

Department of Physics and Astronomy
Heidelberg University

Bachelor Thesis in Physics
submitted by

Yasen Miroslavov Yanev

born in Sofia (Bulgaria)

2026

Characterisation of a Content Addressable Memory Realised on a Photonic Integrated Circuit

This Bachelor Thesis has been carried out by Yassen Miroslavov Yanev at the
Kirchoff Institute for Physics (KIP) in Heidelberg
under the supervision of
Prof. Wolfram Pernice

This work is licensed under a Creative Commons
“Attribution-NonCommercial-NoDerivs 3.0 Unported”
license.



Abstract

The exponential growth of data-intensive workloads, from network packet routing to machine-learning inference, has exposed the fundamental scaling limits of electronic content addressable memories (CAMs), whose lookup throughput is increasingly bottlenecked by interconnect delays and capacitive loading. Microring resonator-based photonic CAMs promise picosecond-scale parallel search by exploiting light’s inherent capacity for massive parallelism, yet existing demonstrations encode entire words within a single optical channel, capping storable memory at the free spectral range of the underlying resonators.

Here we report the design, fabrication and characterisation of a network of four 4-bit silicon-photonic optical CAMs on the imec iSiPP50G platform, in which each microring resonator is tuned to one of two wavelengths to encode a single binary bit, decoupling memory capacity from spectral bandwidth. The 4×4 crossbar integrates microring memory cells, electro-absorption modulators and germanium photodetectors on a single chip, and is interfaced to a custom radio-frequency carrier board. All constituent devices meet or exceed published benchmarks, and end-to-end AC operation of the full optical signal chain is demonstrated. We identify thermal crosstalk, resonance drift and photocurrent non-uniformity as the dominant system-level limitations, and outline closed-loop stabilisation and per-column thresholding as concrete routes to a fully operational query, establishing a scalable architecture for read-intensive photonic memory search.

Kurzfassung

Das exponentielle Wachstum datenintensiver Anwendungen, von der Paketweiterleitung in Netzwerken bis zur Inferenz in maschinellen Lernverfahren, hat die fundamentalen Skalierungsgrenzen elektronischer inhaltsadressierbarer Speicher (CAMs) offengelegt, deren Suchdurchsatz zunehmend durch Verbindungsverzögerungen und kapazitive Lasten begrenzt wird. Auf Mikroringresonatoren basierende photonische CAMs versprechen parallele Suchoperationen auf Pikosekunden-Zeitskalen, indem sie die inhärente Fähigkeit des Lichts zur massiven Parallelisierung ausnutzen, doch kodieren bisherige Demonstrationen ganze Wörter in einem einzigen optischen Kanal, wodurch die speicherbare Kapazität durch den freien Spektralbereich der zugrundeliegenden Resonatoren beschränkt wird.

In dieser Arbeit präsentieren wir den Entwurf, die Fabrikation und die Charakterisierung eines Netzwerks aus vier 4-Bit silizium-photonischen optischen CAMs auf der imec iSiPP50G-Plattform, in dem jeder Mikroringresonator auf eine von zwei Wellenlängen abgestimmt ist, um ein einzelnes Binärbit zu kodieren, wodurch die Speicherkapazität von der spektralen Bandbreite entkoppelt wird. Das 4×4 Kreuzschienenfeld integriert Mikroring-Speicherzellen, Elektroabsorptionsmodulatoren und Germanium-Photodetektoren auf einem einzigen Chip, der mit einer maßgeschneiderten Hochfrequenz-Trägerleiterplatte verbunden ist. Sämtliche Bauelemente erreichen oder übertreffen publizierte Referenzwerte, und der durchgehende AC-Betrieb der vollständigen optischen Signalkette wird demonstriert. Als wesentliche systemische Limitierungen identifizieren wir thermisches Übersprechen, Resonanzdrift und Inhomogenität der Photoströme und skizzieren geschlossene Regelkreise zur Stabilisierung sowie spaltenweise Detektionsschwellen als konkrete Wege zu einer vollständig funktionsfähigen Abfrage, womit eine skalierbare Architektur für leseintensive photonische Speichersuche etabliert wird.

Contents

Abstract	i
Kurzfassung	iii
Contents	v
List of Figures	vii
List of Tables	ix
Introduction	1
1 Theory	3
1.1 Optics	3
1.1.1 Waveguides	3
1.1.2 Couplers	5
1.1.3 Ring Resonators	6
1.1.4 Photodiodes	10
1.1.5 Electro-Absorption Modulators	10
1.1.6 Multimode Interferometers	11
1.2 Electronics	12
1.2.1 Radio-Frequency Signalling	12
1.2.2 Field-Programmable Gate Array	13
2 Design	14
2.1 Content Addressable Memory	14
2.2 Printed Circuit Board	16
2.3 Test Structures	17
2.4 Summary of Relevant Figures	20
3 Experimental Setup and Methods	22
3.1 System setup	22
3.2 Device Fabrication	23
4 Experimental Results	24
4.1 Test Structures	24
4.1.1 Multimode Interferometer and Grating Coupler	24

4.1.2	Electro-Absorption Modulators	25
4.1.3	Ring Resonators	28
4.1.4	Photodiodes	31
4.1.5	Summary of Test Structure Results	33
4.2	Printed Circuit Board	34
4.3	Content Addressable Memory	36
4.3.1	Photodiodes	36
4.3.2	Ring Resonators	37
4.3.3	Electro-Absorption Modulators	42
4.3.4	AC Operation	43
5	Discussion	44
	Bibliography	49
	Acknowledgements	53

List of Figures

1.1	Total internal reflection inside a waveguide.	3
1.2	Simulated intensity distribution of TE mode.	5
1.3	Schematic of a directional coupler.	6
1.4	Add-Drop Ring Resonator	7
1.5	Add-drop and all-pass ring resonator spectra.	9
2.1	Layout of the entire CAM.	14
2.2	Layout of the input section of one CAM row.	15
2.3	Layout of one CAM cell.	16
2.4	3D rendering of the PCB.	17
2.5	Simplified PCB circuit designs.	18
2.6	Layout of passive test structures.	19
2.7	Layout of active test structures.	19
3.1	Experimental setup.	22
4.1	MMI and GC characterisation plots.	25
4.2	Reference and EAM spectra for different bias voltages.	26
4.3	EAM Transmissions	26
4.4	Central performance benchmarks of EAM.	27
4.5	Ring resonator spectrum with Lorentzian fits	28
4.6	Heating response of ring resonator.	29
4.7	Properties of resonances on ring resonators.	30
4.8	Photodiode characterisation results obtained from test chip.	32
4.9	S21 PCB results for different design parameters.	34
4.10	Transient S22 PCB results for different design parameters.	35
4.11	Dark current of photodiodes on CAM chip	36
4.12	Effects of heating cross talk	37
4.13	Thermal Stability of Ring Overnight	38
4.14	Compensated Resonance Peaks from compensated splitting ratios	39
4.15	Heatmaps showing induced photocurrent by individual rings	41
4.16	Scatter Plot with Linear Regression Line for EAM and Photocurrents	41
4.17	Characterization of EAMs on CAM chip.	43

List of Tables

2.1	Published and computed reference values for the characterised on-chip devices.	21
4.1	Measured device performance compared to published reference values. . .	33
4.2	Thermal cross-talk when distant heater is driven at 1 V	38

Introduction

The demand for fast, energy-efficient memory search hardware has grown considerably alongside the expansion of data-intensive workloads such as packet routing in high-speed networks and hardware-accelerated machine learning. A Content Addressable Memory (CAM) is a type of memory circuit that implements a fully parallel search at its core. Rather than retrieving a value at a given address, a CAM accepts an input query and returns the address at which that value is stored, completing the comparison across all stored words simultaneously [1]. This inherent parallelism makes CAMs advantageous in any context where lookup speed is the primary bottleneck.

Electronic CAMs (E-CAMs) based on complementary metal-oxide-semiconductor (CMOS) technology are already mature and widely deployed, benefiting from decades of sustained process scaling. However, as clock speeds approach the limits imposed by interconnect delays and capacitive loading, the performance advantages of conventional electronics begin to erode [2]. Photonic integrated circuits (PICs) present a compelling alternative, due to light's ability to propagate without resistive losses and the massive parallelism enabled by wavelength-division multiplexing (WDM). The combination of these properties results in communication delays on the order of tens to hundreds of picoseconds at chip scale [1]. Microring resonators (MRRs) are attractive candidates for optical memory stores due to their small footprint and low energy consumption. Several photonic CAM architectures on the basis of MRRs have been proposed [1–3]. In simulation, these designs achieve read speeds that are an order of magnitude faster compared to conventional E-CAMs, but many of them encode entire words in a single optical channel, which places an upper bound on storable memory determined by the free spectral range (FSR) of the resonators.

This thesis aims to characterise a silicon-photonic optical CAM (O-CAM) using a novel architecture that avoids the FSR constraint by tuning each MRR to one of two wavelengths, encoding a binary bit. The architecture is realised in a PIC fabricated on the imec iSiPP50G silicon-on-insulator platform, and implements a 4×4 crossbar array of MRRs as memory stores, electro-absorption modulators (EAMs) as query encoders, and germanium photodiodes (PDs) as match detectors, constituting the first demonstration of a network of 4 4-bit O-CAMs. The PIC is designed to be driven by a field-programmable gate array (FPGA) with a custom-designed carrier printed circuit board (PCB) for electrical interfacing. Additional structures are manufactured on different PICs to enable the characterisation of the individual components.

Chapter 1 introduces the necessary theoretical foundations for understanding the operation of the O-CAM architecture, with a special focus on the properties of MRRs. Chapter 2 details the design of the PCB and both O-CAM and test-structure PICs, and provides an overview of the expected performance figures for all components. Chapter 3 describes the experimental setup and fabrication methods employed in our work. Chapter 4 presents all measurement results. Finally, Chapter 5 provides a summary and discussion of the results, and identifies the remaining steps towards a complete demonstration of the O-CAM query operation.

Theory

This chapter introduces the physical principles underlying the optical and electronic components in the employed CAM architecture.

1.1 Optics

The field of optics forms the backbone from which we derive the central working principles of the CAM, thus the majority of the chapter is devoted to its review. The theory described herein can be found in many standard textbooks on the subject, but we base our presentation primarily on [4].

1.1.1 Waveguides

When light travels in homogeneous transparent media, the refractive index n can be used to quantify the increase in phase change per unit distance $k = n2\pi/\lambda$, where λ is the wavelength of light in vacuum. $2\pi/\lambda$ is the free-space wavenumber, sometimes denoted k_0 , thus n represents the factor by which the wavenumber is increased in the medium compared to vacuum.

Light can be guided in a spatially confined region through total internal reflection (TIR). When a propagating wave strikes an interface between two dielectric media at an angle greater than the critical angle $\theta_c = \arcsin(n_2/n_1)$, where $n_1 > n_2$ are the refractive indices of the two media, it is entirely reflected back into the denser medium with refractive index n_1 rather than partially transmitted. By surrounding a high-index core with a low-index cladding on all sides, a waveguide confines and guides light along its length. This process is illustrated in Figure 1.1.




Figure 1.1. Schematic showing light propagating via total internal reflection inside a waveguide. Solid lines represent plane waves. Dashed lines represent the corresponding wavefronts. The ray propagating from A to B is reflected twice while the one from C to D is not reflected at all. The angle θ is greater than the critical angle determined by the refractive indices of the core and cladding, n_{core} and n_{cl} , respectively. Figure adapted from [5].

Guided Modes

Solving Maxwell's equations in a waveguide with the appropriate boundary conditions at the dielectric interfaces yields a discrete set of solutions, each corresponding to a specific transverse field distribution that propagates along the waveguide axis without changing its shape. These solutions are the so-called guided modes of the waveguide. Each mode is characterised by a propagation constant $\beta = 2\pi n_{\text{eff}}/\lambda$, where n_{eff} is the mode's effective refractive index. n_{eff} is a quantity between n_2 and n_1 that captures how strongly the mode is confined. Higher-order modes have smaller n_{eff} and a less confined transverse profile. They approach cut-off and become unguided as the waveguide becomes narrower. Physically, each mode can be understood as a distinct pattern of standing waves across the transverse cross-section, superimposed on propagation along the guide axis.

