41
4.1.1	Modular description of the optical system	43
4.1.2	Nonlinear effects dynamics	46
4.1.3	Parametric solutions and data analysis	48
4.1.4	Benchmarks	49
4.2	Linear Spectra	50
4.2.1	Single microring resonator	52
4.2.2	CROW	53
4.2.3	Taiji microring resonator	53
4.2.4	MRR with Feedback Loop	55
4.2.5	Networks of microring resonators	56
4.3	Dynamics of Nonlinear Systems	58
4.3.1	Single microring resonator	58
4.4	MMI optimiser	61
4.4.1	Field profiles	61
4.4.2	Modal propagation	62
4.4.3	Optimisation	63
4.4.4	Experimental results	65
4.4.5	Optional 3D simulation	67
5	Experimental results	69
5.1	Complex perceptron	69
5.1.1	Operation principles	70
5.1.2	Training and testing procedures	72
5.1.3	Pattern recognition	73
5.1.4	Delayed XOR operation	76
5.1.5	Carrier phase decoding	79
5.1.6	Discussion	79
5.2	SCISSOR-based reservoir networks	80
5.3	Microring resonator as a reservoir	81
5.3.1	Operation principles	81
5.3.2	Experimental measurements	84
5.3.3	Results	85
5.3.4	Discussion	93
6	Conclusion	95
	Publication List	97
A	Experimental Setup	99
A.1	Signal generation	102
A.2	DUT coupling and control	103
A.3	Signal detection	104
A.4	Experimental setup operation	105
A.4.1	Resonance position	106

A.4.2	IR camera aided characterisation	108
A.4.3	Thermal cross-talk	108
A.4.4	Input-output delay calibration procedure	114
B	Photolithographic Masks	115
B.1	Fabrication process	115
B.2	Design process	116
B.3	IR IMEC 2019	118
B.4	Visible IMEC 2019	121
B.5	IR IMEC 2020	122
B.6	IR Applied Nanotools 2021	123
	Bibliography	127

List of Figures

2.1	Representation of a deep FFNN as a graph.	7
2.2	Representation of RNN as a graph. Recurrence can occur at the node level (first hidden layer) or between different layers (second and third layers).	7
2.3	Representation of RCN as a graph. Solid lines represent fixed connections, while dashed lines are the connections that are trained.	9
3.1	Cross-section of a waveguide with a rectangular core. Common sizes and axis orientation are shown.	17
3.2	Examples of field profiles. a Field profiles for the TE and TM modes of a single-mode WG with a rectangular core of size 450 nm × 220 nm. b Field profiles for the lowest TE and TM modes of a multi-mode WG with a rectangular core of size 1500 nm × 220 nm.	18
3.3	Examples of input/output couplers. a Tapered structure for edge coupling, taken from [53]. b Optical fibre with lensed tip. c <i>Apodized</i> grating coupler, taken from [54].	22
3.4	Symmetrical DC drawing and field profile. a Drawing of a symmetrical DC. b Field profiles of the <i>even</i> and <i>odd</i> modes.	24
3.5	Evolution of the field profile in an asymmetrical DC. Reproduced from [61].	25
3.6	Plot of the field profile of the field propagation of an MMI approximated as a 2D structure evaluated through FEM methods. a Symmetrical excitation of an MMI body. b Asymmetrical excitation of ac MMI, producing a mirrored image at the output.	26
3.7	Drawing of MRRs in the a all-pass filter and b add-drop filter configurations. The coupling coefficients κ represent the portion of the field which is exchanged between the bus WGs and the MRR in each coupling region, as shown in c	27
3.8	Spectrum of a MRR in the add-drop filter configuration. Solid and dotted lines represent the transmission spectra of the <i>through</i> and the <i>drop</i> ports, respectively. Arrows indicate the FSR and the full width at half maximum, $\delta\nu$	28

3.9	Drawing of sequences of microring resonators: a CROW and b SCISSOR.	30
3.10	Refractive index of silicon on a wide range of wavelengths [41].	31
4.1	Drawings of fundamental components in PRECISE description of optical systems. a Nodes group together several input/output ports and define the connections from the input to the output ports. In this example there is a node with four ports. b Guides join ports, output to inputs, and describe the phase and amplitude changes associated with the propagation of light. c Input ports in nodes may be designated as <i>input</i> to the whole system; the field of unconnected input ports in nodes are set to zero.	43
4.2	Gaussian pulse propagation: comparison between PRECISE and CMT models. A Gaussian pulse of varying width excites the <i>input</i> port of an ADF with an optical frequency close to the resonance of the MRR. a Peak power in the <i>through</i> port, b peak power in the <i>drop</i> port, c temperature difference, and d FC concentration difference. The vertical lines indicate three times the value of the MRR photon lifetime ($3\tau_p$). Dashed lines (left axis) indicate the results of the two models, while the red lines (right axis) show the residuals.	47
4.3	PRECISE performance analysis on a benchmark suite evaluating the response of a CROW structure: average value and standard deviation of ten consecutive runs. Each new ring in the CROW increases by two the number of elements in the system. a Time spent in the generation of the f_{pfe} function. b Execution time for a single call of the f_{pfe} function; obtained from a batch evaluation of 1001 calls. c Time required to simulate the nonlinear evolution of a CROW with two rings for 51, 101, and 201 frequency points as a function of the number of parallel workers.	50
4.4	Drawings of several structures. a Single MRR in the ADF configuration. b Single MRR with feedback. c Taiji MRR. d CROW composed of four MRRs. Nodes are represented with small squares. Ports in nodes are small circles and are identified by an integer number. The field propagation is illustrated as dashed arrows and input and output to the system are highlighted with red and blue arrows, respectively.	51
4.5	Transmission spectra of the MRR in the ADF configuration. The lines represent the square modulus of the complex amplitude at a given port, which is identified in the legend by the index that refers to fig. 4.4a. For output ports, this corresponds to the power transmission, while, for inner fields, it corresponds to the square modulus of the field enhancement factor. PRECISE evaluates the transmissions associated with the counter-propagating fields (E_1^{out} , E_4^{out} , E_5^{out} , and E_8^{out}) too, however, they are not shown here as they are always zero.	52

4.6	Transmission spectra of the CROW structure in the <i>interferometric band interleaver</i> configuration. Experimental data (black circles) for the <i>through</i> port are compared with the simulation results for the <i>through</i> port (solid cyan line) and <i>drop</i> port (solid red line).	53
4.7	Taiji MRR transmission spectra: comparison of the simulations to the experimental data. <i>Rx</i> stands for reflection, <i>Tx</i> for transmission; <i>bwd</i> stands for backward, as in counter-clockwise field propagation, <i>fwd</i> stands for forward, as in clockwise field propagation.	54
4.8	Maps of the normalised a linear transmission and b internal energy in a MRR cavity with loop feedback as a function of frequency detuning and loop length. The <i>drop</i> transmission is normalised to the <i>input</i> power. The internal energy is normalised to the maximum value in the map and is reported in log scale.	56
4.9	A network is composed of a set of nodes connected by bus WGs. a Each node is a MRR with a given number of coupling regions. b Graph of the network connections between the rings in the network and c associated incidence matrix. Each column in the matrix describes a connection between different nodes (rows), with $-$ ($+$) representing an output from (input to) the node.	57
4.10	The numerical simulation produces the normalised spectra of the energy stored in each MRR in the network.	58
4.11	Self-pulsing phenomenon in a MRR simulated with PRECISE and with the CMT method. The curves show the response of the system observed at the <i>drop</i> port for three different combinations of input power and frequency detuning in respect to the resonance frequency. They have been extracted from longer simulations so that they can be considered stationary states.	59
4.12	Self-pulsing regions are characterised through the period of the most important Fourier transform component, observed in the transmission at the <i>drop</i> port ($ E_6^{out} ^2$). The regions where no self-pulsing is observed are reported having 0 ps period. Comparison between the PRECISE approach a and the CMT approach b . The white horizontal line represents the MRR resonance frequency, i.e. 0 GHz detuning. The pins highlight the curves shown in fig. 4.11.	60
4.13	Drawing of the MMI coupler structure and the associated reference axis.	62
4.14	Effective refractive index n^{eff} for the TE modes as a function of the WG width: a Real and b imaginary parts. Negative imaginary values mean attenuation, either due to absorption or scattering.	63
4.15	Summary plots, displaying the system performance and the geometry parameters across all the tested combinations: a transmission from the input WG to each of the two output WGs, in dB, and b optimal length of the MMI body, in μm	64

4.16	Example of in-depth plot, displaying the field profiles at the input and output of the MMI body as well as the field propagation in the $z - x$ plane. Specifically, the left half of the first plot is the input field profile, while the right half is the excited field profile at the input of the MMI. Both maps represent the top half of the section only, assuming symmetry around the horizontal plane. The second plot is the field propagation on the horizontal plane at $y = 0$. Similarly to the input, the third plot represents the field profile at the output of the MMI in the left half and the excited field profile at the start of the output WGs. The numbers represent the transmission values in dB.	65
4.17	Photographs of matrices of scatterers, used to stimulate neuronal cultures. For more information, see appendix B.4, where the chip that contains these structures is described. a 6×6 matrix addressed by a cascade of 1×2 and 1×3 MMIs. b 5×5 matrix addressed by a cascade of 1×5 MMIs.	67
4.18	Example of 3D simulation of a 2×2 MMI coupler. The colourbar represents the field intensity and is in V/m. Comsol propagation direction is x , and the section is the $y - z$ plane.	68
5.1	Simplified schematic of the complex perceptron [90]. WGs are illustrated as dark paths developing over a substrate. Delay lines are depicted as WG spirals and heaters are the bright volumes covering each WG. Fields are represented as semitransparent plots above the WGs, before the heaters and at the end of the structure. Input and output coupling fibres are shown hovering on top of the input and output grating couplers, respectively.	70
5.2	Sampling times in comparison with bit periods for different bitrates. Black lines show how adjacent bits are sampled by four channels separated by a constant delay of $\Delta T = 50$ ps.	71
5.3	Operation during the recognition task of pattern “1 0”. a Distribution of the input intensity at the best sampling time and b distribution of the output intensity at the best sampling time. The vertical dashed line indicates the optimal threshold.	74
5.4	Time traces for the “1 0” pattern recognition task at 16 Gbps. Input measured optical intensity (blue), output measured optical intensity (red), and ideal target response (green) as functions of the bit number. Line marks represent the best sampling time, while the dashed line represents the optimal threshold.	75
5.5	BER values at the best combination of input power and sampling time for bitrates of 5, 8, 10, and 16 Gbps and different patterns composed of two and three bits. Transparent bars indicate that the value is capped at the statistical limit of the recorded trace. The dashed lines represent the random choice threshold.	75

5.6	BER values at the best combination of input power and sampling time for different bitrates, i.e. 5, 8, 10, and 16 Gbps, and for different bit delays. Transparent bars indicate that the value is capped by the statistical limit of the recorded trace. The solid line represents the random choice threshold.	76
5.7	Map of the BER of the 1-bit XOR operation at 5 Gbps versus the time from the beginning of the bit and input optical power. The best value is obtained at the maximum optical power and 62.5 ps. The BER statistical limit is 10^{-5}	77
5.8	Example of time traces for the 1-bit XOR operation. Input measured optical intensity (blue), output measured optical intensity (red), and ideal target response (green) as functions of the bit number. Line marks represent the best sampling time, while the dashed line represents the optimal threshold.	78
5.9	Distribution of the best sample output intensity for the four cases of the ideal input bits, i.e. “0 0”, “0 1”, “1 0”, “1 1”. The vertical dashed line indicates the optimal threshold.	78
5.10	Time traces for PSK decoding at 10 Gbps. Input measured optical intensity (blue), output measured optical intensity (red), and ideal target response (green) as functions of the bit number.	79
5.11	Normalised transmission spectrum around the resonance, measured at the <i>drop</i> port.	82
5.12	Relation between the number of virtual nodes N_V , the <i>samples per bit</i> S , and the number of observed bits n_R . The bits involved in the task are separated by a delay D	83
5.13	AND 1 with 2 Ridge-bits and $S = 5$. Comparison of a BER map, b <i>RB</i> map, and c power map. Markers identify the combinations that achieved the statistical limit (error-free operation): circles for the input signal and crosses for the output signal.	87
5.14	AND 1 with 1 Ridge-bits and $S = 5$. Comparison of: a BER map, b <i>RB</i> map, and c power map. Markers identify the combinations that achieved the statistical limit (error-free operation): circles for the input signal and crosses for the output signal.	88
5.15	AND 2 with 1 Ridge-bits and $S = 5$. Comparison of: a BER map, b <i>RB</i> map, and c power map.	89
5.16	XOR 1 with 1 Ridge-bits and $S = 5$. Comparison of: a BER map, b <i>RB</i> map, and c power map.	90
5.17	XOR 1 with 2 Ridge-bits and $S = 5$. Comparison of: a BER map, b <i>RB</i> map, and c power map. Markers identify the combinations that achieved the statistical limit (error-free operation): circles for the input signal and crosses for the output signal.	91
5.18	XOR 2 with 1,2, and 3 Ridge-bits and $S = 5$. Comparison of the <i>RB</i> maps for a 1 Ridge-bit, b 2 Ridge-bits, and c 3 Ridge-bits.	92
A.1	Experimental setup.	100

A.2	Sketch of the position of the DUT and of the input/output optical fibres. Vectors show the movements and rotations of the mechanical stages, bundled as axes origins.	103
A.3	a Exploded view of the TEC system with PCB for the packaged die. b temperature fluctuations reported by the PID controller temperature sensor.	104
A.4	Measurement of the resonance of a microring resonator. a Coarse measurement of the positions of all resonances within the broad ASE spectrum of the EDFA1. b , c , and d Fine measurements of the spectral response around a given centre frequency. The first is measured at the <i>drop</i> port, while the latter two are measured at the <i>through</i> port. Both the frequency position and the shape of the resonance are accurately measured.	107
A.5	a Transmission measured at the output grating of the device. b Scattering spectra for each ROI. The vertical dashed lines indicate the detuning at which the IR image has been acquired. c Camera picture with multiple ROIs defined.	109
A.6	Spectrum of the complex perceptron where no currents have been applied to heaters.	110
A.7	Single-side Fourier transform of the spectra: envelope of the 31×4 spectra. The blue line represents the Gaussian filter. Note the presence of three peaks roughly at positions of 6, 12, and 18 nm^{-1}	111
A.8	Timescales involved in the fast measurement prior to thermalisation. a Long measurement. b Zoom around the trigger event. The acquisition happens after 1 ms after the trigger event and usually lasts a few μs . In comparison, 1 ms of measurement contains 10 Mbit at 10 Gbps.	113
B.1	Example of a GDS drawing. Each of the several layers displayed corresponds to a specific step in the fabrication process. Specifically, pink layers work on the device layer, defining the core and the cladding of the WG, while purple and blue are the metallic heater and metallic traces, respectively.	118
B.2	a CAD of the device, highlighting the SCISSOR-based RCNs in green, “swirl” topology RCN in red, the complex perceptron and the FFNN in blue, and the test structures in yellow. b Picture of the manufactured chip.	120
B.3	Picture a and b CAD of the manufactured chip. Highlighted: area (A) on the left and area (B) on the right.	121
B.4	Close-up of the custom structures. a Scatterer. b 1×2 MMI. c 1×3 MMI. d 1×5 MMI.	122
B.5	Picture of the manufactured chip.	123

List of Tables

- 3.1 Typical values for the parameters described in this section for a silicon WG embedded in silica of size $450 \text{ nm} \times 220 \text{ nm}$ measured with IR light at 1550 nm and a standard SMF-28 optical fibre. 20
- 3.2 Summary for the change in effective refractive index for guided modes due to several nonlinear effects. 37
- 4.1 Simulation parameters of the CROW. The coupling region is $10 \mu\text{m}$ long, but is considered a point coupling region and the extra length is associated with the waveguide segment instead. k are the field coupling coefficients. 54
- 4.2 Taiji MRR simulation parameters. L are the lengths of the WG segments and k are the coupling coefficients. Indexes refer to the ports as described in fig. 4.4c. Specifically, indexes on L identify the ports at the ends of the segments and on k identify the ports in the node of the associated coupling region. 55
- 4.3 List of the parameters required to carry out the evolution of the nonlinear dynamics of the MRR. The top half of the table contains values typical of silicon, while the bottom half contains those typical of silicon photonics single-mode channel WG manufactured for 1550 nm light. The asterisks * marks the entries which should be considered as an indication of the correct order of magnitude of the given quantity, rather than its exact value. 61
- 4.4 MMIs geometrical parameters: output lateral shift \bar{x} , width of the input and output WGs, and length L and width of the MMI body. 66
- 4.5 Experimental characterisation results of the MMIs performance. * Note that the “Measured” column reports the insertion loss normalised on the sum of the output power of each channel. This column should be used to compare the channel balance with that of the simulations. 66
- 5.1 Sample per bit N_{spb} for the bitrates employed in the experiments. 71
- 5.2 Truth tables for AND and XOR logical operations. 83

5.3	Best value of the RB and corresponding BER, evaluated at a bitrate of 100 Mbps, a detuning of -25 GHz, and optical power of 16 dBm.	93
A.1	List of instruments and components in the experimental setup. .	101
A.2	Values of the partial derivatives obtained from the Fourier transform spectral analysis.	112

List of Abbreviations

ADC	analog-to-digital converter	EDFA	erbium-doped fibre amplifier
ADF	add-drop filter	EM	electro-magnetic
AI	artificial intelligence	EOM	electro-optic modulator
ANN	artificial neural network	FC	free carrier
APF	all-pass filter	FCA	free carrier absorption
ASE	amplified spontaneous emission	FCD	free carrier dispersion
ASIC	application-specific integrated circuit	FDTD	finite-difference time-domain
AWG	arbitrary waveform generator	FEM	finite element method
BER	bit error rate	FFNN	feedforward neural network
CAD	computer-aided design	FOM	figure-of-merit
CMOS	complementary metal-oxide-semiconductor	FPGA	field-programmable gate array
CMT	coupled mode theory	FSR	free spectral range
CPU	central processing unit	FWHM	full width at half maximum
CRIT	coupled resonators induced transparency	FWM	four-wave-mixing
CROW	coupled resonators optical waveguide	GDS	graphic design system
CW	continuous wave	GDSII	GDSII stream format
DC	directional coupler	GPU	Graphic Processing Unit
DDE	delay differential equation	GVD	group velocity dispersion
DOF	degree of freedom	IC	integrated circuit
DRC	design rule check	ICP-RIE	inductively coupled plasma - reactive ion etching
DUT	device under test	IPU	intelligence processing unit
		IR	infrared

MEMS	micro-electro-mechanical system	RF	radio frequency
ML	machine learning	RNN	recurrent neural network
MMI	multi-mode interference	ROI	region of interest
MPW	multi-project wafer	SCISSOR	side coupled spaced sequence of optical resonators
MRR	microring resonator	SCPI	standard commands for programmable instruments
MZI	Mach-Zehnder interferometer	SEM	scanning electron microscope
NEP	noise equivalent power	SM	single mode
NMZI	nested Mach-Zehnder interferometer	SNN	spiking neural network
NN	neural network	SNR	signal-to-noise ratio
NPU	neural processing unit	SOI	silicon-on-insulator
NRZ	non-return-to-zero	TE	transverse electric
OPD	optical path difference	TEC	thermo-electric cooler
OPU	optical processing unit	TEM	transverse electro-magnetic
OSA	optical spectrum analyser	TIR	total internal reflection
PCB	printed circuit board	TLS	tunable laser source
PIC	photonic integrated circuit	TM	transverse magnetic
PID	proportional-integral-derivative	TOC	thermo-optic coefficient
PM	polarisation maintaining	TOE	thermo-optic effect
PRBS	pseudo-random binary sequence	TPA	two photon absorption
PSK	phase-shift keying	TPU	Tensor Processing Unit
PSW	particle swarm	UV	ultraviolet
RC	reservoir-computing	VOA	variable optical attenuator
RCN	reservoir-computing network	WG	waveguide

Chapter 1

Introduction

This thesis reports the main results I obtained in the development of photonic neuromorphic devices. My research work has covered several aspects: from the installation of a new laboratory equipped with high-frequency instrumentation to the development of new libraries to model linear and nonlinear photonic integrated circuits (PICs), to the experimental characterisation of the devices fabricated.

1.1 The BACKUP Project

BACKUP aim is to investigate the abilities of the biological brain, recreate them *in silico*, and develop a framework that enables the interaction of the two domains. In view of this, BACKUP positions itself at the boundary between biology and technology, between human and machine.

The objective of BACKUP is, in fact, twofold. A first goal is to investigate the activity of biological neurons in the context of information elaboration and, specifically, in the formation of memory in the brain. BACKUP plans to accomplish this by observing neuronal networks from a bottom-up perspective. Interacting with networks of living neurons, i.e. *in vitro* experiments, will allow a more clear and simple, yet hopefully sufficiently insightful view of networks of higher complexity, such as the human brain. The biological techniques of optogenetics enable the interaction through either excitation or suppression of neuronal activity by means of visible light. Specifically, this is allowed by the expression of light-sensitive proteins (channelrhodopsins), which act as light-activated ion channels and can sit in specific cell compartments. By targeting a subset of the network's neurons, specific synapses can be strengthened or weakened, actually changing the way the neurons are connected within the network. The development of an integrated photonics platform, specifically designed to interact with these neurons, will enable the study of the biological network's working principles.

The other goal of the BACKUP project is to create artificial networks of in-

egrated photonic structures that are able to carry out computational tasks. In the field of information technology, artificial neural networks (ANNs) are mathematical models representing a network of nodes in which data is elaborated in parallel. The possibilities shown by these models are increasing rapidly as is the complexity of the tasks they are able to successfully carry out, e.g. image and speech recognition. However, this trend is closely linked to the rise in the computational power required to optimise the internal parameters of these networks. The most recent ANN models may have even millions of parameters. In fact, due to the different structures of ANNs and the computers which emulate them, i.e. the von Neumann architecture, this process is quite inefficient. For the sake of efficiency, researchers are looking at different implementations of ANNs architectures, such as accelerator co-processors (GPUs and TPUs), rather than the application-specific integrated circuit (ASIC), or optics and photonics devices. In this context, BACKUP aims to create recurrent networks of microring resonators, based on the concepts of reservoir-computing networks (RCNs). The reservoir-computing (RC) implemented in BACKUP will be a small network of integrated photonic structures fully exploiting the speed and bandwidth of this technology. Such an approach exploits the complexity of the reservoir dynamic and limits the number of parameters to optimise to accomplish a given task.

The cornerstone of the BACKUP project is the use of integrated photonics both as a means to study the formation of memory in biological neural networks and as a framework to develop a novel paradigm of computation, i.e. artificial neural networks (NNs). In the long-term, BACKUP will eventually bring together the biological and computational efforts to create a hybrid biological-artificial network. Living neurons will interface with a photonic integrated circuit and both will be assisted by electronic circuitry to jointly carry out a given function. This platform will also enable to tackle diseases such as amnesia and epilepsy, by regulating the activity of dysfunctional areas of the brain. Furthermore, it could be envisaged that a hybrid network could behave like a backup for the damaged areas, from which stems the name of the project.

Technological advancements that allow such direct interaction with the human brain, however, require the study of biological neurons. Furthermore, they are certainly a disruptive element for the equilibrium in today's society. Hence, important ethical questions arise both regarding the research and the use of artificial intelligence (AI).

1.1.1 Ethical aspects

BACKUP's neurobiological experiments make use of living neurons, which are obtained from mice embryos. This fact encouraged us to start a discussion regarding the ethical aspects of this kind of research, as well as the development of technologies surrounding the concept of AI. On December 3-4, 2020 a workshop titled "Neuroni artificiali e biologici: etica e diritto"¹ was organised, among others, also under the umbrella of the BACKUP project. It engaged a large

¹ "Artificial and biological neurons: ethics and law"

audience composed of researchers from different fields, such as physics, neurobiology, medicine, law study, and philosophy. In fact, the first aspect of this discussion is that involves a broad range of topics and disciplines, which makes the debate difficult as no language is common to all those fields of research. Nevertheless, many benefits and problems of AI have emerged from the dialogue [1].

Scientists are currently able to investigate the properties of living neurons or the creation of artificial ones, as in the case of AI software tools. These technological advancements are spreading fast in today's world. In the same way as most novel tools, AIs might exacerbate social and/or economic inequalities already existing in our society. Moreover, these developments might be so groundbreaking that they even raise questions about fundamental philosophical definitions, e.g. the concept of intelligent being with respect to artificially created intelligence. For these reasons, the delineation of new regulations is required, which should be created with the involvement of the general public. Rules should also define limits to the research of AIs, depending on the purpose for which they are developed. Finally, encouraging education in regards to these topics is essential to increase the awareness of both the challenging and the promising aspects of AI.

1.2 Thesis outline

At the beginning of this manuscript, the abstract is followed by the acknowledgements, the table of contents, the list of figures, and the list of tables. The main content of this work is organised in the following chapters:

Chapter 1 is the introduction to this work. You are currently reading this.

Chapter 2 introduces the main concepts of neuromorphic photonics, starting from general notions of neuromorphic computing. It provides a review of the most promising neuromorphic architectures in integrated photonics.

Chapter 3 contains the fundamental physical concepts required to model integrated photonics structures and to understand their working principles, with a specific focus on silicon photonics. In the beginning, linear propagation and mode coupling are described, that the description of the vast majority of integrated structures is covered and a few examples are discussed. Eventually, nonlinear effects are introduced and an effective approximation model is presented.

Chapter 4 describes the software for modelling integrated photonics that I developed. The advantage of using numerical models is twofold: first, it enables the evaluation of the optical response even of complex structures, and second, it unifies the method used in the solution of the problems, producing consistent results. The underlying equations are those proposed in the previous chapter. I propose a set of examples, starting from simple linear structures and later adding complexity and nonlinear effects.

Chapter 5 reports the experimental results, starting from those of the *complex perceptron*, which is obtained by the combination of delay lines and an

optical perceptron layer. After, the focus shifts to RC networks composed of microring resonators (MRRs). The primary results in this regard are those of the explorative study of a MRR reservoir.

Chapter 6 conveys the conclusive remarks of the work, emphasizing the objectives achieved as well as the missed opportunities, the unexpected outcomes, and the future perspectives of the BACKUP project.

To avoid broadening the discussion, two appendices collect auxiliary material enriching of additional details concepts and results discussed in the previous chapters.

Appendix A holds a detailed description of the experimental setup in its three main parts: the signal generation stage, the device under test (DUT) coupling and control stage, and the signal detection stage. Moreover, several operative topics are discussed, e.g. remote operation of the instrumentation and the measurement technique in presence of thermal cross-talk.

Appendix B lists the PIC designs to which I contributed, describing for each one the scope of the design, the technology framework used, and the structures which are contained.

Chapter 2

Neuromorphic Photonics

The past few decades have seen the rapid development of AI models, tackling problems and tasks of ever-growing complexity. The optimisation of large networks with massive datasets has been possible thanks to the increase in raw computational power, with a cost accessible by research institutions and technology companies. Each year, processors are paired with new accelerators and ASICs, which are being developed to carry out the specific operations most used by AI algorithms with increased efficiency. The progress in artificial neural network (ANN) models is astonishing and has now become a driving factor for electronics designers and manufacturers.

Researchers are however pushing the boundaries even further and are studying novel computing architectures which implement mechanisms proper of neural networks: neuromorphic computing devices. The platform of choice for the majority of neuromorphic computing devices is integrated electronics, because its mature technology provides high-performance and it is almost ubiquitous throughout the computing landscape. However, integrated photonics is an interesting alternative. In fact, although it is a relatively less established process with respect to electronics, integrated photonics provides far more bandwidth and, in principle, allows an intrinsic parallelisation. Moreover, the backbone of today's digital communications relies almost exclusively on optical communications. Hence integrated photonics computing could be implemented directly within optical communication networks, eliminating the need for opto-electrical and electro-optical conversion layers, which increase the power consumption and the latency.

2.1 Artificial neural networks

ANNs are, in the prevailing sense, a type of computer algorithm which belongs to the larger families of machine learning (ML) and artificial intelligence (AI). Any ANN can be represented as a collection of elementary units, called *nodes*, connected one to the other, loosely resembling the structure of biological neural

networks. Each node is seen as an artificial neuron which, in the conventional McCulloch–Pitts model [2], carries out a *weighted* sum of its inputs x_i and then applies an *activation function* to the result:

$$y = f_a \left(\sum_{i=1}^N w_i x_i + w_0 \right). \quad (2.1)$$

The weights w_i in the sum quantifies the relative importance of the connections and w_0 is the so-called bias. Increasing (decreasing) a weight means reinforcing (weakening) the connection between two nodes. Weights are in fact changed to steer the network’s output toward a certain target. The activation function f_a , on the other hand, determines how the input information is propagated toward the output. Biological neurons activate when the stimuli at their input exceed a given threshold and the information is propagated toward the output of the neuron. In artificial neurons, the information is processed by nonlinear functions. In fact, it is easily demonstrated that if f_a was linear, any number of nodes in cascade could be reduced to a single one.

A single neuron is, however, quite limited in the computational power it can achieve. Indeed, a famous paper [3], demonstrated already in 1969 that such systems cannot solve nonlinearly separable problems, for instance, the XOR logical operation reported in the same scientific work. In the early days of ANNs, this was a limiting factor and hindered their development in favour of the Boolean logic and the von Neumann architecture. This impasse was overcome starting from the 1980s, when networks composed of cascades of many nodes, i.e. *deep* networks, became tractable in an efficient manner with the method of *backpropagation*. In the following decades, other practical problems, such as the *Vanishing (and Exploding) Gradient Problem*, were solved by employing clever expedients, e.g. the choice of an efficient activation function, which enabled the use of deeper and more complex ANNs.

2.1.1 Types of networks

ANNs differentiate one from the others fundamentally because of their topology, i.e. the way each node is connected to the others., In this regard, we can distinguish two categories.

Feedforward neural networks (FFNNs) are organised in *layers* containing a fixed number of nodes, which accepts inputs from nodes in the previous layer and whose output is fed only to nodes in the successive layer. The information flows from the input through each layer toward the output of the network, hence the reason they are called *feed-forward*. The fundamental computing unit of a FFNN is the perceptron. Generally, FFNNs are made by several layers, a network with a few of them render it a *multi-layer* network, while networks with many layers are called *deep*. Typically, the last layer of a FFNN is called the *output* layer and all the others, between that and the inputs, are called *hidden* layers. A layer whose nodes are all connected to all the nodes in the previous layer is called *fully connected*. This enables the most general description of a network, however, it heavily impacts the computation resources and time required to optimise it.

Other choices of connectivity are usually implemented if the knowledge of the problems allows it. For example, in networks used for image recognition tasks, *pooling* and *convolution* layers are frequently employed. In both of these, each node is connected only to a small subset of the input nodes and, most of the time, the weights are shared between all nodes. Non fully connected layers are also called *sparse*. Figure 2.1 shows the representation of a deep FFNN as a graph.

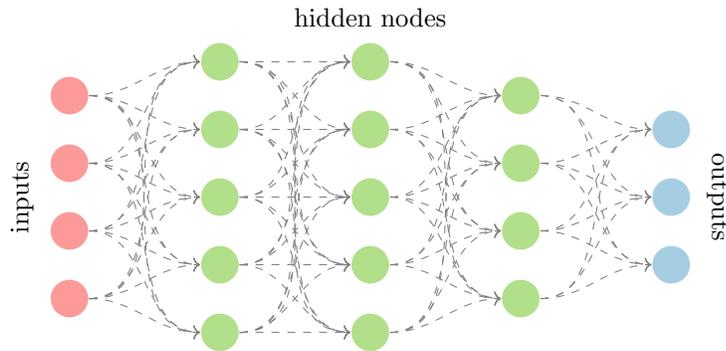


Figure 2.1: Representation of a deep FFNN as a graph.

Recurrent neural networks (RNNs) allow for closed loops, i.e. the information flows back and forth so that the output of a given node depends on the output of other nodes and vice versa in a recurrent fashion. Figure 2.2 shows the representation of a RNN as a graph.

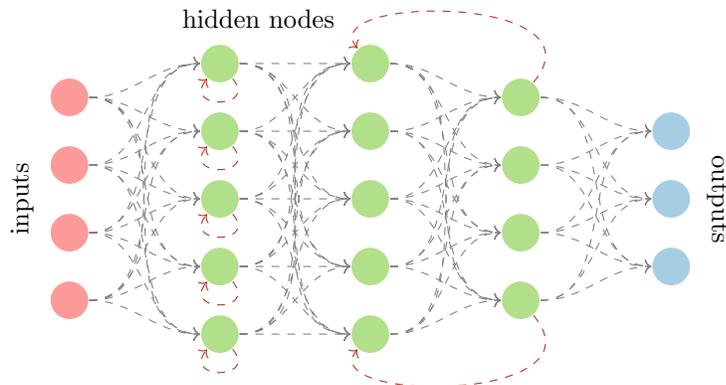


Figure 2.2: Representation of RNN as a graph. Recurrence can occur at the node level (first hidden layer) or between different layers (second and third layers).

Not all nodes in a network must be recurrent or part of a recurrent loop in order to consider it a RNN. These closed loops allow RNNs to exhibit dynamic behaviour. For this reason, RNNs are often employed in the elaboration of temporal sequences, e.g. speech recognition.

2.1.2 Network optimisation

The flexibility of ANNs arises from the huge number of parameters that can influence the response of the network. In order to choose the values of these parameters, an optimisation process is required, which is called *learning* or *training*. During learning, the parameters, or weights, of the network are updated by minimising the output “error” of the network. Usually, this is quantified by a function called *loss* function or *cost* function. The exact way in which the weights are updated and the error is evaluated depends on the learning model used.

The major learning paradigms are three, supervised, unsupervised, and reinforced learning, and they are normally employed depending on the type of network and on the task at hand. Supervised learning evaluates the loss function by comparing the output of a given input to a target value. Datasets for this kind of approach are assembled by pairing each input sample with the corresponding target output, or *label*. The performance of the network is considerably influenced by the quality of the dataset and in particular by the number of samples. On the other side of the spectrum, there are unsupervised learning approaches, whose purpose is to find an underlying structure of the input data, without relying on examples for which the correct answer is known. An example of the application of such methods is cluster analysis, in which the input samples must be grouped by looking at their similarity in respect to one or more criteria. Last but not least, reinforced learning methods are employed mostly in decision-making processes. They do not rely on labelled input samples, however, the network behaviour is influenced by positive rewards, which are assigned by evaluating the state of the network and its environment. Example uses for reinforced learning are AIs for games [4–6] and autonomous driving [7–9].

A very important objective of most ANNs is the ability to generalise a given task to unseen inputs. For this reason, especially when the training is carried out on a finite set of samples, limitations arise. The two most common are overfitting and bias. Overfitting happens when the network’s prediction recognises features that are specific to the training dataset, rather than the more general ones. Similarly, using a dataset containing only a partial description of the model features can lead to a distorted prediction model, i.e. to bias. This effect becomes especially important for ANN algorithms employed in social contexts, e.g. criminal risk assessment [10].

2.1.3 Reservoir computing

A separate description is required to contextualise the reservoir-computing network (RCN) approach. In RC the input is projected to a higher dimensional state by a fixed, nonlinear system called *reservoir*. The reservoir usually consists of a dynamical RNN or can be described as such. The output of the system is obtained by a readout layer performing a simple operation, usually a linear regression, on the state of the reservoir. RCNs are trained by changing the parameters of the readout layer [11]. Figure 2.3 shows the representation of a RCN as a graph.

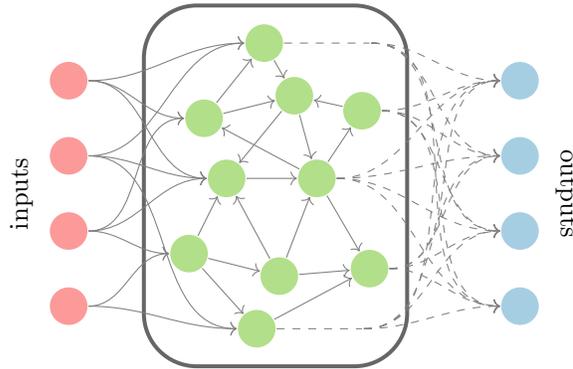


Figure 2.3: Representation of RCN as a graph. Solid lines represent fixed connections, while dashed lines are the connections that are trained.

There are two properties of RCNs which are most important to understand why these kinds of networks are of interest. The first one is that a reduced number of trainable parameters is reflected in a reduced training effort, but also in the diminished flexibility of the network, in respect to training all of the network’s parameters. This compromise becomes more or less appealing depending on the computational cost of the reservoir. Indeed, the second important property of RC systems is that they can harness the complex dynamics of physical systems in the form of a reservoir, drastically decreasing the computational cost of simulating such a network [12].

2.2 Neuromorphic computing

Neuromorphic computing aims at creating computational systems which distance themselves from the prevailing von Neumann architecture and instead employ mechanisms proper of biological neuronal systems. The purposes of such systems are principally two: either precisely reproducing the biological mechanisms, or increasing the computation efficiency of operations normally carried out by CPUs and GPUs. Digital and analog electronics is by far the most used framework, but photonic, spintronic, mechanical, and even biological platforms are investigated as well [11].

Neuromorphic devices aiming at accurately reproducing biological mechanisms were probably the first to appear. Most of these devices employ spiking neural networks (SNNs), which encode the information in isolated impulses, replicating more closely the dynamical behaviour of biological neurons. In this case, the performance is a secondary factor with respect to the accuracy of simulations. Nevertheless, computation efficiency remains very important as a criterion for comparison with respect to biological networks. Moreover, models showing higher performance allow reproducing longer simulations quickly. The most notable examples in this regard are Neurogrid from Stanford University [13], IBM's TrueNorth [14], and HICANN from the University of Heidelberg [15]. These three projects use custom-designed analog and digital electronic circuitry to simulate neural elements and connections between neurons with high efficiency.

Recently, with the increase in computational power spent in ANNs training, improving its efficiency became crucial. For this reason, several companies developed devices that are able to carry out efficiently part of the common mathematical operations normally used during ANNs training, i.e. *accelerators*. Devices of this kind are specialised computation units that trade computation flexibility (with respect to GPUs and CPUs) with computation efficiency, thus consuming less energy for the same task and/or being able to carry out a more intensive task in the same time. Graphcore's IPU [16], Google's TPU [17], and Samsung's NPU [18, 19] are examples of neuromorphic devices implemented with electronics. The first two provide high-performance devices able to drastically decrease the training time of ANNs, while the latter aims at reducing the cost of running ANNs algorithms on the low-power mobile platform. Two examples of commercial optical systems with the goal are Lightmatter's Enviser [20] and LightOn's OPU [21, 22], which implement neuromorphic architectures in integrated photonics and free space optics, respectively.

A third category of devices will constitute the future development of the former two: it will aim to create a standalone device employing neuromorphic techniques to target either general computation or a specific task. While the achievement of general computation is very tough to reach, especially considering the competition of electronics, tackling a single, particular operation could deliver interesting results. The closest analogue is that of ASICs, which are designed ad hoc to carry out a specific task. Of course, the effort is worth it if the performance and/or the efficiency achieved are higher than those obtained with digital electronic devices. Nevertheless, the range of applications for this kind of device is extremely broad, for example, bit header recognition and signal recovery in optical telecommunications, image classification [23], speech recognition [24, 25], chaotic series prediction [26, 27], and biological cell identification [28].

2.3 Neuromorphic photonics

Neuromorphic computing devices implemented in the photonic platform employ a diverse set of approaches to the problem. Both free space and guided optics

techniques have been investigated, each with advantages and disadvantages with respect to the others, but, considering the focus of my work is on the silicon photonics platform, I will discuss only the latter. Each device can be categorised based on a few characteristics. However, this classification creates groups that are often intersecting with each other. I describe the main properties below:

Information Encoding

Optical signals can be differentiated by their static attributes, e.g. frequency (colour), amplitude (power), phase, propagation location (different waveguides), and by their dynamic attributes, e.g. delay, modulation frequency, timing. All of these degrees of freedom (DOFs) can be used to encode information and thus create a network that carries out operations on optical signals.

Dynamics and memory

Neuromorphic devices employ a static design, i.e. a constant input produces a constant output, or they require a time-varying input and they exploit the complex dynamics available in photonic systems. In the former case, the high photonic bandwidth enables fast elaboration rates, usually limited by the bandwidth of the surrounding electronics. In the latter case, the memory of the system determines the duration of the period in which the nonlinear interaction takes place. Integrated photonics reservoirs easily implement short memories and are limited by propagation losses for long memories; on the other hand, short memories require very fast optical modulation/detection, which might be outside the feasibility of current technology [11].

DOFs control

PICs are usually composed of several fundamental structures, each of which adds to the full circuits some DOFs. For example, the light transmitted in a waveguide is associated with an amplitude and a phase, which can be controlled via a Mach-Zehnder interferometer (MZI) and a phase shifter. Some devices will include control structures for most DOFs, while in the case of RCNs many of these DOFs will remain fixed (and only the output, if implemented optically, will contain active controls). The control of the DOFs affects the energy consumption, the losses, the footprint, and the general complexity of the device.

Scalability

Devices containing optically active elements provide the best scalability, as the information is renewed throughout the network, while optically passive devices, on the other hand, show reduced scalability, being limited by propagation and coupling losses. Nevertheless, passive networks remain of interest as they are easier to implement, and could be transparently added to already existing optical infrastructures.

Considering all the properties discussed so far, the possible combinations that could be used to create a neuromorphic photonics device are countless. There are several reviews on the matter, each capturing the state-of-the-art from a

different point of view [11, 28–32]. Nevertheless, given the fact that the work I performed within the BACKUP project is essentially focused on the realisation of RCNs with recurrent networks of microring resonators, it is convenient to focus on the main implementations of RCN with guided optics.

2.3.1 RCN in photonics

RC requirements are that of a fixed reservoir and a simple variable readout layer. These pair very well with the photonic platform because it is possible to create optical reservoirs producing highly nonlinear transformations of the input signal, whereas it proves to be more difficult to create simpler structures controlled efficiently. Photonics reservoirs are commonly divided into two categories: spatially distributed reservoir and delay-based reservoir [29].

Spatially distributed RCNs

The most direct implementation of a RCN is to physically create a number of nodes connected together by waveguides. The delay between the nodes and their nonlinearity enables the system to produce a highly nonlinear response. Then a subset of the nodes can be sampled by the output readout layer. It is also possible to use linear nodes; in this case, the delay and the square modulus operation on the coherent superposition of the light waves assure the required dynamics for the reservoir use. The choice of the connections between the nodes, i.e. the topology of the network, is at the heart of this approach as it must find a trade-off among all the characteristics affecting the performance positively, such as the number of nodes and their connectivity, and negatively, e.g. propagation losses and nonlinear losses.

The earliest, and so far the most successful, approach has been the “swirl” topology, in which nodes are connected in a square mesh with one-way links, resembling the movement of water in a whirlpool. Several variations of this approach have been simulated and manufactured using active or passive elements, and linear or nonlinear nodes [33–36].

Another interesting implementation is the one in which a quarter-stadium cavity is created using a silicon photonic crystal platform [37]. The system was numerically simulated and tested against XOR and AND binary operations and bit header recognition tasks with positive results. This example is also very important from a didactic point of view because it perfectly represents the idea of a reservoir system.

