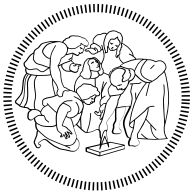

POLITECNICO
MILANO 1863

Department of Electronics, Information and Bioengineering
Doctoral Program In Electronics Engineering

Advanced control techniques for heterogeneous and densely-integrated photonic circuits

Doctoral Dissertation of:
Vittorio Grimaldi

Supervisor:
Prof. Marco Sampietro

Tutor:
Prof. Giuseppe Bertuccio

The Chair of the Doctoral Program:
Prof. Luigi Piroddi

2022 – XXXV Cycle

Abstract

Technologies advances have allowed the fabrication of complex yet reasonably cheap integrated photonic chips. The potentialities of these structures are enormous, ranging from the creation of short-range optical interconnection schemes, to the implementation of artificial intelligence and all-optical signal processors. However, integrated photonic systems have not reached the expected diffusion yet, hindered by the necessity of an electronic layer to be operated. Electronic control algorithms are indeed fundamental to counteract the effect of fabrication mismatches and thermal and wavelength instabilities, that prevent open-loop operations.

In this work, a closed-loop control scheme is proposed, designed to easily scale up for heterogeneous and densely-integrated photonic circuits. The dithering technique is first presented, showing how it can be used to implement a power-independent and calibration-free feedback loop to stabilize the controlled devices on the stationary points of their transfer function. The approach is experimentally validated, showing how it is possible to control multiple cascaded devices with a single external detector by exploiting the concepts of frequency orthogonality and frequency selectivity of the lock-in readout. The dithering technique is then further developed, showing how it is possible to realize a similar control loop for integrated modulators, stabilized

on the maximum-slope working condition. This allows the use of the same control scheme in heterogeneous architectures combining both classes of devices, paving the way to the integration of complex functionalities on the same chip.

The huge flexibility achieved by the implemented design is possible thanks to the adoption of a digital core. To maximize the performance of the control, an FPGA is chosen to operate the system, allowing an unprecedented grade of parallelization. An efficient implementation of all the building blocks of the digital design is presented, allowing further scaling of the system to accommodate future growth of photonic circuits.

Finally, a plasmonic-assisted ultra-compact bolometric sensor is presented, that can be offered as a fundamental building block for the next-generation plasmonic circuits. The effect of the sensor dimensions on the propagation of the plasmonic modes are assessed and the device is designed accordingly. The resulting structure, fabricated in a TiO_2 platform, is fully characterized in terms of sensitivity and time response, and it is eventually used in a first control experiment.

Contents

CHAPTER 1

Introduction

The introduction of optical fiber links for long-range data transmission was a milestone in the history of modern telecommunications. The reliability and efficiency of optical fibers, together with their ease of fabrication, has allowed the creation of a worldwide network that nowadays comprises more than 4 billions km of fibers [1]. While optical-based long-range communications are expected to keep growing in the following years to answer the needs of the growing industry, it is expected that optical interconnections will become the leading paradigm also at at mid-, short- and ultra-short-scale [2]. Indeed, traditional copper-based interconnections are reaching their intrinsic limitations in terms of bandwidth and power consumption, effectively becoming the bottleneck to further increase the computation performance of these systems [3]. To answer these needs, traditional discrete optical components are not a viable alternative, because of their limited flexibility, bulky dimensions, high power dissipated and high cost. Similarly to what happened with microelectronics, a solid-state approach seems to be the answer to bring optics to a

level of miniaturization and price suitable for short range communications. Several integrated optical technologies exist based on different materials, like silicon [4], silicon oxynitride [5], indium phosphide [6] and others, with different strengths, weaknesses and levels of maturity.