In most photonic integrated circuit applications, single-mode operation is strongly preferred, as different modes propagate at different phase velocities, leading to dispersive pulse broadening and unpredictable interference at junctions and couplers. Single-mode operation is achieved by reducing the waveguide cross-section until only the fundamental guided mode remains below cut-off. The precise threshold depends on the wavelength, the core geometry, and the index contrast between core and cladding.

The guided modes of a waveguide are further classified by their polarisation. In a planar dielectric waveguide, two orthogonal families of solutions arise from Maxwell's equations: transverse-electric (TE) modes, in which the dominant electric-field component lies parallel to the plane of the slab, and transverse-magnetic (TM) modes, in which the magnetic-field component takes that orientation instead [6]. In the rectangular strip waveguides used in photonic integrated circuits, both transverse dimensions are finite and the geometry mixes the two polarisations, so the modes are more precisely termed quasi-TE and quasi-TM according to which field component dominates. In silicon-on-insulator (SOI) circuits, the fundamental quasi-TE mode is the standard operating mode. The large refractive-index contrast between the silicon core ($n \approx 3.48$) and the silicon dioxide cladding ($n \approx 1.44$) yields a markedly different effective index and mode-field distribution for the two polarisation families, and most passive and active components are designed specifically for quasi-TE excitation.

Even a perfectly single-mode waveguide introduces some propagation loss due to residual absorption and, more significantly in high-index-contrast systems, scattering from sidewall roughness introduced during fabrication [7].

A related consequence of TIR is the evanescent field. Even when totally internally reflected, the electromagnetic field does not vanish abruptly at the core-cladding interface due to the boundary conditions imposed by Maxwell's equations. Instead, it penetrates a short distance into the cladding as an exponentially decaying wave. The characteristic penetration depth

is set by an extinction coefficient $\gamma = k_0 \sqrt{n_{\text{eff}}^2 - n_2^2}$, confining the evanescent tail to within a fraction of the wavelength. Although the evanescent field carries no net optical power in the transverse direction, it plays a central role whenever two waveguides are placed in close proximity, as will be discussed in the following section. Figure 1.2 shows a simulated field distribution of a TE mode in a waveguide along with the evanescent field which extends into the cladding.




Figure 1.2. Simulated intensity distribution of TE mode inside a strip waveguide. The waveguide is made of Si_3N_4 and surrounded by SiO_2 from below (substrate) and air from above (cladding). The field that reaches beyond the core is the evanescent field. Figure taken from [5].

1.1.2 Couplers

Grating Couplers

A grating coupler (GC) is a periodic corrugation etched into the surface of a waveguide that diffracts light between a guided mode and a free-space beam, typically from an optical fibre held at an oblique angle to the chip surface. Efficient coupling requires that the in-plane component of the incident beam's wavevector match the propagation constant β of the guided mode. A grating of period Λ satisfies this by contributing a discrete spatial-frequency increment. At an incidence angle θ from the surface normal and diffraction order m , the phase-matching condition reads $\beta = k_0 n_2 \sin \theta + m \cdot 2\pi / \Lambda$. Choosing Λ and θ appropriately centres the coupling on a target wavelength. GCs are widely used in photonic integrated circuits because they permit coupling at any point along the chip surface and are compatible with automated fibre-array alignment. Their principal limitation is a wavelength-selective coupling efficiency, due to which the response follows a roughly Gaussian spectral envelope centred on the design wavelength, with a typical 3 dB bandwidth of 30-50 nm for operation on silicon-on-insulator in the conventional spectral range (C-band) of 1530-1565 nm.

Directional Couplers

A directional coupler (DC) is formed by bringing two waveguides into close proximity over a coupling length L_c , so that their evanescent fields overlap. Power is exchanged sinusoidally between the two waveguides as a function of propagation distance. The coupling coefficient per unit length is determined by the degree of field overlap, which depends on the waveguide geometry and the separation between them [6]. Figure 1.3 (A) shows a schematic of a DC along with the field distribution in the coupling region.

At a coupling length corresponding to a quarter of the sinusoidal power-transfer period, half the input power is distributed to each output, constituting a 50:50 splitter. By choosing a different coupling length, any arbitrary splitting ratio can in principle be realised. This point is emphasised by Figure 1.3 (B), which shows simulated transmission values at both outputs as a function of coupling length.




Figure 1.3. (A) Schematic of a directional coupler. The inset heatmap shows the distribution of the electric field along the two waveguides in the coupled region. Light can travel between the waveguides through evanescent coupling. (B) Simulated transmission for a particular waveguide geometry at both outputs as a function of coupler length. A range of splitting ratios can be achieved by varying the length. Figure taken from [5].

1.1.3 Ring Resonators

A ring resonator consists of a closed waveguide loop evanescently coupled to one or more bus waveguides. Light circulating in the ring accumulates phase on each round trip and constructive interference builds up significant circulating power when the round-trip optical path length equals an integer multiple of the wavelength. This resonance condition is

$$n_{\text{eff}}(\lambda) L = m \lambda, \quad m \in \mathbb{N}, \quad (1.1)$$

where L is the total circumference of the ring, and m is the resonance order [7]. In practice, the ring geometry is often extended into a racetrack shape, in which two parallel straight sections replace part of the curvature.

Add-Drop Configuration

A ring resonator coupled to a single bus waveguide constitutes an all-pass filter. When coupled to two separate bus waveguides — one carrying the input (referred to as the through port) and one routing out the resonant light (referred to as the drop port) — the device operates as an add-drop filter. The ring selectively diverts light at its resonant wavelength from the input bus into the drop port, while off-resonant light passes through largely unaffected. An MRR in this configuration can be seen in Figure 1.4.




Figure 1.4. Schematic of an add-drop ring resonator. r_1 and r_2 denote the self-coupling coefficients of the two directional couplers, k_1 and k_2 the cross-coupling coefficients, and a the round-trip field transmission of the ring. Input light is fed into the input port. At resonance, the light exits at the drop port, otherwise it exits at the through port. Figure adapted from [7].

Denoting the field self-coupling coefficients of the two directional couplers as r_1 and r_2 , the cross-coupling coefficients $k_1^2 = 1 - r_1^2$ and $k_2^2 = 1 - r_2^2$, and the round-trip field transmission of the ring as a , the through-port and drop-port power transmissions as a function of the single-pass phase ϕ are [7]

$$T_{\text{through}}(\phi) = \frac{r_1^2 + r_2^2 a^2 - 2 r_1 r_2 a \cos \phi}{1 + r_1^2 r_2^2 a^2 - 2 r_1 r_2 a \cos \phi}, \quad (1.2)$$

$$T_{\text{drop}}(\phi) = \frac{k_1^2 k_2^2 a^2}{1 + r_1^2 r_2^2 a^2 - 2 r_1 r_2 a \cos \phi}. \quad (1.3)$$

Both expressions reduce to Lorentzian lineshapes close to resonance. If the attenuation is negligible ($a \approx 1$), critical coupling, i.e. the coupling at which the transmission at the through port drops to zero, occurs at symmetric coupling ($k_1 = k_2$). For a lossy resonator, critical coupling occurs when the losses match the coupling as $r_2 a = r_1$. At resonance ($\phi = 2\pi m$), T_{drop} is maximised and T_{through} is minimised. The width of each resonance feature is set by the total round-trip loss, which includes both the intrinsic propagation loss, quantified by $1 - a^2$, and the power extracted at the two couplers, given by $k_1^2 + k_2^2$.

As power is transferred from the waveguide to the ring via a DC, the coupling coefficient k for a given coupling region can be estimated by

$$k^2 = \sin^2(\kappa L_c + \kappa_0), \quad (1.4)$$

with coupling length at which the two waveguides are parallel L_c , coupling coefficient per unit distance κ and offset coupling coefficient due to bend sections κ_0 [7]. κ and κ_0 are both wavelength and geometry dependent, and generally constitute empirical parameters.

Free Spectral Range, Quality Factor, and Extinction Ratio

Adjacent resonances are separated in wavelength by the free spectral range (FSR). Because n_{eff} varies with wavelength, a pulse envelope propagating in the waveguide travels not at the phase velocity but at the group velocity $v_g = d\omega/d\beta$. The group index $n_g = c/v_g = n_{\text{eff}} - \lambda dn_{\text{eff}}/d\lambda$ therefore replaces n_{eff} in the FSR expression, since successive resonances are spaced by the round-trip phase accumulated at the group delay rather than the phase delay:

$$\text{FSR} = \frac{\lambda^2}{n_g L}. \quad (1.5)$$

The quality factor Q of a resonance quantifies its spectral sharpness as the ratio of the resonant wavelength to the full width at half maximum (FWHM) $\delta\lambda$:

$$Q = \frac{\lambda_{\text{res}}}{\delta\lambda}. \quad (1.6)$$

The finesse $\mathcal{F} = \text{FSR}/\delta\lambda$ is an equivalent measure on the scale of one FSR. It is related to Q by $Q = \mathcal{F} \lambda_{\text{res}}/\text{FSR}$. For the add-drop ring resonator with low round-trip loss and weak coupling, the finesse can be written as [7]

$$\mathcal{F} = \frac{\pi \sqrt{R}}{1 - R}, \quad R = r_1 r_2 a, \quad (1.7)$$

where R is the product of all amplitude factors accumulated over one round trip and one pass through each coupler. Physically, $\mathcal{F}$ represents the number of round-trips of the light, and Q represents the number of oscillations of the field before the energy in the ring decays to $1/e$ of its initial value.

The extinction ratio (ER) at the drop port, defined as the ratio of the maximum to minimum transmitted power across the spectrum, is [7]

$$\text{ER}_{\text{drop}} = \left(\frac{1 + R}{1 - R} \right)^2. \quad (1.8)$$

This expression grows steeply with R , so a ring with high finesse simultaneously exhibits a high drop-port ER. In practice the achieved ER often falls below the lossless theoretical prediction due to wavelength-dependent coupling coefficients and residual propagation loss. Some of the spectral features discussed up to this point are illustrated in Figure 1.5, along with their corresponding transmission spectra.




Figure 1.5. Transmission spectrum of an all-pass ring and the two outputs of an add-drop ring with indicated ER, FWHM and transmission maximum T . The subscripts n, p, d denote all-pass output, through port and drop port, respectively. Figure taken from [7].

Resonance Tuning

The resonant wavelengths of the ring follow from equation (1.1) and are therefore sensitive to any change in n_{eff} . Two primary mechanisms are used to shift resonances in integrated photonic platforms. Thermal tuning exploits the thermo-optic effect, whereby passing current through a resistive heater element placed near the waveguide raises the local temperature and shifts the resonance through the temperature-dependent change in refractive index [8]. Thermal tuners offer a shift that is linear in dissipated power and are straightforward to implement, but their response time is limited by heat diffusion to the range of microseconds to milliseconds, making them unsuitable for high-frequency modulation [9]. Carrier-based tuning via a reverse- or forward-biased p-n junction modulates n_{eff} through the plasma dispersion effect and can achieve GHz-scale speeds, though at the cost of additional optical loss through free-carrier absorption.

Environmental Sensitivity

The same thermo-optic sensitivity that enables deliberate tuning also makes ring resonators susceptible to ambient temperature fluctuations. A small temperature drift can shift a resonance by a substantial fraction of its FWHM, introducing insertion loss that grows as the ring detunes from the target wavelength [9, 10].