Delay based RCNs

The essential implementation of a delay-based reservoir is a nonlinear node and a feedback loop with delay: the result is a dynamic system described by delay differential equations (DDEs). The property of this system is that its response is not only determined by the current state of the system, but also influenced by the state of the system in the past, thanks to the delay loop. The signal resulting

from the interaction between the input, the feedback, and the state of the node at different times is sampled N times at *virtual nodes*. Additionally, the input signal is usually weighted with a random mask to increase the richness of the data representation. The values of the virtual nodes are then fed to the output layer which carries out the linear readout operation. Thus the virtual nodes, equally distributed in time, represent the input with a high dimensionality; their value is then read by the linear readout layer, whose weights are trained to carry out a given task.

The exceptional advantage of these systems is that they require a single physical node providing nonlinear elaboration. This simple structure relaxes the hardware requirements such that also common (fibre) optical communication components can be used. The drawback is that, by translating the spatial complexity in time, they often require opto-electronics modulation and detection components working at high frequency, i.e. N times that of a spatially distributed RCN. The first examples of this approach were realised with electronic components [38], but they were closely followed by optical implementations [24, 25, 39].

RCNs in BACKUP

The structure of the BACKUP project is based on the use of optically passive networks of MRRs, manufactured on the integrated silicon photonics platform, in order to create a reservoir showing complex dynamics. With this in mind, my principal lines of investigation have been:

- to generate a nonlinear response in the reservoir thanks to the dynamics of the single node, i.e. the MRR;
- to study new network topologies, with special consideration for scale-free networks;
- to develop an all-optical readout layer, in place of the more common approach which carries out the linear operation in the electrical domain.

The first goal requires accurate knowledge of the phenomena occurring inside a MRR. Linear propagation is affected by nonlinear effects involving the temperature and free carrier density in the MRR. Even a single node, provided with a sufficiently powerful input signal, can show extraordinarily rich dynamics. It is therefore important to correctly design such nodes so that the energy required by the nonlinear effects is minimised. The single node is, however, important as much as the way many of them are connected together. In this regard, the literature does not provide much information as only a few options have been explored so far (swirl). For this reason, I developed a simulation framework that allows to build arbitrary topologies and takes into account the dynamics of each node (see section 4.1). The last ambition is somewhat built as an extension of the realisation of the previous two goals. Indeed, even though limited by the propagation losses, matching a passive optical reservoir with an optical readout layer would maintain the power requirement to the minimum and the bandwidth

to the maximum, as the conversion between the electrical and optical domains would happen only at the beginning and end of the network.

Chapter 3

Silicon Photonics Fundamentals

Silicon photonics is the *de facto* standard platform for the development of passive integrated photonics devices. Its main advantages are: the compatibility with the complementary metal–oxide–semiconductor (CMOS) platform, hence exploiting the same manufacturing processes used in the consumer electronics production chain; low propagation losses; high integration density compared to other photonics platforms.

The fundamental structure of integrated photonics is the waveguide (WG), which channels light thanks to the difference in refractive index between its inner part, the core, and its outer part, the cladding. More complex structures can be built by putting together many WGs (splitters, combiners, rings), or by changing the shape of the WG (disks, multi-mode interference (MMI) devices, grating couplers). Furthermore, it is possible to create optical cavities, e.g. microring resonators, in which constructive interference occurs. In this case, the intensity of the electro-magnetic (EM) field may become so large that nonlinear effects are visible.

Optical nonlinear effects are those effects in which the propagation of light in a material, i.e. the polarisation vector P , is influenced by the intensity of light or higher-order terms in the EM field. In general, when light propagates in a material, even in WGs, the linear approximation holds very well. However, if the intensity of the EM field is large enough, nonlinear effects cannot be disregarded anymore. In integrated photonics, the EM fields inside optical cavities such as ring resonators can easily reach intensities at which nonlinear effects can be exploited.

We can separate nonlinear phenomena into two broad categories. The first one contains purely optical effects, that are defined by the expansion order of the refractive index as a function of the EM field. In these effects, the amplitude of the EM field directly affects the value of the refractive index. The second category is the one in which the refractive index change is mediated by another

physical quantity, e.g. the temperature of the material or the free carrier (FC) density in a semiconductor. Each effect acts either on the real or imaginary part of the refractive index, hence changing respectively the phase or the absorption of the propagated light, and on a specific timescale.

Silicon WGs show several nonlinear phenomena. Since silicon is a centrosymmetric material, second-order effects are non-existent. Hence, the most important nonlinear optical effects belong to the third-order and they are the Kerr effect (also known as intensity-dependent refractive index) and the two-photon absorption (TPA). Moreover, given that silicon is a semiconductor, FC-mediated effects, i.e. free carrier dispersion (FCD) and free carrier absorption (FCA), should be considered too. Finally, absorbed photons may generate heat inside the material and hence give rise to the thermo-optic effect (TOE), a mediated effect in which the temperature of the material changes the refractive index.

3.1 Theoretical description

A monochromatic transverse electro-magnetic (TEM) wave of optical frequency ν is represented by a complex wavefunction

$$\mathbf{E}(\mathbf{r}, t) = \mathbf{E}(\mathbf{r})e^{i\phi(\mathbf{r})}e^{-i2\pi\nu t}, \quad (3.1)$$

where $\mathbf{E}(\mathbf{r})$ and $\phi(\mathbf{r})$ are the complex amplitude and phase, respectively [40]. The exact form of both of them is obtained by solving the Helmholtz equation and depends on the specific boundary conditions of the system. The actual value of the electric field \mathbf{e} is obtained by taking the real part of the complex wavefunction

$$\mathbf{e}(\mathbf{r}, t) = \text{Re}\{\mathbf{E}(\mathbf{r}, t)\} = \frac{1}{2}(\mathbf{E}(\mathbf{r}, t) + \mathbf{E}(\mathbf{r}, t)^*), \quad (3.2)$$

where $*$ denotes the complex conjugate operation. The value of the magnetic field is then derived from \mathbf{e} through Maxwell's equations.

The use of the complex wavefunction is particularly convenient. Firstly, the complex notation describes the wave completely as the real field is simply its real part. Secondly, both are solutions of the wave equation and satisfy the same boundary condition. Furthermore, the complex notation simplifies the evaluation of the wave propagation as it can easily represent both the optical oscillation and the slower change in amplitude, due for example to absorption. Finally, the same description applies also to the magnetic field, although for simplicity it is usually omitted from the notation.

3.1.1 Light propagation in waveguides

The most important structure of integrated photonics is the waveguide (WG), as it allows the reliable delivery of optical signals from one part of the photonic integrated circuit (PIC) to another. Integrated WGs are dielectrics, in which light is confined within the core, thanks to total internal reflection (TIR) at

the interface between the core and the cladding. In order for TIR to happen, the core must have a higher refractive index in respect to the cladding, e.g. WGs with a silicon core embedded in a silica substrate, whose refractive indexes are respectively $n_{Si} \simeq 3.46$ and $n_{SiO_2} \simeq 1.44$ [41, 42]. Figure 3.1 shows the cross-section of a dielectric WG with a rectangular core.

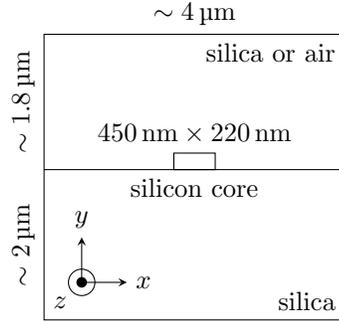


Figure 3.1: Cross-section of a waveguide with a rectangular core. Common sizes and axis orientation are shown.

Light propagates inside a 2D dielectric WG as a TEM wave, described by eq. (3.1), whose complex amplitude and phase can be obtained by solving the Helmholtz equation. Assuming that the refractive index n does not depend on z , the equation can be rewritten via separation of variables as a system of equations:

$$\nabla_z^2 \mathbf{E}_m(x, y) = -\beta_m^2 \mathbf{E}_m(x, y), \quad (3.3)$$

$$\nabla_{xy}^2 \mathbf{E}_m(x, y) = \left[\left(\frac{2\pi\nu}{c_0} \right)^2 n^2(x, y) - \beta_m^2 \right] \mathbf{E}_m(x, y). \quad (3.4)$$

The general solutions to these equations are the *modes* of propagation, indexed with the integer m :

$$\mathbf{E}(x, y, z) = \mathbf{E}_m(x, y) e^{i\beta_m z}, \quad (3.5)$$

TEM waves are associated with a specific transverse field profile $E_m(x, y)$ and a propagation constant, or wavenumber, β_m . Except for the 1D case of slab WGs, no analytical solution exists and both the field profile and the propagation constant have to be found numerically, solving eq. (3.4).

The transverse field profiles are classified initially in transverse electric (TE) and transverse magnetic (TM) modes¹, depending on their polarisation, and furthermore depending on the number of nodes, as shown in fig. 3.2. Moreover, for rectangular WGs, mode indexing is usually carried out with two integers, which represent the number of nodes along the x and y axis respectively, e.g. TE_{2,0} identifies a field profile with two nodes along the x axis, but none along the y axis. WGs with a different core geometry, such as the circular core of

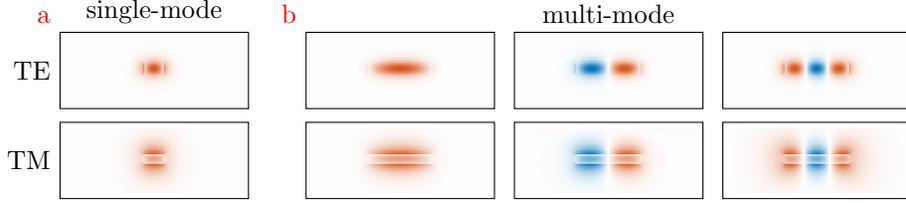


Figure 3.2: Examples of field profiles. **a** Field profiles for the TE and TM modes of a single-mode WG with a rectangular core of size 450 nm \times 220 nm. **b** Field profiles for the lowest TE and TM modes of a multi-mode WG with a rectangular core of size 1500 nm \times 220 nm.

optical fibres, may implement other indexing methods based on the inherent symmetries.

Apart from the classification based on the number of nodes in the field profile, modes can be compared quantitatively by defining the following two interesting quantities: the power confinement factor Γ

$$\Gamma = \frac{\int_{core} n^2 |\mathbf{E}_m(x, y)|^2 dx dy}{\int n^2 |\mathbf{E}_m(x, y)|^2 dx dy}, \quad (3.6)$$

which quantifies how much of the field profile is contained inside the core, and the effective area A^{eff}

$$A^{eff} = \frac{(\int n^2 |\mathbf{E}_m(x, y)|^2 dx dy)^2}{\int n^4 |\mathbf{E}_m(x, y)|^4 dx dy}, \quad (3.7)$$

which is a measure of the area covered by the mode field profile in the transverse plane. Both these quantities can be used to compare two modes travelling in the same or different WGs, without resorting to the complete description of the field profile. They will moreover prove to be useful when dealing with nonlinear effects in WGs (see section 3.3).

The maximum number of modes supported by a dielectric WG with a rectangular section can be approximated with $M \approx \pi(wh/\lambda^2)(n_H^2 - n_L^2)$, where w and h are the dimensions of the WG core, $\lambda = c_0/\nu$ is the wavelength of light, and n_H and n_L are the refractive indexes of the core and the cladding, respectively. Typically WGs are designed such that only the fundamental mode is supported, so that higher-order modes are not excited by defects or bends along the WG. Nevertheless, higher-order modes are commonly used inside individual structures, such as MMIs, or even in full PICs [43, 44]. For the sake of simplicity, for

¹The meaning of the terms TE and TM does actually make sense only in the 1D case of the slab WG, where in TE (TM) modes the electric (magnetic) field non-zero components are transverse to the propagation. This does not happen in 2D WGs, in which both electric and magnetic fields component along the propagation axis is, albeit small, non zero. For rectangular WGs fabricated using lithographic techniques, TE (TM) notation identifies modes whose major electric (magnetic) component is parallel to the substrate layer, i.e. along the x axis.

the rest of this work the use of the fundamental mode $\text{TE}_{0,0}$ is assumed, unless otherwise specified.

While the field profile describes the mode in the transverse plane, the propagation constant β_m defines the evolution of the complex amplitude and phase of the TEM wave along the z axis. Specifically, for each mode β_m can be rewritten in terms of the angular frequency $\omega = 2\pi\nu$ as

$$\beta_m = \frac{\omega}{c_0} n_m^{eff} = \frac{2\pi\nu}{c_0} n_m^{eff}, \quad (3.8)$$

where $n_m^{eff} = c_0/v_{ph}$ is the effective refractive index. The effective refractive index is, in principle, a complex scalar, whose real and imaginary parts are linked to propagation and losses, respectively. Indeed, the module of the complex amplitude decreases along z as a complex exponential $e^{-\alpha z}$, with the attenuation coefficient defined as $\alpha = \frac{4\pi\nu}{c_0} \text{Im} \{n_m^{eff}\}$. On the other hand, the real part of n_m^{eff} defines the wavefront position at each time, hence the phase velocity $v_{ph} = c_0/\text{Re} \{n_m^{eff}\}$. Indicatively, its value lies between the refractive index of the cladding and that of the core, $n_{clad} \leq \text{Re} \{n_m^{eff}\} \leq n_{core}$, and increase with the optical frequency. In fact, the effective refractive index and, thus, also the propagation constant are functions of the optical frequency. This effect is called chromatic dispersion.

Realistic optical signals, having a finite duration in time, are intrinsically non monochromatic: in general, the shorter the wavepacket, the broader the spectrum. For this reason, chromatic dispersion is an important characteristic of light propagation in WGs. If the mode is confined almost completely inside the core, the total chromatic dispersion will be very close to the dispersion of the material in the core [45, 46], such as in large mode area fibres. On the other hand, if the mode confinement and the effective area are low, such as for integrated WGs, the total chromatic dispersion will deviate from that of the bulk material. The extra component of the chromatic dispersion arises from the confinement of the field in the core, and thus from the WG geometry [45–47]. Single mode fibres work in a regime of weak confinement, however their chromatic dispersion can still be tailored by engineering the refractive index profile. In general, different modes show a different chromatic dispersion.

In order to account for the effects of chromatic dispersion, the propagation constant is expanded, usually up to the second-order,

$$\beta_m(\omega) = \beta_m(\omega_c) + \frac{\partial\beta_m(\omega)}{\partial\omega}(\omega - \omega_c) + \frac{1}{2} \frac{\partial^2\beta_m(\omega)}{\partial\omega^2}(\omega - \omega_c)^2, \quad (3.9)$$

where the frequency dependency is expressed in terms of angular frequency $\omega = 2\pi\nu$ instead of optical frequency, to be consistent with wave theory. From this description, several important quantities can be extracted for each mode. The first one is the group velocity v_g , i.e. the velocity at which the envelope of a wavepacket travels, measured in m/s. Additionally, the group index n_g is defined as the ratio of the speed of light c_0 in vacuum to the group velocity v_g .

Both of these are obtained from the first derivative of the propagation constant:

$$\frac{1}{v_g} = \frac{n_g}{c_0} = \frac{\partial \beta_m(\omega)}{\partial \omega}. \quad (3.10)$$

Moreover, the group velocity dispersion (GVD), i.e. the variation of the group velocity for different optical frequencies, is defined as the second-order derivative of the propagation constant

$$\text{GVD} = \frac{\partial^2 \beta_m(\omega)}{\partial \omega^2}, \quad (3.11)$$

having basic unit of measure is s^2/m [48]. Higher-order derivatives are usually referred to as *j-th order dispersion*, but are rarely used. Table 3.1 reports the typical values of α , n_g , and GVD for standard single-mode WGs in silicon and optical fibres.

Parameter	Silicon waveguide	Optical fibre
α	1-2 dB/cm [49]	0.2 dB/km [50]
n_g	3.9 [46]	1.47 [50]
GVD	-1300 ps ² /km [46]	-20 ps ² /km [50]

Table 3.1: Typical values for the parameters described in this section for a silicon WG embedded in silica of size $450 \text{ nm} \times 220 \text{ nm}$ measured with IR light at 1550 nm and a standard SMF-28 optical fibre.

3.1.2 Mode coupling

Since a solution to eqs. (3.3) and (3.4) is still a solution if multiplied by a constant factor, the following normalisation of the field profile is allowed:

$$\int \mathbf{E}_m(x, y) \mathbf{E}_m^*(x, y) dx dy = 1. \quad (3.12)$$

Moreover, it can also be demonstrated that the functions $E_m(x, y)$, representing the field profiles, are orthogonal, i.e.

$$\int_{-\infty}^{\infty} \mathbf{E}_m(x, y) \mathbf{E}_n^*(x, y) dx dy = 0 \quad \text{for } m \neq n, \quad (3.13)$$

at least for lossless WGs [51]. It is, therefore, possible to express an arbitrary field $\mathbf{E}_A(x, y)$ as a superposition of modes:

$$\mathbf{E}_A(x, y, z) = \sum_m a_m \mathbf{E}_m(x, y) e^{i\beta_m z}, \quad (3.14)$$

where

$$a_m = \int_{-\infty}^{\infty} \mathbf{E}_A(x, y, 0) \mathbf{E}_m^*(x, y) dx dy \quad (3.15)$$

is the mode overlap, or degree of similarity, between the original field $\mathbf{E}_A(x, y)$ and the field distribution $\mathbf{E}_m(x, y)$ of the mode m [40]. In other words, an input field with an arbitrary profile produces the excitation of the modes specific to the WG, which will then propagate one independent from the other.

The clearest example of mode excitation is when an arbitrary field is placed at the input of a single-mode WG. The resulting propagation is described in the following way:

$$\begin{aligned} \mathbf{E}_A(x, y, z) &= a_0 \mathbf{E}_0(x, y) e^{i\beta_0 z} + \sum_{m \geq 1} a_m \mathbf{E}_m e^{-\text{Im}\{\beta_m\}z} \\ &\simeq a_0 \mathbf{E}_0(x, y) e^{i\beta_0 z} \quad \text{for } z \gg 0, \end{aligned}$$

assuming $\beta_m = \text{Im}\{\beta_m\} > 0$ for $m \geq 1$ for simplicity. Higher-order modes show a decay in amplitude which is given by the imaginary part of their propagation constant. It is then easy to see that the only mode propagating unperturbed is the fundamental one, while the others quickly disappear. The power transferred from the initial arbitrary field to the fundamental mode will be proportional to the square modulus of the degree of similarity of their profiles $P_{WG} \propto |a_0|^2$. In general, thanks to eq. (3.13) and eq. (3.12), the following equation holds:

$$P_0 \propto \int |E(x, y, 0)|^2 dx dy = \sum_m |a_m|^2 \propto \sum_m P_m.$$

Evanescent coupling

A further case of mode coupling is observed when two WGs are placed one near the other. In this case, the distribution tails of the evanescent field outside the core of one WG may overlap the field profile of the mode propagating in the second WG. The closer the WGs are, the higher the overlap between the two field distributions is, and consequently the larger the exchange of optical power per unit length will be. In the general case, the power exchange can be quite difficult to evaluate and numerical simulations are employed. Structures that employ evanescent coupling in order to split or combine optical power from different WGs are called directional couplers and will be described more in detail in section 3.1.3.

Out-of-chip coupling

One of the main challenges of integrated photonics is delivering light to and reading signals off of PICs, whether they are passive or integrating light sources and detectors. This is due to the fact that the modes in integrated WGs have such a small area in comparison to the modes propagating in optical fibres and even more to free space beams. Indeed, the typical effective mode area A^{eff} of a WG is of the order of $1 \mu\text{m}^2$, whereas for common optical fibres its value is of the order of $100 \mu\text{m}^2$. Direct coupling between an optical fibre and a WG core will result in a loss of even 20 dB of optical power [52]. The only way of increasing the efficiency from that baseline is to introduce, between the fibre

and the WG, a structure that converts the size of one of the two modes to be similar to the size of the other.

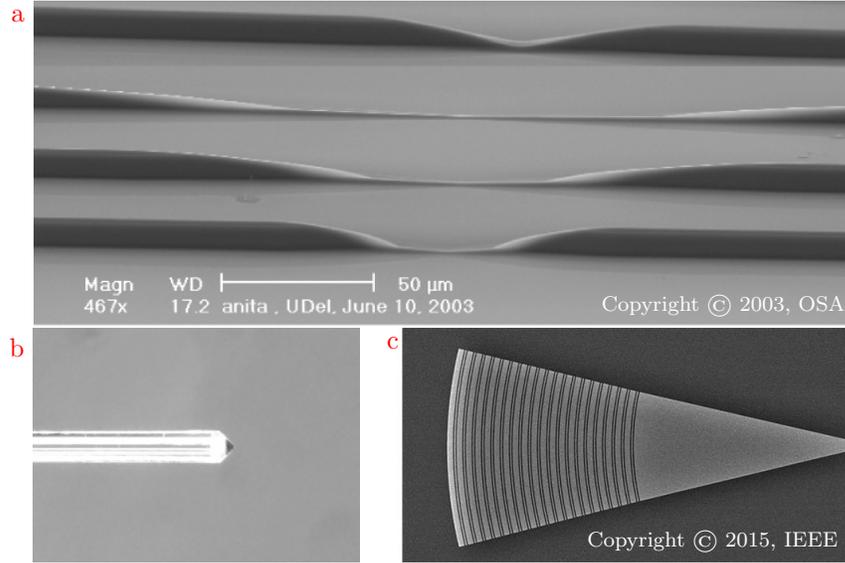


Figure 3.3: Examples of input/output couplers. **a** Tapered structure for edge coupling, taken from [53]. **b** Optical fibre with lensed tip. **c** Apodized grating coupler, taken from [54].

Coupling methods can be separated into two broad categories: edge coupling and surface coupling. Edge coupling happens on the perimeter of the PIC and the input light comes from the same plane of the WG, while surface couplers can be positioned at any place on the surface of the chip and couple to light coming at a certain angle to the normal of the chip surface.

One approach to improving edge coupling is to change the input fibre, e.g. using tapered fibres. Normally, the facets of fibres are cleaved perpendicularly to their axis, so that the transmission from the fibre to the air is maximised, but at the same time, they still suffer from light diffraction due to the small aperture of the core. In tapered fibres, instead, the fibre facet can be tapered in the form of a micro-lens, as shown in fig. 3.3b. The lens can be engineered so that light is focused on a very small spot around the WG core, effectively increasing the mode matching and reducing the insertion losses. The drawback is that tapered fibres are relatively difficult to manufacture and are quite fragile to handle.

A second technique requires the engineering of the input facet of the WG. There are mainly two ways of doing so and both are obtained by adiabatically varying the WG shape. In direct tapering, the size of the WG is increased as much as possible, while in inverse tapering the core is decreased in size, delocalising the mode outside its core. The scanning electron microscope (SEM) image of direct tapering structures is shown in fig. 3.3a. This is an effective way of reducing the

insertion losses to even less than a dB, however, tapered couplers might occupy a considerable surface, as adiabatic tapering require long WGs [53, 55].

Surface couplers, on the other hand, are primarily represented by grating couplers. In these structures, the refractive index of the core is periodically modulated, diffracting an incoming wave into the WG or vice versa. In the simplest case, the angle of this wave is found by satisfying the following equation:

$$\beta_m \Lambda - k_0 n_{clad} \Lambda \sin \theta = 2\pi q \quad \text{with} \quad q = \pm 1, \pm 2, \dots, \quad (3.16)$$

where $k_0 = 2\pi\nu/c_0$ is the wavevector in vacuum, n_{clad} is the refractive index of the cladding material, and Λ is the periodicity of the grating. Grating couplers provide good insertion losses with a relatively small footprint and easy and robust alignment with basic cleaved optical fibres, however, they may be challenging to manufacture due to the small feature size of the periodic modulation and have narrower transmission spectra in comparison to edge coupling methods. More advanced surface coupling is obtained with the addition of micro-machined couplers, reaching insertion losses as low as 1.8 dB on a much broader spectrum in respect to simple grating couplers [56, 57].

Finally, at least in the case of low power signals, the coupling can be improved further by using index matching liquids. For example, glycerol is a colourless, non-toxic liquid with a refractive index of 1.4746 [58], very close to that of glass. The use of glycerol decreases drastically the reflection on the interface facets.

3.1.3 Fundamental structures

A great deal of complexity can be built upon clever use of simple WGs. For example, power splitters are realised by placing two parallel WGs very near, while connecting a WG to itself creates an optical cavity. In the following paragraphs, I will give a brief overview of the most important structures.

Directional coupler

Directional couplers (DCs) are obtained when two WGs are brought together such that in a specific area, called coupling region, they travel parallel and very close one to the other, as shown in the drawing of fig. 3.4a. Thanks to evanescent coupling, optical power is exchanged between the two WGs. Broadly speaking, a smaller gap, i.e. the distance between the guides, provides a faster exchange of the optical power between the WGs. The value of the gap is usually in the order of a few hundreds of nm. Apart from the exact solution, given by solving Maxwell's equations for the whole structure, there are mainly two methods which give good DC approximations: first, considering each of two WGs as a small perturbation to the other, and second, considering the combination of the two WGs as composite structure [40, 59].

The former approach requires the determination of the variation of the field amplitude along z through coupled differential equations. In each differential equation the evanescent field of the other WG will act as a source. From this

method one obtains that only modes whose propagation constants are very similar, $\beta_1 - \beta_2 \simeq 0$, are able to effectively exchange optical power.

The latter method allows the description of the input field in terms of the modes proper of the whole structure, i.e. *normal modes* or *supermodes*. It is found that the mode of an individual WG can be expressed as a superposition of the *even* and *odd* modes [59], an example of which is shown in fig. 3.4b. The resulting field profile at the output of the DC is then given by the interference between the normal modes which have independently travelled along the coupling region with propagation constant β_e and β_o .

The most notable case is where the section of the two WGs is the same and the gap is uniform throughout the length of the coupling region, as shown in fig. 3.4a. In this case, assuming optical power P_0 entering from the first WG, it can be demonstrated that [40, 59]:

$$\begin{aligned} P_1(L) &= P_0 \cos^2(\Delta\beta L) = |\tau|^2 P_0 \\ P_2(L) &= P_0 \sin^2(\Delta\beta L) = |\kappa|^2 P_0, \end{aligned} \quad (3.17)$$

where $\Delta\beta = (\beta_e - \beta_o)/2$ is the phase mismatch per unit length between the *even* and *odd* modes. The transfer length L_0 , that is the distance at which complete power transfer is obtained, is related to the phase mismatch:

$$L_0 = \frac{\pi}{2\Delta\beta} = \frac{\pi}{\beta_e - \beta_o} \quad (3.18)$$

In principle, also the WG bends before and after the coupling region affects the total splitting ratio. This is especially true when $L = 0$, i.e. in the case of point coupling, for which a more refined model is required.

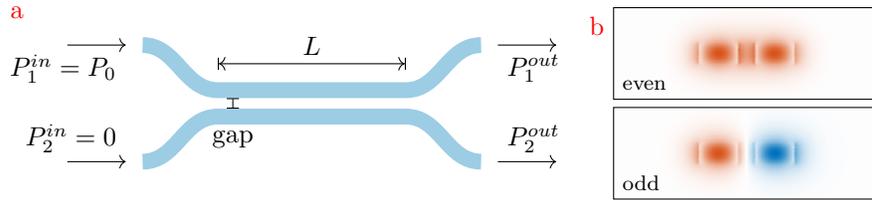


Figure 3.4: Symmetrical DC drawing and field profile. **a** Drawing of a symmetrical DC. **b** Field profiles of the *even* and *odd* modes.

In photonic integrated circuits, DCs are employed for the vast majority as power splitters or combiners. In the former, the optical power of a single input is distributed between two outputs, while in the latter the signal from two inputs is combined and distributed to two output channels, one of which may be designed as a dead-end. The most common splitting ratio is 50%/50%, or 3 dB, as it allows to divide the input power evenly between both output channels. On the other hand, highly unbalanced splitting ratios, such as 99%/1%, or 10 dB, are often used to probe the main signal, by drawing the smallest optical power necessary for a good signal-to-noise ratio (SNR). Other values of splitting ratio

may be used to generate an arbitrary number of channels starting from a single one, i.e. a $1 \times N$ power splitter.

A combiner with a 0%/100% splitting ratio may be used to swap the signals carried by the two WGs, which is useful for the routing of WGs in the PICs design. However, for this purpose specific structures called WG crossings are used instead, because they provide low insertion losses, low cross-talk, and a much smaller footprint [60].

Especially when both input WGs in a coupler are excited with an optical signal, a matricial form is used to represent the relation between the input and output fields:

$$\begin{pmatrix} E_1^{out} \\ E_2^{out} \end{pmatrix} = \begin{bmatrix} \tau & i\kappa \\ i\kappa & \tau \end{bmatrix} e^{i\beta L} \begin{pmatrix} E_1^{in} \\ E_2^{in} \end{pmatrix}, \quad (3.19)$$

where energy conservation condition is obtained for $\tau^2 + \kappa^2 = 1$. This notation is called *transfer matrix* [59].

Asymmetric DCs are employed instead to obtain mode conversion [61, 62]. A single-mode WG is phase-matched with a higher-order mode in a wider WG, effectively transferring optical power from the fundamental mode to a higher one. Figure 3.5 shows this effect through the evolution of the field profile amplitude.

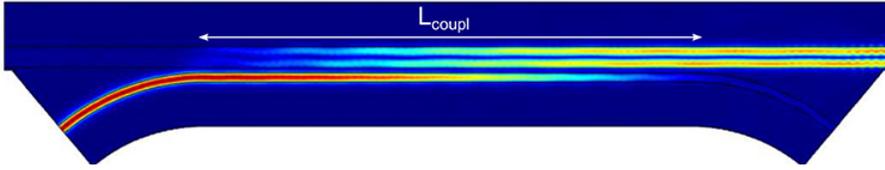


Figure 3.5: Evolution of the field profile in an asymmetrical DC. Reproduced from [61].

Multi-mode interference coupler

A multi-mode interference (MMI) coupler, similarly to the directional coupler, is ordinarily used as a power splitter or combiner. An MMI can be described in general as an $N \times M$ structure, in the sense that N input and M output WGs are connected to the main body. The centre of the MMI is, in its simplest design, a rectangle whose width is (much) larger than that of the WGs. Hence, the body can be thought of as a WG supporting multiple modes.

The optical power coming from the input WG (or WGs) will excite guided modes of the central body. Modal propagation analysis can be used to determine the field at the end of the multi-mode WG. Assuming the multi-mode WG allows M guided modes, then

$$\mathbf{E}(x, y, z) = \sum_{m=0}^{M-1} a_m \mathbf{E}_m(x, y) e^{i\beta_m z}, \quad (3.20)$$

where c_m are the overlap integrals as defined in eq. (3.15). If the input profile is narrow enough to be contained within the MMI body section, then only the guided modes are excited [63].

Moreover, the concurrent propagation of multiple modes in a WG gives rise to the self-image effect. According to the self-imaging principle, the field profile at the input of the WG will be replicated in *single or multiple images* along the path of propagation at periodic locations [63]. The characteristic distance for the generation of images is the *beat length* of the two lowest order modes

$$L_\pi = \frac{\pi}{\beta_0 - \beta_1}, \quad (3.21)$$

where β_0 and β_1 are their propagation constants. The general formula for determining the length at which N images of the input field will be produced is

$$L_N = \frac{p}{N}(3L_\pi) \quad (3.22)$$

where $N \geq 1$ and $p \geq 0$ are integers with no common divisor [63]. For example, a pair of images is created at $L = p(3L_\pi)/2$. However, when the input field respects specific symmetries, the length at which the images will be formed can be greatly decreased. For example, for a symmetrical input field profile only the even modes will be excited and the characteristic length to obtain N images will be reduced by a factor of 3 to $L = p(L_\pi)/N$. This feature is important to keep a small footprint. Figure 3.6a shows the evolution of the field profile in the propagation direction within the multi-mode WG, for a narrow and symmetrical input profile, placed in the middle of the multi-mode WG. It is easy to spot the locations at which single images appear and even a pair of images are spotted effortlessly. Nevertheless, the higher the number of images, the closer they become to the source and each other, and the more difficult their identification is by the naked eye. Spotting the locations where multiple images appear becomes even more difficult for the asymmetric excitation case, reported in fig. 3.6b. In this case at the output a horizontally mirrored image is created; this structure can be used as WG crossing.

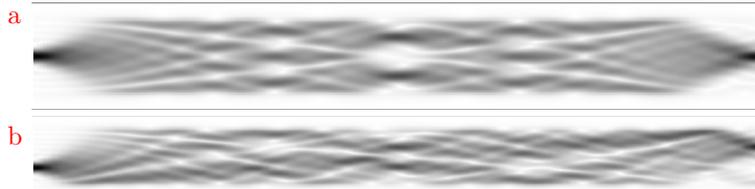


Figure 3.6: Plot of the field profile of the field propagation of an MMI approximated as a 2D structure evaluated through FEM methods. **a** Symmetrical excitation of an MMI body. **b** Asymmetrical excitation of an MMI, producing a mirrored image at the output.

The design of MMIs can be further engineered to reduce the excess loss and the cross-talk between the different channels. For example, tapered input (output) WGs may improve the overlap integrals, increasing the power transferred to

(from) the MMI modes. Moreover, when the power has to be distributed between the input and outputs in a non-even fashion, a general approach has to be used. In section 4.4 I will present a simple, yet effective numerical model, employing modal propagation analysis for the optimisation of MMI couplers geometries.

Microring resonator

A microring resonator (MRR) is an optical cavity composed of a single WG closed on itself. Light is coupled inside the cavity by bus WGs through evanescent coupling in the coupling regions. When a single bus WG is present, the structure takes the name of all-pass filter (APF). In this case, the optical signal enters from the *input* of the bus WG, goes *through* the MRR, and then exits the system. For this reason, the two ports of the bus WG are called *input* and *through*. A MRR with two bus WGs is called add-drop filter (ADF). It is a structure with four ports: the ports on the first bus WG retain the name of *input* and *through*, while the additional ports are the *add* and *drop* ports. Figure 3.7 shows the drawings of both an APF and an ADF where the names of the input/output ports are explicitly reported.

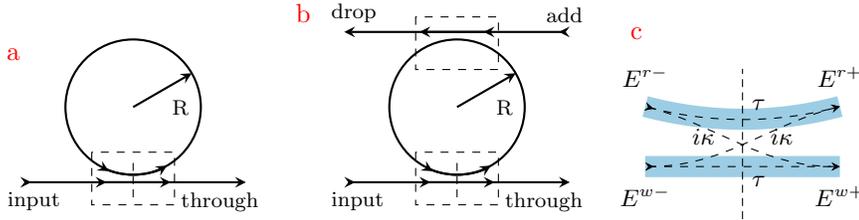


Figure 3.7: Drawing of MRRs in the **a** all-pass filter and **b** add-drop filter configurations. The coupling coefficients κ represent the portion of the field which is exchanged between the bus WGs and the MRR in each coupling region, as shown in **c**.

The spectral response of MRRs is governed by the interference effects occurring inside the cavity and varies depending on the geometry of the ring, the coupling coefficients, the propagation losses, and last but not least the choice of the output port to observe. To simplify the discussion, I will assume that the structure is composed of single-mode WGs only, hence dropping the subscript from the amplitude a_m , and that light propagates from the *input* port to the *through* and *drop* ports only, i.e. no additional excitation is present at the *add* port. The excitation field has the amplitude a^{in} such that $|a^{in}|^2 = 1$ at the *input* port and a^{th} and a^{dr} are the field amplitudes at the *through* and *drop* ports respectively. When light entering the MRR interferes constructively with that already present inside, the energy stored in the cavity increases.

Light propagation inside the ring and to/from the ports at the coupling regions is described by a system of equations. From this system, the ratios

$\eta^{th} \equiv a^{th}/a^{in}$ and $\eta^{dr} \equiv a^{dr}/a^{in}$ can be extracted:

$$\eta^{th} = \tau \frac{1 - e^{i\phi_{rt}}}{1 - \tau^2 e^{i\phi_{rt}}}, \quad (3.23)$$

$$\eta^{dr} = \frac{-\kappa^2 e^{i\phi_{rt}/2}}{1 - \tau^2 e^{i\phi_{rt}}}, \quad (3.24)$$

where $\phi_{rt} = \beta_m L$ is the complex phase added by a round trip inside the cavity and τ and κ are the coupling coefficients, as described in eq. (3.19). The two coupling regions have been assumed identical, i.e. with $\kappa_1 = \kappa_2 = \kappa$. Note that the real and imaginary parts of ϕ_{rt} represent respectively the phase shift and amplitude loss given by a round trip. To simplify the notation, the mode index m will be dropped, assuming fundamental mode propagation, and $n^{eff} = \text{Re}\{n^{eff}\}$ will be assumed to be a real scalar function of the optical frequency ν while and its imaginary part will be described by the attenuation coefficient $\alpha = (4\pi\nu)/c_0 \text{Im}\{n^{eff}\}$:

$$\phi_{rt} = \left(\frac{2\pi\nu}{c_0} n^{eff}(\nu) + i \frac{\alpha}{2} \right) L.$$

This notation reinforces the fact that the phase shift is dependent on the optical frequency of the signal, whereas the amplitude attenuation can be effectively considered as a constant (at least over a restricted frequency span). Figure 3.8 shows the spectral response of a MRR in the ADF configuration, reporting the square modulus of η^{th} and η^{dr} as functions of the phase shift.

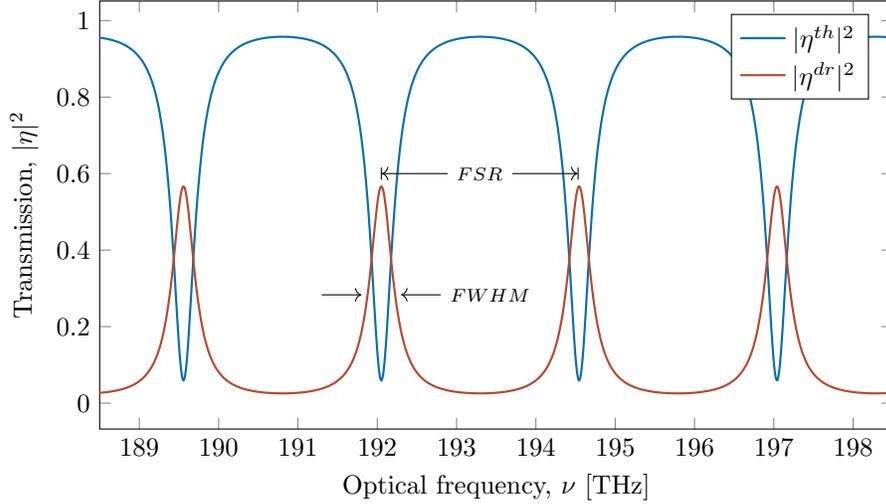


Figure 3.8: Spectrum of a MRR in the add-drop filter configuration. Solid and dotted lines represent the transmission spectra of the *through* and the *drop* ports, respectively. Arrows indicate the FSR and the full width at half maximum, $\delta\nu$.

The cavity resonance condition occurs when the phase shift in a round trip is a multiple of 2π :

$$\text{Re}\{\phi_{rt}\} = \frac{2\pi\nu_n}{c_0} n^{eff}(\nu_n)L = 2\pi n \quad \text{for } n \in N, \quad (3.25)$$

where L is the geometrical length of the optical cavity. Hence, the frequency (wavelength) position of each resonance is given by

$$\nu_n = \frac{c_0}{\lambda_n} = n \cdot \frac{c_0}{n^{eff}(\nu)L}. \quad (3.26)$$

From this equation and employing eq. (3.9), the distance between two consecutive resonances ν_F , i.e. the free spectral range (FSR), becomes

$$\nu_F = \nu_{n+1} - \nu_n \simeq \frac{c_0}{n_g(\langle\nu\rangle)L}, \quad (3.27)$$

where $\langle\nu\rangle$ is the mean value of optical frequency between ν_n and ν_{n+1} . Similarly, employing the same expansion of eq. (3.9) around a given resonance ν_n , the full width at half maximum (FWHM) of the peak (or dip) can be evaluated to

$$\Delta\nu_n \simeq \frac{c_0}{\pi n_g(\nu_n)L} \frac{1 - \tau^2 e^{-\alpha L/2}}{\tau e^{-\alpha L/4}}. \quad (3.28)$$

This quantity is strictly related to the losses in the cavity, i.e. the coupling constant κ and the round trip loss $e^{-\alpha L}$, and in fact is slightly different for an APF. These spectral quantities are actually the characteristics that are easier to measure experimentally. However, a few other interesting properties can be defined. The first one is the quality factor Q , which is defined as

$$Q \equiv 2\pi \frac{\text{energy stored}}{\text{energy loss per cycle}} = \frac{\nu_n}{\Delta\nu_n}. \quad (3.29)$$

The quality factor is also associated with the photon lifetime τ_{ph} in the cavity:

$$\tau_{ph} = \frac{Q}{2\pi\nu}. \quad (3.30)$$

Since the MRR is a passive structure, the optical power at the outputs should be less or at most equal to the power at the input. However, due to constructive interference at and close to the resonance frequencies, the field inside the cavity is higher with respect to the one outside. The (field) enhancement factor EF is defined as the ratio between the field inside the cavity and the field outside and for an ADF microring resonator is:

$$EF = \frac{-\kappa}{1 - \tau^2 e^{i\phi_{rt}}}. \quad (3.31)$$

It is essential to notice that all the principal quantities are strictly related to the value of the n^{eff} at the given frequency or to its derivatives, e.g. n_g . This is

important because the enhanced values of the EM field inside the cavity are able to trigger nonlinear effects, which alter the value of n^{eff} , affecting in turn the coupling of the field to the cavity. This chain of events is able to produce very complex dynamics and will be studied in the following chapters as a physical reservoir in a reservoir-computing network (RCN). This concludes the very brief summary of the most important properties describing MRRs, which allow an easier comparison of their response between each other.

Sequences of microring resonators

Several MRRs can be placed together in order to create a composed structure, where the response is given by the effect of the single cavities and by their interplay. The two most famous types of structures are the coupled resonators optical waveguide (CROW) and the side coupled spaced sequence of optical resonators (SCISSOR).

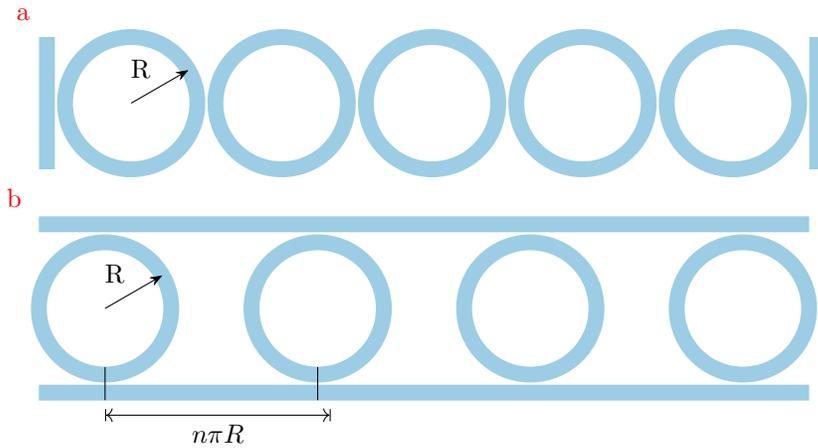


Figure 3.9: Drawing of sequences of microring resonators: **a** CROW and **b** SCISSOR.