In this framework, the most widespread and mature technology platform is indubitably Silicon Photonics, often in the form of Silicon on Insulator (SOI). Silicon has indeed excellent optical characteristics, being transparent to a broad range of wavelengths, spanning from 1100 nm to 7000 nm, hence including the wavelengths of interest in the telecommunication industry (1300 nm and 1500 nm). Moreover, the big difference between the refractive index of the silicon ($n_{\mathrm{Si}} \simeq 3.45$) and of the silicon oxide ($n_{\mathrm{SiO_2}} \simeq 1.45$) allows to integrate compact waveguides where the traveling mode is well confined within the silicon core. High-curvature bending radii are then possible without introducing substantial losses, allowing to further reduce the footprint of the single devices. Hence, by leveraging the consolidated know-how of the CMOS industry, densely-integrated photonic circuits have been fabricated in the past years, featuring different kinds of devices, such as Mach-Zehnder interferometers, ring resonators, arrayed waveguide grating routers [7–9], and more, enabling advanced on-chip manipulation of light beams in a compact footprint.

The choice of silicon also allows to envision the creation of integrated systems that include both optical and electronic components on the same chip. Monolithic integration, that allows to fabricate electronic-photonic integrated circuits on the same chip using a single fabrication process, has already been demonstrated [10–12], and it is expected to be the winning approach once a sufficient level of technological maturity is reached [13, 14]. For the moment, hybrid integration is the most common approach, where the electronic and photonic chips are fabricated with different processes, optimized separately, and then electrically connected through wire or flip-chip bonding [15, 16].

1.1 The need for control electronics

The level of development reached by Silicon Photonics was widely predicted by the experts in the fields [4]. Nonetheless, very few commercial solutions that leverage the potentiality of this technology exist [17]. One reason for this delay can be found in the practical impossibility of having a plug-and-play solution operating these devices in an open-loop fashion. For example, the same high index contrast between Si and SiO_2 that allows dense integration of components can cause a 100 GHz shift of the spectral response of a photonic device due to a length mismatch of just 1 nm [18]. Similarly, the high thermo-optic coefficient of silicon (1.86×10^{-4} K^{-1} at 300 K [19]), makes the integrated devices prone to thermal drift and crosstalk. For instance, a temperature variation of 1 K can shift the resonance frequency of a ring resonator by 10 GHz [20, 21], making it difficult to reliably operate complex photonic architectures in harsh environments like datacenters.

This means that a photonic system has to be supported by an electronic control layer, as schematically presented in Fig. 1.1. To translate the information from the photonic domain into the electronic one, different kinds of sensors can be employed. These sensors can be placed outside the chip, measuring part of (or all) the power emerging from the output ports. However, as the complexity of the assembly increases, it is often necessary to assess the behavior of intermediate points inside the circuit. For these reasons, a lot of research effort has been carried out in developing efficient integrated sensors that can be employed for the task.

Besides extracting the information of interest from the circuit, it is also necessary to be able to change the behavior of the controlled devices. For the purpose, actuators of different kinds have been proposed and integrated over the course of the years.

Finally, the control layer has to be designed as well. The tasks of this layer can be divided into:

- compensate the effect of tolerances and fabrication nonuniformities;
- stabilize the behavior of the circuit compensating the pres-

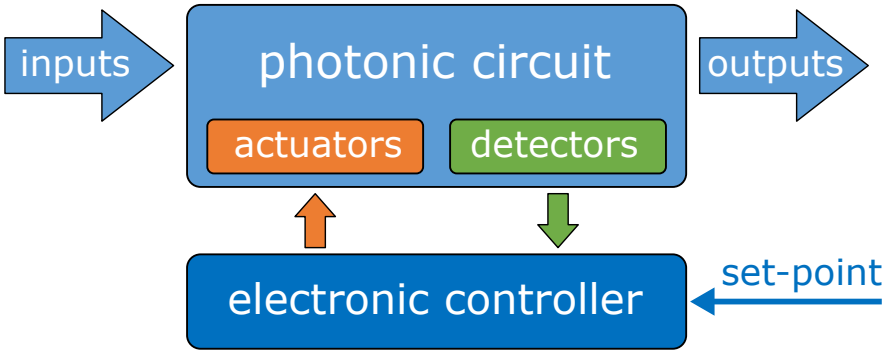

Figure 1.1: *Control scheme of photonic systems. The state of the circuit is monitored with light sensors and stabilized with actuators, driven by an electronic controller.*

ence of temperature variations, aging and mechanical stress;

- operate reconfigurable circuits, that can adapt their behavior to different requirements of the application.