1.1.4 Photodiodes

A photodiode is a semiconductor device that converts optical power into an electrical current. In the band model for semiconductors, the allowed electron energies form two broad bands separated by a forbidden energy gap of width E_g . The two bands are the valence band, which is fully occupied at absolute zero, and the conduction band, which is empty at absolute zero. At finite temperature, a small number of electrons are thermally excited across the gap, leaving behind positively charged vacancies called holes.

When a photon of energy $\hbar\omega \geq E_g$ is absorbed by the semiconductor, it promotes an electron from the valence band into the conduction band, generating an electron-hole pair. Under an applied electric field, the electron and hole are swept in opposite directions, constituting a photocurrent that flows through the external circuit. The magnitude of E_g determines the longest wavelength at which a material can absorb. Operation at telecom wavelengths near 1550 nm (photon energy ≈ 0.80 eV), requires materials with a band gap at or below this value. Germanium, whose direct band gap of approximately 0.8 eV [6] falls precisely at this threshold, is well suited to near-infrared photodetection.

In a practical waveguide-integrated photodiode, the absorbing material is placed between p-type and n-type semiconductor contact regions, forming a p-i-n structure. A reverse bias applied across the p-i-n junction creates a depletion region which is exposed to a strong electric field. Photo-generated carriers inside the depletion region are swept away before they can recombine, improving both the collection efficiency and the response bandwidth. The two central figures of merit are the responsivity $\mathcal{R}$, which relates the photocurrent I_{ph} to the incident optical power P_{opt} ,

$$\mathcal{R} = \frac{I_{\text{ph}}}{P_{\text{opt}}} \quad [\text{A/W}], \quad (1.9)$$

and the dark current I_{dark} , which is the residual leakage current that flows in the absence of any illumination. Dark current arises from thermally generated carriers and sets the noise floor of the detector. It increases with reverse bias and with temperature. High responsivity and low dark current are therefore jointly desirable for sensitive photodetection.

1.1.5 Electro-Absorption Modulators

An electro-absorption modulator (EAM) is structurally a semiconductor diode, utilising the same p-i-n geometry as a photodiode. The underlying physical mechanism is the Franz-Keldysh effect, which dictates that in a bulk semiconductor, the application of an electric field across the absorption region tilts the energy bands, allowing photons with energies slightly below E_g to be absorbed through field-assisted tunnelling into the conduction band [6]. The practical consequence is that the absorption edge shifts to longer wavelengths with increasing reverse bias, so that a material that is transparent at a particular wavelength under

zero field becomes progressively more absorptive as the field is raised. By modulating the bias voltage, the transmitted optical power can be switched between a high-transmission state and a low-transmission state.

The two central performance measures for an EAM are the insertion loss (IL) and the extinction ratio (ER). The IL is the optical attenuation introduced at zero applied bias. The ER is the contrast ratio in transmission $T(\cdot)$ between an applied reverse bias V_1 and 0 V:

$$\text{ER} = 10 \log_{10} \left(\frac{T(0 \text{ V})}{T(V_1)} \right) \quad [\text{dB}], \quad (1.10)$$

A high ER indicates a large voltage-controlled absorption, while a low IL implies little penalty on the optical power budget. These two quantities pull in opposite directions, since a device with a higher ER will generally exhibit a larger IL, e.g. because both scale with device length. [11]. A natural figure of merit (FOM) for capturing this trade-off is therefore

$$\text{FOM} = \frac{\text{ER} [\text{dB}]}{\text{IL} [\text{dB}]}. \quad (1.11)$$

When an EAM is operated with an AC modulation signal of peak-to-peak voltage V_{pp} superimposed on a direct current (DC) bias V_{DC} , the relevant contrast is not between 0 V and a fixed voltage, but between the two extremes of the modulation swing. This quantity, which we call the effective extinction ratio,

$$\text{ER}_{\text{eff}} = 10 \log_{10} \left(\frac{T(V_{\text{DC}} + V_{\text{pp}}/2)}{T(V_{\text{DC}} - V_{\text{pp}}/2)} \right), \quad (1.12)$$

determines the achievable on/off contrast for a given combination of bias point and modulation amplitude.

1.1.6 Multimode Interferometers

A multimode interferometer (MMI) is based on the principle of self-imaging in a wide, multimode waveguide section. When a field is launched into such a section from a single-mode input waveguide, it excites many lateral modes simultaneously. Because each mode carries a different propagation constant, the relative phases between modes evolve along the propagation direction. At specific lengths determined by the beat lengths between mode pairs, the superposition of all modes reconstructs one or more replicas of the input field. For a 1×2 MMI, i.e. an MMI with one input and two outputs, the self-imaging condition produces two images of equal amplitude at half the fundamental beat length, realising a 50:50 power splitter. The same device acts as a 2×1 combiner when driven from both outputs. A key practical advantage over directional couplers is that MMI performance is largely insensitive to the precise waveguide width. The self-imaging length scales with the square of the multimode section width rather than with the exponentially sensitive evanescent overlap, which makes the device more tolerant to fabrication variations [12].

Denoting the input optical power of a 1×2 MMI as P_{in} and the output powers as P_1 and P_2 , the central performance metrics can be found as

$$\text{Excess Loss} = 10 \log_{10} \left(\frac{P_{\text{in}}}{P_1 + P_2} \right), \quad (1.13)$$

$$\text{Power Imbalance} = 10 \log_{10} \left(\frac{P_1}{P_2} \right). \quad (1.14)$$

Excess loss quantifies the total power lost in the device, while power imbalance captures the deviation from an ideal 50:50 split.

1.2 Electronics

Optoelectronic systems require electrical signals both for driving optical modulators and for reading out the generated photocurrents. The following subsections introduce the electronic concepts relevant to high-speed signal generation and acquisition in such systems.

1.2.1 Radio-Frequency Signalling

High-speed electro-optic systems place demanding requirements on the electrical signal path. Electro-absorption modulators can operate at bandwidths well in excess of 50 GHz [13], which means that the electrical transmission lines carrying the modulation signals must themselves preserve signal integrity at these frequencies.

At radio frequencies (RF), the behaviour of a transmission line cannot be described by simple DC resistance alone. The line is better modelled as a distributed network of series resistance R , series inductance L , shunt capacitance C , and shunt conductance G per unit length [14]. Signal attenuation along such a line arises from two dominant mechanisms, namely resistive (ohmic) loss, which increases with frequency through the skin effect because high-frequency currents are confined to a thin surface layer of the conductor, and dielectric loss in the substrate material, which also scales with frequency [14].

When a signal travelling along a transmission line encounters an impedance discontinuity, e.g. at a connector or at the load, part of the signal is reflected back toward the source, corrupting the transmitted waveform. To prevent this, all components in the signal path are typically impedance-matched [14]. This matching condition, which ensures that the characteristic impedance of the transmission line equals the source and load impedances, minimises reflections and maximises power transfer.

Many high-speed signal sources and acquisition systems rely on differential signalling, in which complementary signal pairs V^+ and V^- are transmitted on two separate lines such that the receiver responds to the difference $V^+ - V^-$. Common-mode noise and interference picked up along both lines cancel in the subtraction, offering substantially better noise immunity than single-ended signalling [14]. Where differential signals must interface with

single-ended circuits, a balun (from *balanced-unbalanced*) is used. A balun is a passive device, typically a transformer, that converts between the differential and single-ended domains while preserving signal integrity over the required frequency range.

1.2.2 Field-Programmable Gate Array

A field-programmable gate array (FPGA) is a reconfigurable integrated circuit consisting of an array of programmable logic blocks whose interconnections can be configured after manufacture, allowing the same physical device to implement a wide variety of digital and mixed-signal functions [15]. In photonic measurement and control systems, FPGAs are frequently used as the primary interface between a host computer and high-speed optoelectronic hardware.

Of particular relevance to such systems are RF-grade digital-to-analogue converters (RF-DACs) and analogue-to-digital converters (RF-ADCs). An RF-DAC converts a binary-encoded digital word into a continuous analogue voltage at sample rates that can reach several to tens of gigasamples per second (GSa/s), enabling the synthesis of arbitrary waveforms with spectral content well into the multi-GHz range. An RF-ADC performs the inverse, quantising an incoming analogue signal at comparably high sample rates and thereby capturing wideband signals at GHz frequencies. Both device types are increasingly co-integrated with FPGA logic fabric on a single chip, providing a compact platform for high-speed signal generation and acquisition, as seen in [16].

Design

The principal idea behind the design of the CAM is to encode the memory words in the resonant wavelengths of the MRRs, and to encode the search query in the wavelengths of the input light. If and only if the two are equal will light be transmitted through the drop port of the MRR, where it can be read out by a photodiode. This chapter describes the CAM design in greater detail, along with the designs of the PCB and test structures necessary for characterisation.

2.1 Content Addressable Memory

The CAM is based on a silicon-on-insulator photonic integrated circuit. The PIC consists of a 4×4 crossbar array, where each cell contains an MRR and a heating element used to tune the ring's resonances, along with EAMs for input modulation and on-chip photodiodes for output detection. The light enters the chip at each row through a grating coupler, and the photodiodes convert the result of the optical word search into an electrical current. Figure 2.1 shows the layout of the entire CAM, and Figures 2.2 and 2.3 show the layout of one input section and one cell, respectively.




Figure 2.1. Layout of the entire CAM. Light is coupled into the chip through the GCs on the left, before being selectively absorbed by the EAMs in the modulation section. A 4×4 array of MRRs stores the memory words, encoded through the voltages across the heaters in each cell. Final signal readout happens through the photodiodes at the bottom of the schematic. Electrical wiring of all optoelectronic components is omitted for clarity.

In order to facilitate GHz operation frequencies, each row input sees two lasers, each at a wavelength representing a different bit, either '0' or '1'. The desired bit for a given query is selected via the EAMs, which are modulated at a high frequency to maximally absorb light from the undesired laser input. After passing through the EAM, the light from the two

lasers is combined in a 1×2 MMI splitter and then routed to the crossbar array. Each row uses different wavelengths to represent the search words, in order to avoid interference on the way to the photodiodes.




Figure 2.2. Layout of the input section of one CAM row. Lasers at two different wavelengths representing a '0' and '1' bit couple into the chip through the GCs on the left. The modulating voltage across the EAMs is set with inverted polarity, such that the optical power in one arm is reduced significantly more than in the other, effectively selecting which bit to pass into the crossbar array. The two beams are then combined in a 1×2 MMI splitter before cascading further into the CAM. A 'don't care' bit can be implemented by applying a negative modulating voltage of equal magnitude to both EAMs. Electrical wiring of all optoelectronic components is omitted for clarity.

A horizontal bus waveguide distributes optical power evenly across each column through DCs designed to have a splitting ratio of $1/(4-j+1)$, where j is the column index, ensuring that $1/4$ of the input power reaches each column. Each ring is tuned to resonate at one of the two possible input wavelengths for that row, thus storing either '0' or '1' to memory, and a second DC routes the light from the drop port of the MRR along a vertical bus waveguide to the photodiode. Here, the splitting ratio is designed to be $1/i$ for row index i , thus the total optical power reaching a photodiode in the case of a match from a single input is $1/16$ of the input power, without accounting for the ILs of all components along the optical path. Additionally, inside each cell, there are two more GCs used for characterisation purposes. In total, a 4-bit search word query can be encoded by adjusting the absorption of the EAMs. The resonances of the 16 rings, set by the voltages across the heaters, store 4 different 4-bit words, one along each column. Assuming that each ring induces the same photocurrent I_{ph} , a match between query and memory is detected when the sum of the photocurrents at a column exceeds a predefined threshold current I_{th} , where its magnitude must satisfy

$$(3 + 10^{ER_{eff}/10}) I_{ph} < I_{th} \leq 4I_{ph}. \quad (2.1)$$

ER_{eff} is the effective extinction ratio of the EAMs as defined in equation (1.12), and it appears in (2.1) as a consequence of the non-vanishing contribution of the suppressed optical power at each input. The exact value of I_{th} within the bounds given by (2.1) depends upon




Figure 2.3. Layout of one CAM cell. Electrical wiring of all optoelectronic components is omitted for clarity. **(1)** MRR with thermal resonance tuner. The resonant wavelength of the ring determines the bit stored in the cell. **(2)** Horizontal bus waveguide, connecting the input section of the corresponding row to each of the four columns. **(3)** Horizontal DC designed to couple 1/4 of the input power into the MRR, regardless of the cell's position along the row. **(4)** Vertical bus waveguide, connecting the drop port of the MRR to the photodiode. **(5)** Vertical DC designed to couple 1/4 of the optical power transmitted through the drop port into the photodiode, regardless of the cell's position along the column. **(6)** GCs used only for characterisation.

the noise characteristics of I_{ph} , and the desired error rate.