In the CROW, several microring resonators are coupled side by side in a chain and only the first and last are coupled to a bus WG each. Figure 3.9a shows the drawing of an example of a CROW device with five rings. In the SCISSOR, several MRRs in the add drop configurations are cascaded such that the *through* and *add* ports of each resonator are connected to the *input* and *drop* ports of the next one. Figure 3.9b shows the drawing of a SCISSOR structure with four rings.

When the MRRs that compose these sequences are identical or at least similar, the complete structure has a certain degree of periodicity. The periodicity given by multiple rings gives rise to distributed Bragg reflections. The device transmission will hence be affected by both the resonances of the individual MRRs and those given by the distributed Bragg reflections. Depending on the specific phase condition between the MRRs, which is defined by the geometrical

distance between the MRRs and the dispersion of the material, these families of resonances could occur at the same optical frequency or far apart. In the first case, the coupled resonators induced transparency (CRIT) effect emerges and a transparency window opens within the MRR resonance band of high reflection. CRIT is the integrated photonics analogous to the phenomenon of electromagnetic-induced transparency in atoms. In the second case, new windows of increased reflection appear in between the resonances of the single MRRs.

The optical response of sequences of microring resonators becomes even more intricate when the optical power is increased and optical nonlinear effects come into play. In this case, the spectral analysis is not sufficient to completely understand the response of the structure and the full temporal evolution of the system must be considered to thoroughly comprehend the dynamics involved. Indeed, it has been shown [64] that even a constant input signal might trigger an incredibly complex or even chaotic regime in structures such as SCISSORS.

3.2 Nonlinear Optics in Silicon

Silicon is a semiconductor material, transparent in the infrared (IR) region. Its energy bandgap is $E_g \simeq 1.12$ eV, which corresponds to the energy of a photon of frequency ~ 270 THz or wavelength ~ 1.1 μm .

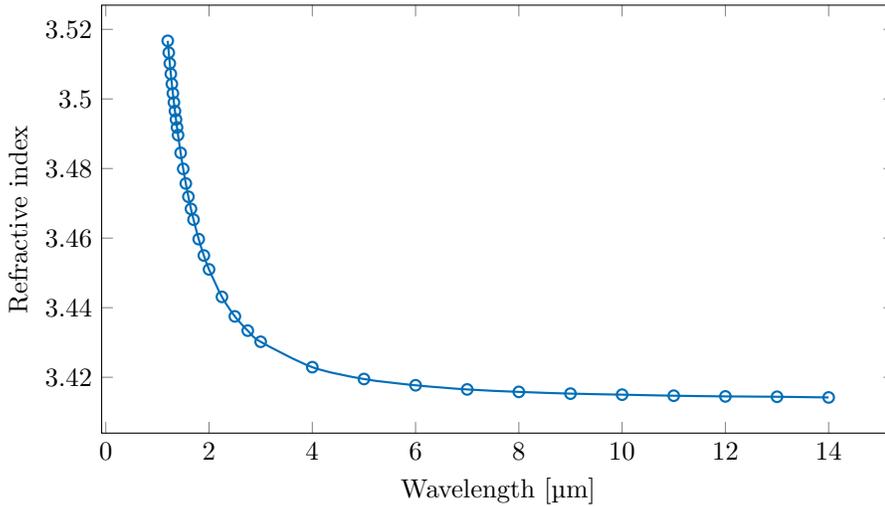


Figure 3.10: Refractive index of silicon on a wide range of wavelengths [41].

For most purposes, silicon can be considered a linear medium. Even in silicon photonics, propagation is usually regarded as linear. However, when dealing with high optical power and/or the high enhancement factor of optical cavities, nonlinear optical effects have to be considered.

There are three sources of nonlinearities in silicon: atomic polarisation, temperature, and the free carrier (FC) population density of the material. The po-

larisation in the material is directly influenced by the EM field of the optical signal (and by other static fields), such that the amplitude of the propagating field depends on itself and those at other optical frequencies. On the other hand, both temperature and FC population density parametrically change light propagation in the medium by altering the refractive index of the material. Nevertheless, they can be considered a nonlinear effect because, in some circumstances, the amplitude of the EM field influences the temperature shift and the FC population density, which, in turn, affect the refractive index, hence the propagation of light.

3.2.1 Optical Effects

The electronic cloud of atoms in a dielectric medium becomes polarised by interacting with an EM field. Usually, such polarisation is approximated linearly with the field amplitude, as the EM field is smaller than the inter-atomic electric fields. Nevertheless, when they reach values from 10^5 to 10^8 V/m [40], e.g. in integrated optical cavities, the former approximation does not hold anymore. The relation between the polarisation vector \mathbf{P} and the field vector \mathbf{E} must be expanded with higher-order terms:

$$\begin{aligned}\mathbf{P} &= \varepsilon_0 \left(\chi^{(1)} \cdot \mathbf{E} + \chi^{(2)} : \mathbf{E}\mathbf{E} + \chi^{(3)} \vdots \mathbf{E}\mathbf{E}\mathbf{E} + \dots \right) \\ &= \varepsilon_0 \chi^{(1)} \mathbf{E} + \varepsilon_0 \Delta\chi \mathbf{E}\end{aligned}\quad (3.32)$$

where the j -th order susceptibility term $\chi^{(j)}$ is a tensor of rank $j + 1$ [48]. For TEM waves travelling in a homogeneous medium, the spatial dependence is commonly simplified and eq. (3.32) is reduced to a scalar equation. From this, we can extend the definition of the refractive index to also include nonlinear effects:

$$n = n_0 + \Delta n = \sqrt{1 + \chi^{(1)} + \Delta\chi} \simeq n_0 + \frac{1}{2n_0} \Delta\chi, \quad (3.33)$$

which is allowed for small perturbations on χ .

Generally, lower-order effects show a larger response with respect to higher order effects. However, given that silicon is a centrosymmetric material, there are no second-order nonlinear effects. Third-order susceptibility, therefore, becomes the lowest effective order. Specifically, they are the Kerr effect (also known as intensity-dependent refractive index) and the two photon absorption (TPA). The response time of both the Kerr effect and the TPA is approximately in the order of femtoseconds and, hence, can be considered effectively an instantaneous response [47]. There are other effects as well, such as those in which multiple optical frequencies are involved, e.g. four-wave-mixing (FWM), or those which arise from the presence of a static electric field E_{DC} . When only a monochromatic signal travels along the WG, these latter effects are usually less significant in respect to Kerr and TPA, so they will be neglected in this work.

The Kerr effect is represented by the real part of the third-order nonlinear susceptibility $\chi^{(3)}$ and it causes a change in the refractive index given by:

$$\Delta n_{Kerr} = n_2 I, \quad (3.34)$$

where I is the intensity related to the EM field and n_2 is the second-order term of the refractive index nonlinear expansion, also known as intensity-dependent refractive index, and is defined by

$$n_2 = \frac{3}{4\varepsilon_0 c_0 n^2} \text{Re} \left\{ \chi^{(3)} \right\}. \quad (3.35)$$

Its value in silicon is found to be $n_2 \simeq 4.5 \times 10^{-9} \text{ } \mu\text{m}^2/\text{mW}$ [65].

Two photon absorption (TPA), on the other hand, is linked to the imaginary part of the third-order nonlinear susceptibility. Its effect is to increase the imaginary part of the refractive index κ , and it is usually represented as a change in the absorption coefficient α , where

$$\Delta \alpha_{TPA} = \beta_{TPA} I, \quad (3.36)$$

where I is again the optical intensity and β_{TPA} is the TPA coefficient defined as

$$\beta_{TPA} = \frac{3\pi\nu}{\varepsilon_0 c_0^2 n^2} \text{Im} \left\{ \chi^{(3)} \right\}. \quad (3.37)$$

Its value in silicon is found to be $\beta_{TPA} \simeq 8 \times 10^{-9} \text{ } \mu\text{m}/\text{mW}$ at wavelengths around $1.55 \text{ } \mu\text{m}$ ($\sim 193.3 \text{ THz}$) [65, 66]. A consequence of the absorption of two photons is that an electron is promoted from the valence band to the conduction band of the material, hence generating a FC that causes nonlinear effects described in the next section.

3.2.2 Free Carrier Effects

When a high power optical beam crosses a semiconductor medium, FCs are generated in the conduction band, changing their concentration population N . The existence of these FCs on one hand produce an excess polarisation of the material, while on the other hand they are able to absorb single photons and subsequently thermalise. The first case is the free carrier dispersion (FCD) effect, which affects the phase of the propagating signal through the real part of the refractive index n , while the second case is the free carrier absorption (FCA) effect, which changes the imaginary part of the refractive index κ , and thus also the absorption coefficient α increasing the propagation losses due to absorption.

These effects are commonly described as a first-order expansion around the equilibrium condition:

$$n(N) = n(N_0) + \left. \frac{dn}{dN} \right|_{N_0} (N - N_0), \quad (3.38)$$

$$\alpha(N) = \alpha(N_0) + \left. \frac{d\alpha}{dN} \right|_{N_0} (N - N_0), \quad (3.39)$$

where the values of the coefficients in silicon are

$$\begin{aligned}\left.\frac{dn}{dN}\right|_{N_0}^{Si} &= \sigma_{FCD}^{Si} = -1.73 \times 10^{-21} \text{ m}^{-3}, \\ \left.\frac{d\alpha}{dN}\right|_{N_0}^{Si} &= \sigma_{FCA}^{Si} = 1.1 \times 10^{-15} \text{ m}^{-2},\end{aligned}\tag{3.40}$$

respectively [65, 66].

In principle, a more refined description could be used to distinguish the effect of electrons and holes on n and α [67]. In this case, the formula for the refractive index becomes [68]:

$$\begin{aligned}\Delta n(N) &= \sigma_{FCD}^{Si}|_e (\Delta N_e)^a + \sigma_{FCD}^{Si}|_h (\Delta N_h)^b, \\ \Delta \alpha(N) &= \sigma_{FCA}^{Si}|_e (\Delta N_e)^c + \sigma_{FCA}^{Si}|_h (\Delta N_h)^d,\end{aligned}\tag{3.41}$$

where each coefficient changes slightly with the wavelength [69]. Hence, the number of coefficients required to describe both effects grows from two to eight. The refinement of eqs. (3.38) and (3.39), however, must be weighted with the increased complexity and the accuracy of the available estimates of each coefficient. Thus, since the literature does not provide sufficient agreement on these values, the first model is used in the rest of this work.

3.2.3 Thermal Effects

The temperature of the material affects the value of its refractive index. This is called thermo-optic effect (TOE) and it is usually modelled as a first-order expansion in temperature T :

$$n(T) = n(T_0) + \left.\frac{dn}{dT}\right|_{T_0} (T - T_0),\tag{3.42}$$

where $\left.\frac{dn}{dT}\right|_{T_0}$ is the thermo-optic coefficient (TOC), which in silicon at $T_0 = 300$ K has value [66]:

$$\left.\frac{dn}{dT}\right|_{300\text{K}}^{Si} = 1.86 \times 10^{-4} \text{ 1/K}.\tag{3.43}$$

When light propagates inside a medium, its optical power is partially absorbed and heats the material, which in turn changes its refractive index. This is especially visible in resonant structures, in which refractive index change is reflected in a change of the resonant condition. For this reason, TOE is commonly included in the list of nonlinear optical effects.

3.3 Nonlinear optics in waveguides

The small effective mode area of integrated WGs allows reaching much higher EM field amplitudes in respect to optical fibres with equal input power. Moreover,

thanks to the field enhancement, resonant structures achieve even higher field amplitudes. This favours the emergence of nonlinear effects, which influence the propagation of light in the WG. It is, therefore, crucial to be able to correctly describe the impact of nonlinear effects on the guided propagation of light.

Due to the lateral confinement of the EM field within the WG core, the change of the effective refractive index is not the same as the change in the refractive index of the bulk material. For this reason, the effects are averaged over the field profile in order to extract a mean effect. The demonstration is long and outside the scope of this work [47, 52], hence I will report here only the noteworthy steps.

Let's first assume the form of the complex field amplitude of eq. (3.1) such that

$$\mathbf{E}_\omega(x, y, z, t) = \sqrt{Z_0 \frac{P_0}{A_0}} \mathbf{E}_m(x, y) e^{i(\beta_\omega z - \omega t)}, \quad (3.44)$$

$$\mathbf{H}_\omega(x, y, z, t) = \sqrt{\frac{1}{Z_0} \frac{P_0}{A_0}} \mathbf{H}_m(x, y) e^{i(\beta_\omega z - \omega t)}, \quad (3.45)$$

where P_0 is the total optical power, $A_0 = w \times h$ is the area of the WG, and $Z_0 = \sqrt{\mu_0/\varepsilon_0} \approx 376.7\Omega$ is the impedance of free space. Afterwards, the following normalisation is imposed

$$\frac{1}{4A_0} \int_{\mathbb{R}^2} (\mathbf{E}_m \times \mathbf{H}_m^* + \mathbf{E}_m^* \times \mathbf{H}_m) \cdot \hat{\mathbf{z}} \, dxdy = 1. \quad (3.46)$$

With this normalisation, the EM field carries the total power P_0 . In fact, the total power flow in the WG is obtained by integrating the component along the propagation direction of the time-averaged Poynting vector over the cross-section:

$$\langle P \rangle_t = \int \langle \mathbf{S} \rangle_t \cdot \hat{\mathbf{z}} \, dxdy = \frac{1}{2} \int \text{Re} \{ \mathbf{E} \times \mathbf{H}^* \} \cdot \hat{\mathbf{z}} \, dxdy = P_0, \quad (3.47)$$

where in the last equality eqs. (3.44) and (3.45) have been substituted into the integral and eq. (3.46) has been used to simplify it.

From these considerations, we can employ the formula [47]

$$\frac{\partial}{\partial z} \int_{\mathbb{R}^2} (\mathbf{E}_1 \times \mathbf{H}_0^* + \mathbf{E}_0^* \times \mathbf{H}_1) \cdot \hat{\mathbf{z}} \, dxdy = i\omega_1 \int_{\mathbb{R}^2} \delta\mathbf{P} \cdot \mathbf{E}_0^* \, dxdy, \quad (3.48)$$

where $(\mathbf{E}_0, \mathbf{H}_0)$ and $(\mathbf{E}_1, \mathbf{H}_1)$ represents the EM fields of the mode propagating in an unperturbed and perturbed WG, respectively, and $\delta\mathbf{P} = \varepsilon_0 \Delta\chi \mathbf{E}_1$ embodies the perturbation of the refractive index caused by nonlinear effects as seen from eq. (3.32). The two fields are defined as follows:

$$\begin{aligned} \mathbf{E}_0 &= \mathbf{E}_0(x, y, z, t) & \text{and} & & \mathbf{E}_1 &= a(z) \mathbf{E}_1(x, y, z, t) \\ \mathbf{H}_0 &= \mathbf{H}_0(x, y, z, t) & \text{and} & & \mathbf{H}_1 &= a(z) \mathbf{H}_1(x, y, z, t) \end{aligned} \quad (3.49)$$

where the slowly varying envelope $a(z)$ has been introduced, such that $\frac{d^2 a}{dz^2} \ll \beta_1^2 a(z)$. If we plug eq. (3.49) into eq. (3.48) we obtain

$$\frac{\partial a}{\partial z} + i(\beta_1 - \beta_0) a(z) = i \frac{\omega_1}{c_0} \left(\frac{\varepsilon_0 c_0 Z_0}{4A_0} \int_{\mathbb{R}^2} \Delta \chi |\mathbf{E}_m|^2 dx dy \right) a(z). \quad (3.50)$$

Moreover, assuming a single optical frequency $\beta_1 = \beta_0$, we can write the evolution of the slowly varying envelope for the guided mode propagation as

$$a(z) = \exp \left(i \frac{2\pi\nu}{c_0} \Delta n^{eff} z \right),$$

where the change in the effective refractive index is expressed as Δn^{eff} , which can be further simplified exploiting the relationship between the optical power and the energy density:

$$P_0 = v_g u_E = \frac{c_0}{n_g} \int \frac{1}{2} \varepsilon_0 n^2 |\mathbf{E}_0|^2 dx dy = \frac{c_0 \varepsilon_0 Z_0}{2A_0} \frac{P_0}{n_g} \int n^2 |\mathbf{E}_m|^2 dx dy \quad (3.51)$$

In the case of the third-order optical nonlinearities, considering only the effects involving the same optical frequency, one obtains $\Delta \chi^{(3)} = \frac{3}{4} \chi^{(3)} |\mathbf{E}|^2$. In this case, the change in the effective refractive index can be expressed as

$$\Delta n^{eff} = \frac{n_g^2}{\varepsilon_0 c_0} P_0 \frac{\int n_0(x, y) \frac{3}{4} \chi^{(3)}(x, y) |\mathbf{E}_m(x, y)|^4 dx dy}{\left(\int n^2(x, y) |\mathbf{E}_m(x, y)|^2 dx dy \right)^2}. \quad (3.52)$$

The upper integral can be restricted to the WG core only, as the mode profile is usually contained for the great part inside of it. Hence, we can write:

$$\Delta n_{Kerr, TPA}^{eff} = \frac{3n_g^2 \chi^{(3)}}{4\varepsilon_0 c_0} \frac{P_0}{A_{eff}} = \left(n_2 + i \frac{c_0}{4\pi\nu} \beta_{TPA} \right) \frac{n_g^2}{n_0} \frac{P_0}{A_{eff}} \quad (3.53)$$

with A_{eff} defined by eq. (3.7).

On the other hand, for parametric effects such as those involving the FCs or the temperature, the change in susceptibility is $\Delta \chi_{N;T} = 2n_0 \Delta n(N;T)$. Hence, the change in the effective refractive index can be expressed as

$$\Delta n_{N,T}^{eff} = n_g \frac{\int n_0(x, y) \Delta n(x, y) |\mathbf{E}_m(x, y)|^2 dx dy}{\int n^2(x, y) |\mathbf{E}_m(x, y)|^2 dx dy}. \quad (3.54)$$

Assuming again that the mode profile is for the great part contained in the core of the WG, we can approximate the change in the refractive index to be restricted to that area. Thus, we can assume that the integrals can be expressed as mean effective value $\langle \cdot \rangle_{x,y}$ multiplied by the power confinement factor, as defined by eq. (3.6). This should be especially true for the temperature profile, in the case of the embedded WG, as the silicon core has a high thermal conductivity with respect to the silica cladding [52]. The results are summarised in table 3.2.

This theory can in principle be further expanded to also include nonlinear phenomena due to static electric fields, such as the linear and quadratic electro-optic effects, or the propagation of light signals at multiple frequencies and their nonlinear interactions, such as in the case of FWM.

Effect	Δn	Δn^{eff}
Kerr	$n_2 I$	$n_2 \frac{n_g^2}{n_0} \frac{P_0}{A_{eff}}$
TPA	$i \frac{c_0}{4\pi\nu} \beta_{TPA} I$	$i \frac{c_0}{4\pi\nu} \beta_{TPA} \frac{n_g^2}{n_0} \frac{P_0}{A_{eff}}$
TOE	$\frac{dn}{dT} \Delta T(x, y)$	$\frac{dn}{dT} \frac{n_g}{n_0} \Gamma \langle \Delta T \rangle_{x,y}$
FCD	$\sigma_{FCD} \Delta N(x, y)$	$\sigma_{FCD} \frac{n_g}{n_0} \Gamma \langle \Delta N \rangle_{x,y}$
FCA	$i \frac{c_0}{4\pi\nu} \sigma_{FCA} \Delta N(x, y)$	$i \frac{c_0}{4\pi\nu} \sigma_{FCD} \frac{n_g}{n_0} \Gamma \langle \Delta N \rangle_{x,y}$

Table 3.2: Summary for the change in effective refractive index for guided modes due to several nonlinear effects.

3.3.1 Free carrier generation

In silicon, at IR wavelengths, FCs are generated through the simultaneous absorption of two photons. Hence, the ratio between the optical power absorbed by TPA (per unit volume) and twice the energy of a single photon $2h\nu$ will directly give the generation rate of free carriers. In this picture, free carrier recombination processes keep the concentration to an equilibrium value. The rate equation, therefore, becomes

$$\frac{\partial}{\partial t} \langle \Delta N \rangle_{x,y} = -\frac{\langle \Delta N \rangle_{x,y}}{\tau_{FC}} + \frac{P_{TPA}}{2h\nu}, \quad (3.55)$$

where τ_{FC} is the characteristic lifetime of FC in the WG. The value of τ_{FC} is of the order of 10 ns for micron-wide WGs, but can decrease to hundreds of ps for smaller structures [70–74]. Nevertheless, higher values up to approximately 50 ns have also been recently reported [75]. The value of P_{TPA} can be estimated from eq. (3.53) to be constant through the WG

$$P_{TPA} = 2 \frac{P_0}{A_{eff}} \left| \frac{\partial a(z)}{\partial z} \right| = \beta_{TPA} \frac{n_g^2}{n_0} \frac{P_0^2}{A_{eff}^2}, \quad (3.56)$$

when it yields a relatively small (linear) absorption over the full WG length L .

3.3.2 Heat generation

When operating at high optical power, the losses due to absorption increase due to nonlinear effects, such as TPA. The temperature of the material rises accordingly to the amount of optical power absorbed. Consequently, the locations in which the EM field is higher, i.e. the centre of the WG in the case of the fundamental mode, reports a localised increase in temperature.

The typical description of the temperature dynamics in a WG is obtained considering an average temperature change in the WG core $\langle \Delta T \rangle_{x,y}$. Hence, the temperature change in a unitary volume is given by:

$$\frac{\partial}{\partial t} \langle \Delta T \rangle_{x,y} = -\frac{\langle \Delta T \rangle_{x,y}}{\tau_{TH}} + \frac{P_{abs}}{\rho C_P}, \quad (3.57)$$

where, similarly to eq. (3.55), τ_{TH} is the characteristic lifetime, here of the thermal relaxation, ρ is the density of the material in the WG core, C_P its specific heat, and P_{abs} is the optical power transformed into heat. Light can be absorbed by the material in many ways and indeed:

$$P_{abs} = \gamma_{lin} P_{lin} + \gamma_{FCA} P_{FCA} + \gamma_{TPA} P_{TPA}, \quad (3.58)$$

where γ_{lin} , γ_{TPA} , and γ_{FCA} coefficients embody the portion of the absorbed power transformed into heat and not lost through scattering or radiative channels. Since silicon is transparent at 1.55 nm, the linear absorption will be negligible in respect to losses due to scattering in the device plane and hence $\gamma_{lin} \ll 1$. On the other hand, since photon emission in silicon is usually a minor effect in respect to thermalisation, $\gamma_{TPA} \sim \gamma_{FCA} \sim 1$. The values of P_{lin} and P_{FCA} can be estimated similarly to eq. (3.56):

$$P_{lin} = \alpha_{lin} \frac{P_0}{A_{eff}}, \quad (3.59)$$

$$P_{FCA} = \sigma_{FCD} \frac{n_g}{n_0} \frac{P_0}{A_{eff}} \Gamma \langle \Delta N \rangle_{x,y}. \quad (3.60)$$

Although eq. (3.57) is usually sufficient to describe accurately the thermal response of a given structure, some cases require a more refined model [75].

3.3.3 Nonlinear effects in microring resonators

The interplay between nonlinear effects and the resonance condition in a MRR is especially of interest as, even for a simple input of constant power, the response of these systems can be quite dynamic. In fact, recalling eq. (3.26), the position of a resonance is linked to the value of the effective refractive index n^{eff} in the cavity.

$$\nu_n = n \cdot \frac{c_0}{n^{eff}(\nu)L}$$

On the other hand, the value of n^{eff} depends on the power circulating in the WG, its FC population density, and its temperature. Therefore, assuming a constant input optical power at a certain fixed frequency close to the resonance, the value of the optical power inside the optical cavity will affect the value of n^{eff} , which in turn will change the position of the resonance, which will modify the optical power in the MRR. This recursive chain of dependencies can establish both positive and negative feedback.

For example, at long timescales, the TOE is the leading effect. The real part of n^{eff} increases with temperature, hence the resonance frequencies diminish. This means that the feedback is positive if the input frequency is lower than the resonance frequency (negative detuning) and negative otherwise (positive detuning). Positive feedback is associated with the effect of *optical bistability* while negative feedback with *optical limiting*. On the other hand, FC reacts faster and at shorter timescales the FCD effect is dominant. Moreover, in this case, the effect of FCD is to reduce the value of n^{eff} proportionally to ΔN . Therefore, the feedback will be positive for positive detunings and negative otherwise.

When the MRR is brought into specific conditions, FC and TOE can interact continuously and produce the *self-pulsing* effect. The position of the resonance oscillates around its initial value due to the periodic interaction of positive and negative feedback due to both FCD and TOE. In order to calculate the specific dynamic of a MRR, numerical evaluation of the coupled differential equation describing the field in the cavity, the temperature of the medium, and the FC population density must be solved. In the next chapter (see section 4.1) I will present a tool that enables this kind of simulation for both single MRRs as well as structures composed of many MRRs.

Chapter 4

Integrated Photonics Modelling

The development of novel integrated photonic devices requires, undoubtedly, a good model of their fundamental components. Analytical solutions are practical only up until the system is simple enough, hence numerical methods are widely used to simulate the response of real devices. This chapter contains a summary of my achievements in this regard.

The largest effort was dedicated to the simulation of arbitrary networks of integrated optical cavities influenced by nonlinear effects. Specifically, the simulation requirements were to correctly predict the complex dynamics of large networks over long period of time. Available tools at the time were not suitable for these needs, so I started the development of `PRECISE: Photonic hybrid EleCtromagnetIc SolvEr`. This software describes the physics of PICs at a high level, allowing the relaxation of the computational requirements. `PRECISE` is available at <https://gitlab.com/erc-backup/precise/> and is distributed under an MIT licence. The first three sections of this chapter (sections 4.1 to 4.3) discuss `PRECISE`'s underlying theoretical base and its implementation, the simulation of linear transmission spectra for several different structures, and the solution of the dynamic evolution of a MRR, respectively.

As a secondary activity, I also developed software that enables the simulation and design of custom MMI couplers. The last section of this chapter (section 4.4) presents this modelling software, which has been used for the design of three MMIs in SiN working at 488 nm (see section 4.4.4 and appendix B.4) and three MMIs in Si working at 1550 nm (see appendix B.6).

4.1 PRECISE

It is certainly possible to describe photonics structures by directly implementing Maxwell's equations with a complete description of the geometry of the materials in which light propagates. Several numerical tools that can rigorously

represent photonics systems in such a way are already available, e.g. Comsol [76], Meep [77], and Lumerical [78]. These tools employ the finite-difference time-domain (FDTD) method or the finite element method (FEM), which subdivide the volume into sections whose sizes are a fraction of the wavelength of the light travelling inside the material, i.e. λ_0/n . Since silicon photonics usually works with light in the IR, either around 1310 nm or 1550 nm, the maximum feature size of these subvolumes must be a few hundreds of nm or less. For this reason, simulating isolated structures of just a few tens of μm is quite difficult without employing clever tricks and modelling whole PICs is usually unfeasible. The computational time and resources required increases rapidly with the size of the structure. Similarly, the response to a long input signal, even for small structures, is simulated with much difficulty.

FDTD and FEM tools are hence typically used to simulate the linear transmission of small, fundamental structures, such as WGs, rings, or MMIs. The response of a full PIC to an input of arbitrary shape is then described as the concatenation of the response of each of these elementary blocks. Unfortunately, these representations hold if the response of each constituent block is linear, i.e. does not change with the optical power circulating inside. When the EM fields increase in optical power, either due to the input signal or to the enhancement factor of resonant cavities contained in the device, the linear response of these systems does not correctly model their response at every input signal. Then, a separate, individual simulation must be carried out for each combination of the input waveform and system parameters specifically, in order to accurately describe the optical signal at the structure's output.

My work aimed to simulate large and complex networks of microring resonators, excited by optical signals of arbitrary shape and power. First of all, the input signals were lengthy sequences of bits, which required likewise long simulations. Secondly, the nature of the networks of ring resonators is that they may produce several resonances, in which the EM field is amplified, possibly triggering nonlinear effects. Hence, these simulations also require considering nonlinear effects in the WGs.

Fortunately, PICs modelling can be considerably simplified under the assumption that light propagates along the WG core in the form of modes. As seen in section 3.1.1, the mathematical representation of a TEM wave travelling in a WG becomes:

$$A(t) \mathbf{E}_m(x, y) e^{i \frac{2\pi\nu}{c_0} n^{eff} z}.$$

This description allows to discretise systems on a WG segment-level, e.g. a MRR in the add-drop filter configuration is composed by two WGs joined at two coupling points. Considering furthermore single-mode WGs, the propagation of the signal in each segment will then be fully described by only two complex amplitudes $A^+(t)$ and $A^-(t)$, representing the two counter-propagating fields. The other parameters of the system are the optical frequency of the fields ν , the geometrical length L of the segment, and the effective refractive index n^{eff} , which might however be affected by nonlinear effects. Different segments are then linked together at coupling regions.

By creating the system of equations that describes the fields in each WG segment and their dynamical quantities, i.e. temperature T and FC population density N , it is possible to evaluate light propagation within the circuit. PRECISE applies the additional assumption (comparable to [79]) that the field propagation reaches its steady-state in a timescale much shorter than the time span in which the nonlinear effects produce a significant change in the refractive index. Such approximation is fulfilled in most devices as the characteristic timescales involved with FC and thermal physical processes are usually several orders of magnitude larger than the timescale of strictly optical phenomena, i.e. Kerr effect, TPA, and the photon lifetime in the cavity.

4.1.1 Modular description of the optical system

In PRECISE, integrated devices are described by decomposing them into fundamental components. Each component will have a number of input/output ports (at least two) where the complex amplitude of the electric field will be evaluated. The system of equations that describes the PIC is symbolically defined by adding each component one by one.

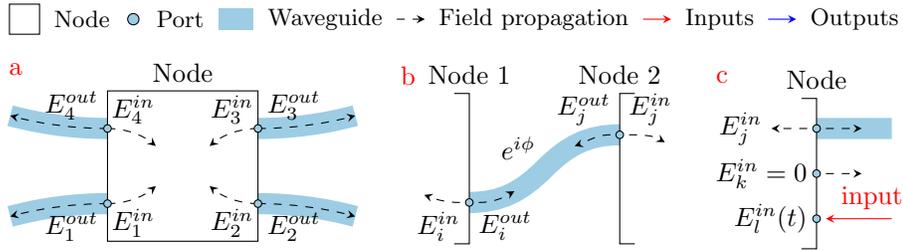


Figure 4.1: Drawings of fundamental components in PRECISE description of optical systems. **a** Nodes group together several input/output ports and define the connections from the input to the output ports. In this example there is a node with four ports. **b** Guides join ports, output to inputs, and describe the phase and amplitude changes associated with the propagation of light. **c** Input ports in nodes may be designated as *input* to the whole system; the field of unconnected input ports in nodes are set to zero.

The most important one is the node: a logical block that groups together a certain number of ports. The output ports in a node are connected to the input ports via

$$\mathbf{A}^{out} = \mathbf{C} \cdot \mathbf{A}^{in}, \quad (4.1)$$

where \mathbf{A}^{out} and \mathbf{A}^{in} are $1 \times n$ column vectors and \mathbf{C} is $n \times n$ matrix. For example, a point coupling region between two WGs is represented as a four-port

node (see fig. 4.1a) and the following matrix:

$$C = \begin{bmatrix} 0 & t & ik & 0 \\ t & 0 & 0 & ik \\ ik & 0 & 0 & t \\ 0 & ik & t & 0 \end{bmatrix}, \quad (4.2)$$

where t and k are the coupling coefficients such that $t^2 + k^2 = 1$. Usually, the elements of such matrices are constants, however, a frequency dependence can easily be introduced as well.

Isolated nodes are joined by WGs segments, which implement the phase and amplitude change associated with the propagation. Specifically, if the WG w connects two ports i and j , then

$$\begin{aligned} E_i^{in} &= E_j^{out} e^{i\beta_w L_w}, \\ E_j^{in} &= E_i^{out} e^{i\beta_w L_w}, \end{aligned} \quad (4.3)$$

where β_w and L_w are the propagation constant and the geometrical length associated with the guide w . Figure 4.1b represents a guide connecting two ports belonging to different nodes. Nonlinear effects are therefore considered by expanding the propagation constant to

$$\beta_w = \frac{2\pi\nu}{c_0} (n_w^{eff} + \Delta n_w^{eff}), \quad (4.4)$$

where Δn_w^{eff} is the change to the linear effective refractive index n_w^{eff} due to nonlinear effects, which was described in section 3.3. In particular, for silicon-on-insulator (SOI) devices, Δn_w^{eff} depends on the power circulating in the WG $P_w \propto |A_w|^2$, the temperature change $\langle \Delta T_w \rangle_{x,y}$, and the FC population density $\langle \Delta N_w \rangle_{x,y}$.

Last but not least, the system is completed by the designation of one or more *input* ports, which allow the excitation of the otherwise passive system. The input ports in nodes are either connected with a guide or labelled as *input*, otherwise, their fields are set to zero. Figure 4.1c represents the three cases.

In summary, an optical system is constructed by grouping input/output ports, adding connections between them, joining them with WGs, and labelling them as *input* ports. In PRECISE these actions are accomplished with the `precise.optical.system` class and its own methods, i.e. `add_node`, `add_link`, `add_guide`, and `add_input`. All this information results in a system of equations based on the symbolic variables ν , A_j^{in} , A_j^{out} , k_c , L_w , n_w^{eff} , Δn_w^{eff} . At this point, the ideal procedure would continue by solving the linear system symbolically and isolating A_j^{out} once and for all. This task, however, becomes progressively more difficult with the increase of elements in the system. For this reason, another approach is followed instead. The symbolic system is converted in the matrix equivalent form with `equationsToMatrix`, which returns a matrix \mathbf{M} and a column vector \mathbf{b} . An internal function, called `matrix_generator`, is then generated through `matlabFunction`, with the task of filling all the symbolic variables with

numerical values. The output of this function, when supplied with all the input values, are the numerical versions of \mathbf{M} and \mathbf{b} . Finally, the stationary propagation of light within the system is solved by feeding these numerical matrices to Matlab's linear solver `linsolve`.

In order to enable the parametric substitution of the symbolic variables, the `generate_pfe` method creates the *parametric field evaluation* function which calls `matrix_generator` and `linsolve` in sequence, when supplied with the correct input arguments. The *parametric field evaluation* function is of the form:

$$A^{out} = f_{pfe}(\nu, \{A^{in}\}, \{k\}, \{L\}, \{n^{eff}\}, \{\Delta n^{eff}\}) \quad (4.5)$$

and its output is

$$A^{out} = [A_1^{out}, \dots, A_P^{out}]^T.$$

The input variables of f_{pfe} are grouped by type with $\{\cdot\}$: A^{in} is the complex field amplitude at the input ports, k are the coupling coefficients, and L , n^{eff} , and Δn^{eff} are respectively the length, the effective refractive index, and its nonlinear change of each WG.

A generic code structure will look like listing 4.1, where all the steps discussed above for the creation of any optical system are highlighted. Furthermore, in the second step, in which the geometry is defined, the creation of nodes, connections, and guides can be incorporated in methods that create more complex building blocks, such as a MRR. A few examples of the methods providing building blocks are `add_ring`, `add_crow`, `add_scissor`, `add_taiji_resonator`, and `add_termination`, which add a MRR, CROW, a SCISSOR, a Taiji MRR, and a WG facet, respectively. These building blocks are defined in the most generic way so that the sub-elements can always be specified by the user. For example, `add_ring` allows an arbitrary number of coupling regions and in `add_crow` and `add_scissor` the user can specify the number of MRR of which they are composed. The response of these structures will be analysed in section 4.2.

```
% 1 - Optical system class initialisation
S = optical_system();

% 2 - Geometry definition: nodes, connections, and guides
S.add_node( n );           % n is the number of ports
S.add_link( p1, p2, K );  % a and b are port indexes,
                          % K is the coupling coefficients matrix
S.add_guide( p1, p2 );    % p1 and p2 are port indexes

% 3 - Designation of one or more 'input' ports
S.add_input( p );

% 4 - Generation of the 'parametric field evaluation' function
f = S.generate_pfe();
```

Listing 4.1: Generic optical system

4.1.2 Nonlinear effects dynamics

The most accurate depiction of the evolution in time of the EM fields propagating in a PIC is given by a system of differential equations describing the relevant physical quantities, i.e. the complex field amplitude, the temperature of the material, and the FC concentration in each element of the integrated structure. The time step used in the solution of the differential equations is upper limited by the smallest timescale, which is due to the dynamic of the complex field amplitude and is in the order of magnitude of 1 to 10 ps. The more steps are required to simulate the system for a given period of time, the more intensive the computation will be. Hence, long simulations with a full description of the system through a system of differential equations may be too expensive.

Conversely, PRECISE assumes that the field instantaneously reaches the stationary state at each time step, while changes in the temperature and FC concentration are negligible. Then, the value of the complex amplitudes can be evaluated considering a linear system where the value of n^{eff} is parametrically changed by the present state, i.e. temperature and FC, as described in eq. (4.4). PRECISE differential step carries out the following operations:

0. The status of the system is described by the complex amplitudes of the counter-propagating fields A_j^{out} and by the average temperature $\langle \Delta T_w \rangle_{x,y}$ and FC concentration $\langle \Delta N_w \rangle_{x,y}$ for each WG.
1. The nonlinear change of the effective refractive index Δn_w^{eff} is evaluated for each WG, as reported in section 3.3
2. The new values of the complex field amplitude are evaluated through the f_{pte}
3. The incremental values of $\langle \Delta N_w \rangle_{x,y}$ are evaluated for each WG, as described in eq. (3.55).
4. The incremental values of $\langle \Delta T_w \rangle_{x,y}$ are evaluated for each WG, as described in eq. (3.57).

It should be noted that, due to the recursive nature, i.e. field amplitude influencing itself, of these effect, an accurate dynamic would be described only with the full description and/or with very short time steps. However, the refractive index change due to the Kerr effect and the TPA is usually negligible in most silicon photonics structures in comparison to the FC effects and even more in respect to the thermo-optic effect. For these reasons, in PRECISE the contributions of the Kerr effect and the TPA are not considered in the evaluation of Δn^{eff} , although the absorbed power due to TPA is necessarily considered for steps 3 and 4.

Essentially, PRECISE aim is to simulate long ($>1 \mu s$) responses in time and is less accurate in the representation of fast dynamics (a few 100 ps). The accuracy of this method is tested against the most common coupled mode theory (CMT) [40] method on a MRR in the ADF configuration (see section 4.3). Specifically,

the system has been excited through the *input* port by a Gaussian pulse of varying width near the resonance frequency. The peak power of the pulse was kept constant, hence the overall total pulse power increases in proportion to its duration. Figure 4.2 compares the peak value - which should represent the error in the worst-case scenario - reached by the power in the *through* and *drop* ports, the temperature, and the FC concentration.

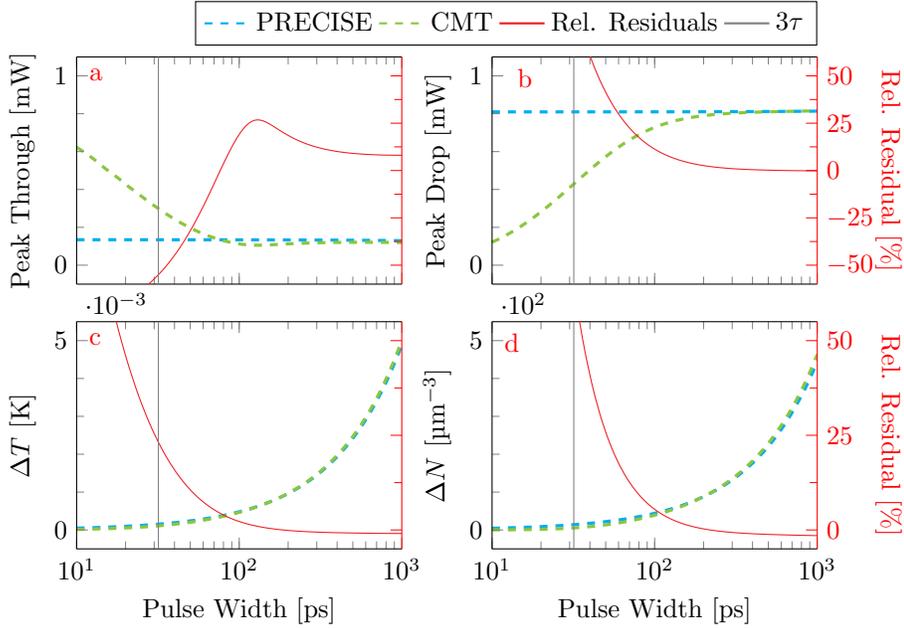


Figure 4.2: Gaussian pulse propagation: comparison between PRECISE and CMT models. A Gaussian pulse of varying width excites the *input* port of an ADF with an optical frequency close to the resonance of the MRR. **a** Peak power in the *through* port, **b** peak power in the *drop* port, **c** temperature difference, and **d** FC concentration difference. The vertical lines indicate three times the value of the MRR photon lifetime ($3\tau_p$). Dashed lines (left axis) indicate the results of the two models, while the red lines (right axis) show the residuals.

The fundamental difference reported by the two models is visible in the loading of the MRR cavity, where the CMT is able to describe the dynamic more accurately than PRECISE, which assumes a steady-state condition. This is most visible in the two top charts (figs. 4.2a and 4.2b). Nevertheless, we can observe that both ΔT and ΔN are correctly represented. Moreover, it can be inferred that PRECISE ability to accurately describe optical systems has a lower limit which is between 1 and 2 orders of magnitude larger than the photon lifetime in the system. Slower dynamics show a difference between the two models around 2%.

4.1.3 Parametric solutions and data analysis

Even in small optical systems, the number of free parameters is very high. For example, in a MRR one might consider its radius, the propagation losses, the coupling coefficients, and many more. Moreover, also the input optical signal brings additional degrees of freedom (DOFs) to the simulation, e.g. the optical frequency, the optical power, and the signal modulation. In view of this, an important aspect of PRECISE is the ability to explore the parameter space of a given optical system methodically. A special class (`data_loader`) stores all the values for each parameter in the system and retrieves any desired combination with a single integer number, spanning from 1 to the total number of combinations N_C . The single combination of parameters can then be used by functions such as the *parametric field evaluation*, for the linear evaluation, or by the function carrying out the differential step.

There are mainly two benefits of this technique. Firstly, the parallel evaluation of many configurations becomes straightforward, employing Matlab's `parfor` loop, hence improving the overall computational efficiency. Secondly, since any given solution can be re-generated at any time by using the associated index, then the time traces can be eliminated, after having evaluated the interesting figures-of-merit from them. This allows to drastically reduce the memory used as roughly every trace, usually composed of several hundreds of thousands of points, can be compressed to a single scalar. These characteristics contribute to a faster overall execution time and a reduced memory footprint, especially advantageous when exploring large parameter spaces.