A review of the main sensors and actuators commonly employed is presented in the following. The proposed architecture for the electronic layer is then introduced and discussed.

1.2 State of the art of on-chip light detectors

The behavior of a photonic circuit can be conveniently monitored by placing light sensors directly in the silicon chip. The use of external detectors, that can be extremely optimized in terms of sensitivity and speed of response, is in fact unpractical and not easily scalable to complex architectures, that would require too many optical input/outputs (I/O) to perform an effective monitoring. Research interest is thus focusing on integrated detectors, that trade the optical I/Os for simpler, more compact and robust electrical connections to the external world. The development of detectors easily compatible with the standard fabrication processes and with good sensitivity is an ongoing activity.

On-chip sensors can be divided into two categories: i) absorbing sensors, that rely on the direct absorption of the traveling photons to convert them into free carriers, and ii) transparent

sensors, that instead measure the changes in some electrical properties of the waveguide due to the photons normally absorbed by the silicon, introducing virtually zero extra losses.

Absorbing sensors are fabricated with non-transparent materials and are often more sensitive than transparent sensors [22]. Germanium is a good candidate for realizing integrating photodetectors, since it is an absorbing material easily introduced in a CMOS fabrication process [23]. Bulk germanium has a good absorption coefficient for wavelengths up to 1600 nm, thus including the full O-band and most of the C and L bands. Several papers in the literature present germanium photo-detectors monolithically integrated on silicon obtaining remarkable performance in terms of responsivity, efficiency, current noise and bandwidth [22, 24–26]. For this reason, many commercial foundries offer nowadays germanium-based photodetectors in their PDK [27].

The in-line integration of this class of devices is usually unfeasible, since the light emerging from the structure would be too heavily attenuated. Instead, a splitter is usually employed to spill a small fraction of light from the main path and reroute it to a dedicated detector. Yet, this approach does not scale to complex systems, since each measurement attenuates and slightly perturbs the propagation of light.

Transparent detectors instead can be integrated directly on the path of the light, without extracting any extra photons. For this reason. even if they are usually characterized by a lower sensitivity, they are a promising candidate to enable a mass integration of sensors in the circuit.

1.2.1 Transparent photodetectors

Although silicon is nominally transparent at the wavelengths of interest in telecommunications (1300 and 1500 nm), the flow of light in a waveguide has been observed to naturally increase the population of free carriers in the core due to sub-bandgap absorption phenomena. The photogenerated free carriers cause a variation of the core electrical resistance that can thus be measured to assess the presence of light. In particular, Defect-Mediated Absorption (DMA) [28], Two-Photon Absorption (TPA) [29] and Surface-State Absorption (SSA) [30] are the main phenomena responsible of the intrinsic propagation losses of silicon waveguides.

The addition of a slab to the core is the most common way to electrically contact the waveguide and measure the photo-generated carriers, since it does not introduce any relevant disturb to the propagation of the light.

Various demonstration can be found in literature of this approach. In [31, 32], Defect-Mediated Absorption is exploited thanks to the fabrication of silicon p-i-n diodes where the number of crystal defects is enhanced by implantation of ions in the core. This approach allows to precisely set the sensitivity of the detector, at the price of additional propagation losses. Other works have reported good performance with $B+$ [33] and $He+$ [34] ion implantations, achieving a responsivity of 3 mA/W in a wavelength range from 1530 nm to 1610 nm with lower losses. An advantage of this approach is that these structures can be annealed to effectively remove the ions and the additional losses [35], allowing the integration of multiple detectors for device testing at wafer scale that can be eventually removed before packaging.

Two-Photon Absorption detectors exploit instead a second-order non-linear phenomenon where two photons of identical or different frequencies are absorbed together to excite an atom to a higher energy state. Since it is proportional to the square of the optical intensity, TPA generally has a weaker contribution than first-order phenomena, but can dominate in presence of intense optical fields. For instance, resonant cavities can be employed to enhance the generated photocurrents [36].