Due to the low modulation speed of the thermal tuners, typically on the order of kHz, the CAM is designed to operate in a 'write-once' mode, whereby the stored memory words are intended to be left unchanged for the duration of the search operation. The high-speed capabilities of the PIC are derived primarily from the short response times of the EAMs.

Though this design constitutes a binary CAM, it can effectively also be used as a ternary CAM. Ternary memories can store an additional 'don't care' bit next to the typical binary bits. This 'don't care' state can easily be implemented here by minimising the absorption of both EAMs at a given input.

2.2 Printed Circuit Board

The CAM is designed to be used together with an FPGA, which requires the PIC to be mounted onto a custom carrier PCB for electrical connectivity. A 3D rendering of the PCB can be seen in Figure 2.4.

The FPGA is equipped with RF-DACs and RF-ADCs, which are used to modulate the EAMs and read out the PD currents, respectively. Additionally, dedicated DAC channels are used to regulate the voltage over the heaters, and to set all bias points. The PCB must therefore provide the appropriate connectors for the RF cables used in the FPGA. Moreover, because the RF-DACs and RF-ADCs rely on differential signalling to reduce common-mode interference, whereas the electrical terminals on the PIC require single ended transmission lines,




Figure 2.4. 3D rendering of the carrier PCB, showing (Left) front and (Right) back views. The PIC is to be glued onto the bare copper area roughly in the centre and wirebonded to the PCB, providing electrical connectivity. The design is optimised for RF together with an FPGA.

the PCB has to convert between the different signal types using baluns.

All transmission lines are length and impedance matched to $50\ \Omega$ to ensure signal integrity and to minimise reflections. The length of the EAM lines is set at ca. 72 mm, and that of the PD lines is ca. 130 mm. The smallest meander bend radius used to allow for length matching within the geometric constraints of the PCB is 0.44 mm. RF lines are left uncovered by solder resist to avoid discrepancies in the impedance. A grid of vias is laid over the entire PCB to improve ground connectivity. Additional vias are placed in the immediate vicinity of all RF lines to provide shielding.

There are in total three different circuits implemented multiple times on the PCB, each individually shown in Figure 2.5. Optoelectronic devices of the same type share a common bias line, with decoupling capacitors placed on the back of the PCB as close to the RF lines as possible to provide a return signal path. The PD circuit in particular is designed to accommodate the use of an off-PCB transimpedance amplifier (TIA) for photocurrent readout. The TIA amplifies the photocurrent and converts it to a voltage before feeding the signal into the RF-ADCs.

2.3 Test Structures

In addition to the CAM, a number of test structures are designed and analysed for the purposes of characterisation. Obtaining data from the test structures requires the design and manufacture of a further PCB without FPGA integration and RF capabilities. The discussion of the test-structure-PCB is omitted from this work due to its simplicity and similarity to the PCB described in section 2.2.




Figure 2.5. Simplified PCB circuit designs. Dashed red lines show portion of circuit that is on PIC. **Top:** Circuit for shifting the resonances of the MRRs via thermal tuners. A single ended line connects a slow DAC channel to one terminal of the heater via bondwire. The other terminal is set to ground. **Center:** Circuit for modulating the EAMs. A balun converts the differential signal from the RF-DAC to single ended. A custom-designed taper reduces the width of the line to the smallest the manufacturers can achieve, before the line connects to the anode of the EAM via bondwire. The bias is applied to the cathode of the EAM with an inverted polarity compared to the signal line. A decoupling capacitor parallel to the bias line ensures a stable bias voltage. On the physical PCB, the capacitor is placed as close as possible to the RF line, providing a return path. **Bottom:** Circuit for photocurrent readout. The basic layout is very similar to the one shown in **Center**. The balun now connects to an RF-ADC channel. Two SMA connectors are placed in the signal path between balun and taper, to accommodate the use of an off-PCB TIA for photocurrent amplification. Black arrows indicate the direction of the signal.




Figure 2.6. Layout of passive test structures. **Left:** Backloop between two GCs. There are two 90 deg bends in the waveguide. **Right:** Structure for MMI characterisation. Light couples into the left-most GC and is split into the two arms of the MMI. Readout happens through the GCs at the end of each arm.




Figure 2.7. Layout of active test structures. **(A)** Structure for EAM characterisation. Light enters through the middle GC and is split into the two arms of the MMI. The light exiting through the top GC is used as a reference, while the light exiting through the bottom GC must first pass through the EAM. The voltage across the EAM is set via electrical contact with the bond pads on the right. **(B)** Structure for PD characterisation. Light enters through the top GC and is split into the two arms of the MMI. The light exiting through the bottom GC is used as a reference, while the light passing through the other arm induces a photocurrent in the PD. There is an additional EAM in the optical path. The voltages across both diodes are set via electrical contact with the bond pads on the right. **(C)** Structure for Ring resonator characterisation. Light enters through the middle GC. Light exiting through the right-most GC must pass through the ring's drop port, whereas light exiting from the left-most GC must pass through the ring's through port. The ring's resonance is tuned by adjusting the voltage across the heater via electrical contact with the bond pads on the top.

There are in total 5 different test structures, each designed to deliver information about a specific component of the CAM. The two structures for characterising passive components are shown in Figure 2.6, and the three structures for characterising active components are shown in Figure 2.7. All test structures consist of a primary and reference optical path, with light coupling in and out of the PIC through GCs. The primary path contains the components of interest, whereas the reference path serves calibration purposes.

2.4 Summary of Relevant Figures

Some exemplary benchmarks for the devices studied in this work can be found in table 2.1. To ensure comparability, the figures were taken from papers which examine either the iSiPP25G or the iSiPP50G imec silicon-on-insulator photonics platform. However, it should be emphasised that these metrics constitute only a qualitative point of comparison, and are not to be understood as expected literature values for the results demonstrated in chapter 4. Some deviation is expected, stemming from variances in device configuration. As an example, the EAM analysed by Srinivasan et al [13] is one made of pure germanium and optimised for the L-band, whereas the EAMs employed in our chips consist of a germanium-silicon alloy, which results in a blueshift of the absorption edge into the C-band. Thus table 2.1 is only meant to serve as a sanity check to our measurements.

To obtain point-of-reference values for the performance metrics of the MRRs, we can use the equations presented in section 1.1.3 along with the geometry of the rings employed in our CAM design.

FSR, Q-factor and ER of MRR Derivation

At a wavelength of $\lambda = 1550$ nm, a typical group index for a single-mode SOI waveguide with TE polarisation is $n_g = 4.3$ [17]. The total circumference of the racetrack resonators we employ is $L = 300$ μm . At the chosen wavelength, from equation (1.5), we find

$$\text{FSR} = \frac{(1550 \text{ nm})^2}{4.3 \times 300 \mu\text{m}} \approx 1.86 \text{ nm}. \quad (2.2)$$

The coupling gap between the bus waveguide and the ring is about 200 nm, which, in SiO_2 cladding, corresponds to a coupling coefficient per unit length $\kappa \approx 0.11/\mu\text{m}$ and an offset coupling coefficient $\kappa_0 \approx 0.17$ [7]. The coupling length is $L_c = 1\mu\text{m}$, which allows us to estimate k^2 per coupler via (1.4)

$$k^2 = \sin^2(0.1 + 0.17) \approx 0.07. \quad (2.3)$$

It follows that $r^2 = 1 - k^2 \approx 0.93$ for each coupler. For negligible attenuation ($a \approx 1$), the round-trip product is $R = r_1 r_2 a = r^2 = 0.93$, giving a finesse from equation (1.7) of

$$\mathcal{F} = \frac{\pi\sqrt{0.93}}{1 - 0.93} = \frac{\pi \times 0.964}{0.07} \approx 43.3.$$

The quality factor is therefore

$$Q = 43.3 \times \frac{1550 \text{ nm}}{1.86 \text{ nm}} \approx 36\,100. \quad (2.4)$$

Finally, the extinction ratio at the drop port can be derived from equation (1.8) as

$$\text{ER}_{\text{drop}} = 10 \log_{10} \left(\frac{(1 + 0.93)^2}{(1 - 0.93)^2} \right) \approx 29 \text{ dB}. \quad (2.5)$$

Table 2.1. Published and computed reference values for the characterised on-chip devices.

Device	Quantity	Reference value	Ref.
PD	I_{dark} at -1 V [nA]	~ 3	[18]
	I_{dark} at -2 V [nA]	~ 8	[18]
	Responsivity R at -1 V , 1 mW [A/W]	> 1.0 (C-band)	[18]
EAM	Maximum FOM (0 V to -2 V)	~ 0.96	[13]
	Wavelength of maximum FOM [nm]	~ 1615	[13]
	IL at maximum FOM at 0 V [dB]	4.8	[13]
	ER at maximum FOM (0 V to -2 V) [dB]	4.6	[13]
Heater	Power required for π phase shift [mW]	23	[8]
Ring	Quality factor Q	36 100	2.4
	Extinction ratio [dB]	29	2.5
	Free spectral range [nm]	1.86	2.2

Experimental Setup and Methods

The goal of this chapter is to elucidate the experimental setup employed for all measurements on the structures introduced in chapter 2, and to provide details on the fabrication process of both PICs and PCBs.

3.1 System setup

The principal method for interacting with the PIC, namely through a fibre array for optical coupling and through a PCB for electrical interfacing, is identical across all measurements. An exemplary setup is shown in Figure 3.1.




Figure 3.1. Experimental setup used for measurements taken on the CAM PIC and PCB. Similar setups used for test-structure measurements as well. **(A)** Fiber array. **(B)** Glued and wirebonded PIC. **(C)** Carrier PCB with connectors for electrical interfacing. **(D)** 3D printed holder for the PCB. **(E)** Motorised optimisation stage. **(F)** Camera to control z-distance between fibre array and PIC. **(G)** Camera to control x-y-position of fibre array relative to PIC.

We employ a SANTEC TSL-770 tunable semiconductor laser as a light source for all optical measurements. The optical fibre from the laser passes through a Thorlabs polarisation controller before connecting to the fibre array, which is aligned to the GCs on the PIC. Alignment is achieved through a Piezo-motor controlled 4-axis optimisation stage. When necessary, output optical power is read out through a New Focus model 2011 photodiode connected to the fibre array.

DC signals are generated and read out using a Keithley Sourcemeter 2400. Vector network analyser (VNA) measurements are performed using a Rhode&Schwarz ZNL6. AC mea-

measurements are performed by modulating the EAM voltage using a Tektronix AFG model 3152C, and by implementing a FEMTO HSA-Y-1-60 TIA to amplify the photocurrent before displaying it on a RIL DS1202 oscilloscope.

The PIC is glued and wirebonded to the PCB, which itself is mounted on the optimisation stage via a 3D printed holder. Contact to the PCB is established through SMA cables and, where necessary, through Samtec ARC6 connectors and a custom breakout board.

3.2 Device Fabrication

The CAM is fabricated on imec's iSiPP50G silicon photonics platform. The design kit provides the monolithically integrated electro-optic devices used and studied in this work. The PCB is fabricated using Multi-CB's defined hybrid impedance pool with four layers, with the top and bottom layers having high frequency Roger material cores, and the two inner layers utilising the standard FR4 material. The PCB is designed to have a characteristic impedance of $50\ \Omega$. Differential lines are transitioned into single ended lines through Mini-Circuits' TCM2-43X+ baluns.