Data analysis

The result of the evaluation of the dynamical system is a number of time traces, representing the complex field amplitudes at each port and the temperature and FC concentration in each WG. For example, even in the relatively simple case of an isolated MRR, the evolution is completely described by 12 time traces: 8 for the complex amplitudes, 2 for the temperature, and another 2 for the FC concentration. The dataset that would contain the time traces for all the combinations evaluated in parallel by the `parfor` loop should be three dimensional, with size $N_C \times N_S \times N_t$, where N_C is the total number of combinations, N_S is the number of time traces per combination (12 for the example above), and N_t is the number of samples in each time trace.

As the complexity of the simulation grows, along with N_C , N_S , and/or N_t , the size of the dataset becomes rapidly unmanageable. However, in consideration that usually $N_t \gg N_C > N_S$, extracting a figure-of-merit (FOM) from each time trace is of paramount importance as it reduces the dataset size by N_t , which may even reach 10^5 to 10^6 . The dataset holding the information of a scalar FOM is then bi-dimensional and has size $N_C \times N_S$.

The most elementary FOMs are the minimum, maximum, average, and last value in the time trace. The mathematical functions that extract these quantities are already found in Matlab as built-in functions. Although they can be

employed to observe more interesting quantities, such as the modulation visibility of the dynamic response, defined as $V = \frac{\max - \min}{\max + \min}$, more significant FOMs require ad-hoc evaluations. I list here three of them:

- `precise.analysis.periodFFT` extracts the most significant component of the Fourier transform of the signal.
- `precise.analysis.extremals` extract the number of unique local minima and maxima, i.e. extremals, in the time trace.
- `precise.analysis.extremals_count` in addition to the same FOM extracted by the previous function, this counts the number a given extremal is observed in the time trace.

Visualisation

Even though the single index approach simplifies the evaluation of the dynamics for each combination, it is usually more comprehensible to restore the original size of the parameters space explored, which is in general n-dimensional. Hence, the resulting dataset requires ad-hoc visualisations. PRECISE offers several tools to help the user in this task. The most important one is `plot_utils.dashboard`, which offers a way to observe a summary of the average, minimum, and maximum values of the complex amplitudes, the temperature, and the FC density for all the combinations. It furthermore plots the time evolution of these quantities for a chosen combination. Another useful tool is the `plot_utils.array_plot` class, which adds the possibility to observe a given FOM as a function of two variables, e.g. displaying the signal period as a function of input frequency and power, for a given set of MRR coupling coefficient and radius. Both these classes have an `update` method that allows the user to interact and change the displayed information.

4.1.4 Benchmarks

In order to understand the performance of the toolkit, I present here a few tests on a CROW structure. The simulation complexity of this structure is proportional to the number of MRRs of which it is composed. All the simulations were run with Matlab 2020a on a laptop equipped with an AMD Ryzen 5 2500U (15W, 4 cores, 8 threads) and 8 GB RAM. Figure 4.3 collects the results, shown in log-log scale, obtained with a statistic of 10 consecutive runs.

The metric shown in the first graph, fig. 4.3a, is the time required to initialise the symbolic system and generate the *parametric field evaluation* function f_{pfe} . We can observe an increase of roughly one order of magnitude in initialisation time for an order of magnitude of increase in the number of elements in the system. Figure 4.3b shows instead the evaluation time of f_{pfe} , where each point is obtained from a batch evaluation of 1001 function calls. In this case, the evaluation time becomes roughly two orders of magnitude longer for one order of magnitude larger number of elements. The third and last metric, shown in

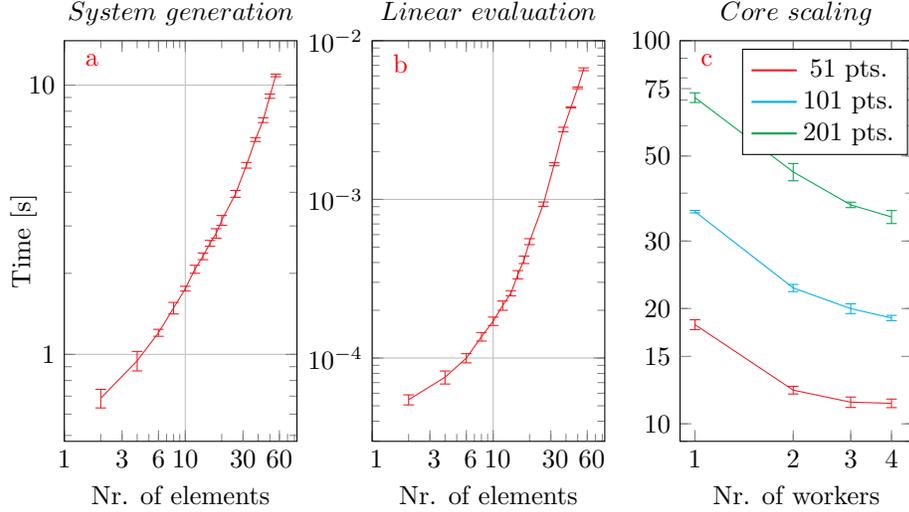


Figure 4.3: PRECISE performance analysis on a benchmark suite evaluating the response of a CROW structure: average value and standard deviation of ten consecutive runs. Each new ring in the CROW increases by two the number of elements in the system. **a** Time spent in the generation of the f_{pfe} function. **b** Execution time for a single call of the f_{pfe} function; obtained from a batch evaluation of 1001 calls. **c** Time required to simulate the nonlinear evolution of a CROW with two rings for 51, 101, and 201 frequency points as a function of the number of parallel workers.

fig. 4.3c, is the scaling of the evaluation time on the number of parallel processes (workers). Although these are relatively small numbers of combinations, we can already see that the total computation time for four workers in respect to a serial code is 38%, 47%, and 51% for 51, 101, and 201 combinations, respectively. A real case scenario would probably involve the simulation of several hundreds or thousands of combinations, for which it is reasonable to expect even higher improvements than those shown in this small benchmark.

4.2 Linear Spectra

Being able to correctly describe the linear spectra, i.e. the linear transfer function of an arbitrary structure was the first step in this work. So, in order to assess PRECISE performance, I modelled several test structures and compared them with the expected results from either theory or experiments. The first device is a simple MRR, then I studied the response of a CROW composed of four rings and a Taiji MRR and compared the results with experimental data. Finally, I explored more complex structures to study their usability as reservoirs in RCN: the SCISSOR, the MRR with a feedback loop between the *through* and *add* ports, and the random network of MRRs. Figure 4.4 groups the drawings of

all the structures used in the following sections.

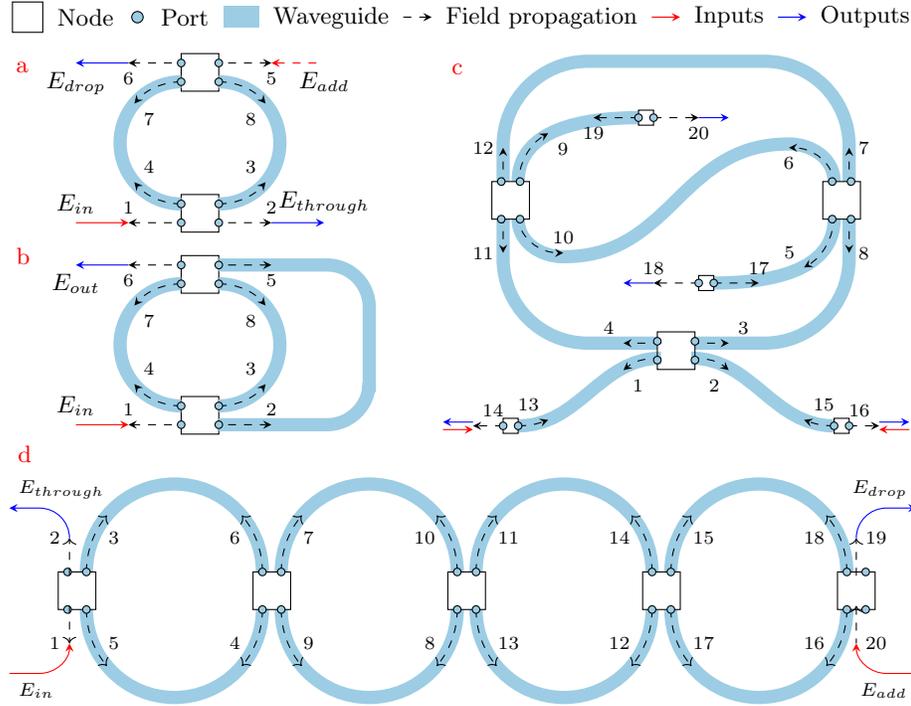


Figure 4.4: Drawings of several structures. **a** Single MRR in the ADF configuration. **b** Single MRR with feedback. **c** Taiji MRR. **d** CROW composed of four MRRs. Nodes are represented with small squares. Ports in nodes are small circles and are identified by an integer number. The field propagation is illustrated as dashed arrows and input and output to the system are highlighted with red and blue arrows, respectively.

Linear spectra are obtained by evaluating the complex amplitude of the fields at each port for an input field with unitary amplitude. The transmission of the system is then evaluated by taking the square modulus of these fields. Indeed, assuming exciting the structure from port 1, then

$$T_j = P_j/P_1 = |E_j|^2/|E_1|^2 = |E_j|^2.$$

In the case of field amplitudes evaluated inside optical cavities, this value represents the square modulus of the field enhancement factor and, as a matter of fact, can be larger than 1. When multiple input fields are considered this description is still useful, although the link between transmission and complex field amplitude is not as straightforward as before. In this case, by default I define the amplitude of secondary fields in respect to the larger one, even though other normalisation methods could be considered, e.g. defining the sum of the square modulus of all the input fields equal to 1.

4.2.1 Single microring resonator

The MRR is probably the simplest structure which shows both a non-trivial spectral response and, due to the optical cavity effects, and nonlinear phenomena. For this reason, it is the first structure that I modelled with PRECISE. As shown in fig. 4.4a, this structure represents a MRR in the ADF configuration and it is composed of two coupling regions and two WG segments. The total number of ports is 8 and port number 1 is specified as the input. This model could represent an APF as well by imposing the coupling coefficient of the second coupling region to be zero.

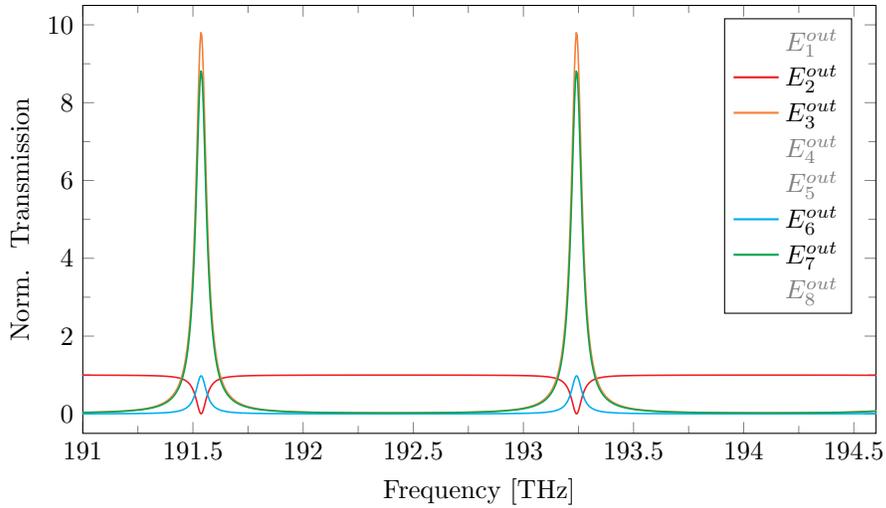


Figure 4.5: Transmission spectra of the MRR in the ADF configuration. The lines represent the square modulus of the complex amplitude at a given port, which is identified in the legend by the index that refers to fig. 4.4a. For output ports, this corresponds to the power transmission, while, for inner fields, it corresponds to the square modulus of the field enhancement factor. PRECISE evaluates the transmissions associated with the counter-propagating fields (E_1^{out} , E_4^{out} , E_5^{out} , and E_8^{out}) too, however, they are not shown here as they are always zero.

The results shown in fig. 4.5 report the transmission spectra, i.e. the square modulus of the fields, of the MRR model with a set of default parameters defined in PRECISE example script `precise.examples.linear_ring`. The shape of the fields at the *through*, E_2^{out} , and *drop*, E_6^{out} ports show the typical Lorentzian form. On the other hand, the field inside the optical cavity, E_6^{out} and E_7^{out} , provide insight on the field enhancement factor achieved by the device around the resonance condition. Since PRECISE actually simulates both counter-propagating fields in a given structure, also E_1^{out} , E_4^{out} , E_5^{out} , and E_8^{out} are evaluated. However, given that the MRR is excited only from port 1, they are always zero.

4.2.2 CROW

A CROW is composed of several MRRs, with two bus WGs coupling to the first and last ring in the chain, respectively [80]. Specifically, the one I simulated in this case is made of four rings and both the *input* and *add* ports are excited with the same input signal. The phase difference between the two ports establishes if the output power is directed toward the *through* port or the *drop* port. This arrangement is called *interferometric band interleaver* [81, 82]. Figure 4.4d shows the drawing of this CROW in PRECISE.

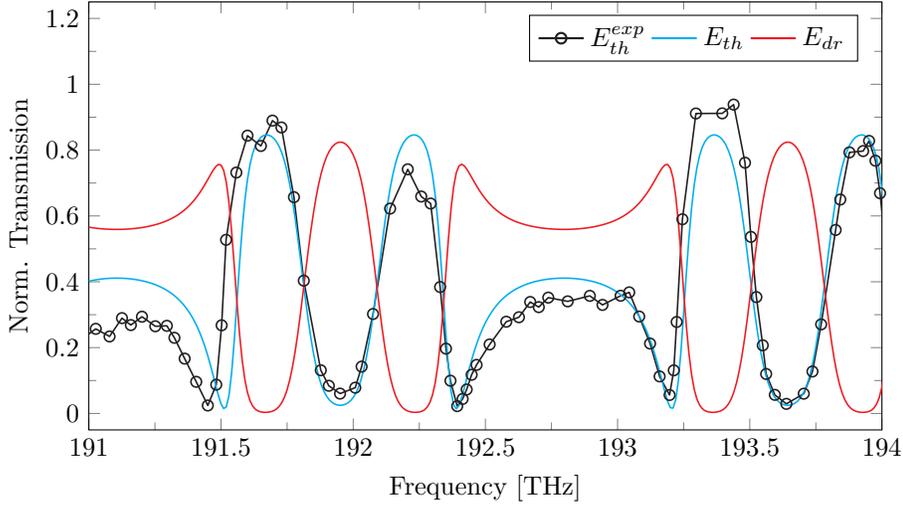


Figure 4.6: Transmission spectra of the CROW structure in the *interferometric band interleaver* configuration. Experimental data (black circles) for the *through* port are compared with the simulation results for the *through* port (solid cyan line) and *drop* port (solid red line).

Figure 4.6 shows the linear transmission spectra of the CROW, comparing PRECISE description to the experimental measurements reported in [82, 83]. PRECISE is able to correctly predict the resonance positions and the free spectral range as well as the complicated profile due to the interaction of the multiple MRRs. In order to choose the correct set of parameters describing the real device, Matlab's `particleswarm` has been used on PRECISE model, with the assumption that all coupling coefficients and MRRs are equal. The final set of parameters is reported in table 4.1.

4.2.3 Taiji microring resonator

The Taiji MRR is a modification of the APF configuration which enables the excitation of the counter-propagating field. Specifically, as shown in fig. 4.4c, two coupling regions are connected between their counter-clockwise ports. Hence, the field propagating counter-clockwise excites the field propagating clockwise,

Fixed Parameter	Value	Fit Parameter	Value
loss	2 dB/cm	n^{eff}	2.6391
number of rings	4	n_g	4.3416
WG segments, L	$\pi R + 10 \mu\text{m}$	radius, R	3.2544 μm
		k	0.83645

Table 4.1: Simulation parameters of the CROW. The coupling region is 10 μm long, but is considered a point coupling region and the extra length is associated with the waveguide segment instead. k are the field coupling coefficients.

but not the other way around. The total transmission of the structure is the same for both propagation directions, however, the reflection is asymmetrical. Further information on this structure can be found in [84].

In addition to the isolated structure, the input and output WGs have been added to the mode in order to include also the effect of the WG facets reflectivity. In fact, a simple WG enclosed between its facets is a Fabry-Perot optical cavity, which produces distinctive oscillations in the spectrum, due to frequency-dependent interference effects.

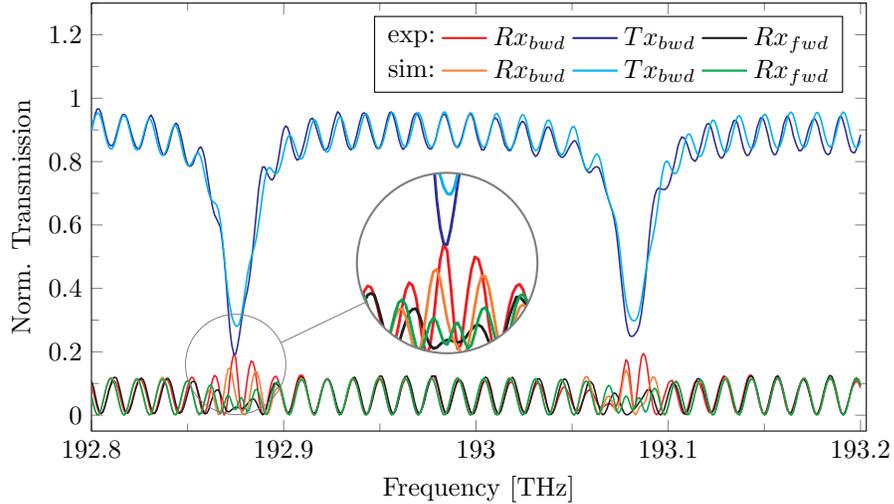


Figure 4.7: Taiji MRR transmission spectra: comparison of the simulations to the experimental data. Rx stands for reflection, Tx for transmission; bwd stands for backward, as in counter-clockwise field propagation, fwd stands for forward, as in clockwise field propagation.

Figure 4.6 shows the transmission from port 14 and port 16 and the reflection at port 14 and at port 16. In the first two cases, the input signal excites port 14, while in the latter excites port 16. The simulated spectral response is very close to the measured experimental one obtained from the real device. In-

deed, the model correctly predicts all the main features of the structure. In the transmission spectra, the resonances are in the expected positions and the oscillations due to the Fabry-Perot cavity between the two entry facets are reported. Moreover, the reflection spectra show again the Fabry-Perot oscillations as well as the difference between the clockwise and counter-clockwise reflections. The good agreement between the experimental results and the numerical simulations confirms that PRECISE is able to correctly handle counter-propagating fields.

Similarly to the CROW spectra, Matlab’s `particleswarm` algorithm has been employed to find the optimal parameters that reproduce the experimental results of the real device. These parameters are reported in table 4.2 and, if given as input the script `precise.examples.linear_taiji`, they allow to easily reproduce these results.

Fixed Parameter	Value	Fit Parameter	Value
linear loss	2.16 dB/cm	n^{eff}	1.6623
L_{1-13}	3.3034 mm	n_g	1.8224
L_{2-15}	2.6715 mm	$k_{1,2,3,4}$	0.29751
$L_{3-8} = L_{4-11}$	0.202 mm	$k_{5,6,7,8} = k_{9,10,11,12}$	0.38599
L_{7-12}	0.391 mm	k_{F1}	0.20305
L_{6-10}	0.611 mm	k_{F2}	0.20328
$L_{5-17} = L_{9-19}$	0.118 mm		

Table 4.2: Taiji MRR simulation parameters. L are the lengths of the WG segments and k are the coupling coefficients. Indexes refer to the ports as described in fig. 4.4c. Specifically, indexes on L identify the ports at the ends of the segments and on k identify the ports in the node of the associated coupling region.

4.2.4 MRR with Feedback Loop

A further interesting case is that of a MRR where the *through* port is connected with a feedback loop WG to the *add* port. In this case, the transmission at the *drop* port, i.e. the only output port of the system, is determined by the resonance condition of the MRR and by that of the external loop.

The radius of the MRR is $R = 7 \mu\text{m}$ and the default loop length is fixed to three times the ring semi-perimeter length $L = 3 \cdot \pi R$. The coupling coefficient are $k_1 = k_2 = k = \text{sqr}(0.1) \simeq 0.32$, the effective index and the group index are respectively $n^{eff} = 2.4$ and $n_g = 4$. The propagation losses were fixed to the standard 2 dB/cm for the two segments composing the ring and to 200 dB/cm for the loop. This extremely high propagation loss is implemented only to exaggerate the effect of the feedback loop in respect to the normal ADF operation. The linear spectrum of the device has been simulated in a 400 GHz range around the resonance ($f_0 = 193.2413 \text{ THz}$) and by changing the length of the feedback loop WG, and thus the interference with the MRR, in the range between -0.3 to $0.7 \mu\text{m}$.

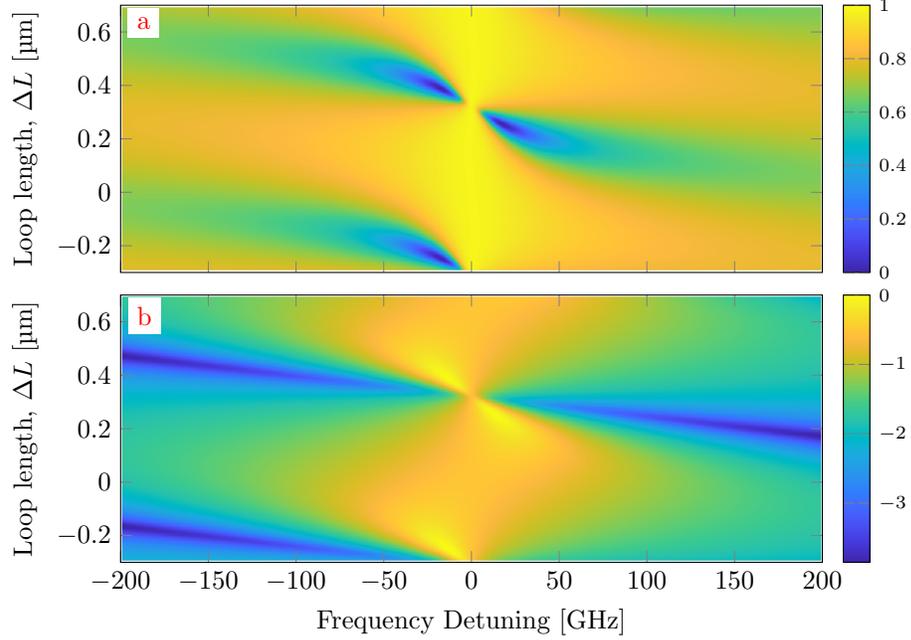


Figure 4.8: Maps of the normalised **a** linear transmission and **b** internal energy in a MRR cavity with loop feedback as a function of frequency detuning and loop length. The *drop* transmission is normalised to the *input* power. The internal energy is normalised to the maximum value in the map and is reported in log scale.

Figure 4.8 shows the linear transmission to the *drop* port and the normalised internal energy (in \log_{10}) as functions of the detuning and the loop length change. From the first map, it is clear that when the input signal is in perfect resonance, all the energy deviates directly to the *drop* port. However, moving a little away from the resonance the transmission decreases, slowly when the signal from the feedback loop is in phase with the *drop* signal and abruptly when it is out of phase. In fact, the minimum transmission is obtained close to the resonance, but only for restricted ranges of feedback loop length. On the other hand, the normalised internal energy shows a different picture. The maximum internal energy is observed close to but not perfectly on resonance, close to the conditions of minimum *drop* transmission. In conclusion, the overall device response is non-trivial in respect to that of the isolated MRR, but is relatively simple to simulate.

4.2.5 Networks of microring resonators

Networks of integrated optical MRRs are a promising physical implementation of reservoir-computing (RC) [35, 85]. The complex topology and the strong

recurrence of RCN provide a higher dimensional representation of the input, enabling its linear separability. Then, a subset of the nodes is used as output signals on which a simple linear regression is carried out. The performance of these methods is hence dependent on the network topology. Considering that nanophotonics cannot rival electronics in terms of integration scale, to fully exploit its capabilities, i.e. speed and parallel computation, finding the correct procedure to design these networks is a decisive factor.

The modelling of recurrent neural networks (NNs) is a formidable task and it becomes even more complex for nanophotonic because the dynamics of the network nodes is described by systems of coupled nonlinear equations. PRECISE is able to model such networks and to evaluate their stationary and dynamic properties. It therefore enables the study of the effect of the network topology on its spectral and dynamical response.

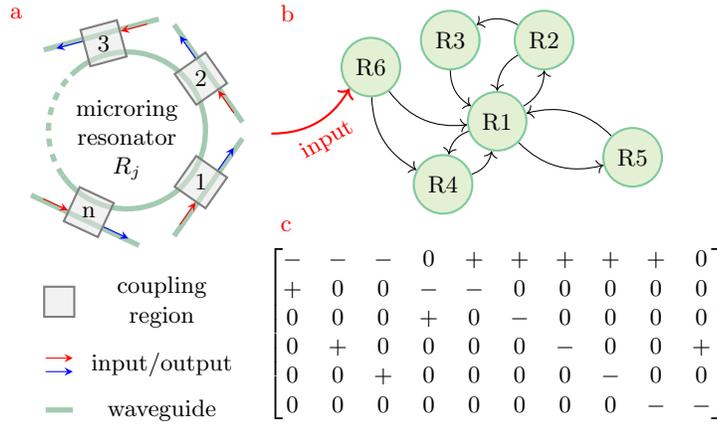


Figure 4.9: A network is composed of a set of nodes connected by bus WGs. **a** Each node is a MRR with a given number of coupling regions. **b** Graph of the network connections between the rings in the network and **c** associated incidence matrix. Each column in the matrix describes a connection between different nodes (rows), with $-$ ($+$) representing an output from (input to) the node.

As the response of large nanophotonic systems strongly depends on fabrication tolerances, we start by investigating RCNs with scale-free networks [86], which provide robustness against manufacturing defects. The incidence matrix of the mathematical network (both shown in fig. 4.9), which defines the number of nodes and their connections, is given as input to PRECISE, which generates the required system of equations. Once the model characteristics are fixed, i.e. the coupling constants and the length of the WG segments, the system of equations is integrated and the model's stationary and dynamic properties can be evaluated. As an example, fig. 4.10 reports the spectra of the square modulus of the field enhancement factor in each ring cavity, which is directly linked to the optical energy stored inside.

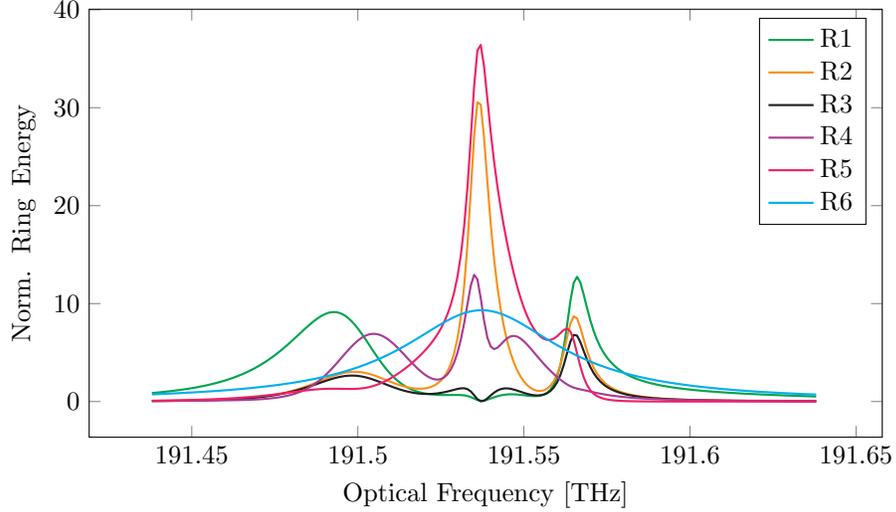


Figure 4.10: The numerical simulation produces the normalised spectra of the energy stored in each MRR in the network.

4.3 Dynamics of Nonlinear Systems

In order to test the dynamical response of a system, I simulated that of a MRR. In fact, in an optical cavity, the n^{eff} change due to nonlinear effects affect the optical power propagating inside the cavity itself, which in turn influences the nonlinear effects and the n^{eff} change they produce. This recursive causality develops in a non-easily predictable way. As a comparison, I run the same simulations with a CMT method, which do not incorporate the stationary state approximation.

4.3.1 Single microring resonator

MRR is the simplest system in which the consequences of nonlinear phenomena arise visibly in the *self-pulsing* effect. As discussed in section 3.3.3, during *self-pulsing* the positive and negative feedback mechanisms generate a dynamical (periodic) response to an input of constant optical power. The unique shape and periodical components of the output signal are determined in a non-trivial way by numerous parameters, e.g. the optical power and frequency of the input signal, but also the geometry of the MRR, as well as the material physical characteristics.

The MRR in consideration is an ADF with a radius of $7\ \mu\text{m}$ and coupling coefficients $k_1 = k_2 = k = 0.17$. The drawing is the same as for the linear propagation and is reported in fig. 4.4a. The system response is simulated in a 50 GHz interval around the resonance frequency and for input powers between 1.7 mW and 14 mW. All the other parameters required to carry out the system

evolution are reported at the end of the section in table 4.3.

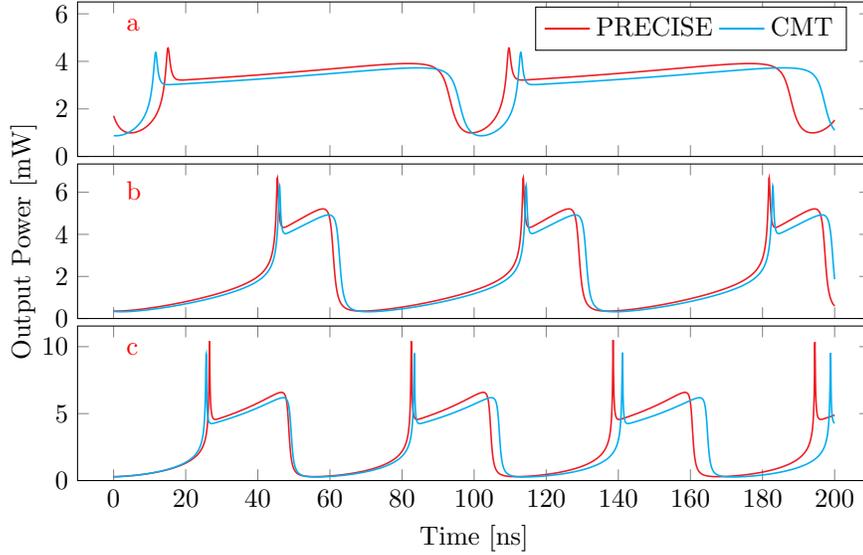


Figure 4.11: Self-pulsing phenomenon in a MRR simulated with PRECISE and with the CMT method. The curves show the response of the system observed at the *drop* port for three different combinations of input power and frequency detuning in respect to the resonance frequency. They have been extracted from longer simulations so that they can be considered stationary states.

Figure 4.11 compares the nonlinear response at the *drop* ($|E_6^{out}|^2$) port of the MRR for three combinations of input optical power and frequency obtained from PRECISE and from the CMT model. PRECISE is able to correctly describe the *self-pulsing* oscillations both in frequency and amplitude, with only small differences compared to the CMT model. These small discrepancies arise from the different approximations carried out by the two models. The most important is surely the fact that, while in CMT the detuning from the resonance is given as an exact variable, i.e. $\Delta\nu = \nu - \nu_0$, in PRECISE this quantity naturally emerges from the propagation throughout the geometry. A very small misalignment of the detuning frequencies between the two models is most likely the culprit of the difference in the results, given the fact that the system is extremely sensitive to the simulation conditions. Further differences are the overall dispersion of n^{eff} and α , and the fact that the optical cavity is considered as a whole in CMT, while is divided into two segments in PRECISE, enabling in principle to have a different value for the physical quantities in each section of the MRR.

To further examine the description of the system given by PRECISE and by the CMT model, it is interesting to generate a map of the response of the MRR as a function of optical frequency and power. Figure 4.12 shows, side by side, the colourmaps representing the period of the most important non-zero component of the Fourier transform of the output signal. The period has been extracted

using `precise.analysis.periodFFT` on the second half of all the traces, to analyse only the stationary response.

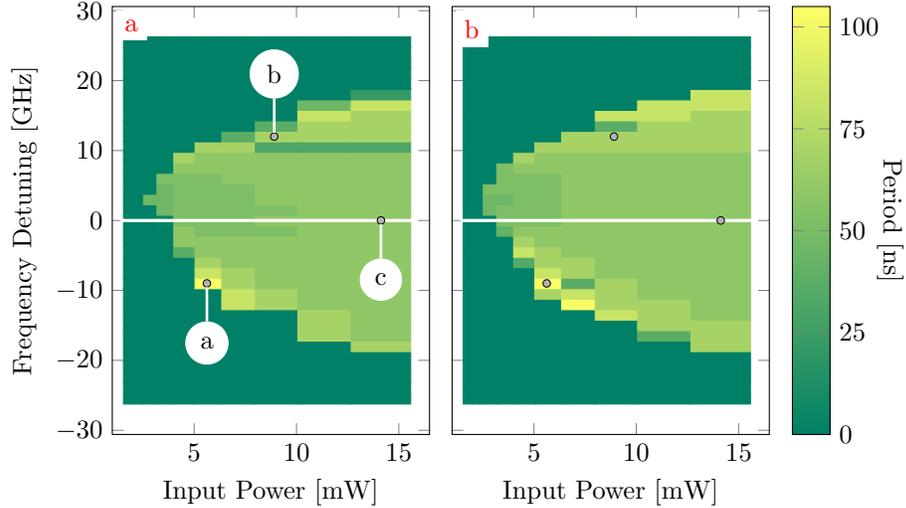


Figure 4.12: Self-pulsing regions are characterised through the period of the most important Fourier transform component, observed in the transmission at the *drop* port ($|E_6^{out}|^2$). The regions where no self-pulsing is observed are reported having 0 ps period. Comparison between the PRECISE approach **a** and the CMT approach **b**. The white horizontal line represents the MRR resonance frequency, i.e. 0 GHz detuning. The pins highlight the curves shown in fig. 4.11.

As anticipated, the *self-pulsing* behaviour occurs in a limited region close to the resonance frequency and for input powers above a certain threshold. The boundary between the stable and unstable regions resembles a parabolic curve in both models, although the actual shape is strictly correlated to the specific system parameters. Likewise, the individual values of the *self-pulsing* period are close in the two different models. Also from this point of view, the similarity between the results of the CMT and those of PRECISE is clear, with both models obtaining the same behaviour when exciting a MRR with constant input power.

Quantity	Description	Value	Unit	Source
n_0	silicon refractive index	3.485		[66]
n_2	Kerr effect	4.5×10^{-9}	$\mu\text{m}^2/\text{mW}$	[65]
β_{TPA}	two photon absorption	8.1×10^{-9}	$\mu\text{m}/\text{mW}$	[65, 66]*
n_{TOE}	thermo-optic effect	1.86×10^{-4}	1/K	[66]
σ_{FCD}	FC dispersion	-4.0×10^{-9}	μm^3	[66]*
σ_{FCA}	FC absorption	1.4×10^{-9}	μm^2	[65, 66]*
C_P	silicon specific heat	0.710	J/g K	[66]
ρ	silicon density	2.329×10^{-12}	g/mm ³	[87]
n^{eff}	eff. refractive index	2.4		FEM*
n_g	group index	4		FEM
α	linear losses	2×10^{-4}	dB/ μm	[88]*
τ_{FC}	FC lifetime	3.333	ps	[89]*
τ_{TH}	thermal lifetime	22	ps	[89]*
Γ	field conf. factor	0.99		FEM
A_0	WG area	0.099	μm^2	FEM
A_{eff}	effective area	0.141	μm^2	FEM

Table 4.3: List of the parameters required to carry out the evolution of the nonlinear dynamics of the MRR. The top half of the table contains values typical of silicon, while the bottom half contains those typical of silicon photonics single-mode channel WG manufactured for 1550 nm light. The asterisks * marks the entries which should be considered as an indication of the correct order of magnitude of the given quantity, rather than its exact value.

4.4 MMI optimiser

Multi-mode interference (MMI) couplers are an important building block in the design of PICs, as they are passive structures that enable the exchange of optical power between several WGs. Since not all integrated photonics manufacturers provide MMI couplers as ready to use building blocks, I developed software that allows the design of rectangular MMIs in any arbitrary material.

This software is composed of a set of Matlab scripts, which exploit Comsol's FEM engine to evaluate the lateral field profiles. For this purpose, an empty Comsol template for a waveguide with a rectangular section is provided. The numerical evaluation of the performance of a given MMI geometry is based on the modal propagation theory presented in section 3.1.3, where the field profiles of the modes are simulated using FEM methods. Figure 4.13 describes the system and the associated reference axis.

4.4.1 Field profiles

The first operation to carry out is to generate a collection of field profiles for different WG widths. With this objective in mind, a Comsol template for a WG

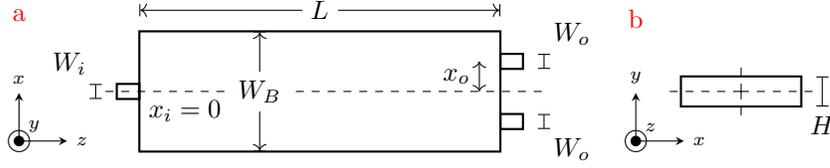


Figure 4.13: Drawing of the MMI coupler structure and the associated reference axis.

with a rectangular section is provided. The default materials are silicon for the core and silica for the cladding, but they can be easily changed. The FEM model defines a geometry where the width of the WG is swept on a set of values, usually specified in steps between a minimum and a maximum values, while the height is fixed.

The simulation, which can be carried out for a single or a range of optical frequencies, evaluates the field profile of the first N modes. The resulting field profiles are analysed by the Matlab script, discriminating between TE and TM polarisation. In general, not all geometries support a high number of modes, in fact, narrow WGs may guide the fundamental mode only. The modes showing a value of the $\text{Im}\{n^{eff}\}$ higher than a given threshold (default 1e-6) are considered lossy or radiating modes and are discarded from the collection. To check the results for each geometry, the script creates one figure containing the profiles of the allowed TE modes and one for the TM modes.

At the end of the simulation and sorting procedures, the value of the real and imaginary parts of n^{eff} are plotted for each mode number and polarisation as a function of the WG width. Figure 4.14a and fig. 4.14b show an example of such charts for TE polarisation modes.

4.4.2 Modal propagation

The field profile of the fundamental mode is shifted horizontally by x_i for each input channel, from 1 to I , and similarly by x_o for each output channel, from 1 to O . This sets up input field profiles E_i and output field profiles E_o at the right positions in respect to the centre of the MMI. Then, for each input channel, the overlap integral C_{im} is evaluated between E_i and the MMI field profiles E_m and normalised over $\iint |E_m|^2 dx dy$. Then each mode in the MMI is phase-matched in the centre of the section $(x, y) = (0, 0)$, and they are propagated to a given length $z = L$ each with their respective propagation constant β_m :

$$E_i^{end} = \sum_m C_{im} E_m(x, y) e^{i\beta_m L}.$$

For each output channel, the overlap integral C_{io} is evaluated between E_i^{end} and the output profiles E_o , again normalising over the field $\iint |E_o|^2 dx dy$. The transmission of each channel then becomes

$$T_{io} = |C_{io}|^2 \frac{\iint |E_o|^2 dx dy}{\iint |E_i|^2 dx dy}. \quad (4.6)$$

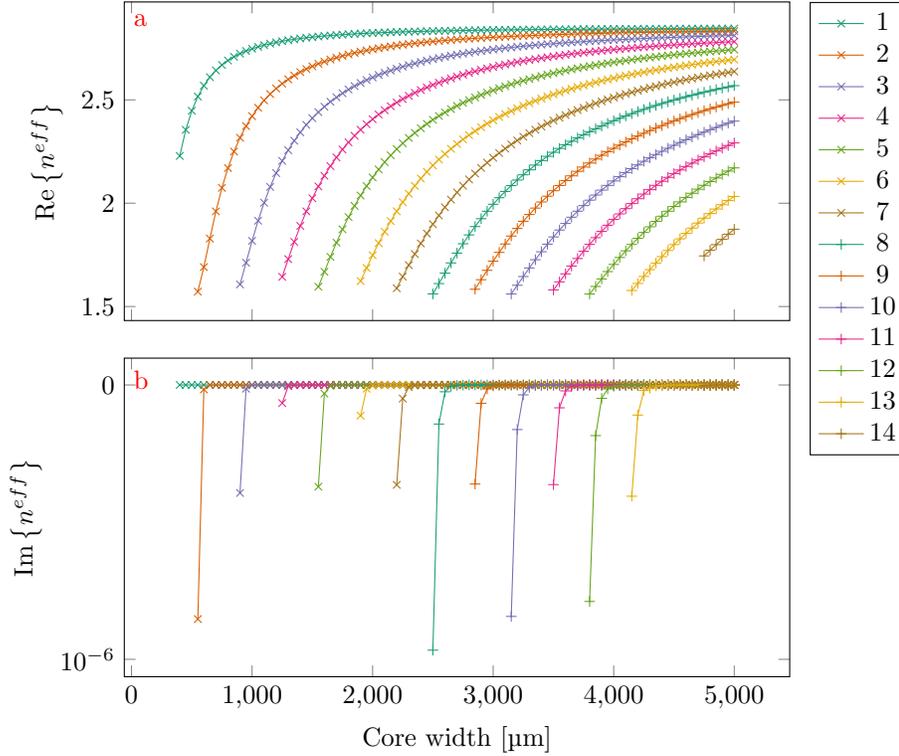


Figure 4.14: Effective refractive index n^{eff} for the TE modes as a function of the WG width: **a** Real and **b** imaginary parts. Negative imaginary values mean attenuation, either due to absorption or scattering.

The result is a matrix of size $I \times O$ containing the transmission T_{io} from each input channel to each output channel. Finally, the performance of the geometry is evaluated against a target \hat{T}_{io} and a scalar figure of merit is evaluated as $p = \sum |\hat{T}_{io} - T_{io}|$.

4.4.3 Optimisation

In order to explore a vast number of combinations of the system parameters, a hybrid optimisation approach is employed. The full procedure consists of an optimisation algorithm, specifically a genetic algorithm, nested in a **for** loop.

The external **for** loop is employed to evaluate the performance for each combination of the system parameters, e.g. the wavelength, the input WG width, the width of the MMI, and output WG width. The width of the output WGs is usually fixed to be equal to that of the input WGs, so that the number of DOFs is reduced by one. For each combination, the input field profiles, the MMI field

profiles, their corresponding propagation constants, and the output field profiles are retrieved from the collection of data.

The genetic algorithm, Matlab's own `ga`, is employed to find the best length L and the best lateral displacements x_i and x_o of the input and output WGs, respectively. Again, in order to reduce the number of DOFs a single input is usually considered, thanks to the exploitation of the symmetries in the system. The transmission from the inputs to the outputs is then evaluated, as described in section 4.4.2, and transformed in dB to be able to clearly appreciate smaller variations. If some error occurs in the `ga` evolution, the results are filled with NaN values, so that they are easily recognised and discarded. At the end of the evaluation, summary and in-depth plots can be generated.