Finally, Surface-State Absorption relies on the generation of free carriers in the waveguide thanks to the interaction of photons with intra-gap energy states present at the interface between core and cladding [37, 38]. This phenomenon is particularly relevant in integrated waveguides, as they feature a very high surface-to-volume ratio and because the interface between crystalline silicon and amorphous oxide is a preferential site where intra-gap traps are easily generated [39, 40]. SSA is enhanced by surface roughness of the waveguide walls but would also happen in perfectly smooth devices, due to the different stoichiometry of the two materials of core and cladding [41]. Both p-i-n diodes [42] and photoconductors [43–46] have been demonstrated exploiting this effect.

The sensitivity of transparent detectors reported in literature

is around $-35\,\mathrm{dBm}$. This number depends both on the small amount of free carriers photogenerated in the sensors and on how their readout is performed. The resolution of the electronic front-end determines in fact the minimum light variation that can be detected. To obtain the best possible performance, the development of the sensor must thus be complemented by a corresponding refinement of the readout electronics.

1.2.2 The CLIPP sensor

The most common transparent detectors still affect the waveguide losses, purposely or not. Moreover, they often require either the integration of non-CMOS materials, or the ion implantation of doped silicon, complicating the technological stack. A different approach comes from the CLIPP sensor, a truly non-invasive detector developed at Politecnico di Milano in the past years that is compatible with passive silicon technological platforms [47].

The CLIPP (ContactLess Integrated Photonic Probe) is a sensor that allows to measure the quantity of light inside a silicon waveguide by measuring the variations of conductance of the core induced by the light. However, in order not to perturb the signal, no direct contact to the core is established. The measurement is performed in a contactless way: two electrodes are placed on the cladding of the waveguide, at a sufficient distance so to have no measurable effect on the optical field, and the core of the waveguide is accessed capacitively.

The electrical equivalent model of the sensor is depicted in Fig. 1.2(a), where $\mathrm{C_A}$ is the access capacitance due to the coupling between the electrode and the waveguide and $\mathrm{R_{WG}}$ is the resistance of the portion of the waveguide between the two electrodes. To probe the waveguide resistance R_{WG}, the access capacitance C_A has to be bypassed, so an AC stimulation of the sensor is needed. To do so, an alternated voltage V_{AC} is applied to one electrode (called *force*), and the generated current I_{clipp} is collected from the other electrode (called *sense*) and read with a Trans-Impedance Amplifier (TIA) [48]. Voltages up to $10\,\mathrm{V}$ are usually applied without causing any relevant perturbation of the light propagating in the waveguide [47].

To properly bypass the access capacitance and probe the waveguide resistance, the stimulation frequency should be around

	Germanium PD [22]	Silicon PD [31]	Photoconductor [45]	CLIPP [49]
Footprint	15 µm × 15 µm	40 µm × 5 µm	100 µm × 5 µm	70 µm × 20 µm
Operating frequency	DC	DC	DC - 150 kHz	100 kHz - 10 MHz
Maximum detection bandwidth	40 GHz	1 GHz	500 kHz	10 kHz
Responsivity	750 mA/W	3.3 mA/W* 4 µA/W	-	-
Detectable light @10 Hz BW	−80 dBm**	−60 dBm* −35 dBm	−35 dBm	−50 dBm

* Inside a ring resonator ** Theoretical, shot noise limited.

Table 1.1: *Comparative table of state-of-the-art on-chip light detectors.*

10 times bigger than the access frequency set by the $\frac{C_A}{2} \cdot R_{WG}$ product. Since limiting the working frequency is beneficial both to reduce the noise contribution of the front-end [48] and to facilitate the design of the following electronics, it is recommended to design an access capacitance C_A as large as possible. For this reason, that state-of-the-art implementation of this sensor relies on the use of a slabbed waveguide beneath the access electrodes, as shown in Fig. 1.2(b) and (c). This stratagem allows to increment the coupling between the electrode and waveguide, hence increasing the access capacitance. In this way, it is possible to shorten the CLIPP length L, increasing the conductance of the waveguide G_{WG} and the variation of conductance associated with the luminous signal, increasing, in turn, the sensitivity of the sensor. In Fig. 1.2(d), a microscope picture of a fabricated CLIPP is shown, and in Fig. 1.2(e) the sensitivity curve of the sensor is reported, showing how it is possible to assess an electrical signal associated to −50 dBm of optical power, a 10-fold improvement with respect to previous CLIPP detectors [47], outperforming the resolution of other transparent sensors found in literature [31,45].