Experimental Results

In this chapter, we present the characterisation results of all devices from the test structures and then examine the performance of the devices when used in unison on the final CAM chip.

4.1 Test Structures

The results discussed in this section are derived from measurements made on the structures presented in section 2.3. We extract central performance figures and put them in relation to the metrics shown in Table 2.1. All data presented herein is collected by wire bonding the PICs containing the relevant structures to a custom built PCB, which is mounted onto a 4-axis calibration stage and optically contacted through a fibre array, as described in section 3.1.

4.1.1 Multimode Interferometer and Grating Coupler

The MMI splitters and grating couplers are discussed together in this subsection, as isolating the MMI performance from the raw measured spectrum requires an independent estimate of the GC insertion loss.

The GC spectrum was measured on an optical backloop test structure seen in Figure 2.6 **Left**. A backloop is a waveguide that enters and exits through two GCs without any intervening component, so that the measured loss can be attributed to the two couplers, discounting any losses from bends in the waveguide.

The MMI spectrum was obtained using a test structure containing a 1×2 MMI followed by a bent waveguide that routes one output arm back to a second GC, shown in Figure 2.6 **Right**. Gaussian fits are applied to both spectra, and the central performance figures such as IL extracted at 1550 nm.

Figure 4.1 **Left** shows the raw spectra together with their Gaussian fits. It can clearly be seen that the fits only approximate the true spectral shapes well when close to the GC maximum. This is also reflected in the higher normalised root mean square deviation of the fits. The GC spectrum is centred around 1549 nm.

The strip chart in **Right** summarizes the extracted figures across three chips. The mean MMI excess loss is -0.4 dB, and the mean power imbalance is -0.5 dB. The mean GC insertion loss is -2.17 dB.

All metrics are sufficiently low. In particular the mean GC insertion loss is well below the commonly assumed -3 dB, though the inter-chip variance is noticeable.




Figure 4.1. MMI and GC characterisation plots. **Left:** MMI and GC spectra with fitted Gaussians. All fits achieve a normalised root mean square deviation of $\leq 6.7\%$. Red cross shows GC peak. **Right:** Stripchart showing excess loss and power imbalance of the MMI alongside the insertion loss of the GC at 1550 nm. The reported GC insertion loss has been divided by two to account for the two GCs in the backloop optical path. Data collected across three chips.

4.1.2 Electro-Absorption Modulators

The results presented in this section are obtained from the test structures shown in Figure 2.7 (A). Light couples into a GC and is split into two arms, of which one serves as a reference spectrum and the other passes through an electrically contacted EAM. Both arms leave the structures through additional GCs.

Due to the wavelength dependence of the GCs, the optical transmission spectra seen in Figure 4.2 exhibit a Gaussian shape. This results in very poor signal-to-noise ratio (SNR) at wavelengths distant from 1550 nm. As in the previous section, to obtain metrics in a noise-robust way, we fit Gaussians to the measured spectra. The resulting fits can be seen against the raw data in Figure 4.2.

One may notice that the EAM spectrum exhibits a red shift with respect to the reference spectrum, which becomes more pronounced at higher bias voltages. This is a direct consequence of the red shift of the absorption edge via the Franz-Keldysh effect. As we are only interested in modulating the intensity of light, this effect is only significant insofar as it impacts the wavelength-dependent extinction ratio of the EAMs.

Figures 4.3 (A) and (B) depict the relative transmission obtained from both measurement and fit at various bias voltages. The depicted transmissions are obtained by dividing the data from the EAM arm of the MMI by that of the reference arm. There is a good visual concordance between fit and data spectra in the wavelength region of interest, roughly between 1540 and 1590 nm. The highest transmission occurs at wavelengths between the peaks of the individual spectra, and decreases with increasing reverse bias voltage.

Figure 4.3 (C) offers further insights into the EAM behaviour. The previously mentioned




Figure 4.2. Reference and EAM spectra for different bias voltages. The grey points represent the raw data. The colourful lines show the resulting Gaussian fits. All fits achieve a normalised root mean square deviation of $\leq 3.8\%$.




Figure 4.3. EAM Transmissions (A) Transmission obtained from measured spectrum in Figure 4.2. (B) Transmission obtained from Gaussian fits in Figure 4.2. (C) Normalized transmission as a function of wavelength and bias voltage. (D) Effective extinction ratio as a function of DC bias voltage for different modulation amplitudes V_{pp} .




Figure 4.4. Central performance benchmarks of EAM. Data marked with red crosses is used for comparison to reference values. **Top:** Extinction ratio and insertion loss obtained from Gaussian fits in 4.2 as a function of wavelength for various bias voltages. Red crosses show IL and ER values between 0V and 2V reverse bias at the wavelength corresponding to maximum FOM. **Bottom:** Resulting figure of merit as a function of wavelength. Red cross shows maximum value between 0V and 2V reverse bias.

red shift can be seen as a lighter spot at a wavelength of about 1580 nm. As before, the figure shows that the transmission decreases progressively with increasing reverse bias. One should note the importance of the transmission values not only at the bias point itself, but at the adjacent voltage levels as well.

Due to the nature of our application of the EAMs, in which all inputs pass through a biased diode where they are effectively turned "on" or "off" by an AC modulation signal, it is precisely this relationship, captured by the effective extinction ratio, which is most significant. Figure 4.3 (D) shows precisely ER_{eff} for $\lambda = 1550$ nm at different bias voltages and modulation amplitudes V_{pp} . Increasing V_{pp} opens the door to greater effective ERs, but also requires higher reverse bias voltages in order for the ER to be maximal. Moreover, despite the monotonic decrease of transmission with reverse bias, using a bias that is too high can hamper the ER_{eff} . Thus the choice of DC bias must necessarily account for the V_{pp} setting. Insertion loss, extinction ratio with respect to 0V, and the FOM as defined by equation (1.11) obtained from the fitted spectra are depicted in Figure 4.4.

The ER and IL in the **Top** figure demonstrate the same spectral response as those reported

by Srinivasan et al [13], however the values extracted at the specified points are dissimilar. Their ER is 0.8 dB higher, which may be attributable to the fact that our EAM consists of a GeSi alloy rather than pure Ge. At the same time, our IL is nearly 2 dB lower, which likely stems from run-to-run variations in the fabrication process [19]. Particularly due to our small IL, the FOM shown in 4.4 **Bottom** is over two times larger than the reference FOM in table 2.1. This, however, is not by itself an argument for a superior device, as the most important benchmark for our application remains the ER.

4.1.3 Ring Resonators

The ring measurement is carried out on the structure shown in Figure 2.7 (C). Light enters through a GC into the ring and exits either through the drop or the through port, depending on the wavelength. A wavelength sweep generates a spectrum, shown for the drop port in Figure 4.5. The spectrum is additionally calibrated by dividing by the grating coupler envelope to compensate for the wavelength-dependent coupling efficiency. While this procedure yields clean resonance peaks in the vicinity of the GC centre wavelength, it introduces artefacts at the spectral edges where the Gaussian fit diverges from the true GC response.




Figure 4.5. Ring resonator spectrum with Lorentzian fits. All fits share common offset C , determined from the 10th percentile of the spectrum. FWHM is determined from the fits, the ER is estimated from the marked true peaks and the noise floor. All fits achieve a normalised root mean square deviation of $\leq 5.9\%$. The FSR is roughly 1.85 nm.

The resonance peaks vary noticeably in amplitude. This is a known consequence of the wavelength dependence of the coupling coefficient between ring and bus waveguide, which is itself sensitive to fabrication tolerances [7]. Since the CAM readout depends on the trans-




Figure 4.6. Heating response of ring resonator. **Left:** Resonance shift as a function of power over heater. Red cross shows required power for π phase shift extracted from linear fit. **Right:** Resonance shift as a function of applied voltage. Red cross shows required voltage for π phase shift extracted from quadratic fit.

mitted power at resonance, such amplitude variations directly affect the photocurrent uniformity across columns.

Figure 4.7 collects the central performance metrics for the ring resonators, extracted from measurements across three separate chips. The figure shows mean transmission, FWHM, Q-factor, ER and IL as functions of wavelength, along with the standard deviation across chips. The inter-chip spread across all metrics narrows considerably near the GC centre wavelength around 1550 nm, where the calibration is most reliable, and fans out towards the edges of the observed spectrum. Still, it is important to note that the inter-chip variability in transmission is at times significant even close to the GC peak.

In the region around 1550 nm, the mean Q-factor settles above 35 000, in reasonable agreement with the computed reference value of 36 100 from table 2.1. The mean extinction ratio fluctuates around 16 dB, achieving a strong contrast but still falling noticeably short of the computed expectation of 29 dB. This discrepancy, however, very likely stems from the idealized assumptions for attenuation and is not indicative of poor device performance.

The measured FSR of approximately 1.85 nm matches the computed value of 1.86 nm very well. On the whole, these results confirm that the ring resonators perform as designed, with deviations attributable partially to calibration limitations and wrongful model assumptions rather than device deficiencies.

The integrated heaters are used to tune the positions of the rings' resonances. The heating response can be quantified by performing wavelength sweeps over the drop port while applying different levels of power to the heater, and then tracking the shift of the peaks from 0V. This shift $\Delta\lambda$ can be converted into units of π by normalising as $\Delta\lambda_\pi = 2\Delta\lambda/\text{FSR}$. Figure 4.6 shows this heating response as a function of both dissipated power and applied




Figure 4.7. Properties of resonances on ring resonators. Data is collected from three separate chips. From top to bottom: Transmission at resonance, FWHM, Q-factor, extinction ratio, and insertion loss. Only FWHM and Q-factor are obtained from Lorentzian fits, remaining figures are estimated directly from the ring spectrum.

voltage over the heater. The resonance shift scales linearly with dissipated heater power. This is expected from the proportionality between change in temperature and refractive index. A linear fit yields a π -phase-shift power of $P_\pi = 19.11$ mW, which performs similarly to the 23 mW reported by Liu et al. [8].

When plotted against voltage, the resonance shift follows the expected quadratic dependence ($P \propto V^2$), and a quadratic fit extracts $V_\pi = 1.44$ V. This figure constitutes the voltage range the FPGA DACs must cover in order to tune any ring across a full FSR.

4.1.4 Photodiodes

The structure measured for obtaining the results presented in this subsection can be seen in Figure 2.7 (B). Light enters through a GC into an MMI, which diverts half of the optical power out through a further GC, which constitutes a reference spectrum. The other half passes through an EAM and then excites a photocurrent in a photodiode, which is read out through an electrical contact.

Figure 4.8 shows the central performance metrics of the PD, namely dark current and responsivity. The former is read out by applying varying DC reverse bias voltages across the diode without illumination. The latter follows from sweeps across both bias and optical power.

Subfigure (A) shows dark current as a function of voltage. It can be seen that the PD exhibits a typical IV characteristic. The desired region for reverse bias operation is between 0V and 4.5V. The dark current in the photodetection operation region is low, on the order of magnitude of a few nanoamperes.

Subfigure (D) shows the responsivity as a function of bias voltage for various optical powers. The plot demonstrates an increase in responsivity with bias voltage up to a plateau at roughly 2V reverse bias. Though the values marked with red crosses match the expected magnitude for high performing photodiodes (responsivity at 1V reverse bias > 1 A/W), it is notable that they also demonstrate a strong dependence on optical power. This can be seen from the separation of the three curves.

The severity of this dependence on optical power is made especially clear in subfigure (B), which shows responsivity as a function of optical power for different bias voltages. The subfigure depicts a monotonic decrease in responsivity with increasing optical power. At sufficiently low powers, shown in the inset plot in subfigure (B), the observed power-dependent responsivity results in a small variability of about 0.05 A/W, or roughly 4% (drop-off between 0.3 and 1 mW input power at 2V reverse bias).