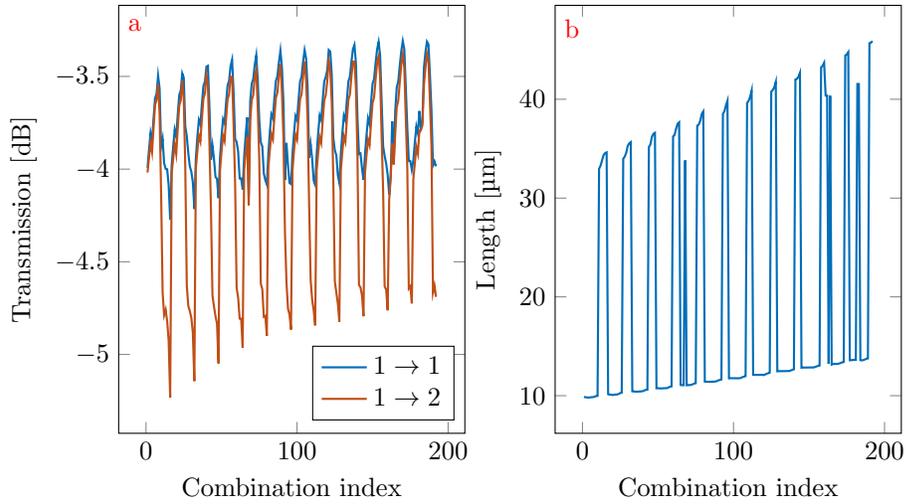


Figure 4.15: Summary plots, displaying the system performance and the geometry parameters across all the tested combinations: **a** transmission from the input WG to each of the two output WGs, in dB, and **b** optimal length of the MMI body, in μm .

The first charts summarise the performance of the MMI geometries optimised for each combination of the system parameters at once. The most important information is given by the transmission values between every pair of input and output channels, as shown in fig. 4.15a. The second most significant characteristic is the MMI body length L , because reducing the footprint frees space that can be used by other structures. Figure 4.15b shows the length L for each of the tested combinations. Finally, additional plots report other significant information, such as the relative position x_i and x_o .

The second, more in-depth plot that can be generated displays information of one combination at a time, as fig. 4.16 shows. In this case, a more complete overview of the field propagation is displayed. Specifically, the plot shows the input field profile in both the input WG and the MMI body, the evolution of the

field profile inside the MMI volume along the $z - x$ plane, and the output field profile in both the MMI body and the output WGs. Moreover, a filter can be applied to produce a single image for the chosen combination or a set of images for each of those combinations which verify a given condition, e.g. those such that the transmission of the first input channel to the second output channel is above a certain threshold.

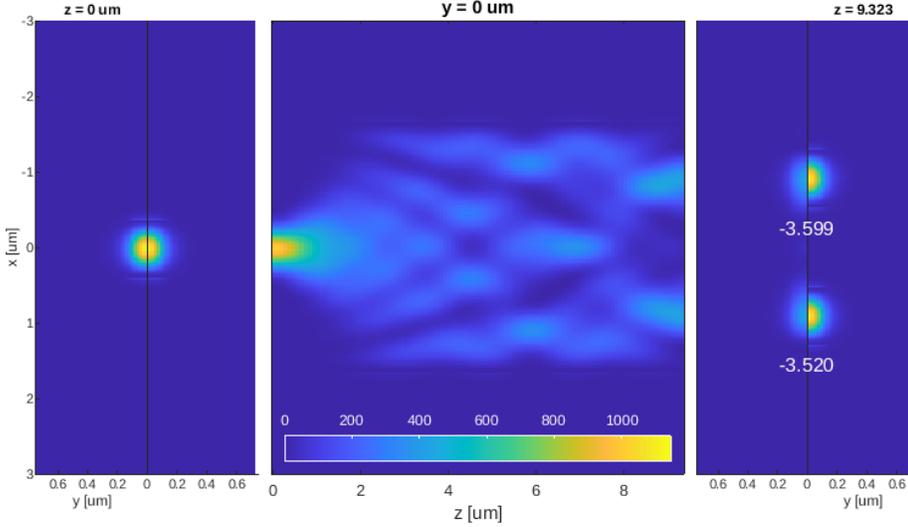


Figure 4.16: Example of in-depth plot, displaying the field profiles at the input and output of the MMI body as well as the field propagation in the $z - x$ plane. Specifically, the left half of the first plot is the input field profile, while the right half is the excited field profile at the input of the MMI. Both maps represent the top half of the section only, assuming symmetry around the horizontal plane. The second plot is the field propagation on the horizontal plane at $y = 0$. Similarly to the input, the third plot represents the field profile at the output of the MMI in the left half and the excited field profile at the start of the output WGs. The numbers represent the transmission values in dB.

4.4.4 Experimental results

I report here the experimental measurements of three MMI designed using this optimisation procedure. These devices have been implemented in the design reported in appendix B.4. The target wavelength was 488 nm, with SiN rectangular WGs embedded in SiO₂ cladding. I run the optimisation routines for a 1×2 , a 1×3 , and a 1×5 couplers, which evenly distribute the input optical power to the output channels. In all three cases, the input WG position x_i was fixed at the centre of the MMI body. For the 1×2 coupler, the positions x_o of the two channels were symmetrical in respect to the centre, so it was required the optimisation of a single DOF in addition to the MMI length L . Similarly, in

the 1×3 coupler the three positions x_o were defined symmetrically in respect to the centre: one shifted by $x = \bar{x}$, another at $x = 0$, and the last at $x = -\bar{x}$. Again, this geometry required the optimisation of two DOFs: the lateral shift \bar{x} and the length L . Finally, for the 1×5 coupler, we resorted for simplicity to the same method, using a single variable \bar{x} to describe all the five lateral shifts x_o : $2\bar{x}$, \bar{x} , 0 , $-\bar{x}$, and $-2\bar{x}$. The length L and the lateral shift \bar{x} were once more the only DOFs to optimise. In general, a second DOF could be substituted to the $2\bar{x}$ lateral shift. Table 4.4 reports the sizes of the selected MMI geometries.

Structure	Output WGs		MMI body	
	\bar{x} [nm]	width [nm]	L [μm]	width [μm]
1×2	606	500	10.335	2.350
1×3	1058	550	11.801	3.100
1×5	700	550	16.939	3.400

Table 4.4: MMIs geometrical parameters: output lateral shift \bar{x} , width of the input and output WGs, and length L and width of the MMI body.

Table 4.5 reports the results of both the simulations and the experimental characterisation of the three MMIs. Even though the structures are symmetrical in the transverse direction, there are some asymmetries in the values of the expected insertion loss. These are due to numerical error in the lateral shift of the field profiles and the limited resolution of the rectangular mesh they are represented with. On the other hand, the real structures are affected by the limits of the manufacturing processes, especially for designs for visible light.

Structure	Output channel	Insertion loss [dB]		
		Expected	Target	Measured*
1×2	1	-3.11	-3.0103	-2.83
	2	-3.10		-3.18
1×3	1	-4.89	-4.7712	-4.95
	2	-4.90		-4.44
	3	-4.82		-4.81
1×5	1	-6.88	-6.9897	-9.21
	2	-7.62		-5.38
	3	-6.42		-10.00
	4	-7.63		-4.32
	5	-6.92		-9.58

Table 4.5: Experimental characterisation results of the MMIs performance. * Note that the ‘‘Measured’’ column reports the insertion loss normalised on the sum of the output power of each channel. This column should be used to compare the channel balance with that of the simulations.

The first two structures appear to be reasonably balanced both in the simulations and in the experiments. The third structure, however, shows a moderate unbalance of the odd channels in respect to even channels in the simulations. Conversely, in the experimental data the values are in favour of the even channel in respect to the odd ones. The discrepancy between the prediction of the model and the experimental measurements underlines the necessity to take into consideration the limits of the manufacturing processes. Figure 4.17 confirms the (un-) balance of the designed structures by observing matrices of scatterers. In these structures, light from a single input channel is distributed over a number of scatterers, thanks to cascades of the MMIs described in this section.

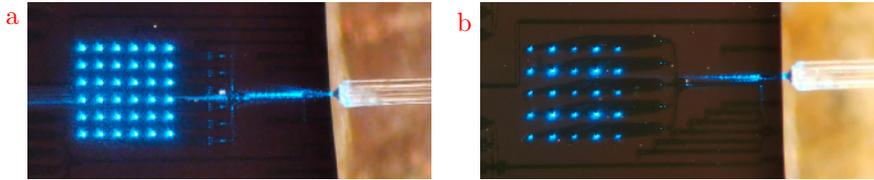


Figure 4.17: Photographs of matrices of scatterers, used to stimulate neuronal cultures. For more information, see appendix B.4, where the chip that contains these structures is described. **a** 6×6 matrix addressed by a cascade of 1×2 and 1×3 MMIs. **b** 5×5 matrix addressed by a cascade of 1×5 MMIs.

4.4.5 Optional 3D simulation

This optimisation procedure has the objective of rapidly exploring a large number of geometries. The method of modal propagation analysis provides a good approximation of the real device and the computational cost is quite low. In comparison, a full 3D model of the MMI evaluated through FEM methods would certainly be more accurate, but also slower and computationally more expensive.

Nevertheless, after the conclusion of the optimisation through modal propagation, it is easy to identify a few of the geometries with the best performance, e.g. low insertion loss, low cross-talk, and high channel balance. Optionally, these geometries could be studied with a full 3D FEM simulation to verify the results. Moreover, these geometries could be optimised further, although in a small range around the initial values, by employing full 3D FEM simulations. Indeed, restricting the range in which the system parameters are optimised makes the use of full 3D FEM simulations more approachable, computationally speaking.

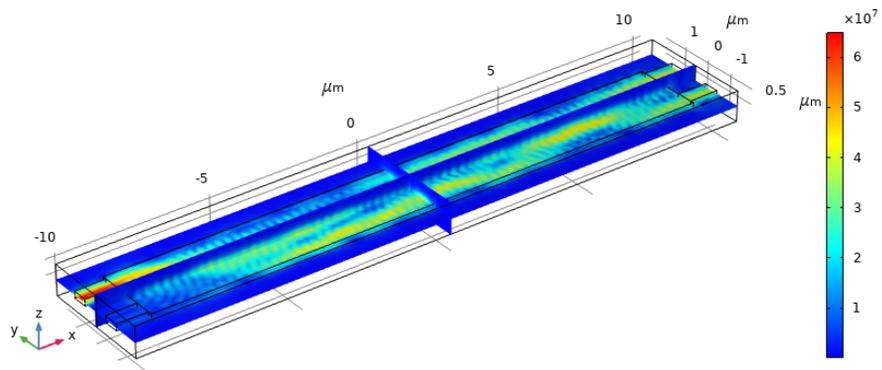


Figure 4.18: Example of 3D simulation of a 2×2 MMI coupler. The colourbar represents the field intensity and is in V/m . Comsol propagation direction is x , and the section is the $y - z$ plane.

Chapter 5

Experimental results

I experimentally characterised two types of structures: the complex perceptron and reservoir networks based on MRRs. The former is a complex-valued feed-forward neural network, where linear memory is provided by integrated delay lines, while the activation function is the square modulus implemented by the photodetector. In the latter kind, it was planned to use the SCISSOR structure, in which the interplay of each MRR provides the complex dynamics for the reservoir. The field is probed between the rings by several channels, whose signals are then coherently summed together, similarly to the complex perceptron. Due to the high propagation loss in these structures, I instead concentrated my effort on the investigation of the dynamical response of a single MRR. Here I report the explorative yet extensive study of a single MRR used as a reservoir with a time multiplexing technique.

5.1 Complex perceptron

The complex perceptron is a single node artificial neural network (ANN) whose inputs and weights are complex-valued and is designed for operation on high-speed optical bitstreams [90]. The input signal is multiplexed over a number of channels, each delaying the signal by a multiple of a given period, i.e. the first channel has no delay, the second channel has ΔT delay, and the n -th channel has $n \cdot \Delta T$ delay. After, complex weights are applied independently to each copy and then they are coherently summed together. The resulting signal is then fed to a photodetector, which carries out the nonlinear activation function, i.e. the square modulus $|\cdot|^2$.

Our implementation has been fabricated with a multi-project wafer (MPW) scheme on the iSiPP50G technology process by imec. The devices are manufactured on SOI wafers with a device layer thickness of 220 nm. Waveguides are 450 nm wide to make sure that both TE and TM polarisations support the first mode only. The propagation loss of about 6 dB/cm has been measured with test delay lines of different lengths on the same chip. This value is quite higher than

the 2 dB/cm one given as a standard performance for imec’s iSiPP50G process. A minimum value of 3.8 dB is the measured insertion loss through a single grating coupler. Localised thermal tuning is enabled by TiN tracks, deposited on top of the silica cladding, directly above waveguide segments. See appendix B.3 for further details.

I tested the device with pseudo-random binary sequence (PRBS) input signals at 5, 8, 10, and 16 Gbps on four different tasks: the XOR operation, 2-bit and 3-bit pattern recognition, and the phase decoding. The results are promising: we see a bit error rate (BER) less than 10^{-5} on both the XOR task and the 2-bit pattern recognition, whereas the 3-bit pattern task showed BER almost as low as 10^{-3} . The phase decoding task was carried out at 10 Gbps reaching error-free operation.

5.1.1 Operation principles

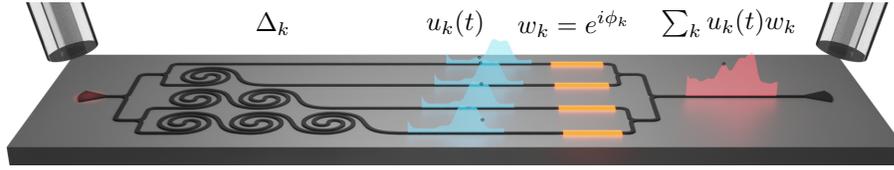


Figure 5.1: Simplified schematic of the complex perceptron [90]. WGs are illustrated as dark paths developing over a substrate. Delay lines are depicted as WG spirals and heaters are the bright volumes covering each WG. Fields are represented as semitransparent plots above the WGs, before the heaters and at the end of the structure. Input and output coupling fibres are shown hovering on top of the input and output grating couplers, respectively.

The simplified design of the complex perceptron is shown in fig. 5.1. The device is a structure with four channels, providing delays of 0, 50, 100, and 150 ps, respectively. Hence, the four copies of the original signal are $u_k(t) = u(t - k\Delta T)$ for $k = 0, \dots, 3$. The weights applied to the copies are defined as $w_k = a_k e^{i\phi_k}$, where the phase is controlled by the thermal heater on top of the waveguides, while the amplitude is due to the propagation losses of the delay lines. The four channels are hence unbalanced and their amplitudes, relative to the zero-delay channel, are $a_k^2 = \{1, 0.58, 0.34, 0.2\}$. The four channels are coherently summed together and their square modulus is physically evaluated by a photodetector. The signal at the output of the device is

$$y(t) = \left| \sum_k w_k u_k(t) \right|^2. \quad (5.1)$$

The input signal $u(t)$ is given by the continuous repetition of a PRBS of length $2^8 = 256$ containing words with 8 bits. The input sequence is always encoded in the optical intensity through a non-return-to-zero (NRZ) code, with the exception shown in section 5.1.5. The binary 0 and 1 are mapped respectively to the

lowest and the highest transparency level of the electro-optic modulator (EOM), with an extinction ratio of 7 dB and a SNR of 14 dB. The input bitrate is either 5, 8, 10 or 16 Gbps so that the oscilloscope with its 80 GSa/s can acquire an integer number of samples per bit.

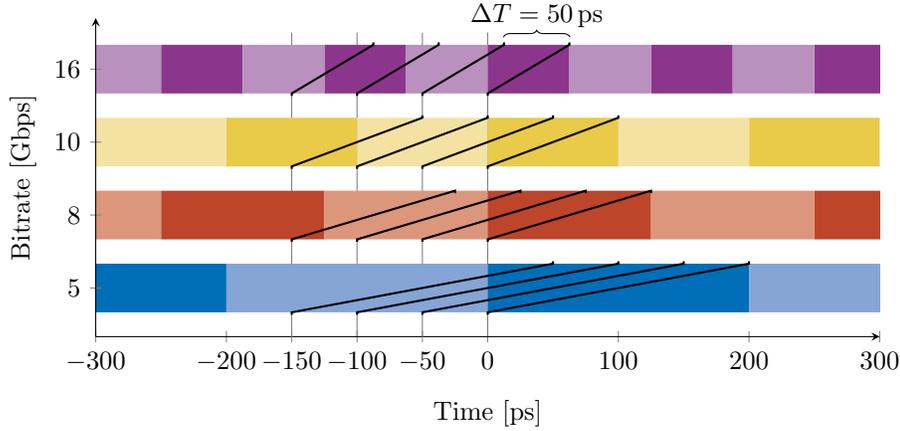


Figure 5.2: Sampling times in comparison with bit periods for different bitrates. Black lines show how adjacent bits are sampled by four channels separated by a constant delay of $\Delta T = 50$ ps.

The possibility to map each bit with multiple samples permits to learn detail of the sequence such as the bit-to-bit transient. Figure 5.2 shows the possible sampling times of each of the four channels in relation to the bit periods of the bitrates considered. For example, at 5 Gbps the four channels are able to observe one or two consecutive bits; at 8 and 10 Gbps only two and three bits are observed at the same time; at 16 Gbps the complex perceptron can either observe three or four consecutive bits. Notice that sampling one bit more than the other might counteract the channel unbalance due to the propagation losses.

The input/output signals are sampled by the oscilloscope at 80 GSa/s for a duration of 2 μ s and a total of 160 kSa. Assuming that the input signal has a given bit rate f_b , then the samples in each bit are numbered with $n = 1, \dots, N_{spb}$, where $N_{spb} = 80 \text{ GSa}/f_b$ are the sample per bit.

bitrate [Gbps]	5	8	10	16
N_{spb}	16	10	8	5

Table 5.1: Sample per bit N_{spb} for the bitrates employed in the experiments.

The acquisition is synchronised with the operation of the current source generator, which defines the phase weights ϕ_k by flowing electrical current in the thermal heaters. The acquisition starts 1 ms after the trigger event, in order to take into account both the response time of the current source generator and

the dynamic of heat diffusion in the device. Due to the path difference for the optical signal to reach the input and the output detector, the output trace has an intrinsic delay with respect to the input one. The delay is evaluated by a calibration procedure every time the bitrate of the input signal is changed (see appendix A.4.4 for details).

The device performance is quantified by evaluating the bit error rate (BER) on experimental traces measured with the experimental setup discussed in appendix A. A generic task involves a number of consecutive bits and the last one is called the reference bit. Causality dictates that the response of the perceptron to a given task should happen at the same time as the reference bit. Hence the target trace is generated from the input one by carrying out the task operation and by placing the result in the same time slot of the reference bit. The BER is then evaluated by comparing bit-wise the output trace to the target trace.

5.1.2 Training and testing procedures

The training procedure consists of the minimisation of the BER changing, for each channel, the value of the current flowing in the heater. Since the number of DOFs is low, I employed a particle swarm (PSW) algorithm in the training routine [91]. Actually, only three out of four phases are trained, since the system is sensitive to the relative phase difference of the channels and not to their absolute phase. The PSW ensures a relatively fast, gradient-free, and, most importantly, robust to noise minimisation of the cost function, i.e. the value of the BER. Structures with more DOFs may benefit from the use of algorithms such as gradient descent, which, however, suffer more from the experimental noise and the presence of local minima. The trade-off between faster minimisation and result accuracy must be analysed case-by-case.

During the training routine, the PSW algorithm evaluates a given cost function, which eventually estimates the BER of the output trace produced by the system in the current state. The cost function is composed of the following steps:

1. A new set of currents are applied to the device by the current controller, which also generates the trigger signal for the oscilloscope.
2. The oscilloscope acquires the input $\{x_j\}$ and the output $\{y_j\}$ traces, with 160 kSa each.
3. The two traces are aligned and extra samples are discarded, obtaining traces of length $M \times N_{spb}$. They are then rescaled between 0 and 1.
4. The input trace $\{x_j\}$ is digitised to $\{X_j\}$, comparing the level of the middle sample ($\lfloor N_{spb}/2 \rfloor$) of each bit to the average level of the trace, i.e. 0.5.
5. The target trace $\{T_j\}$ is evaluated by applying the task operation to $\{X_j\}$.
6. The output trace $\{y_j\}$ is digitised for each sample n and using a variable threshold r , obtaining a set of output sequences $\{Y_j\}_{n,r}$. The values of r are evenly spaced between 0 and 1, extremes excluded.

7. The error number $E_{n,r} = \sum_{j=1}^M \|Y_j - T_j\|$ is evaluated comparing each sequence in the set $\{Y_j\}_{n,r}$ to the target sequence. The BER is given by the minimum value $C = \min_{n,r}(E_{n,r})$, found for the best sampling position n_b and the best threshold value r_b .

The PSW keeps iterating until either error-free operation or the max iteration condition is reached. The optimal currents resulting from this procedure are used to test the final performance of the system. The test procedure repeats the measurement of the cost function, usually 10 times, while applying the optimal currents. The final value of the BER is the average of these BER measurements.

It is important to note that the response of the complex perceptron is not affected by sequences longer than the observed bits, due to its linear memory. Moreover, all the possible input combinations are included in the PRBS, as its order is higher than the maximum number of observed bits. Hence, the optimisation procedure has sufficient information to describe the complete task and is only affected by the experimental noise. The most significant sources of noise on the perceptron are: phase noise on the weights, due to temperature fluctuation of the delay lines and to noise of the current source generator; detection apparatus noise, e.g. detector noise on the voltage signal and jitter of the oscilloscope sampling time. The noise slows the optimisation procedure, however, it also produces more true-to-life results with respect to simulations. The use of adaptive thresholds on both signal intensity and sampling time is normally used in real case scenarios for telecommunication applications to minimise the effects of slow offsets and drifts owed to the environmental parameters, such as ambient temperature.

5.1.3 Pattern recognition

The pattern recognition task requires the output of the structure to be high, if and only if a given patterns appears at the input. I analysed the device response to 2-bit and 3-bit pattern. “0 0” and “0 0 0” patterns are not considered, because their output would be identically null. The probabilities of giving the correct answer with a random choice, i.e. the random choice thresholds, for 2-bit and 3-bit patterns are 25% $\approx 10^{-0.6}$ and 12.5% $\approx 10^{-0.9}$, respectively.

We can observe more in detail the operation of the trained device during the recognition task of the pattern “1 0” at 16 Gbps. Figure 5.3a shows the distribution of the normalised amplitude of the current sample for the four possible patterns in the input. The four distributions are centred close to the value of the reference bit, i.e. 0 for “0 0” and “1 0” and 1 for “0 1” and “1 1”, with just small offsets due to the inter-bit interference. With these distributions, it would be fairly difficult to separate the first two patterns and almost impossible to separate the last two, as their distributions overlap. However, the distribution of the output signal, similarly illustrated in fig. 5.3b, is completely different. In fact, the distribution for the target pattern is almost completely separated from the distributions of the other three patterns. Hence, the output of the device is converted into a binary sequence with a simple threshold operation.

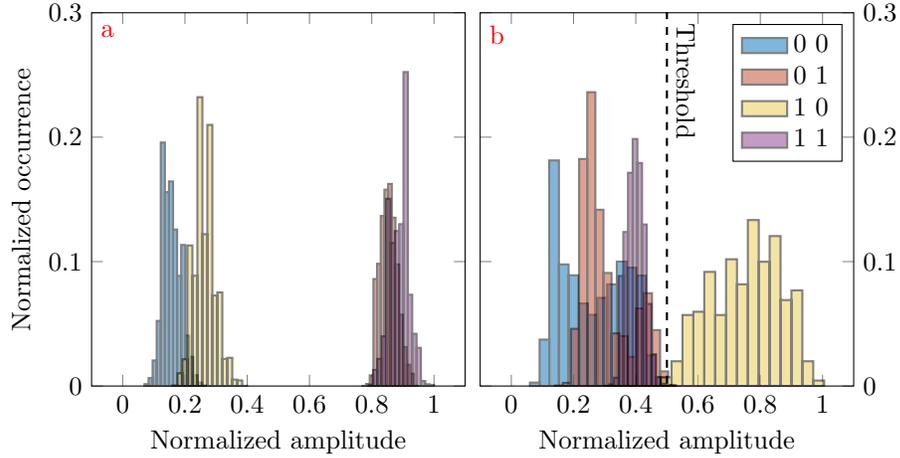


Figure 5.3: Operation during the recognition task of pattern “1 0”. **a** Distribution of the input intensity at the best sampling time and **b** distribution of the output intensity at the best sampling time. The vertical dashed line indicates the optimal threshold.

Figure 5.4 shows a fragment of the typical time traces of the input and output optical amplitudes compared to the target sequence. Notice that the target trace is high only when the reference bit is low and the bit in the past is high. The output level is able to replicate this behaviour by choosing the best sampling time and the threshold value.

Figure 5.5 shows the performance reached by training the device to recognise different patterns at four bitrates: 5, 8, 10, and 16 Gbps. Pattern recognition of 2-bit sequences shows outstanding results, with most cases achieving a BER well below 10^{-4} , many of which even reaching the statistical limit. The best case is obtained at 5 Gbps, achieving its statistical limit for each of the three patterns.

On the other hand, 3-bit pattern recognition tasks display mixed results. In this case, the performance is heavily influenced by the bitrate of the input signal. In fact, for 5 Gbps, the bit period is 200 ps, longer than the maximum observable window of 150 ps. Performance at this bitrate is marginally better than the random choice threshold. At 8 and 10 Gbps, the bit periods decrease to 125 and 100 ps, respectively, and the device can observe three bits, by cleverly positioning its sample just before/after the bit transients. The performance improves in both cases, except for the “1 0 0” pattern. At 16 Gbps, the bit period becomes 62.5 ps, which means that the device observes a window of four bits. This specific configuration is able to reach very low values of BER for patterns in which the first bit is high, i.e. “1 1 1”, “1 1 0”, and “1 0 1”, except for “1 0 0”, while has more trouble with patterns that begin with a 0. The worst performing pattern, for all bitrates, is the “1 0 0” one, which requires the perceptron to make principally use of the channel with the longest delay line and highest propagation loss. The high propagation loss, leading to higher channel unbalance, and the absence of amplitude control on the weights are the reasons why the device is able to perform well in some cases but not in all.

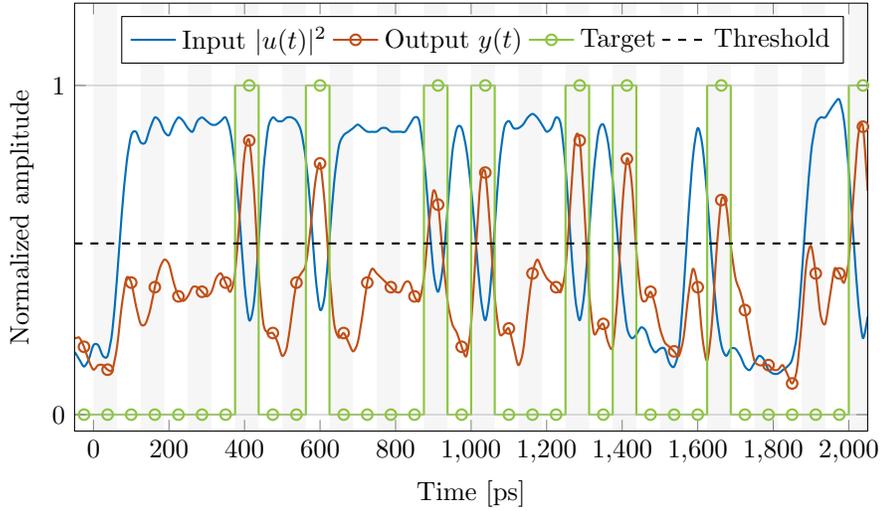


Figure 5.4: Time traces for the “1 0” pattern recognition task at 16 Gbps. Input measured optical intensity (blue), output measured optical intensity (red), and ideal target response (green) as functions of the bit number. Line marks represent the best sampling time, while the dashed line represents the optimal threshold.

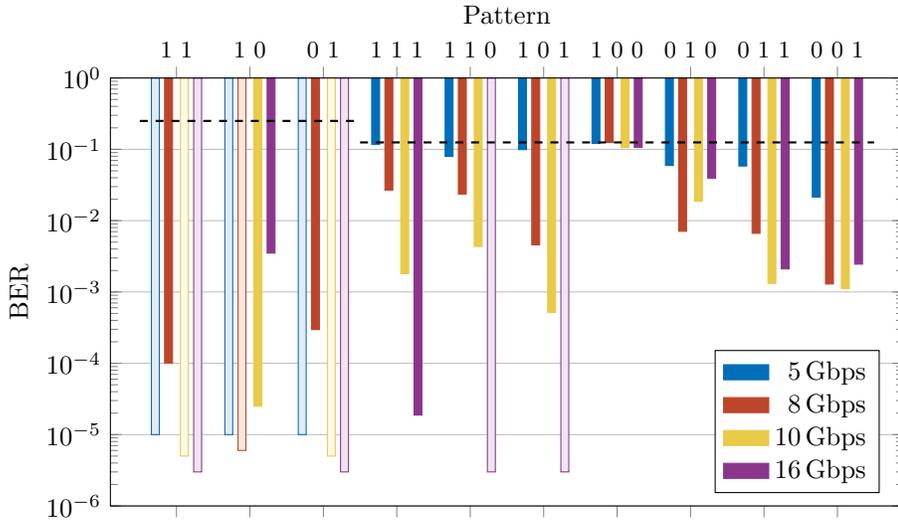


Figure 5.5: BER values at the best combination of input power and sampling time for bitrates of 5, 8, 10, and 16 Gbps and different patterns composed of two and three bits. Transparent bars indicate that the value is capped at the statistical limit of the recorded trace. The dashed lines represent the random choice threshold.

5.1.4 Delayed XOR operation

The n-bit delayed XOR operation consists of applying the XOR binary operation between the reference bit and the n-th previous bit, e.g. the 1-bit delayed XOR is carried out on consecutive bits. I studied the perceptron performance in the case of 1 to 3-bit delays. The XOR binary operation produces outputs that are non linearly separable, so a real-valued perceptron would not be able to solve successfully this kind of operation. The linear separability threshold achievable by a real-valued perceptron is 25% or $10^{-0.6}$.

Figure 5.6 shows a summary of the results of the delayed XOR operation at different bitrates. In the case of 1-bit delay, the perceptron achieves error-free operation (statistical limit of $\approx 10^{-5}$) at the bitrates of 5 and 8 Gbps, while at 10 and 16 Gbps we see BER values close to 10^{-2} and 10^{-3} , respectively, in any case well below the linear separability threshold. At the two lowest bitrates, the device is able to sample the input signal across the transient between adjacent bits, extracting information from both. On the other hand, at 10 and 16 Gbps the device struggles to obtain good results. This may be due to the fact that at higher bitrates the levels of the samples are more affected by the bit transient, leading to an increase in signal noise and reduced performance.

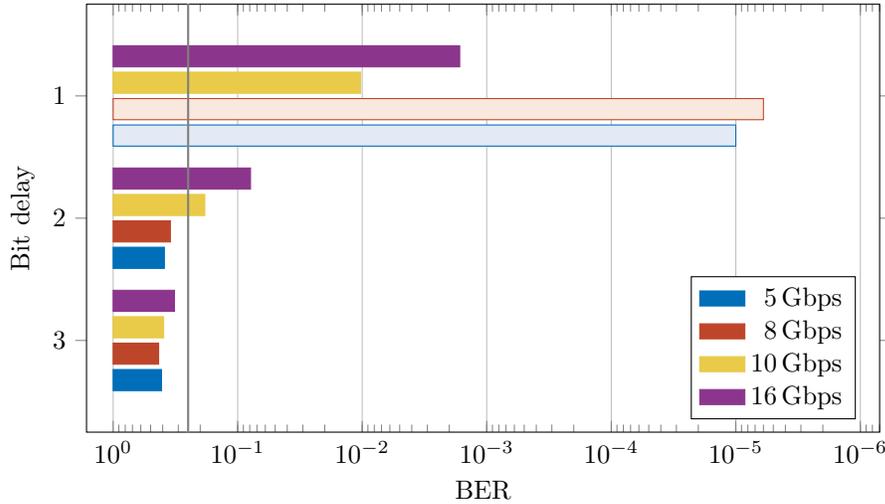


Figure 5.6: BER values at the best combination of input power and sampling time for different bitrates, i.e. 5, 8, 10, and 16 Gbps, and for different bit delays. Transparent bars indicate that the value is capped by the statistical limit of the recorded trace. The solid line represents the random choice threshold.

Bit delays of 2 and 3 show a completely different picture. In fact, with 2-bit delay, only 10 and 16 Gbps show a BER slightly better than the separability threshold, while with a 3-bit delay no bitrate is able to overcome the threshold.

To better understand the role of the sampling time, let's examine its effect on the value of BER. Figure 5.7 shows the value of BER as a function of the

sampling position n and for several different values of input optical power, in the case of a 1-bit delay at 5 Gbps. Regarding the input power, the best results are obtained at the highest input power. The obvious conclusion is that the SNR of the output signal decreases drastically at lower input optical powers.

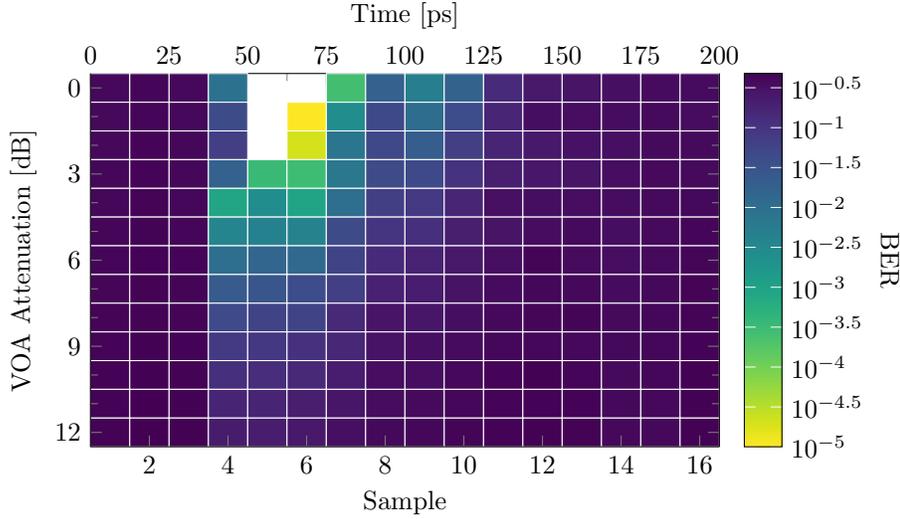


Figure 5.7: Map of the BER of the 1-bit XOR operation at 5 Gbps versus the time from the beginning of the bit and input optical power. The best value is obtained at the maximum optical power and 62.5 ps. The BER statistical limit is 10^{-5} .

The sampling positions are spaced 12.5 ps one from the other, i.e. the inverse of the sampling frequency of the oscilloscope. At this bitrate, the best sampling time is clearly around 62.5 ps. The channel with the shortest delay (and lowest attenuation) is sampling the reference bit, the second shortest is on the transient, while the remaining two channels are observing the value of the past bit. By choosing the correct sampling time the device is selecting which features of the bit shape to observe and which to avoid, e.g. the bit value in the middle of the bit or the bit transient at the beginning/end of the bit. Furthermore, since the device has no control over the amplitudes of the weights, the perceptron tries to equalise the channel unbalance by sampling twice the bit in the past and once the reference bit. An example of the time traces of the input and output optical amplitudes as well as the target sequence for the 1-bit delayed XOR operation carried out at 16 Gbps is shown in fig. 5.8. The horizontal line represents the best threshold and the line markers indicate the best sampling time. Figure 5.9, on the other hand, reports the distributions of the samples taken at the best sampling time, for the input bitrate of 16 Gbps. It is clear that the effect of the perceptron decreases the value of the “1 1” while maintaining high the value of both “1 0” and “0 1”.

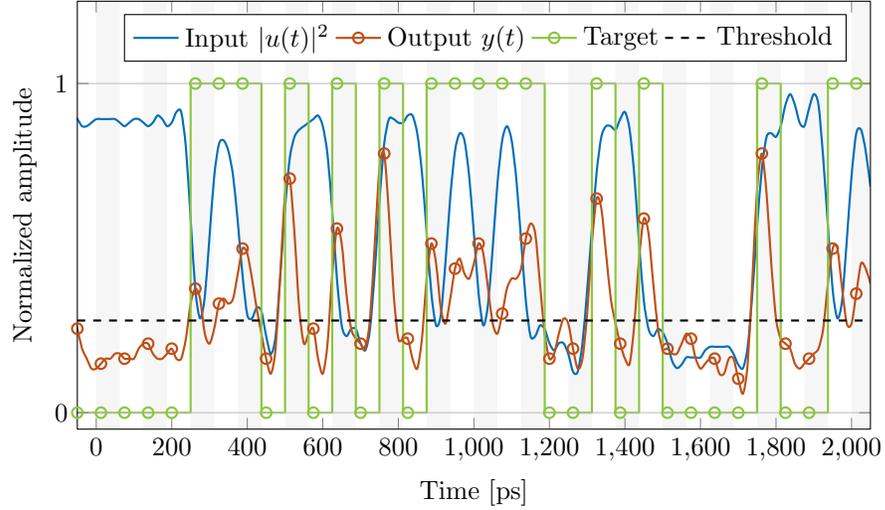


Figure 5.8: Example of time traces for the 1-bit XOR operation. Input measured optical intensity (blue), output measured optical intensity (red), and ideal target response (green) as functions of the bit number. Line marks represent the best sampling time, while the dashed line represents the optimal threshold.

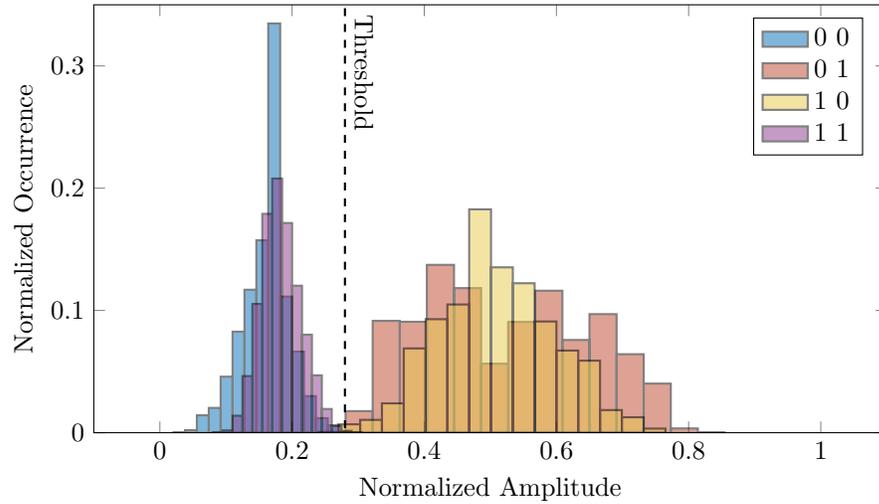


Figure 5.9: Distribution of the best sample output intensity for the four cases of the ideal input bits, i.e. “0 0”, “0 1”, “1 0”, “1 1”. The vertical dashed line indicates the optimal threshold.

5.1.5 Carrier phase decoding

While all the previous tasks have been carried out on NRZ signals encoded with amplitude modulation, it is possible to employ different encodings. An interesting example is given by phase-shift keying (PSK), a digital modulation format in which each symbol is associated with a specific phase of the carrier wave. In this case, I used a binary encoding where binary 0 and 1 are mapped to phase 0 and π , respectively.

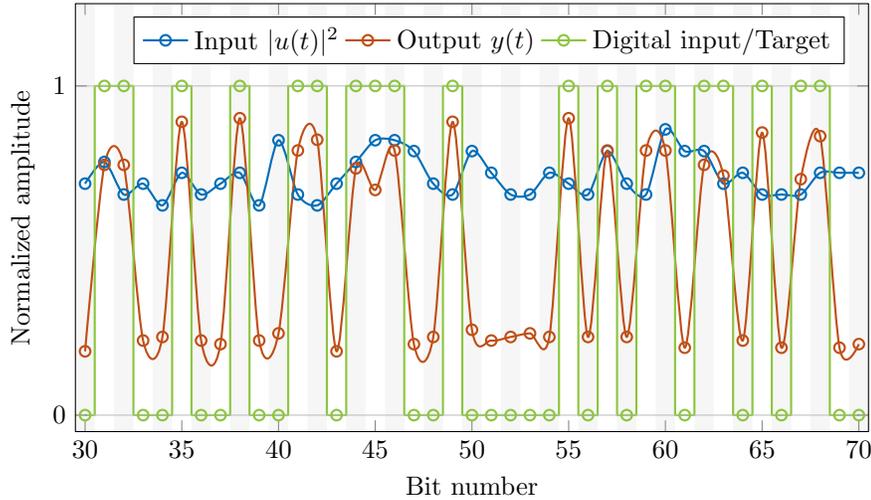


Figure 5.10: Time traces for PSK decoding at 10 Gbps. Input measured optical intensity (blue), output measured optical intensity (red), and ideal target response (green) as functions of the bit number.

Figure 5.10 shows an example of the time traces for the PSK decoding operation carried out at 10 Gbps. The ideal input sequence is encoded in the input optical signal, changing only its phase. After the training procedure, the complex perceptron is able to continuously decode pseudo-random binary sequences without errors, as seen from the trace of the output intensity. The Pearson correlation coefficient between the ideal input sequence and the measured optical intensity of the input trace is 0.002. From such a low value, we can infer that the perceptron response is not due to the input intensity, but rather to its optical phase.

5.1.6 Discussion

The device performance is influenced by several factors. The length of the delay lines is the most important characteristic of the perceptron, as it sets the boundaries of the observation window and the number of optical samples per bit, for a given bitrate of the input signal. These features alone restrict the tasks that the device can perform. For example, with four channels evenly spaced by

$\Delta T = 50$ ps, our perceptron cannot observe more than two bits when the input signal is modulated at 5 Gbps.

Moreover, the presence of delay lines introduces propagation losses, which unbalance the optical power delivered at the end of each channel and decrease the SNR at the output. Propagation losses are proportional to the length of the waveguide, hence a perceptron designed for higher bitrates would suffer less from performance degradation due to channel unbalance and lower SNR. Since in our case each weight in the perceptron acts on the phase of the channel but not on its amplitude, the choice of the best sampling time becomes rather important, due to channel unbalance. In fact, by choosing the right sampling time the device can sample a certain bit more than the others, acting against the channel unbalance. Replacing the stage in which MMIs evenly split the input power to all the channels with a cascade of addressable Mach-Zehnder interferometers (MZIs) would allow to control the amplitude weight and to correct the channel unbalance at the same time. This would increase both the power draw and the device performance.

Notably, our measurement setup contains the most important noise sources typical of real-world optical transmission lines in telecommunication networks. Moreover, the task investigated, e.g. pattern recognition, XOR operation, or PSK decoding, are general and do not limit the application spectrum of the complex perceptron.

Unlike systems employing software linear regression, which require offline data elaboration, this device could work instantaneously, once the threshold level and the sampling time are set. However, these variable parameters are commonly updated online in optical transmission lines. A major improvement of our complex perceptron would be to minimise the thermal cross-talk, which is currently limiting the continuous operation of the device.

Finally, comparable opto-electronic systems performing similar operations, usually (super)sampling the input optical signal, are affected by the limited bandwidth of electronics and they require energy-consuming components such as high-speed analog-to-digital converters (ADCs) and field-programmable gate arrays (FPGAs). In comparison, the complex perceptron is a transparent device, whose only power draw comes from the heaters and a photodiode, and it is easy to be installed on existing hardware.

5.2 SCISSOR-based reservoir networks

A SCISSOR is a structure made by a sequence of MRRs whose bus waveguides are shared with each other [83]. The idea at the base of this device is to create a resonant cavity that could be used as a reservoir. Depending on the relative position of the resonances, the output of these structures changes dramatically. At low input power, two resonances could couple together forming a resonance with a higher Q-factor. At higher input power, nonlinear effects are triggered in different resonators simultaneously. The output signal is then affected by the specific dynamics of each resonator. Additionally, delay lines can be interposed

on the bus WGs between one ring and the next. This introduces in the systems new dynamics which work at timescales fixed by the length of the delay lines and thus enabling an even more complex signal mixing.

In the first BACKUP photo-lithographic mask (see appendix B.3) I designed a set of SCISSOR-based networks with different design parameters, such as the ring coupling coefficients, the presence of a delay line between the rings and the value of the delay provided. Unfortunately, the devices delivered by the manufacturer had two important problems: the linear losses were much higher than the supposed nominal value (about $3\times$) and the time scale of the FC nonlinear effects was $\tau_{FC} = 45$ ns [75, 92], i.e. between one and two orders of magnitude larger than the one expected [70–74]. The former hindered the use of the devices with delay lines because the signal was attenuated too much. The latter moved the nonlinear dynamics of the MRRs towards lower frequencies, hampering the performance of reservoir networks with short or no delays.

5.3 Microring resonator as a reservoir

Due to the aforementioned manufacturing problems, I took a step back and studied the dynamics of a single MRR manufactured with better quality.

The structure is a MRR in the ADF configuration with a radius of $7\ \mu\text{m}$ and WGs with section $450\ \text{nm} \times 220\ \text{nm}$. There are two point-coupling regions with a gap of $200\ \text{nm}$. Figure 5.11 shows the normalised transmission spectrum obtained by measuring the *drop* port. The quality factor has been estimated to be $Q = 6 \times 10^3$ at a frequency of $193.5\ \text{THz}$. Consequently, the photon lifetime in the cavity is of the order of $\tau_{PH} = 5\ \text{ps}$ [93].

Differently from the structure of section 5.2, which showed a FC lifetime $\tau_{FC} = 45\ \text{ns}$ [75, 92], this MRR has a FC lifetime of about a few ns, which is at least one order of magnitude shorter. Because of this, transient optical phenomena related to FC occur at a temporal scale ten times faster than in the other devices. This value is consistent with the common values reported in the literature [70–73]. Moreover, the thermal relaxation time for similar structures is expected to be in the order of $100\ \text{ns}$ [27, 94].

Here I propose to study this single MRR employed as a reservoir in a RCN [95], in which its optical response, driven by the nonlinear dynamics taking place in the cavity, is sampled by an electronic readout layer (implemented offline for simplicity). My resolve is to explore the performance of the system as a function of the parameters of the input optical signal.

5.3.1 Operation principles

The technique I propose in this section is similar to the one based on the delay loop [38], however, in this case, no delayed feedback is present. The reservoir is composed of a single MRR and the linear readout layer. The input signal is transformed by the dynamics of the MRR, which arise from the enhanced EM

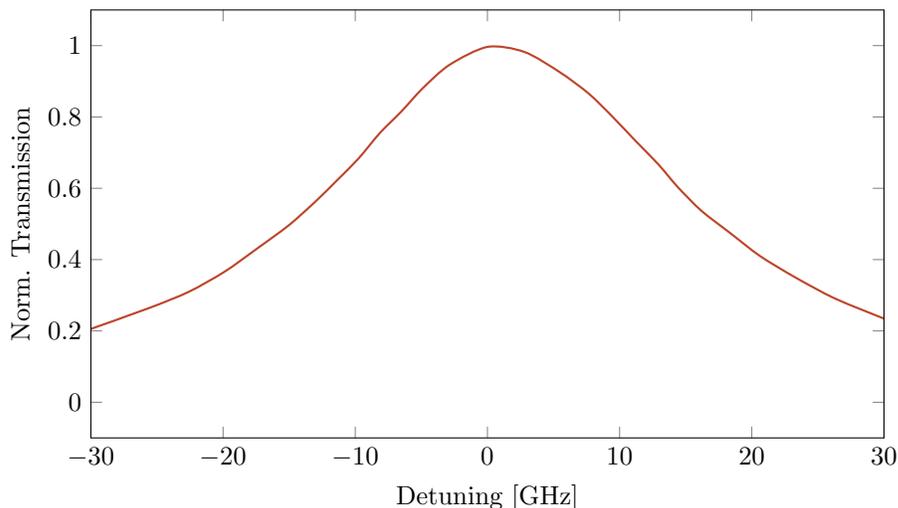


Figure 5.11: Normalised transmission spectrum around the resonance, measured at the *drop* port.

field in the cavity generating nonlinear effects. The output signal is fed to the readout layer, which is carried out offline.

The input data is a binary sequence

$$D^I = [b_1, \dots, b_j, \dots, b_N], \quad (5.2)$$

encoded in the amplitude of the optical signal X^I which is given as input to the MRR. The output is then sampled multiple times for each bit, creating the trace of bits:

$$X^O = [x_1^O, \dots, x_j^O, \dots, x_N^O], \quad (5.3)$$

where each element contains the samples from 1 to the number of *samples per bit*, S ,

$$x^O = [s_{j,1}^O, \dots, s_{j,k}^O, \dots, s_{j,S}^O]. \quad (5.4)$$

The input optical signal is sampled in the same way, creating a trace X^I composed of the bits x_j^I , each containing S samples $s_{j,k}^I$.

The number of virtual nodes N_V is given by the number of *samples per bit* S and the number of bit n_R in the observation window, with the temporal separation among nodes being $\delta T = T_B/S$ linked to the bit duration T_b . Figure 5.12 graphically explains the same concept.

Then the readout layer produces the output Y by multiplying the virtual nodes in each bit with the weights. The total number of virtual nodes depends on how many bits are observed by the readout layer. Both input and output sequences can be reshaped into matrices X_M^I and X_M^O , whose columns are the vectors $[x_{j-k}^O, \dots, x_j^O]^T$ (and similarly for X^I), containing all the samples of

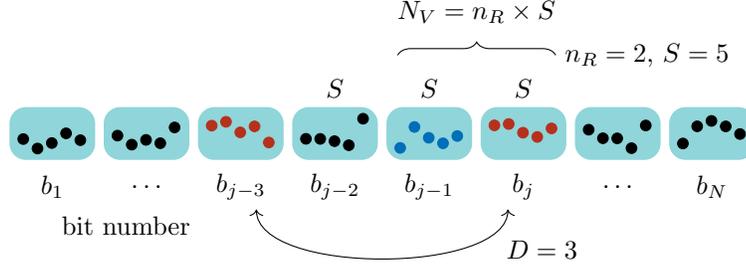


Figure 5.12: Relation between the number of virtual nodes N_V , the *samples per bit* S , and the number of observed bits n_R . The bits involved in the task are separated by a delay D .

a sequence of $k + 1$ bits. For example, the column in X_M^O for an observation window of $n_R = 3$ bits would be $[x_{j-2}^O, x_{j-1}^O, x_j^O]^T$. Finally, the output of the RCN is produced as:

$$Y = WX_M^O. \quad (5.5)$$

In general, W is a $Q \times N_V$ matrix where Q is the number of classes of the output and $N_V = n_R S$ is the total number of virtual nodes observed. In this case, the tasks are Boolean, hence a single class $Q = 1$ is sufficient.

The training, i.e the optimisation of the elements in W , is accomplished by regularised least squares (ridge regression) [92, 96], in which the λ parameter is decided with 5-fold cross-validation. The result of this optimisation procedure is the matrix W that minimises the regularised least square error $\sum_{i=1}^N \|Y_T - WX\|^2 + \lambda^2 \|W\|$, where Y_T is the target output sequence obtained from the task function as $Y_T = f_{task}(D^I)$. For this operation, Matlab’s `fitrlinear` algorithm is employed.

AND	0	1	XOR	0	1
0	0	0	0	0	1
1	0	1	1	1	0

Table 5.2: Truth tables for AND and XOR logical operations.

This RCN was trained to solve AND and XOR logical operations, whose truth tables are presented in table 5.2. Each task is carried out as a moving logical operation (LO) on the sequences of bits, involving the current bit and another bit in the past. The delay between the two bits is D , where $D = 1$ represents consecutive bits. The total number of virtual nodes fed as samples to the ridge regression is represented by a window of n_R Ridge-bits, starting from the current bit and going backwards, each containing S samples. To simplify the discussion, I introduce the notation “ $LO D$ with n_R Ridge-bits and S ”, whose general schematic is reported in fig. 5.12. As an example, “AND 3 with 2 Ridge-bits and $S = 5$ ” is the AND operation carried out between the current

bit and the one three bits in the past, while the virtual nodes are 10, belonging to the current bit and the one before only.

The following combinations were investigated:

- “LO” AND and XOR
- for each LO , D from 1 to 3
- for each D , n_R Ridge-bits from 1 to $D + 1$
- S equal to 3, 4, 5, 10, 15, 20, and 30

The fundamental characteristics that these tasks are meant to explore are the fading memory of the system [97–99] and its nonlinear transformation of the input. The AND operation is a linear task and therefore the input data is already linearly separable, hence the linear regression is always able to correctly predict the target. However, if the input data to the linear regression do not comprise both bits, i.e. $n_R < D + 1$, that missing information must be recovered from the nonlinear transformation imprinted on the output optical signal by the MRR dynamics. Furthermore, by changing the delay D associated with the logical operation and the number of Ridge-bits n_R supplied to the regression, the limiting values of the fading memory of the system can be observed.

On the other hand, the XOR operation is not linearly separable. In this case, when the input of the ridge regression contains information on both the bits involved, the operation can be solved if the MRR dynamics imprints on the output a sufficiently important nonlinear transformation. Additionally, by reducing the number of Ridge-bits for a fixed delay, the MRR must supply both a nonlinear transformation and fading memory in order for the ridge to correctly predict the target output.

5.3.2 Experimental measurements

The MRR is stimulated with a PRBS of order 8 and length $2^8 = 256$, repeated indefinitely. The setup, discussed in detail in appendix A, allows changing the bitrates, the frequency, and the power of the signal (measured at the input fibre). The response of the MRR has been recorded for each combination of the following variables:

- **bitrate:** 20, 40, 50, 80, 100, 200, 250, 400, 500, 800, 1000, 2000, and 4000 Mbps;
- **detuning:** -30 , -25 , -20 , -15 , -10 , -5 , 0 , 5 , 10 , 15 , 20 , 25 , and 30 GHz;
- **power:** 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, and 18 dBm.

To ensure a fair comparison, the SNR at the output photodetector must be the same for all measured traces. For this reason, the optical power is maintained at a constant level by a VOA and an EDFA placed in the optical line between the sample and the output photodetector.

The opportunity to vary both the frequency and the optical power of the input signal allows studying the effect of the MRR nonlinear response for a broad region of the parameter space of the system [75]. Indicatively, the MRR response can be considered linear at all frequencies at 8 dBm, while at the highest values of input optical power self-pulsing is observed. For example, an input signal resonant with the MRR, with an average power of 16 dBm, triggers the self-pulsing effect. In between, where no self-pulsing is observed, the nonlinear effects still influence the MRR's response. Furthermore, the detuning of the input signal from the resonance frequency can have several effects. At the timescale of the bit, fast dynamics (FC effects) are stimulated. In this case, negative detuning produces optical limiting, while positive detuning leads to optical bistability. Additionally, at longer timescales, such as those at and below 10 MHz, the TOE influences the response of the MRR. In this case, also the average optical power incident on the device drive the thermal effects. The relation is opposite to that of the FC effects: negative detuning causes optical bistability and positive detuning cause optical limiting.

In order to maintain an integer number of samples per bit, the input bitrate has been set to submultiples of the sampling rate of the oscilloscope (20 GSa/s). Therefore, each bitrate has a different number of S . Specifically, the number is as low as 5 at 4000 Mbps and as high as 1000 at 20 Mbps. To be able to compare the results of the ridge regression optimisation algorithm, the same number S should be used. For this reason, before proceeding with the training, the traces are *re-binned* to obtain a fixed number S . As an example, suppose $S = 10$, then the samples for each bitrate are separated into 10 bins in the following way:

- from 20 to 1000 Mbps, each of the ten bins contains an equal (± 1) number of samples, e.g. at 800 Mbps the ten bins contain respectively 3, 2, 3, 2, 3, 2, 3, 2, 3, and 2 samples.
- at 2000 Mbps, each bin contains exactly one sample
- at 4000 Mbps, the first five bins contain one sample each and the remaining five are filled with zeros.

5.3.3 Results

The input and output traces recorded for each combination of bitrate, detuning, and power are analysed by training a ridge regression algorithm for several tasks. The study has been carried out for several values of virtual nodes N_V , observed bits n_R , and samples per bit S , but I report here only the case where $S = 5$, which is the number of S for the highest bitrate. In general, higher values of S showed superior results but they would also translate into a requirement of faster hardware to carry out the readout operation in real-time.

The performance of the network has been evaluated both on the input and output signals, as the BER for each combination of the input optical signal properties, resulting in two three-dimensional arrays. Then, for each combination

of bitrate and detuning, the best value of the BER and the corresponding optical power are extracted. Eventually, three important quantities are obtained from the analysis. First, the best BER evaluated on the output signal as a function of the bitrate and detuning of the input signal. Second, the ratio $RB = \min(\text{BER}^I) / \min(\text{BER}^O)$ which highlights the combinations in which the performance increases thanks to the presence of the MRR as a reservoir. Third, the optical power of the input signal associated with the best BER.

For the sake of clarity, these three quantities are reported as contour maps where the colourmap is represented with a logarithmic (\log_{10}) scale. The BER maps have a blue (good) to yellow (bad) colourmap and fixed range $[-3.5, -0.5]$. Markers pinpoint the combinations achieving the statistical limit of the BER. The power maps report the power in dBm, with a colourmap going from 8 dBm (black) to 18 dBm (bright yellow). The RB maps have a double colourmap, going from white (bad) to black, then blue, and finally yellow (good). Specifically, the colour white marks the combinations for which the performance of the MRR is worse than that of the input signal, the colour black identifies the combinations where the two values are the same, and the colours blue/green/yellow designate the combinations for which the MRR performance shows and improvement over the one evaluated on the input signal. In this case, the combinations where the statistical limit is achieved on the input signal are reported as well, with cross markers.

Linear operation, logical AND

Linear operations between delayed bits are used to test the fading memory of the MRR reservoir.

AND 1 with 2 Ridge-bits

The baseline for the results regarding the logical operation AND is given by the task “AND 1 with 2 Ridge-bits and $S = 5$ ”. In this case, both the current and past bits are given as input to the ridge regression. Being a linear task, the training both on the input and on the output should be able to correctly predict the target output. Indeed, this is the case, as shown in fig. 5.13.

The statistical limit (markers) is reached almost everywhere on the map, with a minimum BER value of $10^{-3.4}$. Furthermore, no evident dependence of the BER on the detuning or the bitrate is observed. The power map reports principally the lower value (8 dBm), which means that the MRR is driven in the most linear regime within the observed range. Finally, a slight deterioration of the performance is observed in a small region at higher bitrates, at 2000 and 4000 Mbps, and it is visible also in the RB map.

AND 1 with 1 Ridge-bit

The previous task “AND 1” is repeated, this time providing only the current bit to the ridge regression instead of both. In this case, in order to correctly carry out the task, the ridge regression must extract from the current bit the information regarding the value of the past bit, imprinted on the output signal

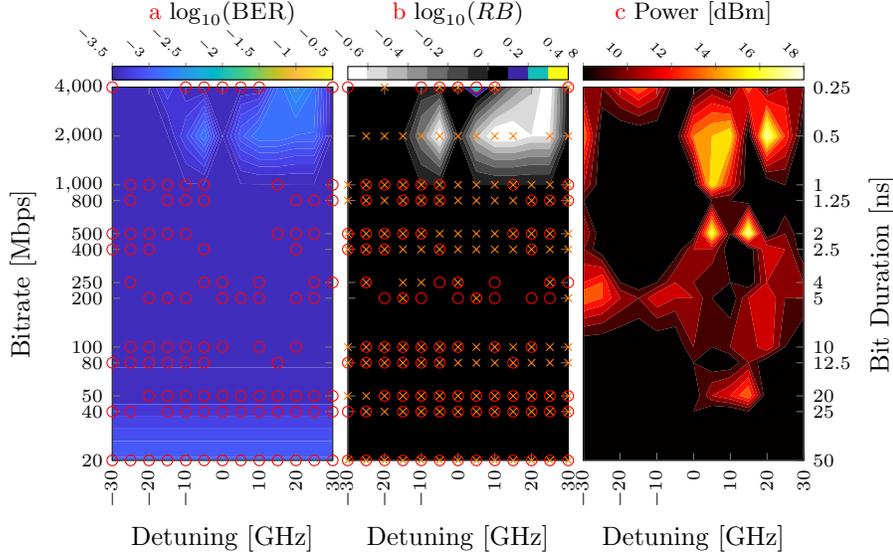


Figure 5.13: AND 1 with 2 Ridge-bits and $S = 5$. Comparison of **a** BER map, **b** RB map, and **c** power map. Markers identify the combinations that achieved the statistical limit (error-free operation): circles for the input signal and crosses for the output signal.

by the nonlinear transformation due to the MRR dynamics. Figure 5.14 shows a situation diametrically opposed to the one observed earlier. The performance by analysis splits the BER and RB maps into two distinct regions. From 20 Mbps to 500 Mbps the values of the BER are very low (min. $10^{-3.4}$), usually reaching the statistical limit, and there is a marked improvement in respect to that evaluated on the input, with values of RB between 10^1 and 10^2 . From 500 Mbps to 4000 Mbps, instead, the values of the BER increase by up to two orders of magnitude and, more importantly, the RB map shows mixed results, from limited improvements ($10^{0.5}$) to decisively adverse effects (10^{-2}). In the power map, a similar division arises, with a separation around 500 Mbps. At low bitrates, low optical power is sufficient to obtain good performance around zero detuning, while it smoothly increases when moving away from the resonance frequency. This is a clear dependence of the system performance on the optical frequency. At high bitrates, on the other hand, there is no clear dependence on the frequency for the power map.

The important aspects that emerged from the analysis of this task are two. The first is that the system is able to perform the “AND 1” task with virtual nodes belonging only to the current bit, hence providing a memory of 1 bit. The second and more critical result is that at high bitrates, the input signal already contains distortions which allow the ridge regression to carry out the linear task.

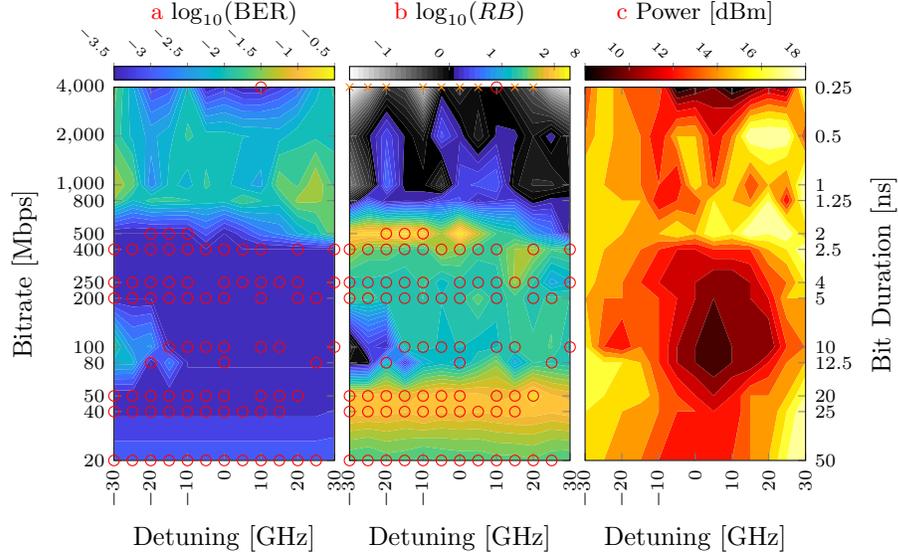


Figure 5.14: AND 1 with 1 Ridge-bits and $S = 5$. Comparison of: **a** BER map, **b** RB map, and **c** power map. Markers identify the combinations that achieved the statistical limit (error-free operation): circles for the input signal and crosses for the output signal.

The bitrate of 4000 Mbps is the extreme case in which the BER evaluated on the input achieves the statistical limit. This confirms the presence of undesirable memory in the signal, which the generation and detection stages alone are able to exploit to solve the task.

AND 2

Figure 5.15 shows the analysis for the “AND 3 with 1 Ridge-bit and $S = 5$. In this case, there is a marked region where the RCN can solve the task with a relatively low BER of 10^{-2} . This result is confirmed by the RB map, which shows for the same area an area of RB greater than unity, whereas the power map reports optical powers mostly between 13 dBm and 16 dBm, i.e. in the nonlinear regime. Outside of this sector, the performance is definitely not noteworthy (around 10^{-1}) and roughly on par with that evaluated on the input, if not worse. At the same time, the power map shows predominantly high values (16 to 18 dBm), except for a small region at 250 Mbps where the power drops values proper of the linear regime (8 dBm). In this confined area, the system nonlinearities seem detrimental to the solution of the task.

The best value is obtained at a bitrate of 100 Mbps and a detuning of -20 GHz and an optical power of 15 dBm, with the value of RB approximately $10^{1.5}$. This suggests a link to the nonlinear effects related to the FC dynamics,

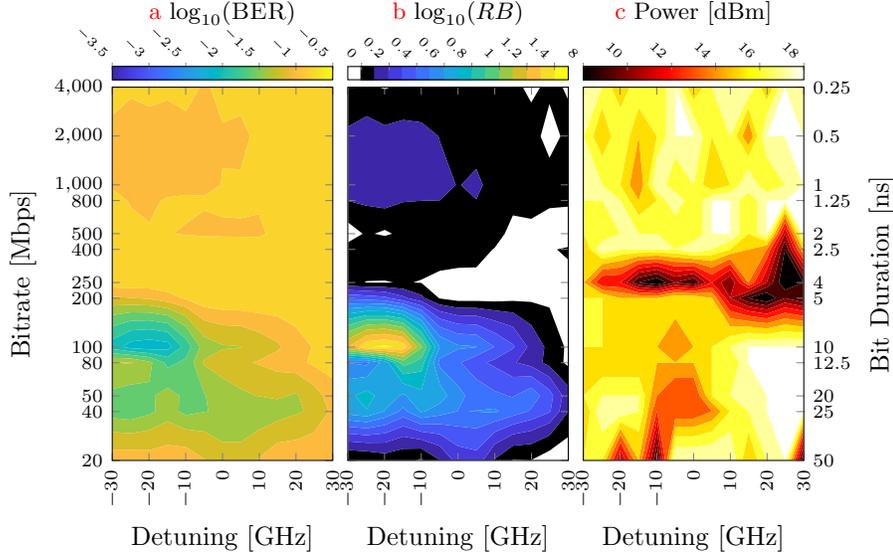


Figure 5.15: AND 2 with 1 Ridge-bit and $S = 5$. Comparison of: **a** BER map, **b** RB map, and **c** power map.

whose lifetime is about a few ns. From these results, a fading memory of two bits due to the MRR dynamics is observed.

AND 3

The “AND 3 with 1 Ridge-bit and $S = 5$ ” task reaches an absolute minimum BER of 10^{-1} at a bitrate of 100 Mbps and -20 GHz of detuning. Clearly, the MRR is not able to supply the fading memory sufficient to carry out this 3-bit delay task.

Nonlinear operation, logical XOR

The delayed XOR operation requires the MRR dynamics to provide a nonlinear transformation of the input when supplying both input bits to the ridge regression. Additionally, when less data is given, MRR dynamics must supply the fading memory too.

XOR 1 with 1 Ridge-bit

The “XOR 1 with 1 Ridge-bit” task require both nonlinearity and memory. However, since the two input bits are adjacent, memory may be already present in the input signal in the form of inter-symbolic interference. The RB map should be able to clarify this aspect.

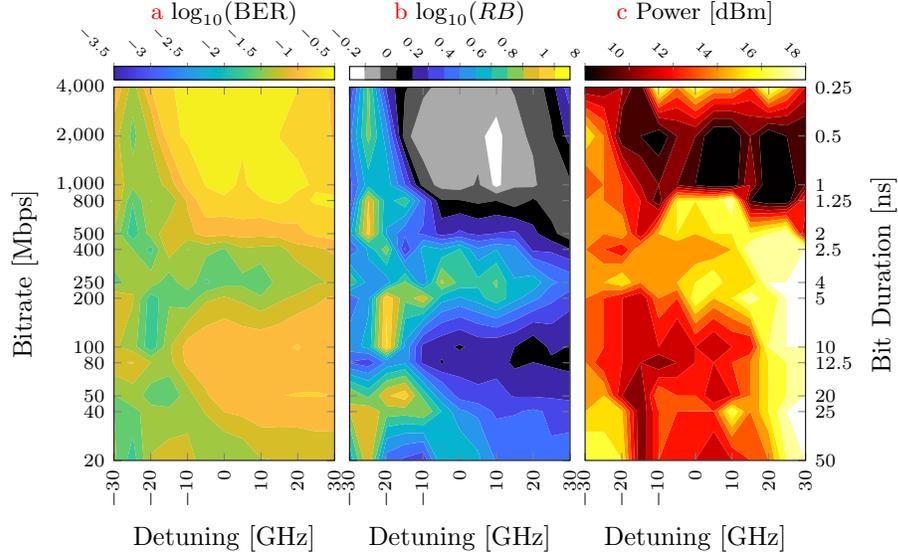


Figure 5.16: XOR 1 with 1 Ridge-bits and $S = 5$. Comparison of: **a** BER map, **b** RB map, and **c** power map.

Figure 5.16 contains the results of the analysis. This task shows a mediocre performance, with the absolute minimum value of BER of $10^{-1.7}$. The lower values of the BER are obtained either at all bitrates and negative detunings, roughly from -15 GHz to -30 GHz, or at bitrates of 200 and 250 Mbps and most detunings. There are two regions, one at detunings greater than -10 GHz and bitrates smaller than 200 Mbps and another one at detunings greater than -20 GHz and bitrates higher than 800 Mbps, which reports BER larger than 10^{-1} . Similarly, also the RB map separates the detuning-bitrate space in the same three regions, with the highest performance improvement just shy of $10^{1.1}$ and worst degradation around $10^{-0.3}$. The power map does not add much information, except by showing a steady increase in power from negative detuning and low bitrates to positive detuning and bitrates around 250 to 500 Mbps. Above 800 Mbps and for detuning larger than -20 GHz, there is a sharp decrease in the best incident optical power.

In conclusion, the system is not able to provide good results when both the fading memory and the nonlinear transformation are required at the same time. However, in all three maps, there is a clear preference for negative detunings, reporting lower BER, lower incident power, and higher performance improvement in respect to the training on the input signal.

XOR 1 with 2 Ridge-bits

The same task is repeated, this time supplying both input bits to the ridge

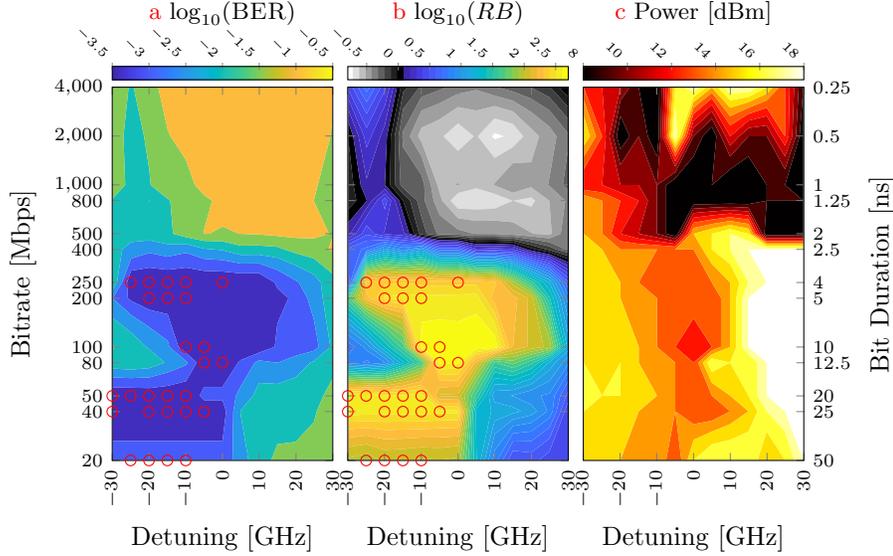


Figure 5.17: XOR 1 with 2 Ridge-bits and $S = 5$. Comparison of: **a** BER map, **b** RB map, and **c** power map. Markers identify the combinations that achieved the statistical limit (error-free operation): circles for the input signal and crosses for the output signal.

regression. The burden of providing the memory is then removed from the MRR dynamics. In this case, much better results are expected in comparison to the previous task. Indeed, the results shown in fig. 5.17 confirm this expectation.

The performance analysis divides the detuning-bitrate space essentially into two distinct regions: from 20 Mbps to 400 Mbps and from 500 Mbps 4000 Mbps. In the first region, the performance reported is optimal, as a large number of combinations reach low values of BER and several achieve even the statistical limit at a BER of $10^{-3.5}$. The same combinations are those which exhibit the highest performance improvement in the RB map, with a peak ratio of $10^{2.7}$. The power map in this region describes a situation in which lower power is required for detuning close to 0 and increases with larger (in absolute value) detunings. In particular, positive detunings on the edge of the map (30 GHz) are associated with twice the incident optical power in respect to those on the other side.

On the other hand, in the second region at high bitrates, the values of the BER achieved are only 10^{-1} to 10^{-2} . Moreover, in this region, the RB map confirms that the performance of the system without the MRR, i.e. evaluated on the input signal, is better than that evaluated on the output signal. Conversely, the incident optical power lean with few exceptions toward the lowest values (8 to 11 dBm). It is clear that in this area the system is able to exploit the

distortions associated with the generation and detection of the optical signal. Nevertheless, no combinations enable the input signal to achieve the statistical limit for this task.

In summary, by providing both input bits to the ridge regression, the system can correctly carry out the XOR logical operation for bitrates below 400 Mbps. Optimal performance is obtained between 40 to 250 Mbps with a preference for negative detunings.

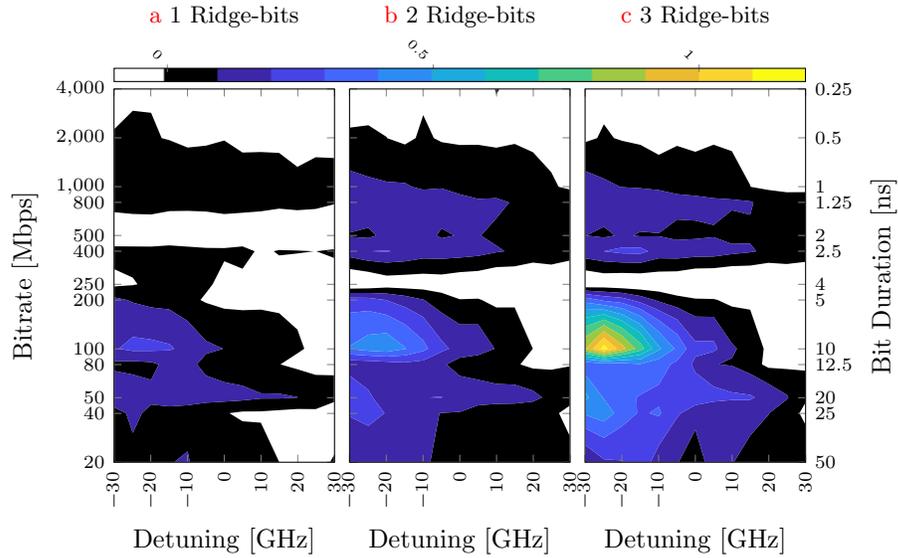


Figure 5.18: XOR 2 with 1, 2, and 3 Ridge-bits and $S = 5$. Comparison of the RB maps for **a** 1 Ridge-bit, **b** 2 Ridge-bits, and **c** 3 Ridge-bits.

XOR 2 with 1, 2, and 3 Ridge-bits

As a final example, I report here the results for the “XOR 2” task obtained by giving 1, 2, and 3 Ridge-bits to the ridge regression algorithm. The expected result is, of course, a performance improvement by the subsequent addition of Ridge-bits due to the associated decrease in the memory requirement. Figure 5.18 shows the RB maps for the three variations of the XOR 2 task.

Providing only the current bit to the ridge regression algorithm, the performance is poor: the best RB is $10^{0.2}$ at a bitrate of 100 Mbps and a detuning of -25 GHz. A small surrounding region shows a negligible improvement, but the remaining combinations of bitrate and detuning show only a detrimental effect of the MRR. The picture does not change much when 2 Ridge-bits are provided. In fact, the maximum RB is $10^{0.4}$ for a BER of $10^{-0.9}$, again at 100 Mbps and -25 GHz. The area of improvements is slightly broader and another area of minimal improvement arises for bitrates between 400 and 1000 Mbps. Finally,

supplying 3 Ridge-bits to the readout, both the best RB and BER improve drastically, even though for an isolated combination of parameters, achieving respectively $10^{1.1}$ and $10^{-2.6}$. The areas in which the presence of the MRR is beneficial for the performance remain almost the same, although they see an increase in the RB values, as is observed by brighter colours in the RB map. The results have been summarised in table 5.3.

XOR 2	1 Ridge-bit	1 Ridge-bits	1 Ridge-bits
$\min(BER)$	$10^{-0.7}$	$10^{-0.9}$	$10^{-2.6}$
RB	$10^{0.2}$	$10^{0.4}$	$10^{1.1}$

Table 5.3: Best value of the RB and corresponding BER, evaluated at a bitrate of 100 Mbps, a detuning of -25 GHz, and optical power of 16 dBm.

5.3.4 Discussion

This is an exploratory study of the performance of a single MRR in the ADF configuration. The MRR is used as a reservoir in which the field enhancement proper of the optical cavity is used to amplify nonlinear optical effects, mainly associated with FC, and to change the dynamics of the reservoir. Its ability to provide either fading memory, a nonlinear transformation of the input signal, or both, is tested with linear and nonlinear tasks for a considerable number of combinations of bitrate, frequency detuning, and incident power of the input optical signal. An offline readout layer exploits the MRR’s dynamics to execute a given function.

The logical task AND, carried out between bits at several delays, is used to test the fading memory property. “AND 1 with 2 Ridge-bits” is used to verify the whole system, as the presence of the MRR, in this case, is redundant. Indeed even the readout layer alone is able to complete the task and the MRR adds at most some noise. “AND 1 with 1 Ridge-bit” and “AND 2 with 1 Ridge-bit” show that the MRR reservoir is able to supply the readout layer with a fading memory of one bit for most combinations of the input parameters for bitrate below 500 Mbps and with two bits memory for a smaller set of parameters. Tasks with delay 3 or longer show that the MRR is not able to provide sufficient memory to the readout. The RB maps confirm that this fading memory comes in fact from the presence of the MRR and is caused by the distortion already present in the input.

The second kind of task studied is that of the logical XOR operation, whose outputs are not linearly separable. For this reason, this function requires the MRR to provide both memory and nonlinearity. “XOR 1 with 1 Ridge-bit” shows that the reservoir is not able to provide both. “XOR 1 with 2 Ridge-bits”, on the other hand, shows that if external memory is given to the readout layer, then the dynamics of the MRR successfully transforms the optical input and the nonlinear task can be solved. In this case, several combinations with bitrates

up to 250 Mbps and negative detunings achieve error-free operation (statistical limit of $10^{-3.4}$). Furthermore, the results for the XOR 2 operation detail a similar situation, in which external memory greatly improve the performance of the system. Also in this case, the *RB* maps verify that these results are due to the presence of the MRR.

Considering both the linear and the nonlinear tasks, a certain trend can be observed. Most combinations of input parameters showing good performance have bitrates from 100 to 400 Mbps. Therefore, it seems that the system is harnessing the nonlinear FC effects, whose timescale is close to those bitrates: the inverse of the FC lifetime $1/\tau_{FC} \simeq 200$ to 500 MHz. A similar situation is reported in [92], where the much longer FC lifetime $\tau_{FC} = 45$ ns gives rise to the peak performance nearby 20 Mbps. Additionally, the timescale of the thermal effects can be estimated to be around $1/\tau_{TH} \simeq 10$ to 20 MHz and therefore falls outside the observed bitrate range. Similarly, also the photon lifetime does not affect the system performance, because the charge and discharge effect of the optical cavity occurs at a frequency of $\simeq 200$ GHz.

In conclusion, this study highlights the importance of isolating the effect of the reservoir on the performance of the system. This is especially true for those approaches which rely on super-sampling and offline analysis. Indeed, in this case, the input optical signal could contain distortions, usually due to the finite bandwidth of electronic and electro-optic components in the signal generation system, which can be exploited by the detection and readout stages. The most evident is the inter-symbol interference, which creates a smoothed transient between contiguous bits. Consequently, approaches based on virtual nodes could benefit from the analysis of the optical input as well as the optical output signals, so that the *RB* values can be estimated. Only this can provide concrete proof of the reservoir performance, which otherwise could be mistaken with the performance of the whole experimental apparatus.

Chapter 6

Conclusion

In this thesis, I investigated the use of the silicon photonics platform for the development of novel neuromorphic computing architectures. In particular, I investigated the use of a network of MRRs to create an all-optical reservoir.

During the thesis, I became familiar with several aspects of the research topic. Initially, I developed a flexible computational library able to reliably simulate arbitrary PICs. The result is PRECISE [100], an open-source Matlab library able to describe PICs using a modular approach, decomposing the geometry in waveguide segments, and taking into account several nonlinear effects, in particular FC and thermal phenomena. Additionally, PRECISE effective propagation model implements a few approximations to decrease the computational burden so that an extensive analysis of the parameter space is easily obtained.

Moreover, I also participated in the installation of the new experimental setup, specifically designed to characterise PICs with optical signals modulated at high frequency. Toward this goal, a comprehensive set of routines has been developed to enable the remote operation of the instruments in the setup. This allowed measurements with high stability and repeatability. Likewise, I also took part in the design of the PICs that were subsequently fabricated, focusing on the SCISSOR-based RCN and the complex perceptron schemes. I was later able to study and measure the response of the manufactured devices. Unfortunately, the samples showed high propagation losses and a long lifetime associated with the FC effects. The deviation of these two properties from the standard values resulted in the degradation of the SCISSOR reservoir performance and in the shift of the best rate of operation toward longer timescales. This, in conjunction with a technical limitation of the experimental setup at low modulation frequencies, delayed the study of the structure. Nevertheless, despite the high propagation losses, the analysis on the complex perceptron was accomplished and showed promising results for a fast, yet simple all-optical network. The device was tested against tasks such as pattern recognition, XOR logical operation, and phase decoding and it was able to achieve error-free operation at bitrates as high as 16 Gbps. Finally, I investigated the possibility to use a single MRR as a dynamic nonlinear system to perform linear and nonlinear operations and to

study the possibility to store information within an integrated optical resonator.

In conclusion, in this work, I illustrated in particular two aspects: the modelling of networks of MRRs, whose inherent complexity is enhanced by the presence of nonlinear optical effects, and the experimental study of PICs with high-frequency signals. I was able to develop a flexible and powerful simulation tool and I carried out the experimental measurement of several devices. The simulation tool is a solid starting point, which can be certainly improved and adapted to the needs of modelling the future RCNs. Moreover, the knowledge gained from the study of the early devices allowed us to design new structures, which have been recently manufactured and will therefore be soon characterised. I believe that the results achieved are a promising first step in the direction of the development of novel RCNs that employ the full potential of all-optical devices.

Publication List

In chronological order

[75] Massimo Borghi, Davide Bazzanella, Mattia Mancinelli, and Lorenzo Pavesi. “On the modeling of thermal and free carrier nonlinearities in silicon-on-insulator microring resonators”. In: *Optics Express* 29.3 (2021), pp. 4363–4377. DOI: 10.1364/oe.413572

[1] Davide Bazzanella, Giorgia Bincoletto, Monica Consolandi, Marta Fasan, Francesca Gennari, Federico Carmelo La Vattiata, Luca Rinaldi, David Roccaro, and Clara Zaccaria. “Artificial and Biological Neurons: Interdisciplinary Issues and Future Perspectives. White Paper”. In: *BioLaw Journal - Rivista di BioDiritto* 1 (Mar. 2021), pp. 353–379. ISSN: 2284-4503. DOI: 10.15168/2284-4503-756

[90] Mattia Mancinelli, Davide Bazzanella, Paolo Bettotti, and Lorenzo Pavesi. “A photonic complex perceptron for ultrafast data processing”. In: *Scientific Reports* 12.1 (Mar. 2022), p. 4216. ISSN: 2045-2322. DOI: 10.1038/s41598-022-08087-2

[100] Davide Bazzanella, Mattia Mancinelli, Massimo Borghi, Paolo Bettotti, and Lorenzo Pavesi. “PRECISE Photonic Hybrid Electromagnetic Solver”. In: *IEEE Photonics Journal* 14.3 (2022), pp. 1–10. DOI: 10.1109/JPHOT.2022.3168157

[95] Davide Bazzanella, Stefano Biasi, Mattia Mancinelli, and Lorenzo Pavesi. “A Microring as a Reservoir Computing Node: Memory/Nonlinear Tasks and Effect of Input Non-ideality”. In: *arXiv preprint arXiv:2203.07045* (2022)

Appendix A

Experimental Setup

I have devoted a large portion of my efforts to the design and assembly of the laboratory's experimental setup. The main scope of the setup is to measure passive photonic integrated circuits (PICs) and to characterise their ability to elaborate information encoded in the optical input. Specifically, the setup can generate signals modulated in both amplitude and phase up to 20 GHz and record the optical response of the device under test (DUT).

Since many experiments are carried out close to transition thresholds of nonlinear optical phenomena, such as optical bistability and self-pulsing, the setup should provide high accuracy in setting the system state. In fact, even small noise in the system DOFs could trigger the transition from a stable state to another.

Another important property of the experimental setup is the ability to explore wide parameter spaces with consistent results. For example, analysing the response of the DUT for a set of optical power and optical frequencies of the input signal. The number of single measurements required to cover the parameter space grows as the number of parameters increases. As a result, the total measurement time will be longer. In order for each measurement to be comparable to the others, the setup must provide high stability in time.

The components of the experimental setup can be grouped in:

- **signal generation stage**, where the input signal is generated with a specific optical frequency, average power, polarisation, and modulation;
- **DUT coupling and control stage**, in which input/output optical fibres are coupled to the DUT edge or grating couplers and integrated electrically active elements such as thermal heaters are contacted and controlled;
- **signal detection stage**, where the signal collected from the output of the DUT is amplified and filtered in the optical domain, before being detected by a photodetector, and sampled by the oscilloscope.

The schematic of the experimental setup is shown in fig. A.1 and its components are listed in table A.1.