Ideally, responsivity should be constant up to saturation levels. However the small variability at low powers suggests that the observed decline may be negligible for weak photocurrents. The linearity plot in subfigure (C), which depicts photocurrent as a function of optical

power along with residuals of a linear fit, further supports this claim. The plot demonstrates a good fit up to $70 \mu A$ photocurrent, which corresponds to an incident optical power of roughly 0.7 mW at 2 V reverse bias.

Combining all insights, the dark current and responsivity demonstrate a good device performance. The decline in responsivity with power may be attributable to early-onset saturation effects, thus one should remain cognizant of the incident optical power during applications. The diminishing returns of bias voltages beyond 2 V provide a justification for setting the operating point at precisely that voltage level.




Figure 4.8. Photodiode characterisation results obtained from test chip. Incident optical power was estimated by accounting for the losses in the optical path. Data marked with red crosses is used for comparison to reference values. All data collected at a wavelength of 1550 nm . (A) Dark current as a function of bias voltage, where the coloured regions represent different operation modes of the diode. (B) Responsivity as a function of optical power for different bias voltages. Inset plot shows region of approximate linearity between photocurrent and optical power. (C) Photocurrent as a function of incident optical power at a reverse bias of 2 V . Inset plot shows residuals of linear fit. Assumption of linearity breaks at currents over $70 \mu A$. (D) Responsivity as a function of bias voltage for different optical power levels. Increasing bias over 2 V offers diminishing returns.

4.1.5 Summary of Test Structure Results

Table 4.1 collects the key figures extracted from the test structure measurements and places them alongside the reference values from table 2.1. On the whole, the measured devices meet or exceed the expected performance. The photodiodes exhibit low dark currents and high responsivity, though the power-dependent responsivity warrants additional considerations. The EAMs provide a useful extinction ratio with remarkably low insertion loss, yielding a FOM more than double the reference value. The ring resonators show stable Q-factors and extinction ratios, and the heaters tune efficiently with a P_π below the reference benchmark. The MMI splitters introduce minimal excess loss and acceptable power imbalance.

Table 4.1. Measured device performance compared to published reference values.

Device	Quantity	Reference	Measured
PD	I_{dark} at -1 V [nA]	~ 3	~ 3
	I_{dark} at -2 V [nA]	~ 8	~ 7
	Responsivity at -1 V, 1 mW [A/W]	> 1.0	~ 1.1
EAM	Max. FOM (0 to -2 V)	~ 0.96	~ 2.1
	IL at max. FOM [dB]	4.8	~ 2.9
	ER at max. FOM [dB]	4.6	~ 3.8
Heater	P_π [mW]	23	19.1
Ring	Quality factor Q ($\times 10^3$)	36	~ 35
	Extinction ratio [dB]	29	~ 16
	Free spectral range [nm]	1.85	1.86
MMI	Excess loss [dB]	—	-0.38
	Power imbalance [dB]	—	-0.48
GC	Insertion loss [dB]	—	-2.17

4.2 Printed Circuit Board

To validate the design decisions made during the PCB layout, a set of dedicated test traces is measured using a vector network analyser. The data are collected via S21 and S22 VNA analysis. The test traces systematically vary three parameters: the bend radius used in the meander sections, the presence or absence of solder resist over the signal traces, and the overall trace length.




Figure 4.9. S21 PCB results for different design parameters. **Top:** Comparison of attenuation for different meander bend radii. **Center:** Comparison of attenuation with and without solder resist. **Bottom:** Effect of trace length on attenuation.

Figure 4.9 presents the forward transmission for each parameter variation in frequency space. Subfigure **Top** shows that changing the meander bend radius has a negligible effect on attenuation, particularly below 1 GHz, which is the frequency range relevant to the intended FPGA modulation speeds. Subfigure **Center** reveals a similarly insignificant dif-

ference between traces with and without solder resist, with no notable effects up to several GHz. In contrast, subfigure **Bottom** demonstrates that trace length is by far the dominant contributor to signal attenuation. The longer trace suffers appreciably more loss at all frequencies, which is the expected behaviour for a lossy transmission line where attenuation scales with length.




Figure 4.10. Transient S22 PCB results for different design parameters. **Left:** Comparison of reflections for different meander bend radii. **Right:** Comparison of reflections with and without solder resist.

The transient reflection data presented in figure 4.10 shows that while the solder resist again has no measurable impact on the reflected signal in **Right**, the bend radius does produce visible differences in **Left**. After 1.5 ns, that is, after the signal has travelled through the cable connecting the VNA to the PCB and back, the curves for different radii begin to separate, with the smallest radius exhibiting the largest reflected power. This motivates the use of the largest feasible bend radius in the final PCB layout to minimise reflections.

These results confirm that the main design levers for maintaining signal integrity at the frequencies of interest are the meander geometry and trace length.

It is worth noting that while the solder resist has no discernible effect on either attenuation or reflections, leaving the traces uncovered carries practical downsides. Bare copper traces can act as unintentional antennas, both radiating energy from the driven lines and picking up ambient electromagnetic interference. While these effects can be mitigated through via fencing, due to the fact that the VNA measurements show no electrical benefit to omitting the solder resist, covering the traces remains the preferable choice.

4.3 Content Addressable Memory

Having established the individual device performance through the test structures, we now turn to the final CAM chip, where all components operate together. On the CAM chip, the ring resonators, EAMs, DC couplers, and photodiodes are connected as outlined in chapter 2, so the behaviour of each device is no longer isolated but coupled to its neighbours through optical crosstalk, thermal interactions, shared electrical bias lines, and RF electrical crosstalk. This integrated context reveals challenges that are not apparent from the test structure characterisation alone.

4.3.1 Photodiodes




Figure 4.11. Characterization of photodiodes on CAM chip. **Left:** Dark current as a function of bias voltage for all four column photodiodes. The dashed purple curve shows the current measured with no shorted terminals. **Right:** Dark current with the purple-line contribution subtracted.

On the CAM chip, all four column photodiodes share a common bias supply, so individual diodes must be selected by shorting one terminal of the desired PD to ground while leaving the others open. Figure 4.11 **Left** shows the thus extracted dark current as a function of applied voltage for each photodiode, and reveals that a substantial current flows even when no photodiode is shorted. This open-circuit current exhibits a diode-like IV characteristic of its own. This is indicative of a parasitic leakage path on the PCB.

At 2 V reverse bias, the open-circuit contribution exceeds 30 nA, which is larger than the dark current of any individual photodiode at the same bias. When the open-circuit contribution is subtracted, as in **Right**, the four PD curves recover a shape that is broadly consistent with the test structure results, with dark currents in the expected single-digit nanoampere range up to about 2 V.

4.3.2 Ring Resonators

As there are 16 rings on the CAM chip, we shall disambiguate them by referring to them with their row and column indices, e.g. ring 41 refers to the cell at row 4, column 1.

We will begin by analysing the effects of thermal crosstalk on the ring resonators. We do this by coupling into the intracell GC seen in Figure 2.3 and performing wavelength sweeps whilst simultaneously applying voltage to one or multiple other heaters on the chip. We can further compute the resulting detuning loss L_{det} if the thermal crosstalk were left unaccounted for by assuming a Lorentzian lineshape $T(\lambda; \lambda_0, \Gamma, A)$ of the resonance

$$T(\lambda; \lambda_0, \Gamma, A) = \frac{1}{\pi\Gamma} \frac{2A}{1 + 4((\lambda - \lambda_0)/\Gamma)^2}, \quad (4.1)$$

with wavelength λ , resonance wavelength λ_0 , FWHM Γ , and peak amplitude A . The detuning loss can then be computed as the ratio of the transmission at the resonance wavelength with and without the thermal crosstalk-induced shift $\Delta\lambda$:

$$L_{\text{det}} = 10 \log_{10} \left(\frac{T(\lambda_0; \lambda_0 + \Delta\lambda, \Gamma, A)}{T(\lambda_0; \lambda_0, \Gamma, A)} \right) = \frac{1}{1 + 4(\Delta\lambda/\Gamma)^2}. \quad (4.2)$$




Figure 4.12. Effects of heating cross talk on ring resonator. **Left:** Phase shift of resonance wavelength as a function of voltage applied to other heaters. **Right:** Phase shift of resonance wavelength as a function of power applied to other heaters.

Figure 4.12 shows the detuning effects over ring 41 with heating at rings 13, 31, 43 and 21. The **Left** subfigure shows a monotonic increase in detuning with increasing voltage. It can also be seen that the simultaneous heating of multiple rings produces an effect that is smaller than the sum of the individuals. This is very likely caused by thermal dissipation across the chip.

The **Right** subfigure shows the same data plotted against the power dissipated in each heater. The effect of the crosstalk is significant. At $V_\pi \approx 1.4\text{V}$, heating a single neighbour produces a detuning loss of roughly -10 dB, while heating all four simultaneously causes a loss greater than -20 dB. Heating all 16 rings simultaneously would produce an even stronger detuning.




Figure 4.13. Resonance shifts of ring due to environmental temperature changes on two different days. **Left:** Resonance wavelengths at different days for different heater voltages. **Right:** Violin plot of the combined population of shifts across all voltages. The right-hand axis is calibrated in the resulting detuning loss that would occur if the shift were not compensated. The white dot represents the median, the box represents the interquartile range, and the whiskers extend to the extremas of the distribution.

Detuning of this magnitude would render the ring effectively unreadable, as an unaccounted detuning of only -1.25 dB across each ring in a column, corresponding to a loss of a whole I_{ph} , already violates the threshold condition (2.1).

Table 4.2. Thermal cross-talk measured at Ring 41 when each distant heater is driven at 1 V. The detuning loss is calculated from the resonance shift assuming a Lorentzian lineshape.

Ring Index	Distance to Ring 41 (μm)	Applied Voltage (V)	Resonance Shift (π)	Detuning Loss (dB)
13	930	1.0	0.0404	-6.21
21	580	1.0	0.0479	-7.37
31	280	1.0	0.0398	-6.10
43	260	1.0	0.0413	-6.35

Table 4.2 shows the resonance shift, detuning loss and distance to ring 41 with an applied voltage of 1V. There is no discernible effect of distance on the detuning, suggesting that a global heating of the chip is occurring rather than a local effect. This is consistent with the findings of Biasi et al. [20], who report a dominating global heating effect for rings separated by more than 200 μm . The slightly stronger crosstalk from ring 21 may originate from residual heat from previous measurements or environmental factors.

These findings underscore the necessity for either active compensation of the ring resonance wavelength or a tuning strategy that accounts for the cross-heating between rings.

Next, we examine the stability of the ring resonances over time under ambient conditions. This is done by performing wavelength sweeps over the same structure that was used for the




Figure 4.14. Peak magnitudes from compensated splitting ratios. Measurement was taken using the grating coupler pair adjacent to each ring, resulting in an inverted splitting ratio at each DC compared to the design. Each peak was multiplied by a correction factor to compensate for the inversion. Peaks corresponding to a 1:0 splitting ratio are omitted, as the correction factor diverges. Error bars show the mean $\pm$ one standard deviation.

crosstalk analysis at different heater voltages on two different days and tracking the shifts in resonance wavelengths.

Figure 4.13 **Left** shows the exact resonance positions for ring 41 at different heater voltages. The day-to-day shifts are small but consistently present and not uniform. The violin plot in the **Right** subfigure aggregates the shifts observed at all heater voltages into a single distribution and translates them into the resulting detuning loss. The median detuning loss is modest, around -0.5 dB, but the distribution extends up to approximately -3 dB in the worst case.

The non-uniformities may arise as a consequence of unstable voltage supply to the heaters, but as a shift is present even without heating, it is evident that environmental effects must also be compensated for during operation. This also underscores the importance of using a stable voltage supply.

It should be noted that these measurements reflect the ambient conditions of only two particular days. Larger temperature excursions would produce correspondingly larger shifts.