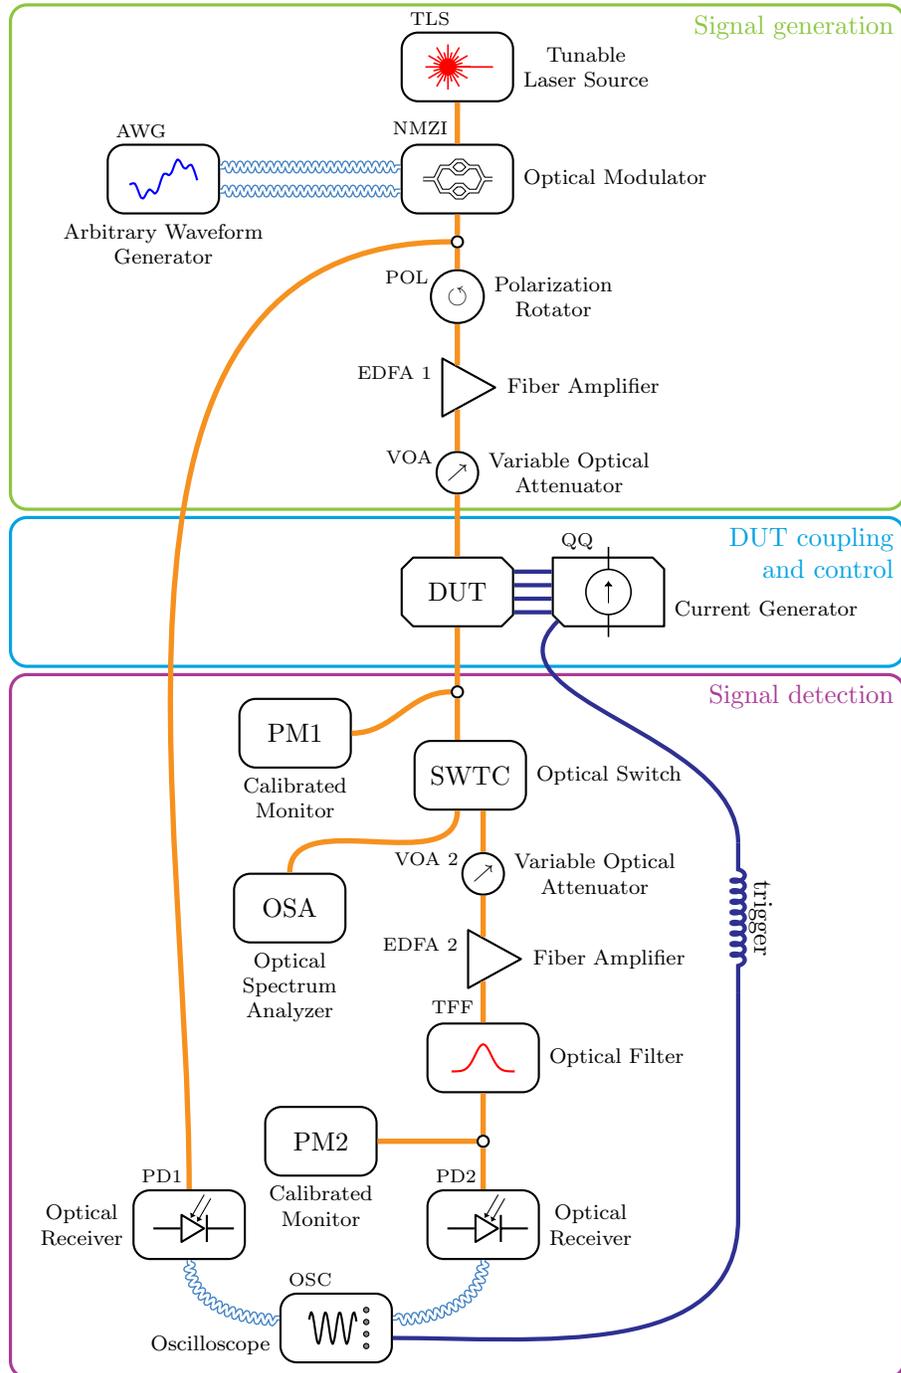


Figure A.1: Experimental setup.

Table A.1: List of instruments and components in the experimental setup.

Label	Model and Description
TLS, TLS2	PurePhotonics PPCL200 Tunable laser source (TLS) working between 191.5 and 196.25 THz and with an output power from 6 to 17 dBm.
NMZI	iXblue MXIQ-LN-30 Lithium niobate EOM with analog bandwidth of 30 GHz working between 1530 and 1580 nm. It is an IQ modulator, which modulates both intensity and phase thanks to the nested Mach-Zehnder interferometer (NMZI) configuration.
AN1, AN2	iXblue DR-AN-28-MO Wideband radio frequency (RF) amplifier characterised by a low noise figure and a linear transfer function.
AWG	Keysight M8195A Two channel arbitrary waveform generator (AWG) with a sample rate of 65 GSa/s, an analog bandwidth of 25 GHz, and 8 bit vertical resolution.
POL	Newport F-POL-APC Babinet-Soleil compensator for polarization control based on an all-fibre design. Enable manual control of the output polarisation state
EDFA1	IPG Photonics EAD-1K-C High power erbium-doped fibre amplifier (EDFA) providing a maximum output power of 30 dBm and a typical noise figure of 5 dB. Works in current control, power control, and gain control.
LS1, LS2	Thorlabs Nanomax MAX312D/M Precision 3-axis linear stage with sub-micron resolution and 4 mm travel.
QQ	Qontrol Q8iv modules and BP8 backplane 3 × Q8iv modules act as a 24 channel current generator.
VOA1	Viavi mVOA-C1 Calibrated variable optical attenuator (VOA) with high input power capability, high accuracy and repeatability, and low insertion loss.
PM1, PM2	Viavi mOPM-C1 Calibrated photodetector with 110 dB dynamic range and 13 mW maximum input power.
SWTC	Viavi mOSW-C1 Calibrated optical switch providing a 2 × 2 configuration.
OSA	Yokogawa AQ6360 Optical spectrum analyser (OSA) for the IR band in the range 1200 to 1650 nm with an adjustable wavelength resolution between 0.1 and 2 nm and a measurement range from −80 to 20 dBm.
EDFA2	Thorlabs EDFA100S Erbium-doped fibre amplifier (EDFA) providing small signal gains greater than 30 dB, noise figure smaller than 5 dB, and polarisation-dependent gain smaller than 0.2 dB.

Continue on the next page

Label	Model and Description
VOA2	Thorlabs V1550 MEMS-based voltage controlled variable optical attenuator (VOA) with attenuation greater than 25 dB, single mode (SM) fibre, maximum optical input power of 100 mW, and modulation up to 1 kHz.
TFF	OZ-Optics TF-100-11 Motor driven tunable optical pass-band filter with 0.3 nm 3 dB line width. Tunable range is from 1470 to 1620 nm.
PD1, PD2	Thorlabs DXM20AF Ultrafast fibre optic InGaAs photodetector with 20 GHz bandwidth and a responsivity of 0.90 A/W, giving a conversion gain of 22.5 V/W. Its NEP is 28 pW/ $\sqrt{\text{Hz}}$.
OSC	Teledyne LeCroy SDA816Zi-A Digital oscilloscope with 16 GHz analog bandwidth and up to 80 GSa/s in either 4×20 GSa/s channel or 2×40 GSa/s channel configurations.

A.1 Signal generation

The output of the signal generation stage is an optical signal with a constant optical frequency, polarisation, and average power, as well as time-varying amplitude and phase modulations. To maximise stability and accuracy, while minimising noise and preserving flexibility, each property of the optical signal is generated with a specific instrument.

The tunable laser source (TLS) produces the carrier signal with an optical frequency in the range from 191.5 to 196.25 THz, with a constant power of at least 10 dBm. Its polarisation is linear and is maintained through polarisation maintaining (PM) optical fibres. Subsequently, the signal is modulated both in amplitude and phase through a nested Mach-Zehnder interferometer (NMZI), whose two electrical driving signals are generated by a dual-channel arbitrary waveform generator (AWG) and amplified by two RF amplifiers (AN1 and AN2). When the modulation is only on the amplitude, the two electric driving signals are identical. At this point, the signal optical frequency and its modulation have been set. A fibre splitter (50/50) is used to probe the generated optical signal with an ultrafast photodetector (PD1).

A polarisation controller (POL) compensates against the polarisation rotation given by the optical fibre sections that follow, in order to reach the DUT with a TE polarization. The last two elements of this stage, an erbium-doped fibre amplifier (EDFA1) and a calibrated variable optical attenuator (VOA1) provide the signal with an accurate and stable average optical power. EDFA1 amplifies the signal up to the maximum average optical power, which in our case is usually 20 dBm, but it can be set as high as 30 dBm. The amplifier operates in the automatic power control mode, producing a stable average optical power, thanks to the feedback of an internal photodiode. Finally, VOA1 attenuates with high accuracy the power to the value required by the measurement.

A.2 DUT coupling and control

The generated signal is coupled to the DUT, where mechanical stages align the input/output SM fibres or fibre array to the DUT. In our case, the typical configuration comprises two 3-axis (x, y, z) linear stages with open-loop piezoelectric actuators (LS1 and LS2) on top of which custom mounts hold the input/output optical fibres at a specific angle (12°) from the normal to the DUT surface. Figure A.2 shows a sketch of the position of input/output optical fibres and the DUT and their movements and rotations.

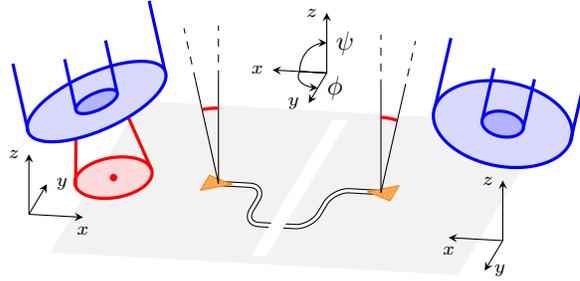


Figure A.2: Sketch of the position of the DUT and of the input/output optical fibres. Vectors show the movements and rotations of the mechanical stages, bundled as axes origins.

The DUT is either a naked die or a packaged chip and is placed on top of a pillar consisting of mechanical stages and the thermo-electric cooler (TEC) system. The mechanical stages are two: a 1-axis (y) linear stage and a 5-axis (x, y, z , pitch, and yaw) stage. Their movement is coarser than those holding the fibres and are used to facilitate the handling of the DUT rather than to align the fibres to the sample. The DUT temperature is stabilised by a Peltier cell, driven by a proportional-integral-derivative (PID) controller (PID). The Peltier cell dissipates or collects heat through an aluminium heatsink and is in contact with the DUT through a copper mass, which houses a $10\text{ k}\Omega$ thermistor used as feedback by the PID controller. Figure A.3a shows the exploded view of the TEC system, while fig. A.3b shows the thermal stability of the PID readings.

The TEC system is built to accommodate a printed circuit board (PCB) with a socket for the packaged version of the DUT. The PCB was designed in conjunction with a breakout board for the multi-channel current source generator (QQ). The current source generator has a modular structure that is able to feed 8, 16, or 24 channels independently. Each channel is driven as a constant current source so that changes in the resistance due to the common ground are automatically taken care of. The breakout board allows connecting the current channels to the pin of the packaged PIC in any combination, through the use of jumper wires.

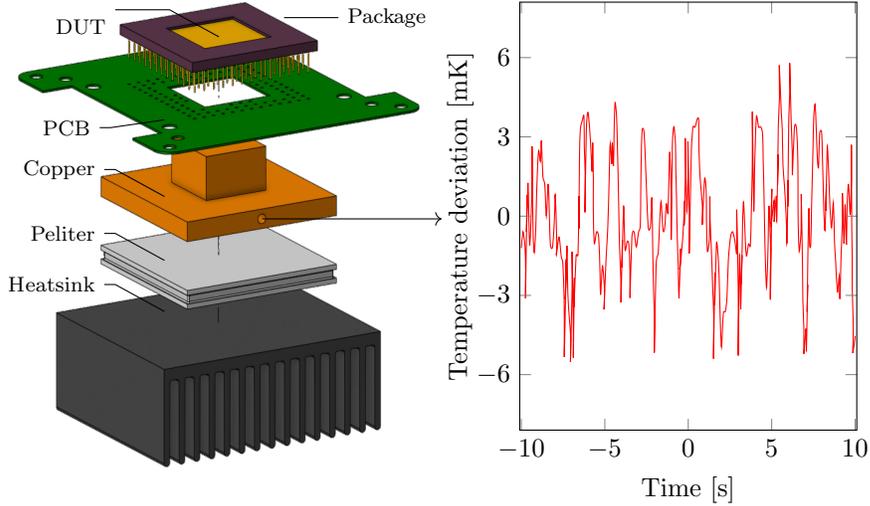


Figure A.3: **a** Exploded view of the TEC system with PCB for the packaged die. **b** temperature fluctuations reported by the PID controller temperature sensor.

A.3 Signal detection

A fibre splitter (99.9/0.1) is used to probe the output signal with one of the two calibrated photodetectors (PM1), obtaining an accurate measurement of its average optical power. This allows to evaluate the insertion losses of the DUT and thus the actual optical power travelling in the input WG of the structure.

The next element of the setup is the optical switch (SWTC), which directs the signal either to the optical spectrum analyser (OSA), by default, or to the second ultrafast photodetector (PD2). The OSA is a precision instrument designed to carry out spectral measurement of optical signals. It has high sensitivity, but it is relatively slow, taking from a portion of a second to several seconds to complete a measurement scan, depending on the instrument settings. It is especially useful in order to find the approximate position of microring resonances (see appendix A.4.1).

The signal collected from the DUT output ports has usually low optical power, due to the non-negligible coupling losses, which results in poor SNR for direct detection with a photodetector. In our case, an erbium-doped fibre amplifier (EDFA2) operating with a constant pump current amplifies the output optical signal. As a rule of thumb, electrical amplification is cheaper and more practical for sub-GHz signals, whereas for ultrafast signals optical amplification becomes more effective because it is not limited to the bandwidth of electronics components. EDFA2 is preceded by a calibrated variable optical attenuator (VOA2), whose task is to reduce the optical power at the input of EDFA2 so that its working point remains approximately constant. The spectral noise introduced by optical amplification is filtered out by a pass-band tun-

able optical filter (TFF). The average optical power is accurately measured by the second of the two calibrated photodetectors (PM2), through another fibre splitter (99.9/0.1). The resulting signal is hence amplified with precision up to 0–3 dBm, resulting in a good SNR within the linear range of the detector.

A multi-channel oscilloscope (OSC) samples both PD1 (input signal) and PD2 (output signal) with up to 80 GSa/s as well an external trigger signal coming from the current source generator. The analog bandwidth of the oscilloscope is 16 GHz.

A.4 Experimental setup operation

Considering the primary objective of the experimental setup, i.e. measurement accuracy and exploration of the system DOFs, the obvious approach is remote operation. Indeed, it enables the reliable control of the settings for each instrument, the synchronous operation of multiple instruments, and the automation of serial procedures. All these properties, in turn, improve the measurement speed, hence the stability in time, and its repeatability, reducing the probability of human error.

The experimental setup is controlled by a single computer (PC) and the instruments are connected to it with a variety of interfaces: mostly TCP/IP, USB, and RS-232, but seldom also the “dreaded” GPIB. Many of them employ the SCPI communication protocol, in which commands are human-readable strings of text defined by a common syntax, while some of them offer C or python libraries. Furthermore, custom instruments can be assembled with micro-controllers, e.g. Arduino based boards, and sensors and actuators. In this case, also the low-level programming, implementing the logic of the instrument, must be taken care of as well as the communication part.

On the PC, a Matlab environment is used to communicate with the instruments. However, the number of commands associated with a single instrument can easily be in the hundreds. For this reason, an object-oriented approach allows collecting all the properties and functions of an instrument in a *class*. This description allows creating instances of these classes, representing each a single instrument. Then their properties and functions are easily accessed with the dot notation, e.g. `instance.property` to access the property value and `instance.method(arguments)` to call a given function, or *method*, with a set of arguments. There are at least three big advantages of this approach that improve how the code is managed and used: first, being based on plain text, versioning can be easily implemented, e.g. with git, facilitating the debug and review of the code; second, many components of the code can be reused, minimising both effort and errors; third, the commands can be structured hierarchically, where the functions associated to a class are collected in a single place, simplifying both its use and development.

A versioned library aggregates the *classes* that I developed for all the instruments, enabling easy access to their functionalities. Ideally, each class has the same set of core functions and properties, although there are some exceptions.

There are two standard properties, `con` which contains the instrument interface object, i.e. the connection, such as Matlab's `serialport` object, and `verbose`, which is a Boolean flag used to enable or disable writing to the command prompt what each command does (useful during debugging). Furthermore, there are usually three standard methods: `send`, `listen`, and `query`, which implements sending a command, receiving a command, and the two in sequence. Additional commands are `enable` and `disable`, common for active devices such as the TLSs or the AWG. In the case of python and C libraries, it is usually convenient to define a Matlab *class* wrapping the original library.

Finally, an additional class groups all the instruments of the setup and provides a set of functionalities involving two or more instruments together, e.g. a transmission spectrum can be obtained by operating the TLS and the PM1 concurrently. Scripts and functions can easily access all the instruments and their functions by addressing this class. The setup is usually started with a sequence of four steps:

1. `init` carries out the instrument initialisation, instantiating all the classes inside the system class and adjusting the parameters of the connections when needed.
2. `pre_alignment` activates the signal generation stage, starting up its instruments with the correct sequence, typically producing a continuous wave (CW) input signal at a frequency around 193 THz and 0 dBm of incident power. A shutter within the VOA1 blocks the signal from entering the DUT. Normally the NMZI is calibrated during this phase.
3. manual alignment of the input/output optical fibres on the surface of the DUT. Their position is manually changed with the three micrometer actuators (thumbscrews), by looking at a live feed of the PM1 reading.
4. `post_alignment` carries out operations that require the DUT to be aligned, such as the spectral analysis or the alignment of the TFF

At this point, the setup is ready to carry out the measurements.

A.4.1 Resonance position

The spectral response of microring resonators with radii of about 10 μm usually shows just a few resonances in the frequency range covered by the TLS, i.e. between 191.5 and 196.25 THz. It is therefore important to know where these resonances appear. The experimental setup allows for two measurement procedures: the first one measures simultaneously the coarse position of all resonances over a broad range of frequencies, while the second is able to accurately measure the position and shape of a resonance on a narrow range of frequencies.

The coarse position measurement is carried out by using the OSA to acquire the transmission of the microring resonator when stimulated with a broad light source. In principle, the resonance positions can be isolated by comparing the spectra of the light transmitted through the DUT to that of the broad light

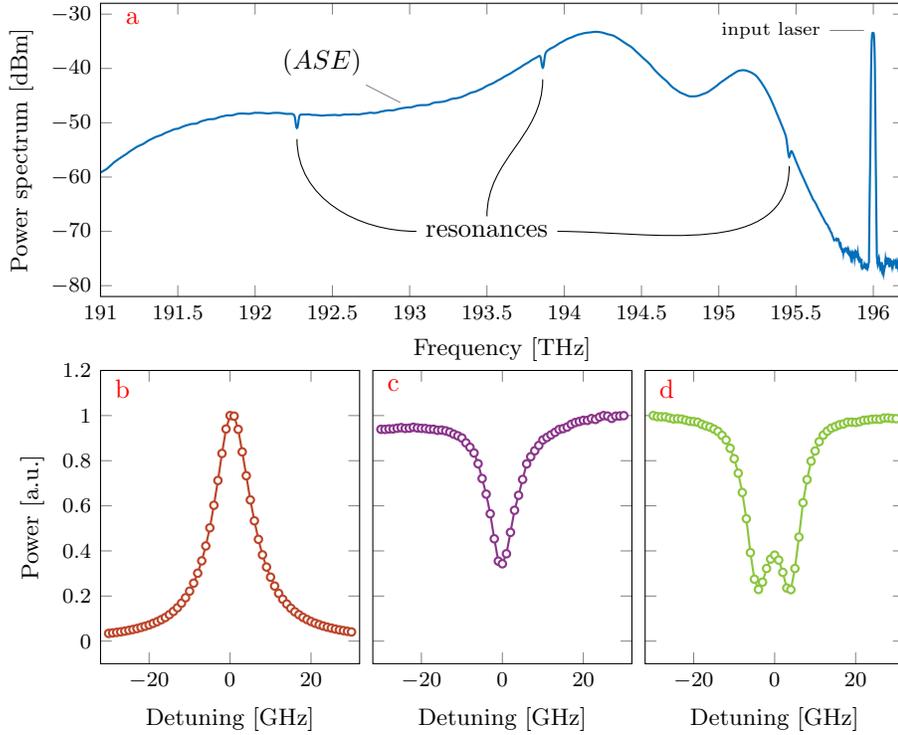


Figure A.4: Measurement of the resonance of a microring resonator. **a** Coarse measurement of the positions of all resonances within the broad ASE spectrum of the EDFA1. **b**, **c**, and **d** Fine measurements of the spectral response around a given centre frequency. The first is measured at the *drop* port, while the latter two are measured at the *through* port. Both the frequency position and the shape of the resonance are accurately measured.

source alone. However, since the spectrum of the broad light source is usually smoothly shaped, the measurement of the transmitted light is sufficient in order to extract the resonance position with simple assumptions on the transmission shape, i.e. the transmission in resonance is a maximum or a minimum. Figure A.4a shows the spectrum of the transmitted light, highlighting the spectrum of the source and position of resonances.

In order to obtain a broad light source, the TLS and the NMZI are configured to produce a CW signal at 196 THz or higher, which is then fed to the EDFA1. Since this frequency is outside the EDFA1 amplification range, amplified spontaneous emission (ASE) occurs instead. In ASE, the input signal is amplified very little while spontaneous emission of the gain medium is amplified more. Given that the EDFA1 is maintaining the average output optical power constant to 20 dBm, this power is distributed more or less evenly on the whole amplification

range.

The narrow, accurate measurement is carried out by exploiting the fine-tuning mechanisms allowed by the TLS itself. In fact, the TLS is able to tune its frequency between -30 GHz and 30 GHz from a given centre frequency, with a resolution of 0.1 GHz. The centre frequency of the TLS is set to the approximate position of a certain resonance, e.g. one found with the coarse position measurement. By fine-tuning the laser frequency around the centre and measuring the corresponding transmitted intensity with PD1, an accurate spectrum is obtained. The smaller the steps in frequency, the higher the resolution with which the transmission spectrum is defined. A few examples of the resulting spectrum are shown in figs. A.4b to A.4d.

A.4.2 IR camera aided characterisation

Scattering of light from integrated structures outside the plane of the chip can be collected by an IR camera observing the chip from the top. Hence, it is possible to observe the optical response of a device by observing its scattered light.

Usually, light scattered from a specific structure is somewhat proportional to the light circulating within. For example, in the case of a microring resonator, it is possible to measure the shape of its resonances as shown in Figure A.5c: by integrating over the pixels inside the region of interest (ROI), a quantitative estimation of the light circulating within the ring can be obtained. Figure A.5b shows the spectra measured by repeating the integration of the ROIs for different frequency detunings of the TLS. Similarly, a fixed input frequency can be used to characterise the resonance shape as a function of the heater current.

Bear in mind that this method requires a clear line of sight between the structure and the camera. Therefore any obstruction above the chip, such as the fibre holders and the mechanical stages, or on the chip, such as metal tracks and dust, could make the measure more difficult, if not impossible altogether.

A.4.3 Thermal cross-talk

Integrated devices featuring thermal heaters on top of WGs may show thermal cross-talk. Cross-talk happens when a controller influences a system DOF which should be exclusively managed by another controller. Ideally, heaters change the temperature only of the WG directly below them, however, they might also change the temperature of nearby WGs.

In order to avoid thermal cross-talk, several options are at hand. The first and most obvious one is to design a device without thermal cross-talk, i.e. heating is localised to the WG underneath the heater. This is achieved either by placing the WGs distant enough from each other or by reducing heat diffusion in the volume between nearby WGs, e.g. with thermal isolation trenches. This approach is limited by design constraints, such as trivially the total design area available. The second option is to accurately model the effect of each heater of the DOFs of the system and to invert this relation. However, isolating the

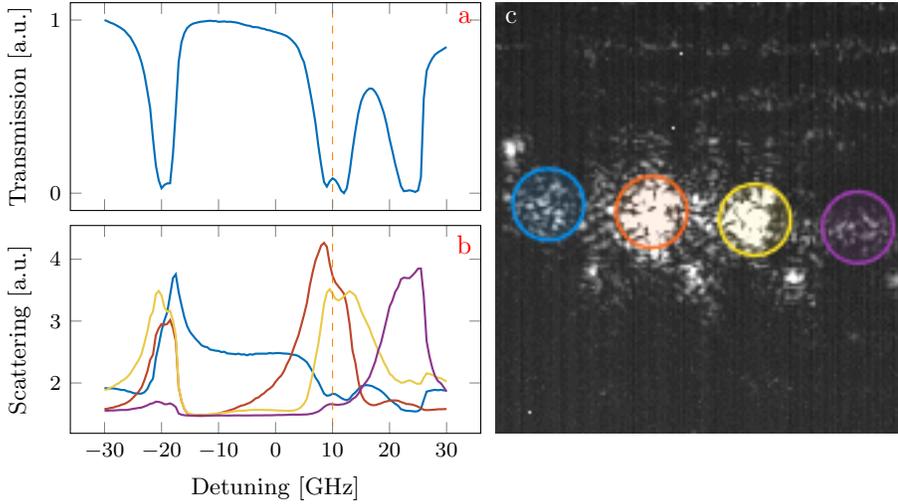


Figure A.5: **a** Transmission measured at the output grating of the device. **b** Scattering spectra for each ROI. The vertical dashed lines indicate the detuning at which the IR image has been acquired. **c** Camera picture with multiple ROIs defined.

effect of every actuator on each DOF may prove quite challenging, as described later in this Section. The last technique is to work at timescales in which the local thermalisation has reached a steady state, but the heat diffusion has still to reach nearby WGs, causing cross-talk. Although this strategy is quite simple and effective, the drawback is the impossibility of operating the device in a continuous fashion.

Cross-talk matrix

Here I describe the results obtained trying to correct the thermal cross-talk effect by applying the inverse operation that mathematically describes it. For example, the phase change induced by a thermal heater on an optical signal is proportional to the power dissipated by the heater and hence to the square of the electrical current flowing, i.e. $\Delta\phi \propto P \propto i^2$. Assuming there are a total of four thermal heaters in the device, we can initially approximate the cross-talk linearly:

$$\begin{pmatrix} \phi_1 \\ \phi_2 \\ \phi_3 \\ \phi_4 \end{pmatrix} = \mathbf{M} \begin{pmatrix} i_1^2 \\ i_2^2 \\ i_3^2 \\ i_4^2 \end{pmatrix}, \quad \text{where } \mathbf{M} = \begin{bmatrix} m_{11} & m_{12} & m_{13} & m_{14} \\ m_{21} & m_{22} & m_{23} & m_{24} \\ m_{31} & m_{32} & m_{33} & m_{34} \\ m_{41} & m_{42} & m_{43} & m_{44} \end{bmatrix}. \quad (\text{A.1})$$

Depending on the system symmetries, dependence between the matrix \mathbf{M} coefficients can be assumed. For example, if the WG and heaters corresponding to the index 1 and 2 were the same length and parallel, we could assume $m_{12} = m_{21}$.

Complex perceptron cross-talk matrix measurement

The complex perceptron can be considered as a four arms MZI. Each of the four arms has a delay line (0, 50, 100, and 150 ps) and a 50 μm long heater that is used as a phase shifter. The heaters share the electrical ground channel. Since they are quite close one to the other (40 μm), flowing current in one heater will inevitably affect also the other arms. For simplicity, from here onward the four arms and heaters will be addressed by an integer number from 1 to 4, starting from the one with the smallest delay.

The complex perceptron does not allow to measure the effect of cross-talk independently. There is a single four optical output in which the fields of the four different arms interfere. Figure A.6 shows the spectrum of the complex perceptron excited with white-light (ASE from EDFA1) measured with the OSA. The trace is obtained for a restricted wavelength range (1547–1557 nm) so that both the light source emission and the grating coupler transmission is easily taken into consideration with the measurement of the reference spectrum. In order to study the effect of cross-talk, I repeated the measurement of the spectrum changing the current flowing in the four heaters one at a time. The data collected is composed of roughly 250 000 transmission points for:

- 2001 points in wavelength
- 31 points in current
- four heaters

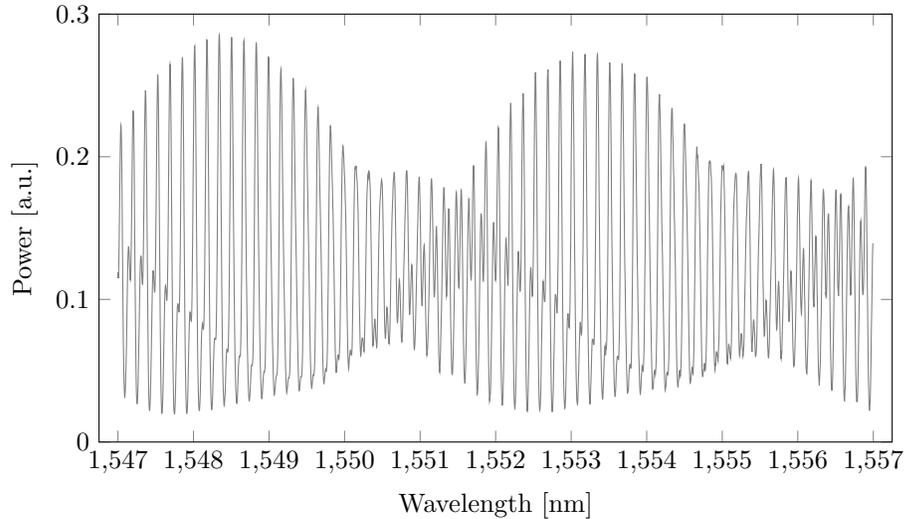


Figure A.6: Spectrum of the complex perceptron where no currents have been applied to heaters.

The shape of the transmission spectrum is due to the interference effect of the six paths in the MZI. Specifically, a longer optical path difference (OPD)

between two arms is translated in a spectrum characterised by more oscillations in the same wavelength range. There are three values of OPD among the four arms, 3.6, 7.2, and 10.8 mm, which give rise to roughly 6, 12, and 18 oscillations per nm. Hence, it is possible to isolate the various components via Fourier transform, filtering, and Fourier anti-transform. Figure A.7 shows an example in which a Gaussian filter is being applied to the first peak. Notice how the first and second peaks are split, due to minimal deviations from the nominal value between paths with the same OPD.

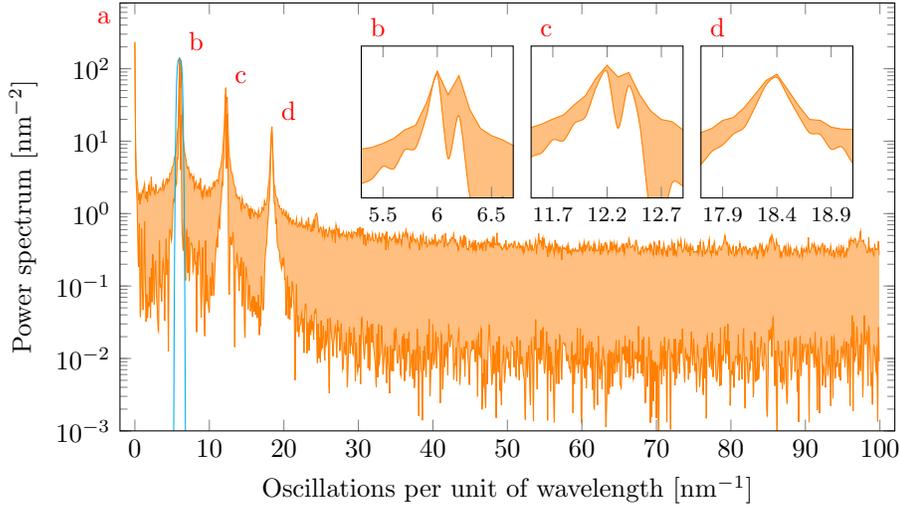


Figure A.7: Single-side Fourier transform of the spectra: envelope of the 31×4 spectra. The blue line represents the Gaussian filter. Note the presence of three peaks roughly at positions of 6, 12, and 18 nm^{-1} .

The result is a set of 31×4 spectra, in which the effect of a given OPD is isolated. The interference pattern at zero power for each heater (4 spectra) is fitted with a cosine wave for each pair of paths causing the interference. Specifically, the longer OPD has one cosine (filter on the third peak), the one in the middle has two cosines (filter on the second peak), and the shortest OPD has three cosines (filter on the first peak). The fit takes into account the dispersion of the effective index with a second-order polynomial. After this, the influence of each heater is evaluated as a first-order derivative on the phase change for each current sweep:

$$\Delta\phi_{a,b} = \frac{\partial\Delta\phi_{a,b}}{\partial P_i} P_i = \left(\frac{\partial\phi_a}{\partial P_i} - \frac{\partial\phi_b}{\partial P_i} \right) P_i, \quad (\text{A.2})$$

where $\Delta\phi_{a,b} = \phi_a - \phi_b$ is the phase difference measured on the interference of the paths a and b . The cross-talk matrix is then obtained by isolating each derivative as $m_{ij} = \frac{\partial\phi_i}{\partial P_j}$. In order to do this, I have to make the (strong) assumption on the system symmetry, such that the cross-talk matrix can be

described as:

$$\mathbf{M} = \begin{bmatrix} \phi_0 & \phi_I & \phi_{II} & \phi_{III} \\ \phi_I & \phi_0 & \phi_I & \phi_{II} \\ \phi_{II} & \phi_I & \phi_0 & \phi_I \\ \phi_{III} & \phi_{II} & \phi_I & \phi_0 \end{bmatrix}, \quad (\text{A.3})$$

and furthermore $\phi_{III} = 0$. The roman numerals indicate the distance between the heater and the channel affected, e.g. in ϕ_0 the heater changes the phase of the WG underneath, in ϕ_I the heater affects the phase of the WG next to it, and so on. From this, the value of ϕ_0 can be extracted from the analysis of the third peak, in fact, the effect of P_1 on the longest OPD gives:

$$\Delta\phi_{1,4} = \frac{\partial\Delta\phi_{1,4}}{\partial P_1} P_1 = (\phi_0 - \phi_{III}) P_1 = \phi_0 P_1, \quad (\text{A.4})$$

and similarly analysing the effect of P_4 gives the same result $\frac{\partial\Delta\phi_{1,4}}{\partial P_4} = -\phi_0$. Then, studying the OPD of 100 ps, also ϕ_I and ϕ_{II} can be evaluated, by inverting the following relations

$$\frac{\partial\Delta\phi_{1,3}}{\partial P_1} = \frac{\partial\Delta\phi_{2,4}}{\partial P_2} = \phi_0 - \phi_{II} \quad (\text{A.5})$$

$$\frac{\partial\Delta\phi_{1,3}}{\partial P_3} = \frac{\partial\Delta\phi_{2,4}}{\partial P_4} = \phi_{II} - \phi_0 \quad (\text{A.6})$$

$$\frac{\partial\Delta\phi_{1,3}}{\partial P_4} = \phi_{III} - \phi_I = -\phi_I \quad \frac{\partial\Delta\phi_{2,4}}{\partial P_1} = \phi_I - \phi_{III} = \phi_I \quad (\text{A.7})$$

and employing the value of ϕ_0 already found. Finally, the value of the partial derivatives of the paths $\Delta\phi_{1,4}$, $\Delta\phi_{1,3}$, and $\Delta\phi_{2,4}$ are reported in table A.2.

	$\frac{\partial}{\partial P_1}$	$\frac{\partial}{\partial P_2}$	$\frac{\partial}{\partial P_3}$	$\frac{\partial}{\partial P_4}$
$\Delta\phi_{1,4}$	0.0651	-0.0547	-0.0606	-0.1999
$\Delta\phi_{1,3}$	0.0746	-0.0508	-0.1654	-0.0551
$\Delta\phi_{2,4}$	-0.0094	0.1008	-0.0237	-0.1624

Table A.2: Values of the partial derivatives obtained from the Fourier transform spectral analysis.

Unfortunately, the values obtained by this analysis do not reflect the assumption on the symmetry employed to carry out their evaluation. The outcome, hence, is that the system is not symmetrical and this does not allow to properly isolate the effect of the different heaters in order to evaluate the cross-talk matrix.

Measurement prior thermalisation

Exploiting the different timescale at which the temperature of the WG directly under the heater reaches the steady-state (around 100s of μs) and that of the

surrounding volume in which the nearby WGs are (around 100s of ms), it is possible to obtain a measurement not affected by thermal cross-talk. Indeed, as fig. A.8 shows, the current applied by QQ reaches a stationary state in a few ms, after which there is a plateau where the structure can be considered without thermal cross-talk, and only after several tens to hundreds of ms the rest of the material thermalise, showing the full extent of the cross-talk.

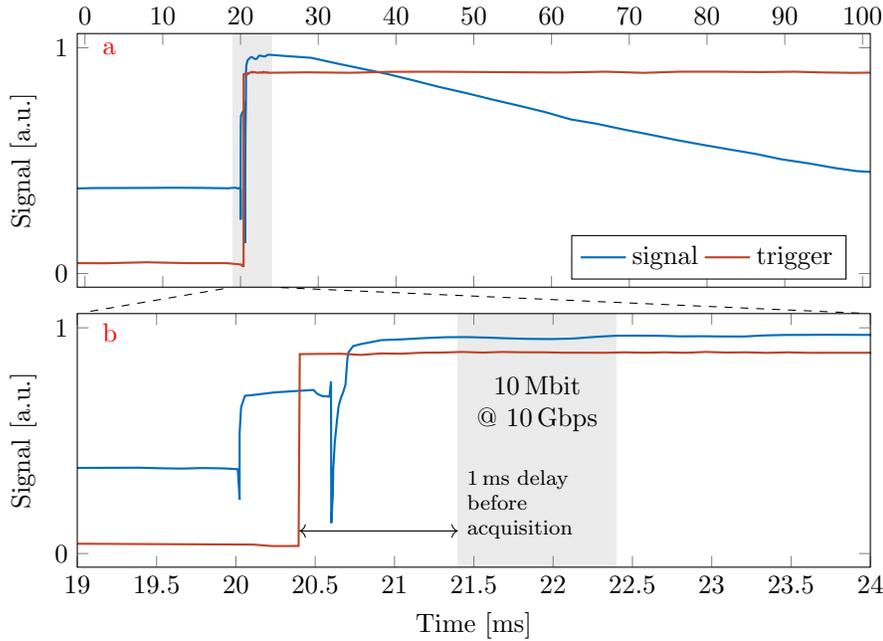


Figure A.8: Timescales involved in the fast measurement prior to thermalisation. **a** Long measurement. **b** Zoom around the trigger event. The acquisition happens after 1 ms after the trigger event and usually lasts a few μs . In comparison, 1 ms of measurement contains 10 Mbit at 10 Gbps.

In order to obtain such measurement, one of the channels of the QQ was connected to the auxiliary port of OSC, with a $50\ \Omega$ load. When a current flowing in the channel is changed from 0 to 10 mA, OSC observes the increase in the voltage difference from 0 to 500 mV, which is easily detectable and used as a trigger for the OSC. So the acquisition from the oscilloscope is synchronised with the flowing of current in the channels of QQ and thus in the DUT. By fine-tuning the acquisition delay on OSC, measures about 100 μs long can be acquired without seeing particular underlying drifts due to thermal cross-talk. Considering an input bitrate of 10 Gbps, a measure of just 10 μs contains 10^5 bits, which are sufficient to obtain reliable statistics.

A.4.4 Input-output delay calibration procedure

Due to the fact that the two fast photodetectors are separated by a considerable optical path, the same optical signal is seen by PD2 with a delay in respect to PD1. For this reason, in order to synchronise the traces acquired by OSC, it is necessary to measure the value of this delay. Since also the integrated structure can add a noticeable delay, this measure should not change the optical line. Indeed, this is an online procedure based on the evaluation of the correlation between the input and output traces, so no change is required.

In order to provide the absolute value of the delay, this operation should be carried out with an input signal whose periodicity is longer than the delay. If the period of the input signal is shorter than the delay, then only the relative delay can be evaluated. However, as it is in our case, this is sufficient when no remarkable drifts in the response of the system occur over several periods. Moreover, the DUT heavily transforms the waveform of the input signal, so that it becomes difficult to recognise it in the output signal. For this reason, the `particleswarm` is employed to find the optimal values (current) for the heaters in the structure, such that the output signal becomes as close as possible to the input one. Specifically, the algorithm maximises the peak correlation between the input and output.

The measurement process is composed of the following steps:

- (optional) Apply currents to heaters.
- The input and output traces are acquired.
- The correlation between the input and output is evaluated.
- The position of the maximum represents the delay.

These steps are repeated several times in order to obtain reliable statistics. The final delay is the mode of all these values.

The delay evaluated is then employed in all the following measurements. The calibration must be repeated whenever there is a change in the optical line or the period of the modulation in the input signal is modified, in the case of the relative delay.

Appendix B

Photolithographic Masks

I designed several integrated optical devices for the BACKUP project and/or for related projects, such as ALPI (optical signal recovery by Photonic neural network Integrated in a transceiver) and PELM (Photonic Extreme Learning Machine).

The first two designs have been carried out using the `gdspy` python library [101], which I customised in order to include helpful features, such as the connection ports coordinates of building blocks. Successive designs have been accomplished by using `nazca` python library [102], which is a modern open-source framework for the design of PICs.

Custom structures have been usually simulated with Comsol, which provides a powerful and flexible FEM solver. However, in some cases, python or Matlab scripts have been used as well.

B.1 Fabrication process

The development of integration processes and high-performance components have the consequence of reducing the production cost of the single component. High performance and low cost are, in fact, the principal characteristics which brought the CMOS technology to be in the lead of the integrated electronics industry.

In integrated circuit (IC) manufacturing, structures are created by shaping several layers of new materials over a substrate, e.g. a silicon wafer, through several removal or additive steps. This can be obtained through some lithographic techniques, such as photo- and electron beam lithography, and nanoimprinting. All these techniques share the fundamental working principle of creating a patterned thin film over the substrate, protecting particular areas from and exposing others to a given implantation, deposition, or etching operation.

Photo-lithography is the conventional technique for manufacturing electronic and photonic integrated circuits. It uses light, typically ultraviolet (UV) light, and a photo-lithographic mask to project the pattern on a thin film of photore-

sist, a light-sensitive substance. Irradiated areas may either expose or protect the substrate underneath. Its advantages are given by the high patterning resolution, currently enabling the manufacturing of features sizes of the order of 100 nm [62], in conjunction with the possibility to process entire wafers at a time, hence reducing the cost per operation. This technique is used in the vast majority of the IC manufacturing plants.

Electron beam lithography, on the other hand, gives access to higher resolution in comparison to photo-lithography [62, 103], however, its throughput is low and cannot be used to pattern at the wafer-scale level. In this technique, the thin film composed of an electron-sensitive material is patterned via direct writing with an electron beam.

Although these methods allow high precision, a certain degree of manufacturing uncertainty should always be considered. For example, in photo-lithography, two nominally identical structures in the same chip may differ in dimension of a few nm, while between two chips on the same wafer this difference might be 5–10 nm [89].