What remains to be seen is whether the photocurrents generated by rings at different positions on the chip are uniform. This depends on the variable ILs of the rings, but also on the splitting ratios of the DCs meant to evenly distribute optical power across the CAM, as explained in chapter 2, the responsivities of the photodiodes, and on the absorption of the

EAMs.

We begin by coupling into the intracell GCs of each ring and performing a wavelength sweep over the drop port. This methodology results in an inversion of the splitting ratios at each DC compared to the design: A ring that would nominally receive a fraction α of the input power in normal operation now receives $(1 - \alpha)$. Similarly, a fraction β of the drop port signal ordinarily couples to the vertical bus waveguide leading to the photodiode, but the GC in the cell now sees $(1 - \beta)$. To account for this, each measured peak is multiplied by a compensation factor $((1 - \beta)(1 - \alpha))^{-1}$, derived from the nominal splitting ratios. Peaks corresponding to a nominal ratio of 1:0 (i.e. zero power through the test port) are omitted, as the compensation factor diverges.

The thus compensated magnitudes can be seen in figure 4.14. The red bars show standard deviations in both resonant wavelength and magnitude. Though their means fall within each other's tolerance regions, the magnitudes exhibit a large spread at every resonance wavelength.

Ideally, all rings should produce the same transmitted power at resonance, since the photocurrent generated in the column photodiode scales directly with this quantity. The observed variability — with a standard deviation on the order of 20–30% of the mean — means that some rings contribute substantially more photocurrent than others, complicating the definition of a clean detection threshold. As the spread in transmission observed here is greater than the one obtained from the test structures near the GC peak (see Figure 4.5), the most likely cause is the splitting ratios.

However, it should also be noted that the resonance wavelengths themselves are tightly clustered. The horizontal scatter at each FSR period is small, indicating consistent ring radii and waveguide widths across the chip.

The question of how this variability expresses itself in the real photocurrents can be explored by measuring them after entrance through the intended inputs, seen in Figure 2.2. To make this measurement possible, a photodiode at a desired column is shorted, thus enabling the signal readout, and an input wave is sent into a desired row. The laser wavelength is iteratively tuned to maximise the photocurrent, which is equivalent to finding the resonant wavelength of the ring at that position.

Repeating this procedure for all 16 rings and documenting the induced photocurrents produces the heatmap shown in Figure 4.15 **Left**. The bar charts on the margins show the column-wise and row-wise means of the photocurrents. We observe a maximal difference of $1.1 \mu\text{A}$, which is over half of the highest recorded photocurrent at $2.0 \mu\text{A}$. Furthermore, despite the homogeneous row-wise means, there is an apparent systematic tendency for the first column to produce higher photocurrents.

These observations are complicated by the fact that the EAM currents, which were also




Figure 4.15. Heatmaps showing the photocurrent induced by individual rings in the 4×4 grid. Marginal bar charts show the column-wise (top) and row-wise (right) means. **Left:** Raw photocurrent readout at the column photodiode after light passes through the selected ring. **Right:** Photocurrent normalised by the EAM current measured at the corresponding row input.




Figure 4.16. Scatter plot with linear regression line for EAM and PD photocurrents across all 16 ring-column combinations. The Pearson correlation coefficient $r = 0.70$ ($p = 0.003$) confirms a statistically significant positive correlation.

recorded, are not uniform across rows, despite the fact that each row shares a common EAM and input waveguide. This can be seen in Figure 4.16, which plots the EAM currents against the photocurrents — under identical coupling to the input GC, the EAM currents, even if different for each row, should be stacked on top of one another in the plot. Instead, they are scattered across the plane, exhibiting a highly significant positive correlation with the photocurrent readout ($p \ll 0.05$).

The most probable explanation is that the coupling into the input GC was not, in fact, identical across rows. Despite efforts to realign the stage before each measurement, the process of swapping shorts between the photodiode SMA connectors and moving the laser light along the fibre array likely led to differences in calibration which could not be compensated. Naturally, a weaker coupling would then produce a smaller current at the EAM and a smaller photocurrent. To get an idea of how the photocurrent distribution may have looked like under identical calibration conditions, we can normalise the photocurrent heatmap to the same EAM current. The result can be seen in Figure 4.15 **Right**. The column-wise variability is reduced, but the row-wise variability is increased, with row 4 demonstrating the highest photocurrents.

This is undesirable: while differing column-wise photocurrents can be accounted for by defining different detection thresholds for each PD, row-wise variability produces systematic biases toward certain bits. As an example, imagine a 2-bit column, in which the maximal photocurrent generated by the first cell is $1\mu\text{A}$, and that of the second cell is $2\mu\text{A}$. Allowing for a noise margin of 2 dB per cell, the threshold current becomes $I_{\text{th}} \approx 1.8\mu\text{A}$. As $2\mu\text{A} > I_{\text{th}} > 1\mu\text{A}$, the column could detect a match even if only the second cell corresponds to the input, while the first is completely off-resonance. The inverse could not happen (a match can still be detected while the second cell is below the 2 dB margin, but not if it produces $0\mu\text{A}$ photocurrent). In a column with identical photocurrent generation, the probability of a mismatch is symmetric and independent of the cell.

Taken together with the compensated peak amplitudes in Figure 4.14, these heatmap results offer a strong argument for deviations of the DC splitting ratios from the design. However, a more controlled measurement environment is necessary to disentangle the intrinsic device non-uniformity from the extrinsic calibration noise.

4.3.3 Electro-Absorption Modulators

Although, unlike with the MRRs, the EAMs do not have dedicated GCs on the CAM chip for testing, the ER can still be extracted from the photocurrents, provided the wavelength sent into the input GC matches the resonance of the ring at the same row and in the same column as the shorted photodiode.

Figure 4.17 **Left** shows the thus determined extinction ratio as a function of reverse bias for




Figure 4.17. Characterization of EAMs on CAM chip. **Left:** Extinction ratio derived from photocurrents as a function of reverse bias voltage. Red line shows the mean, shaded region indicates $\pm 1\sigma$. **Right:** Comparison of extinction ratios between test chip EAMs (from optical spectra) and CAM chip EAMs (from photocurrents) at several wavelengths.

all eight EAMs on the CAM chip. The spread between devices is small, with the 1σ band remaining within roughly 1 dB of the mean at all bias points. This is important for the CAM application, where every EAM must provide a comparable on/off contrast. The mean ER reaches approximately 8 dB at 2.5 V reverse bias, which exceeds the test structure values at the same voltage.

This apparent improvement is examined more closely in the **Right** subfigure, which overlays the CAM chip ER (derived from photocurrents) with the test chip ER (derived from optical spectra) at several wavelengths. The CAM chip data points sit consistently above the test chip curves. While this might at first suggest superior EAM performance on the final chip, a more likely explanation is offered by ambient temperature fluctuations and the resulting ring detuning.

4.3.4 AC Operation

A simple AC measurement is carried out as a proof of concept for the full signal chain. An arbitrary waveform generator (AWG) is connected to one of the EAMs, the laser wavelength is tuned to the resonance of the corresponding ring, and the photodiode output is routed off-chip through a transimpedance amplifier and read out on an oscilloscope.

The waveform applied by the AWG is successfully reproduced in the photodiode signal, demonstrating that the end-to-end optical path, from EAM modulation through ring filtering to photodetection, functions as intended under AC drive conditions. While this test does not yet constitute a full CAM query, which would require simultaneous modulation of multiple EAMs and threshold-based evaluation of all four column photodiodes, it confirms the principal feasibility of AC operation.

Discussion

The results presented in chapter 4 demonstrate that all individual components of the CAM photonic integrated circuit perform at or above the expected benchmarks and that the end-to-end optical signal path functions as intended under AC drive conditions. This chapter synthesises the principal findings and identifies the challenges that must be addressed before a full CAM query can be demonstrated. We conclude with an outlook on the steps required to advance from the present proof-of-concept of this architecture towards a practically deployable photonic CAM system.

Device Performance

The test structure characterisation confirms that each component meets its design targets. The photodiodes exhibit low dark currents in the expected single-digit nanoampere range up to 2 V reverse bias, and a responsivity exceeding 1 A/W at 1 V, in agreement with the reference values from table 4.1. The power-dependence of the responsivity of the PDs is negligible at sufficiently small power levels, and linearity between photocurrent and optical power is maintained up to about 70 μ A. The EAMs achieve a figure of merit more than twice the reference value, driven primarily by their remarkably low insertion loss of approximately 2.9 dB. The ring resonators show Q-factors around 35 000 and extinction ratios of approximately 16 dB, both consistent with the computed reference values, with the discrepancy in ER attributable to the idealised lossless assumption in the theoretical model. The heaters deliver a π -phase-shift power of 19.1 mW, thus achieving a higher efficiency compared to the 23 mW reference. The MMI splitters introduce only -0.38 dB of excess loss and -0.48 dB of power imbalance, and the grating couplers achieve a mean insertion loss of -2.17 dB, comfortably below the commonly assumed -3 dB. Taken together, these results confirm a well-functioning photonic platform.

An important cross-cutting remark is that upgrading from the current iSiPP50G process design kit to a more recent imec platform would be expected to yield across-the-board improvements in device performance, in particular higher EAM extinction ratios and reduced waveguide propagation losses, without requiring architectural changes to the CAM design.

Thermal Stability and Crosstalk

The dominant operational challenge identified by the measurements is thermal instability of the ring resonances. The violin plot in figure 4.13 shows that the day-to-day resonance shift can translate into a detuning loss of up to approximately -3 dB in the worst case, with

a median around -0.5 dB. This sensitivity to ambient temperature is a known property of silicon microring resonators and is a direct consequence of the thermo-optic effect [9].

A practical mitigation strategy is to initialise the system by driving all heaters to a fixed elevated operating power before commencing any search operation. Because the photonic CAM is operated in a write-once mode for the duration of a query sequence, only relative resonance positions matter during operation, not absolute resonance wavelengths. Raising all heaters to a common baseline power brings the chip to a stable elevated temperature, from which individual resonances can be shifted in both directions by increasing or reducing the heater voltage. Finding the target resonance for each ring then necessitates some calibration process, an example of which is described below.

First, the EAM voltages are set to transmit only the ‘0’-bit laser on the first row. The readout photocurrent at each column photodiode is then compared to the expected value: it should fall below $1/4$ of the full-match threshold if the first row of the stored word is not a ‘0’ bit, and above it if it is. If the current does not match the expectation, the heater voltages for each column are tuned iteratively until the photocurrent condition is satisfied. The procedure is then repeated for the ‘1’-bit laser of the first row, and then for subsequent rows. After all rows have been calibrated, all lasers are simultaneously switched on and the photocurrents are verified to lie above the full-match threshold in the expected columns. If any column fails, the calibration restarts from the first row. Crucially, a stable, low-noise voltage source must be used to supply the heater currents, as voltage noise directly translates into resonance jitter and degrades the detection margin.

A further complication is thermal crosstalk between adjacent heaters, which has been identified as a central challenge in photonic systems of this type [20]. As shown in figure 4.12, applying voltage to a single distant heater can shift the resonance by close to 0.05π at 1 V, corresponding to a detuning loss of over -7 dB. This sort of detuning is incompatible with CAM operation, and must be carefully managed.

Two broad classes of solution to thermal instability have been proposed in the literature. The first is athermalization, in which the temperature dependence of the resonance wavelength is reduced by incorporating materials with negative thermo-optic coefficients into the waveguide cladding, or by thermally isolating the ring from the substrate through etching [9]. While effective in principle, this approach is costly and difficult to integrate with standard CMOS fabrication processes, as polymer cladding materials tend to degrade at the elevated temperatures used in back-end processing, and etching is geometrically constrained by a given architecture. The second approach is control-based stabilisation, in which a feedback controller continuously monitors the resonance position and adjusts the heater voltage to compensate for drift [10]. A lightweight implementation suitable for the present system would inject a periodic calibration query — for example, once every million search queries

— in which all laser inputs are activated simultaneously. The FPGA then verifies programmatically that the photocurrent on each column exceeds the threshold corresponding to a full match. If any column falls below threshold, the previously described ring calibration sequence is triggered automatically before search operations resume. The cost of this approach is a small periodic interruption and modest additional power consumed by the calibration query.