Silicon photonics brings to the table several advantages over other technologies, namely relatively low cost and the well-known electric and optical properties as well as its robust technological platform. Moreover, it is optically transparent both at 1.3 and 1.55 μm and its high refractive index enable high confinement factors and thus small devices. Furthermore, its thermal conductivity allows high tolerance to power dissipation and optical modulation is obtained either via carrier injection or thanks to the thermo-optic effect. However, silicon also has some drawbacks. First of all, it has an indirect bandgap, which means that native optical sources are not possible. Secondly, it does not provide second-order nonlinearities, used in optical modulation, whereas allowed modulation mechanisms tend to be moderately slow.

Hybrid integration of low-cost silicon photonics with optically active components belonging to other material platforms is nowadays gaining interest. A number of technologies have been proposed and are currently being investigated, such as micro-transfer printing [104], engineering of epitaxial growth [105, 106], and photonic wire-bonding via additive manufacturing [107]. The perspective is to enable a low-cost platform with the benefits of both silicon photonics and other III-V platforms.

B.2 Design process

The principal objective of the design is to manufacture PICs functioning at the target performance. The ideal device functionality should be transformed into a circuit composed of integrated structures by taking into account their fundamental operation mechanisms, i.e. mainly linear propagation and interference effects. However, the characteristics of the real device deviate from those of the ideal model. The manufacturing process is obviously not able to reproduce the ideal geometry of the structures with infinite precision and the modelling itself often lacks the complexity to include all of the characteristics of the real

device. In fact, the causes of deviation from the ideal operation are many and the design should attempt to counteract these effects or at least include them in such a way that the device performance is not hampered by them. For example, loss due propagation is strongly affected by surface roughness on the walls of the core of the WG and is usually taken into account by measuring the propagation losses experimentally on test structures. Additionally, even infinitesimal local variations of the effective refractive index or small changes to the geometrical path in a MRR can shift the resonance position. For this reason, electrical heaters are employed to locally control the temperature of the material, effectively changing the resonance position. The use of active elements (heaters), however, introduces other issues, such as the thermal cross-talk between nearby structures. Therefore, it is also important to be able to characterise single parts of a complex PIC, so that deviations from the ideal functionality can be measured and hopefully accounted for either directly in the interaction with the PIC or in a post-processing phase following the measurements.

Interaction with the PIC should be ensured in both the optical and electrical domains. The layout should take into consideration the optical coupling mechanisms and ensure access to the input/output couplers. For example, a design employing side couplers should keep one or two edges clear while those making use of surface couplers should free the area around such couplers from obstacles. More advanced techniques, such as optical fibre arrays or photonic wire-bonding, may be employed to improve insertion loss and ease of use. Electrical connections, on the other hand, are obtained at specific areas called contact pads. Here the connection is secured by either metallic contact with a temporary tungsten tip or by permanent bonding of wires. Electrical wire-bonding provides good, but delicate, connections, so their length should be minimised. Hence, contact pads for wire-bonding are typically placed only on the chip edges, reducing the travel from the pads on the chip to those outside as much as possible. Temporary optical coupling and electrical contact solutions only require the naked chip and are thus easier to carry out, for this reason, they are usually employed in the early stages of research and development. Nevertheless, late development stages and commercial products benefit from permanent coupling which is coupled with the packaging of the device. The packaging ensures reliable optical and electrical coupling as well as, in some cases, protection from the environment and temperature control.

In conclusion, the design of a PIC should deliver the target performance, starting from the ideal operation mechanisms, but also taking into account the effects which influence the real device. Furthermore, the layout should be developed so that the device footprint is minimised while assuring the required optical coupling and electrical connections.

For the foundries to be able to produce a PIC, they have to be supplied with its design. The *de facto* standard file format for this task is the GDSII stream format (GDSII), where the acronym stands for graphic design system (GDS). Figure B.1 gives an example of GDS drawing. The fundamental aspect of this file format is that it contains sets of 2D polygons, each representing a piece of the final PIC, organised in layers. The different layers are then combined through

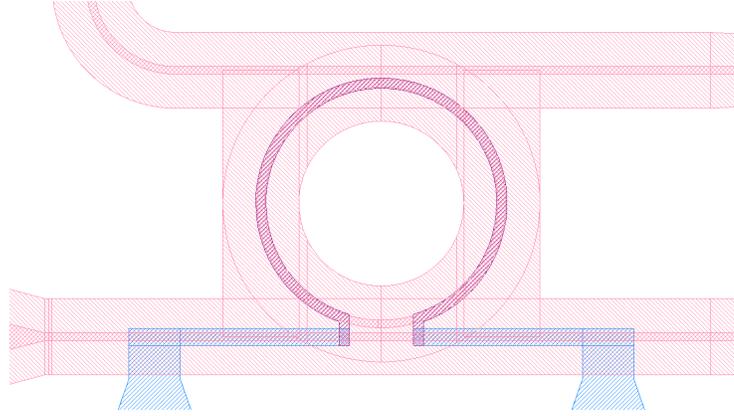


Figure B.1: Example of a GDS drawing. Each of the several layers displayed corresponds to a specific step in the fabrication process. Specifically, pink layers work on the device layer, defining the core and the cladding of the WG, while purple and blue are the metallic heater and metallic traces, respectively.

Boolean operations, e.g. union and difference, and mapped to specific physical steps of the manufacturing process. The design is typically further organised in a hierarchical structure of cells, containing any given subset of polygons even belonging to different layers, which improve the clarity of the design. There are several tools that help in the design, e.g. Klayout [108], gdspy [101], Nazca Design [102], Lumerical [78]. Some are open and free, others are proprietary.

Each foundry is able to carry out a specific set of operations on the base silicon wafer. The set of allowed operations defines a manufacturing technology, which has some intrinsic limitations embodied by the *design rules*. In order to maximise the manufacturing yield and the device performance, these rules are enforced before the beginning of production through the design rule check (DRC) on the GDSII file.

B.3 IR IMEC 2019

This has been the first photo-lithographic mask designed for the BACKUP project. The design was carried out by me and Mattia Mancinelli at the beginning of 2019.

The main focus of this design was to study RCNs based on MRRs, although other types of structures are present. The two most innovative features of this mask are the integration of an optical perceptron layer directly on-chip and the use of SCISSORs as reservoirs. Integrating the optical perceptron layer moves the burden of the weighted sum of the reservoir attributes in the optical domain, in principle enabling faster operation, as optical components have higher bandwidth than electrical ones, but at the same time increasing the complexity of structure and its training procedure. This layer is used both standalone (see

the description of the complex perceptron in section 5.1) and in the RCNs as an optical readout layer. Five variations of SCISSOR-based networks, each with four MRRs, but differing in inter-ring delay and the MRR quality factor were placed on the design. In addition to the SCISSOR-based networks, the chip also contains one RCN based on the “swirl” topology [36], although modified with MRR nodes and an optical readout layer, and a feedforward neural network with two inputs, one hidden layer composed of two nonlinear nodes, and a single output nonlinear, in order to study the use of MRR nonlinear response as activation function. Finally, several test structures are placed all around the mask, in between the main devices, where there is unused space.

This design is based on imec’s Passive+ technology platform, which is a subset of the iSiPP50G technology, enabling the design of passive structures and heaters working in the C-band [109]. The devices employed two of the foundry’s own building blocks: the grating coupler and the 1×2 MMI. The design area is half-block, i.e. $2.5 \text{ mm} \times 5.15 \text{ mm}$, and contains the following structures:

- A feed-forward neural network with three layers: 2 inputs linear nodes, four weights, two nonlinear nodes, two weights, one output nonlinear node. There are a total of 19 heaters on the structure.
- The complex perceptron (see section 5.1)
- Five SCISSOR structures with four MRRs and feedback loop between the through and the add ports. 8 taps positioned between each MRR collect $\sim 3\%$ of the optical signal. Their signal is optically summed through a cascade of 2×1 MMIs until a single output is obtained. These structures differentiate themselves based on the inter-ring delay, the coupling region length, and the gap between the WG and the MRR:
 - Three copies with an inter-ring delay of 12.5 ps, 25 ps, and 50 ps. The microring resonators are racetracks with radius $\sim 7 \mu\text{m}$ and coupling regions of length $5 \mu\text{m}$ and gap $0.25 \mu\text{m}$.
 - Two copies without inter-ring delays. The microring resonators in the first structure are racetracks with radius $\sim 7 \mu\text{m}$ and coupling regions of length $2 \mu\text{m}$ and gap $0.2 \mu\text{m}$. The microring resonators in the second structure are rings with radius $\sim 7 \mu\text{m}$ and point coupling regions with gap $0.18 \mu\text{m}$.
- A “swirl” structure with a 4×4 ring matrix, fully electrically addressed, and a complex perceptron as readout.
- test structures organised as follows:
 - ring resonators (6 copies)
 - directional coupler (6 copies)
 - grating-to-grating

- 1×2 MMI, fab design
 - 1×4 MMI, new design [110]
 - waveguide crossings
 - delay lines of length 12.5, 25, 50, 75, and 100 ps
- two areas on the long edges of the chip were reserved for structures regarding quantum technologies.

Figure B.2 reports the final computer-aided design (CAD) of the mask, highlighting the different structures discussed above, and the picture of a manufactured sample.

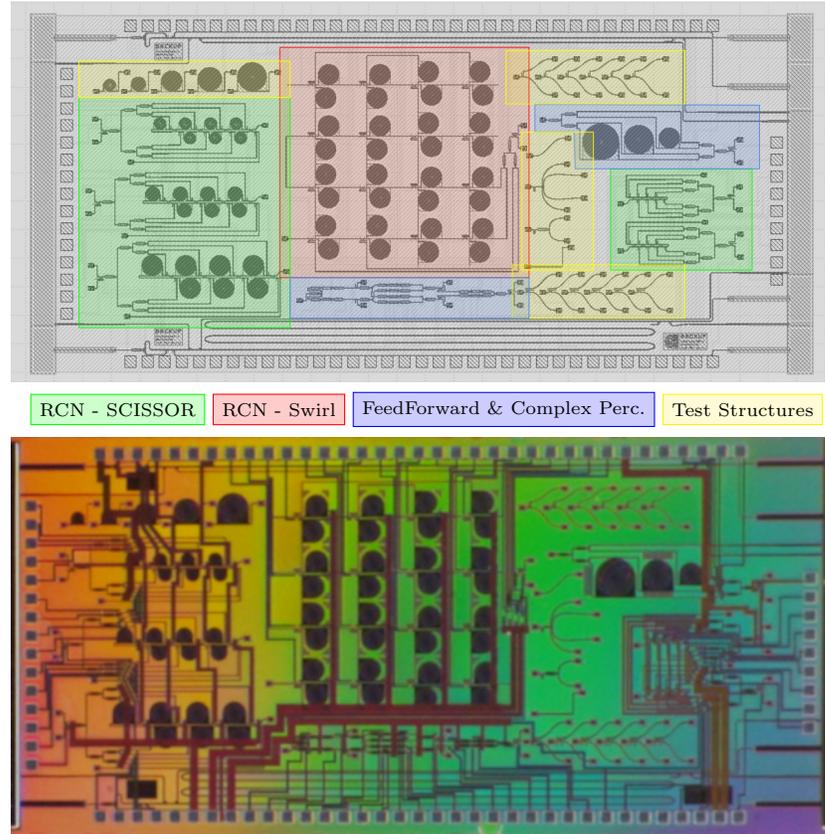


Figure B.2: **a** CAD of the device, highlighting the SCISSOR-based RCNs in green, “swirl” topology RCN in red, the complex perceptron and the FFNN in blue, and the test structures in yellow. **b** Picture of the manufactured chip.

B.4 Visible IMEC 2019

The second photolithographic mask designed for the BACKUP project aimed at creating a PIC able to hold living neurons on its surface and to stimulate a subset of the neurons with visible (blue) light. In order to target individual neurons, light must be scattered from the PIC at precise locations over its surface. The design was carried out by myself, Mattia Mancinelli, and Clara Zaccaria during September and October in 2019 and the manufactured samples have been delivered in May 2020. Figure B.3a shows a picture of the manufactured chip.

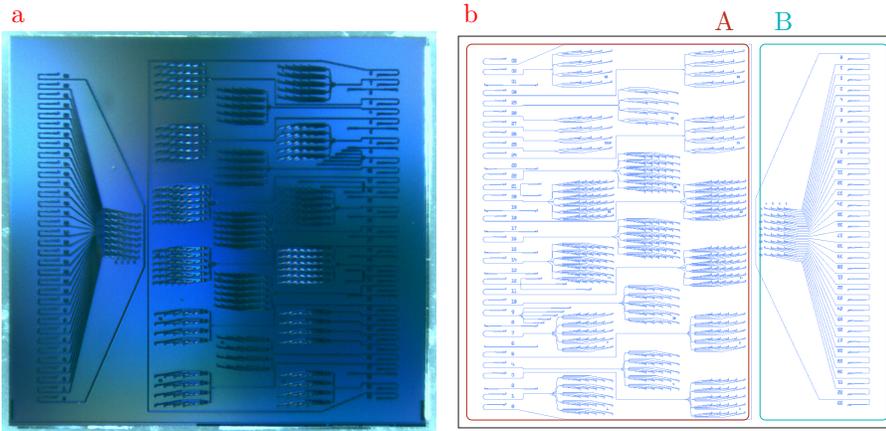


Figure B.3: Picture **a** and **b** CAD of the manufactured chip. Highlighted: area (A) on the left and area (B) on the right.

In order to be able to guide visible light at 488 nm, the wavelength required for the activation of the channel-rhodopsin transfected in the living neurons, a silicon nitride (SiN) platform has been chosen, specifically, imec's BioPIX150 technology. It employs standard photo-lithographic CMOS fabrication processes and can manufacture 150 nm high SiN WGs embedded in SiO₂ cladding, whose target operation lies between 400 to 650 nm.

The design contains only passive components and has a total surface of 5.300 mm × 4.750 mm. The total area is divided into two parts with full height and reduced width: (A) containing several groups of scatterers and (B) containing a single group of scatterers. The input grating couplers are disposed on a single column of 34 elements, placed on the left side of each area at a distance of 127 μm one from the other, which is the standard core-to-core pitch for the production of fibre arrays. Figure B.3b shows the CAD of the device, with the two halves highlighted.

Area (A) contains a total of 17 groups, each of which consists of either 4 × 4, 5 × 5, or 6 × 6 scatterers disposed on square grids. All groups are addressed by a single optical input, except for of a 4 × 4 grid in which each row is connected to an independent optical input: inputs 1, 3, 4, 5, 7, and 10 are connected to 5 × 5 grids, inputs 11, 14, 16, 18, 20, and 22 are connected to 6 × 6 grids,

inputs 24, 29, 30, and 32 are connected to 4×4 grids, and inputs 25 to 28 are connected each to one row of a 4×4 grid. The inputs 0 and 33 are connected one to the other in order to facilitate the alignment of the fibre array to the chip and the remaining inputs are connected to the structure used for test and characterisation purposes.

Area (B) contains a single group of 32 scatterers, disposed on an 8×4 square grid. Each one of them is connected to an independent input, from number 1 to 32. Similarly to the area (A), inputs 0 and 33 are connected one to the other to simplify the alignment procedure.

The design makes use of imec's grating couplers used as inputs, custom grating couplers used as scattering structures, and custom 1×2 , 1×3 , and 1×5 MMIs for distributing one input equally over many scatterers. The MMIs have been designed using the procedure described in section 4.4. Figure B.4 shows an enlarged view of the custom structures.

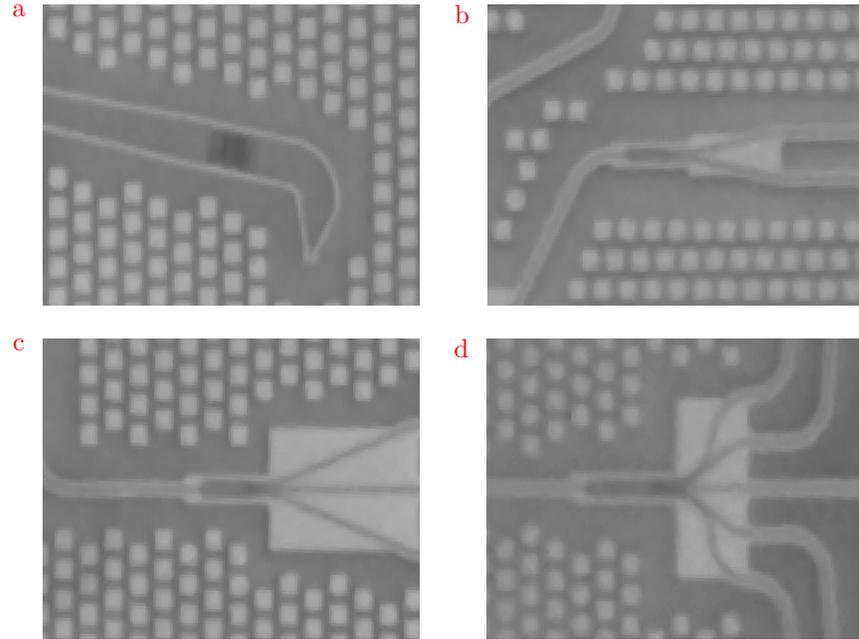


Figure B.4: Close-up of the custom structures. **a** Scatterer. **b** 1×2 MMI. **c** 1×3 MMI. **d** 1×5 MMI.

B.5 IR IMEC 2020

The third mask of the BACKUP project was developed in collaboration with the PELM project. It is based on the same technology platform of the first mask, imec's Passive+, and, similarly, targets the study of neuromorphic structures.

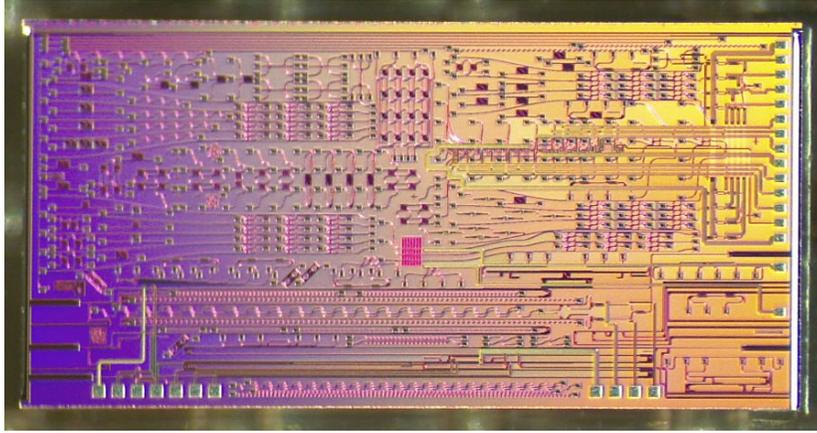


Figure B.5: Picture of the manufactured chip.

The design area is again half-block, i.e. $2.5 \text{ mm} \times 5.15 \text{ mm}$.

This device contains principally two types of structures, one containing a matrix of MRRs to be driven with CW signals and another one based on a revised version of CROW and SCISSOR reservoirs employed for their dynamics. This mask contains also several test structures, designed to allow the characterisation of the relevant parameters and test the single components of the bigger structures. This time the main structures were developed to exchange light with a fibre array. So, instead of using isolated grating couplers, all the input/output grating couplers of the main structures have been aligned and distanced by $127 \mu\text{m}$. Figure B.5 shows a pic of the manufactured chip and gives the idea of the complexity of the design and the density of the structure placed on such tiny surface.

B.6 IR Applied Nanotools 2021

From Tuesday 9th November 2021 to Friday 19th November 2021, I designed together with Mattia Mancinelli a supplementary design for the ALPI project, based on the complex perceptron. The total design area is $9 \text{ mm} \times 9 \text{ mm}$ and is divided into E sections following a 3×3 rectangular grid. Each chip of smaller area is approximately $2.7 \text{ mm} \times 2.7 \text{ mm}$ large. The left and right sides are used for edge coupling and are hence manufactured with deep trenches. Input and output coupling is obtained with edge couplers based on inverse tapering, except for monitoring points and some test structures which may employ grating couplers. Each one of the nine sections contains one of two variations of the main structure, implemented with a specific combination of MMIs and characterised by a specific delay.

A total of four different MMIs have been used in the design: three, (a), (b), and (c), have been designed and simulated in-house, while a fourth, (d), similar

to a Y-junction, is taken from the “SiEPIC EBeam PDK” [111]. Two of them, (a) and (d), are 1×2 MMIs, while the other, (b) and (c), are 2×2 MMIs. The first three structures were designed using the procedure described in section 4.4.¹

The main structures have been organised as follows:

- variation 1 (active, edge couplers)
 - 25 ps delay, (a) and (b) MMIs (2 copies)
 - 25 ps delay, (d) MMIs only (2 copies)
 - 15.625 ps delay, (a) and (b) MMIs (1 copy)
 - 20.833 ps delay, (a) and (b) MMIs (1 copy)
- variation 2 (active, edge couplers)
 - 25 ps delay, (a) and (c) MMIs (2 copies)
 - 15.625 ps delay, (a) and (c) MMIs (1 copy)

Where the remaining space allowed it, the following test structures have been inserted in the chips:

- edge coupler to edge coupler (6 copies)
- delay lines of 1, 2, 3, 4, 5, 6, 7, and 8 mm (3 copies each)
- grating to grating, two versions (3 copies each)
- Mach-Zehnder interferometers (active, grating couplers)
 - heater length of 300 and 600 μm (3 copies each)
- cascades of MMIs for insertion loss and channel unbalance (grating couplers)
 - (a), (b), (c), and (d) MMIs (1 copy each)
- 7 μm radius ring resonators (grating couplers)
 - in both all-pass filter and add-drop filter configurations
 - gap of 150, 180, 210, 240, 270, and 300 nm
- rings with “Euler-bends” (grating couplers)
 - all pass with a gap of 289 and 310 nm
 - CROW with two rings, external gap 289 nm, inner gap 260 nm

¹To my surprise, after two years of inactivity, the code I wrote for MMI simulation was used by a different person with successful outcomes.

The active structures employ one or more WG sections over which $4\ \mu\text{m}$ wide heaters, made of a Ti/Walloy, have been placed. Active elements are connected to two sets of Albond pads placed on the perimeter of the chip, along two opposite sides. There is a total of 24 pads on each chip.

The design has been manufactured on a SOI platform using a 100 keV electron beam lithography process developed by Applied Nanotools. The silicon layer, the buried oxide layer, and the silicon substrate layer are 220 nm, $2\ \mu\text{m}$, and $725\ \mu\text{m}$ thick, respectively. Structures are fully-etched and are created using an e-beam mask material and anisotropic inductively coupled plasma - reactive ion etching (ICP-RIE) process. The standard WG section is a trapezium with sidewall angles of approximately 88° , with an area of $500\ \text{nm} \times 220\ \text{nm}$. The expected propagation loss for the TE mode is 1.5 dB/cm on straight WG and 3.5 dB/cm on curved WG. The minimum feature size is 60 nm and the minimum spacing is 70 nm, excluding proximity effects. A protective oxide cladding layer $2\ \mu\text{m}$ thick is deposited on top of the other layers.

Bibliography

- [1] Davide Bazzanella et al. “Artificial and Biological Neurons: Interdisciplinary Issues and Future Perspectives. White Paper”. In: *BioLaw Journal - Rivista di BioDiritto* 1 (Mar. 2021), pp. 353–379. ISSN: 2284-4503. DOI: 10.15168/2284-4503-756.
- [2] Warren S McCulloch and Walter Pitts. “A logical calculus of the ideas immanent in nervous activity”. In: *The bulletin of mathematical biophysics* 5.4 (1943), pp. 115–133.
- [3] Marvin Minsky and Seymour Papert. “Perceptrons.” In: (1969).
- [4] John Markoff. “Computer wins on ‘Jeopardy!’: Trivial, it’s not”. In: *New York Times* 16 (2011).
- [5] Steven Borowiec. “AlphaGo seals 4-1 victory over Go grandmaster Lee Sedol”. In: *The Guardian* 15 (2016), p. 6.
- [6] *Train a software agent to behave rationally with reinforcement learning*. Accessed 2022-03-06. 2017. URL: <https://developer.ibm.com/articles/cc-reinforcement-learning-train-software-agent/>.
- [7] Wang Tong et al. “Artificial intelligence for vehicle-to-everything: A survey”. In: *IEEE Access* 7 (2019), pp. 10823–10843.
- [8] Sorin Grigorescu et al. “A survey of deep learning techniques for autonomous driving”. In: *Journal of Field Robotics* 37.3 (2020), pp. 362–386.
- [9] Ammar Haydari and Yasin Yilmaz. “Deep reinforcement learning for intelligent transportation systems: A survey”. In: *IEEE Transactions on Intelligent Transportation Systems* (2020).
- [10] Solon Barocas and Andrew D Selbst. “Big data’s disparate impact”. In: *Calif. L. Rev.* 104 (2016), p. 671.
- [11] Gouhei Tanaka et al. “Recent advances in physical reservoir computing: A review”. In: *Neural Networks* 115 (2019), pp. 100–123.
- [12] André Röhm and Kathy Lüdge. “Multiplexed networks: reservoir computing with virtual and real nodes”. In: *Journal of Physics Communications* 2.8 (2018), p. 085007.

- [13] Ben Varkey Benjamin et al. “Neurogrid: A mixed-analog-digital multi-chip system for large-scale neural simulations”. In: *Proceedings of the IEEE* 102.5 (2014), pp. 699–716.
- [14] Paul A Merolla et al. “A million spiking-neuron integrated circuit with a scalable communication network and interface”. In: *Science* 345.6197 (2014), pp. 668–673.
- [15] Johannes Schemmel et al. “A wafer-scale neuromorphic hardware system for large-scale neural modeling”. In: *2010 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE. 2010, pp. 1947–1950.
- [16] Zhe Jia et al. “Dissecting the graphcore ipu architecture via microbenchmarking”. In: *arXiv preprint arXiv:1912.03413* (2019).
- [17] Yu Emma Wang, Gu-Yeon Wei, and David Brooks. “Benchmarking TPU, GPU, and CPU platforms for deep learning”. In: *arXiv:1907.10701* (2019).
- [18] Jinook Song et al. “An 11.5 TOPS/W 1024-MAC butterfly structure dual-core sparsity-aware neural processing unit in 8nm flagship mobile SoC”. In: *2019 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE. 2019, pp. 130–132.
- [19] Jun-Woo Jang et al. “Sparsity-aware and re-configurable NPU architecture for samsung flagship mobile SoC”. In: *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. IEEE. 2021, pp. 15–28.
- [20] Carl Ramey. “Silicon photonics for artificial intelligence acceleration: HotChips 32”. In: *2020 IEEE hot chips 32 symposium (HCS)*. IEEE. 2020, pp. 1–26.
- [21] Jonathan Dong et al. “Optical reservoir computing using multiple light scattering for chaotic systems prediction”. In: *IEEE Journal of Selected Topics in Quantum Electronics* 26.1 (2019), pp. 1–12.
- [22] Daniel Hesslow et al. “Photonic co-processors in HPC: using LightOn OPUs for Randomized Numerical Linear Algebra”. In: *arXiv preprint arXiv:2104.14429* (2021).
- [23] Xing Lin et al. “All-optical machine learning using diffractive deep neural networks”. In: *Science* 361.6406 (2018), pp. 1004–1008.
- [24] Yvan Paquot et al. “Optoelectronic reservoir computing”. In: *Scientific reports* 2.1 (2012), pp. 1–6.
- [25] Laurent Larger et al. “Photonic information processing beyond Turing: an optoelectronic implementation of reservoir computing”. In: *Optics express* 20.3 (2012), pp. 3241–3249.
- [26] Miguel C Soriano et al. “Optoelectronic reservoir computing: tackling noise-induced performance degradation”. In: *Optics express* 21.1 (2013), pp. 12–20.

- [27] S Ortín et al. “A unified framework for reservoir computing and extreme learning machines based on a single time-delayed neuron”. In: *Scientific reports* 5.1 (2015), pp. 1–11.
- [28] Andrew Katumba et al. “Neuromorphic computing based on silicon photonics and reservoir computing”. In: *IEEE Journal of Selected Topics in Quantum Electronics* 24.6 (2018), pp. 1–10.
- [29] Guy Van der Sande, Daniel Brunner, and Miguel C Soriano. “Advances in photonic reservoir computing”. In: *Nanophotonics* 6.3 (2017), pp. 561–576.
- [30] Febin P Sunny et al. “A survey on silicon photonics for deep learning”. In: *ACM Journal of Emerging Technologies in Computing System* 17.4 (2021), pp. 1–57.
- [31] Hsuan-Tung Peng et al. “Neuromorphic photonic integrated circuits”. In: *IEEE Journal of Selected Topics in Quantum Electronics* 24.6 (2018), pp. 1–15.
- [32] Bhavin J Shastri et al. “Photonics for artificial intelligence and neuromorphic computing”. In: *Nature Photonics* 15.2 (2021), pp. 102–114.
- [33] Kristof Vandoorne et al. “Toward optical signal processing using photonic reservoir computing”. In: *Optics express* 16.15 (2008), pp. 11182–11192.
- [34] Kristof Vandoorne et al. “Parallel reservoir computing using optical amplifiers”. In: *IEEE transactions on neural networks* 22.9 (2011), pp. 1469–1481.
- [35] Kristof Vandoorne et al. “Experimental demonstration of reservoir computing on a silicon photonics chip”. In: *Nature communications* 5.1 (2014), pp. 1–6.
- [36] Florian Denis-Le Coarer et al. “All-optical reservoir computing on a photonic chip using silicon-based ring resonators”. In: *IEEE Journal of Selected Topics in Quantum Electronics* 24.6 (2018), pp. 1–8.
- [37] Floris Laporte et al. “Numerical demonstration of neuromorphic computing with photonic crystal cavities”. In: *Optics express* 26.7 (2018), pp. 7955–7964.
- [38] Lennert Appeltant et al. “Information processing using a single dynamical node as complex system”. In: *Nature communications* 2.1 (2011), pp. 1–6.
- [39] François Duport et al. “All-optical reservoir computing”. In: *Optics express* 20.20 (2012), pp. 22783–22795.
- [40] Bahaa EA Saleh and Malvin Carl Teich. *Fundamentals of photonics*. John Wiley & sons, 2019.
- [41] HH Li. “Refractive index of silicon and germanium and its wavelength and temperature derivatives”. In: *Journal of Physical and Chemical Reference Data* 9.3 (1980), pp. 561–658.

- [42] Ian H Malitson. “Interspecimen comparison of the refractive index of fused silica”. In: *Josa* 55.10 (1965), pp. 1205–1209.
- [43] Stefano Signorini et al. “Intermodal four-wave mixing in silicon waveguides”. In: *Photonics Research* 6.8 (2018), pp. 805–814.
- [44] Stefano Paesani et al. “Near-ideal spontaneous photon sources in silicon quantum photonics”. In: *Nature communications* 11.1 (2020), pp. 1–6.
- [45] Eric Dulkeith et al. “Group index and group velocity dispersion in silicon-on-insulator photonic wires”. In: *Optics express* 14.9 (2006), pp. 3853–3863.
- [46] Xiaogang Chen, Nicolae C Panoiu, and Richard M Osgood. “Theory of Raman-mediated pulsed amplification in silicon-wire waveguides”. In: *IEEE journal of quantum electronics* 42.2 (2006), pp. 160–170.
- [47] RM Osgood et al. “Engineering nonlinearities in nanoscale optical systems: physics and applications in dispersion-engineered silicon nanophotonic wires”. In: *Advances in Optics and Photonics* 1.1 (Jan. 2009), pp. 162–235. DOI: 10.1364/AOP.1.000162.
- [48] Govind Agrawal. *Nonlinear Fiber Optics*. Optics and Photonics. Elsevier Science, 2013.
- [49] Long Zhang et al. “Ultralow-Loss Silicon Photonics beyond the Single-mode Regime”. In: *Laser & Photonics Reviews* (2022), p. 2100292.
- [50] Corning[®] SMF-28 Ultra Product Specification. Accessed 2022-03-06. 2021. URL: <https://www.corning.com/media/worldwide/coc/documents/Fiber/PI-1450-AEN.pdf>.
- [51] William S. C. Chang. *Fundamentals of Guided-Wave Optoelectronic Devices*. Cambridge University Press, 2009. DOI: 10.1017/cbo9780511807114.
- [52] Massimo Borghi. “Linear, nonlinear and quantum optics in Silicon Photonics”. PhD thesis. University of Trento, 2016.
- [53] Anita Sure et al. “Fabrication and characterization of three-dimensional silicon tapers”. In: *Optics express* 11.26 (2003), pp. 3555–3561.
- [54] Yun Wang et al. “Apodized focusing fully etched subwavelength grating couplers”. In: *IEEE Photonics Journal* 7.3 (2015), pp. 1–10.
- [55] Na Fang et al. “Three-dimensional tapered spot-size converter based on (111) silicon-on-insulator”. In: *IEEE Photonics Technology Letters* 21.12 (2009), pp. 820–822.
- [56] H Gehring et al. “Low-loss fiber-to-chip couplers with ultrawide optical bandwidth”. In: *APL Photonics* 4.1 (2019), p. 010801.
- [57] Helge Gehring et al. “Broadband out-of-plane coupling at visible wavelengths”. In: *Optics letters* 44.20 (2019), pp. 5089–5092.
- [58] Glycerine Producers’ Association et al. *Physical properties of glycerine and its solutions*. Glycerine Producers’ Association, 1963.

- [59] Chin-Lin Chen. *Foundations for guided-wave optics*. John Wiley & Sons, 2006.
- [60] Sailong Wu et al. “State-of-the-art and perspectives on silicon waveguide crossings: A review”. In: *Micromachines* 11.3 (2020), p. 326.
- [61] Stefano Signorini et al. “Silicon photonics chip for inter-modal four wave mixing on a broad wavelength range”. In: *Frontiers in Physics* (2019), p. 128.
- [62] Yikai Su et al. “Silicon photonic platform for passive waveguide devices: materials, fabrication, and applications”. In: *Advanced Materials Technologies* 5.8 (2020), p. 1901153.
- [63] L.B. Soldano and E.C.M. Pennings. “Optical multi-mode interference devices based on self-imaging: principles and applications”. In: *Journal of Lightwave Technology* 13.4 (Apr. 1995), pp. 615–627. DOI: 10.1109/50.372474.
- [64] M Mancinelli et al. “Chaotic dynamics in coupled resonator sequences”. In: *Optics express* 22.12 (2014), pp. 14505–14516.
- [65] Shaowu Chen et al. “Bistability and self-pulsation phenomena in silicon microring resonators based on nonlinear optical effects”. In: *Optics Express* 20.7 (2012), pp. 7454–7468.
- [66] Thomas J Johnson, Matthew Borselli, and Oskar Painter. “Self-induced optical modulation of the transmission through a high-Q silicon microdisk resonator”. In: *Optics express* 14.2 (2006), pp. 817–831.
- [67] Richard A Soref and Brian R Bennett. “Electrooptical effects in silicon”. In: *IEEE journal of quantum electronics* 23.1 (1987), pp. 123–129.
- [68] Q Lin, Oskar J Painter, and Govind P Agrawal. “Nonlinear optical phenomena in silicon waveguides: modeling and applications”. In: *Optics express* 15.25 (2007), pp. 16604–16644.
- [69] Milos Nedeljkovic, Richard A Soref, and Goran Z Mashanovich. “Free-Carrier Electrorefraction and Electroabsorption Modulation Predictions for Silicon Over the 1–14- μm Infrared Wavelength Range”. In: *IEEE Photonics Journal* 3.6 (2011), pp. 1171–1180.
- [70] Wolfram HP Pernice et al. “Carrier and thermal dynamics of silicon photonic resonators at cryogenic temperatures”. In: *Optics express* 19.4 (2011), pp. 3290–3296.
- [71] Qianfan Xu and Michal Lipson. “All-optical logic based on silicon microring resonators”. In: *Optics express* 15.3 (2007), pp. 924–929.
- [72] Vilson R Almeida et al. “All-optical control of light on a silicon chip”. In: *Nature* 431.7012 (2004), pp. 1081–1084.
- [73] Lian-Wee Luo et al. “Power insensitive silicon microring resonators”. In: *Optics letters* 37.4 (2012), pp. 590–592.

- [74] Libin Zhang et al. “Multibistability and self-pulsation in nonlinear high-Q silicon microring resonators considering thermo-optical effect”. In: *Physical Review A* 87.5 (2013), p. 053805.
- [75] Massimo Borghi et al. “On the modeling of thermal and free carrier nonlinearities in silicon-on-insulator microring resonators”. In: *Optics Express* 29.3 (2021), pp. 4363–4377. DOI: 10.1364/oe.413572.
- [76] COMSOL AB, Stockholm, Sweden. *COMSOL Multiphysics* [®]. 2022. URL: <https://www.comsol.com>.
- [77] Ardavan F Oskooi et al. “MEEP: A flexible free-software package for electromagnetic simulations by the FDTD method”. In: *Computer Physics Communications* 181.3 (2010), pp. 687–702.
- [78] Lumerical Inc. *Lumerical* [®]. 2022. URL: <https://www.lumerical.com/products/>.
- [79] Thomas Van Vaerenbergh et al. “Simplified description of self-pulsation and excitability by thermal and free-carrier effects in semiconductor microcavities”. In: *Physical Review A* 86.6 (2012), p. 063808.
- [80] Joyce K. Poon et al. “Transmission and group delay of microring coupled-resonator optical waveguides”. In: *Optics Letters* 31.4 (Feb. 2006), p. 456. DOI: 10.1364/ol.31.000456.
- [81] Andrea Melloni and Mario Martinelli. “Synthesis of direct-coupled-resonators bandpass filters for WDM systems”. In: *Journal of Lightwave Technology* 20.2 (2002), p. 296.
- [82] Mattia Mancinelli et al. “Interferometric switching in coupled resonator optical waveguides-based reconfigurable optical device”. In: *Optics letters* 38.2 (2013), pp. 217–219.
- [83] Mattia Mancinelli et al. “Optical characterization of a SCISSOR device”. In: *Optics Express* 19.14 (2011), pp. 13664–13674.
- [84] A Calabrese et al. “Unidirectional reflection from an integrated “taiji” microresonator”. In: *Photonics Research* 8.8 (2020), pp. 1333–1341.
- [85] Mantas Lukoševičius and Herbert Jaeger. “Reservoir computing approaches to recurrent neural network training”. In: *Computer Science Review* 3.3 (2009), pp. 127–149.
- [86] Réka Albert and Albert-László Barabási. “Statistical mechanics of complex networks”. In: *Reviews of modern physics* 74.1 (2002), p. 47.
- [87] Robert Hull. *Properties of crystalline silicon*. 20. IET, 1999.
- [88] Yurii A Vlasov and Sharee J McNab. “Losses in single-mode silicon-on-insulator strip waveguides and bends”. In: *Optics express* 12.8 (2004), pp. 1622–1631.
- [89] Mancinelli Mattia. “Linear and non linear coupling effects in sequence of microresonators”. PhD thesis. University of Trento, 2013.

- [90] Mattia Mancinelli et al. “A photonic complex perceptron for ultrafast data processing”. In: *Scientific Reports* 12.1 (Mar. 2022), p. 4216. ISSN: 2045-2322. DOI: 10.1038/s41598-022-08087-2.
- [91] Jagdish Chand Bansal, Pramod Kumar Singh, and Nikhil R Pal. *Evolutionary and swarm intelligence algorithms*. Springer, 2019.
- [92] Massimo Borghi, Stefano Biasi, and Lorenzo Pavesi. “Reservoir computing based on a silicon microring and time multiplexing for binary and analog operations”. In: *Scientific Reports* 11.1 (2021), pp. 1–12.
- [93] Stefano Biasi et al. “Time response of a microring resonator to a rectangular pulse in different coupling regimes”. In: *Journal of Lightwave Technology* 37.19 (2019), pp. 5091–5099.
- [94] Thomas Van Vaerenbergh et al. “Cascadable excitability in microrings”. In: *Optics express* 20.18 (2012), pp. 20292–20308.
- [95] Davide Bazzanella et al. “A Microring as a Reservoir Computing Node: Memory/Nonlinear Tasks and Effect of Input Non-ideality”. In: *arXiv preprint arXiv:2203.07045* (2022).
- [96] Andrei Nikolaevich Tikhonov et al. *Numerical methods for the solution of ill-posed problems*. Vol. 328. Springer Science & Business Media, 1995.
- [97] Joni Dambre et al. “Information processing capacity of dynamical systems”. In: *Scientific reports* 2.1 (2012), pp. 1–7.
- [98] Masanobu Inubushi and Kazuyuki Yoshimura. “Reservoir computing beyond memory-nonlinearity trade-off”. In: *Scientific reports* 7.1 (2017), pp. 1–10.
- [99] Matthew Dale et al. “A substrate-independent framework to characterize reservoir computers”. In: *Proceedings of the Royal Society A* 475.2226 (2019), p. 20180723.
- [100] Davide Bazzanella et al. “PRECISE Photonic Hybrid Electromagnetic Solver”. In: *IEEE Photonics Journal* 14.3 (2022), pp. 1–10. DOI: 10.1109/JPHOT.2022.3168157.
- [101] Lucas H Gabrielli. *Python module for creating GDSII stream files, usually CAD layouts*. URL: <https://github.com/heitzmann/gdspys>.
- [102] Ronald Broeke and Xaveer Leijtsens. *Nazca design, photonic IC design framework*. URL: <https://nazca-design.org/>.
- [103] Richard J Bojko et al. “Electron beam lithography writing strategies for low loss, high confinement silicon optical waveguides”. In: *Journal of Vacuum Science & Technology B, Nanotechnology and Microelectronics: Materials, Processing, Measurement, and Phenomena* 29.6 (2011), 06F309.
- [104] Jing Zhang et al. “III-V-on-Si photonic integrated circuits realized using micro-transfer-printing”. In: *APL photonics* 4.11 (2019), p. 110803.

- [105] Yuting Shi et al. “Optical pumped InGaAs/GaAs nano-ridge laser epitaxially grown on a standard 300-mm Si wafer”. In: *Optica* 4.12 (2017), pp. 1468–1473.
- [106] Yuting Shi et al. “Novel adiabatic coupler for III-V nano-ridge laser grown on a Si photonics platform”. In: *Optics Express* 27.26 (2019), pp. 37781–37794.
- [107] Muhammad Rodlin Billah et al. “Hybrid integration of silicon photonics circuits and InP lasers by photonic wire bonding”. In: *Optica* 5.7 (2018), pp. 876–883.
- [108] Matthias Köfferlein. *Klayout*. URL: <https://klayout.de/>.
- [109] *Photonic integrated circuit prototyping and small volume production*. Accessed 2022-03-06. 2019. URL: https://www.imec-int.com/drupal/sites/default/files/2019-03/IC-link%20integrated%20si-photonics%20platform_web.pdf.
- [110] Runkui Yao et al. “Compact and Low-Insertion-Loss $1 \times N$ Power Splitter in Silicon Photonics”. In: *Journal of Lightwave Technology* 39.19 (2021), pp. 6253–6259.
- [111] Lukas Chrostowski et al. “Design and simulation of silicon photonic schematics and layouts”. In: *Silicon Photonics and Photonic Integrated Circuits V*. Vol. 9891. International Society for Optics and Photonics. 2016, p. 989114.