Ideally, a combination of both approaches would be employed, with mild athermalization reducing the magnitude of thermal drift, and an active feedback loop to correct the residual instability. Of the two, the control-based approach is more immediately compatible with the existing FPGA infrastructure and is therefore the more feasible near-term solution.

Detection Threshold and Signal Integrity

Setting a reliable detection threshold requires the photocurrent produced by a full match to be clearly distinguishable from that of a near-miss. This condition is governed by equation 2.1, which states that the threshold I_{th} must satisfy $(3 + 10^{ER_{eff}/10}) I_{ph} < I_{th} < 4 I_{ph}$, where I_{ph} is the photocurrent contribution of a single matching ring. A usable threshold window therefore exists only for sufficiently large ER_{eff} .

In practice, the photocurrents measured across all 16 ring–column combinations as well as the compensated magnitudes have been shown to be variable with ring position along the CAM grid. While this effect necessitates further study, it would be reasonable to conclude that a single, globally defined threshold current is unlikely to perform optimally for all columns simultaneously. Instead, a per-column threshold would better accommodate the device-to-device variability. Even if the photocurrent generated by each ring differs from cell to cell, these values are characteristic of the specific device. The threshold for each column can therefore be defined as the sum of the characteristic single-ring photocurrent excitations. Whilst this approach would bias the error rate toward particular bits, it would account for the non-uniformity without requiring redesign.

All photocurrents presented in the heatmap 4.15 lie well below $70 \mu A$, the saturation onset identified in the photodiode characterisation. The photodiodes therefore operate in the linear regime throughout all measurements, and non-linearity is not likely to have contributed to the observed photocurrent variability.

Path to Full CAM Operation

The AC measurement presented in section 4.3.4 establishes that the end-to-end optical signal chain operates as intended under high-frequency drive conditions. However, this demonstration involved only a single EAM modulating a single row, with a single photodiode reading out a single column. A complete 4-bit CAM query requires simultaneous modula-

tion of all four row EAMs with complementary bit-encoding voltages, parallel readout of all four column photodiodes, and a threshold comparison in the FPGA that identifies which column, if any, registers a full match. Realising this in the FPGA requires synchronised RF-DAC channels for EAM drive and simultaneous RF-ADC sampling of all four photocurrent channels, with digital threshold logic operating on the captured values, not unlike the ones suggested in the previous sections.

Outlook

The present work demonstrates the feasibility of a silicon-photonics content-addressable memory on a standard fabrication platform and identifies the principal challenges that must be overcome to realise its practical operation. Several directions for future work emerge naturally from the results.

The most immediate next step is the integration of the full CAM query into the FPGA firmware, enabling simultaneous multi-row modulation and parallel photodiode readout with real-time threshold comparison. This would allow a direct measurement of query latency and error rate under controlled conditions. Closely coupled to this is the implementation of the closed-loop thermal stabilisation scheme discussed above, which is a prerequisite for stable, unattended operation in a realistic environment.

Scaling the present 4×4 architecture to larger arrays is the central longer-term challenge. The optical power reaching each photodiode scales as $1/(MN)$ for an $M \times N$ array, since each successive column tap and row tap removes a fixed fraction of the circulating power. Maintaining a sufficient signal-to-noise ratio at larger scales will require either higher input laser power, lower-loss splitting elements or on-chip optical amplification, or all three in combination, depending on the size of the array.

Finally, it is worth placing this work in the context of the broader landscape of photonic and electronic memory technologies. Electronic CMOS ternary CAMs are mature, densely integrated, and widely deployed, but their power consumption and interconnect delays become increasingly problematic as clock frequencies scale [2]. Photonic CAMs operating at gigahertz query rates with low per-bit energy costs represent a compelling alternative for the specific case of read-intensive parallel lookup, provided that the challenges of thermal stability and optical power budget can be resolved. The results presented in this thesis confirm that the individual building blocks of such a system are already available on a standard silicon-on-insulator platform, and that the principal remaining obstacles are at the system integration and control level rather than at the device level.

Bibliography

- [1] Yousra Alkabani, Mario Miscuglio, Volker J. Sorger, and Tarek El-Ghazawi. OE-CAM: A Hybrid Opto-Electronic Content Addressable Memory. *IEEE Photonics Journal*, 12(2):1–14, April 2020.
- [2] Md Abdullah-Al Kaiser, Ankur Singh, Shay Snyder, Hamed Poursiami, Maryam Parsa, and Akhilesh Jaiswal. OptiCAM: An optical content-addressable memory architecture for ultra-fast pattern matching. *APL Photonics*, 10(12):126116, December 2025.
- [3] Yanir London, Thomas Van Vaerenbergh, Luca Ramini, Antoine Descos, Luca Buonanno, Jinsung Youn, Can Li, Catherine E. Graves, Marco Fiorentino, and Raymond G. Beausoleil. Multiplexing in photonics as a resource for optical ternary content-addressable memory functionality. *Nanophotonics*, 12(22):4137–4155, November 2023.
- [4] Bahaa E. A. Saleh. *Fundamentals of photonics I: Optics/ Bahaa E.A. Saleh, Malvin Carl Teich*. Wiley series in pure and applied optics. Wiley, Hoboken, NJ, third edition edition, 2019.
- [5] Michael Te Vrugt, editor. *Artificial Intelligence and Intelligent Matter: Nanoscience, Soft Matter, Philosophy*. Machine Intelligence for Materials Science. Springer Nature Switzerland, Cham, 2026.
- [6] Jia-ming Liu. *Photonic Devices*. Cambridge University Press, 1 edition, April 2005.
- [7] W. Bogaerts, P. De Heyn, T. Van Vaerenbergh, K. De Vos, S. Kumar Selvaraja, T. Claes, P. Dumon, P. Bienstman, D. Van Thourhout, and R. Baets. Silicon microring resonators. *Laser & Photonics Reviews*, 6(1):47–73, January 2012.
- [8] Shengping Liu, Junbo Feng, Ye Tian, Heng Zhao, Li Jin, Boling Ouyang, Jiguang Zhu, and Jin Guo. Thermo-optic phase shifters based on silicon-on-insulator platform: state-of-the-art and a review. *Frontiers of Optoelectronics*, 15(1):9, April 2022.
- [9] Kishore Padmaraju and Keren Bergman. Resolving the thermal challenges for silicon microring resonator devices. *Nanophotonics*, 3(4-5):269–281, August 2014.
- [10] Kishore Padmaraju, Johnnie Chan, Long Chen, Michal Lipson, and Keren Bergman. Thermal stabilization of a microring modulator using feedback control. *Optics Express*, 20(27):27999–28008, December 2012.

- [11] Gordon B. Morrison, James W. Raring, Chad S. Wang, Erik J. Skogen, Yu-Chia Chang, Matt Sysak, and L.A. Coldren. Electroabsorption modulator performance predicted from band-edge absorption spectra of bulk, quantum-well, and quantum-well-intermixed InGaAsP structures. *Solid-State Electronics*, 51(1):38–47, January 2007.
- [12] L. B. Soldano and E. C. M. Pennings. Optical multi-mode interference devices based on self-imaging: principles and applications. *Journal of Lightwave Technology*, 13(4):615–627, April 1995.
- [13] Srinivasan Ashwyn Srinivasan, Marianna Pantouvaki, Shashank Gupta, Hong Tao Chen, Peter Verheyen, Guy Lepage, Gunther Roelkens, Krishna Saraswat, Dries Van Thourhout, Philippe Absil, and Joris Van Campenhout. 56 Gb/s Germanium Waveguide Electro-Absorption Modulator. *Journal of Lightwave Technology*, 34(2):419–424, January 2016.
- [14] David M. Pozar. *Microwave engineering*. Wiley, Hoboken, NJ, 4th ed edition, 2012.
- [15] S. Brown and J. Rose. FPGA and CPLD architectures: A tutorial. *IEEE Design & Test of Computers*, 13(2):42–57, June 1996.
- [16] Lennart Meyer, Jelle Dijkstra, Simon Tebeck, Liam McRae, Niklas Bahr, Daniel Steinmeyer, Sergey Koptyaev, Johana Bernasconi, Nikolay G. Pavlov, Maxim Karpov, John D. Jost, Wolfram Pernice, and Frank Brücknerhoff-Plückelmann. Deep neural network inference on an integrated, reconfigurable photonic tensor processor. *Nature Communications*, 17(1):3396, April 2026.
- [17] Eric Dulkeith, Fengnian Xia, Laurent Schares, William M. J. Green, and Yurii A. Vlasov. Group index and group velocity dispersion in silicon-on-insulator photonic wires. *Optics Express*, 14(9):3853–3863, May 2006.
- [18] Hongtao Chen, Peter Verheyen, Peter Heyn, Guy Lepage, Jeroen Coster, Philippe Absil, Gunther Roelkens, and Joris Van Campenhout. High-Responsivity Low-Voltage 28-Gb/s Ge p-i-n Photodetector With Silicon Contacts. *Journal of Lightwave Technology*, 33:820–824, February 2015.
- [19] Filippo Ferraro, Peter Heyn, Minkyu Kim, Natarajan Rajasekaran, Mathias Berciano, Grigorij Muliuk, Dieter Bode, Guy Lepage, Sofie Janssen, Rafal Magdziak, Jeroen Coster, Hakim Kobbi, Sebastien Lardenois, Negin Golshani, Leili Shiramin, Chiara Marchese, Selva Rajmohan, Shankr Nadarajan, Neha Singh, and Philippe Absil. Imec silicon photonics platforms: performance overview and roadmap. page 9, March 2023.

-
- [20] Stefano Biasi, Riccardo Franchi, Davide Bazzanella, and Lorenzo Pavesi. On the effect of the thermal cross-talk in a photonic feed-forward neural network based on silicon microresonators. *Frontiers in Physics*, 10, December 2022.

Acknowledgements

First and foremost, I would like to thank professor Wolfram Pernice for affording me the opportunity to grow and learn in his research group. It has been without a doubt one of the most enriching experiences I've had to date.

I would also like to thank Philipp and Lennart for assuming joint custody in the tedious task of my supervision. You were always friendly and approachable, and I felt much more an equal contributor to the project than I had anticipated. Thank you also for taking the time to review my thesis even so terribly close to the deadline.

It would be amiss not to thank Falk as well, for answering my million questions regarding PCB design, never with any delay or reproach.

I wish to extend the warmest of gratitudes to my girlfriend, Zhuli, for her terrific patience, love and compassion. I surely would not have succeeded, or, at least, not without losing my sanity, if it wasn't for you. I have incurred a debt I know not how to pay, besides to say 'I love you'.

I should thank Derya, for asking me to take care of her dog, Pamuk, whose company I grew increasingly fond of during my weeks of writing, and whose companionship and frequent need for walks kept me tethered to reality even at the most abysmal of times.

Thank you also to Grisha, Bobi, Lozi, Marto and Sasho, for your friendship and readiness to hang out with me even at the shortest of notices.

And of course, thank you mum and dad for enabling me to come here and study in the first place, and for all the support you have given me throughout the years.

Erklärung

Zur Überprüfung der Grammatik und sprachlichen Richtigkeit wurde teilweise ChatGPT verwendet. Alle Analyse-Codes wurden in VS Code geschrieben, mit Inline-Vorschlägen des integrierten Copilot-Assistenten. Sämtliche von KI generierten Code-Vorschläge wurden vor der Implementierung sorgfältig überprüft und überarbeitet.

Ich versichere, dass ich diese Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, den 28.04.2026,

A handwritten signature in black ink, appearing to be 'Andreas', written in a cursive style.