1.3 Actuators

In order to actively shift the working condition of the device of interest, actuators have to be employed. The key actuation

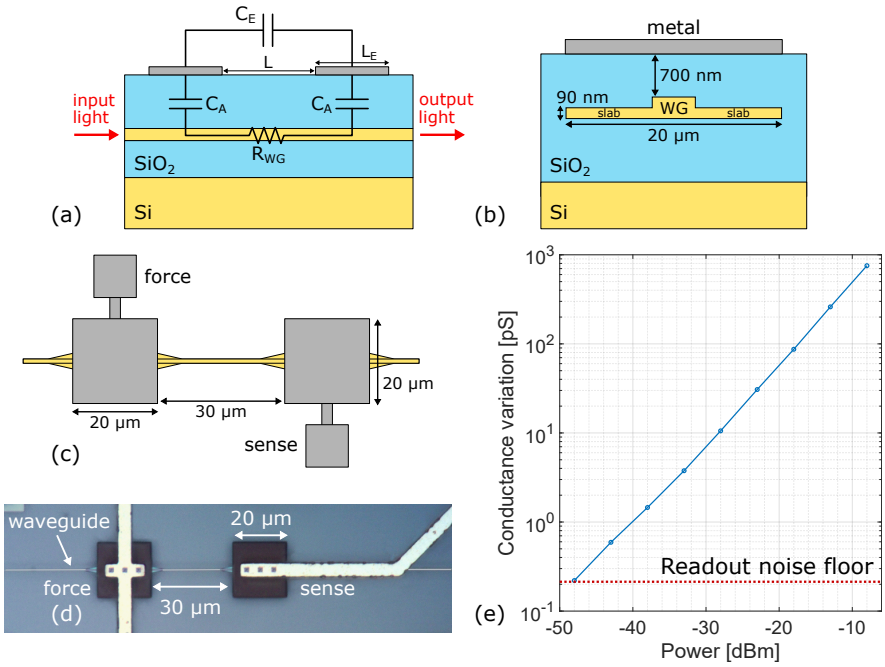

Figure 1.2: *(a) Equivalent electrical model of the CLIPP. (b) Cross-section and (c) top view of the layout in the state-of-the-art implementation. (d) Microscope photograph of the fabricated sensor. (e) Sensitivity curve.*

mechanism is to induce an optical phase shift, possibly without introducing loss. In order to fully operate the device in any working condition, an actuator should be able to induce a phase shift that amounts to at least 2π, defined as the period of the transfer function after which it repeats itself.

The most common approach today is through the use of heaters, that are easily integrated and introduce no losses in the optical field. Yet, they can be very power hungry, so they cannot be relied upon to scale up the dimension and the complexity of a circuit. Very popular are also electro-optical actuators that exploit the free-plasma carrier dispersion effect, that however exhibit losses that depends on the applied phase shift, and can hardly reach an overall 2π phase shift in normal operating conditions. Both these classes of devices have reached a sufficient level of technological maturity so that they are often offered in the PDK of commercial foundries [27]. To overcome the criti-

calities described, a lot of research effort is being spent towards the development of new actuators [50]. In particular, optical micro electro-mechanical systems (MEMS) are becoming more and more popular, but they have not reached the level of maturity to be adopted in large scale yet.

1.3.1 Thermal actuators

The most commonly-employed actuators are the thermal ones. These devices rely on the fact that the refractive index n of a certain material depends on its temperature by a quantity $\frac{dn}{dT}$ called thermo-optic coefficient (TOC). A common way to integrate these class of devices is to fabricate a sheet of metal in the proximity of the device that has to be controlled. By injecting current into this metal, that is usually called heater, power is locally dissipated, and the temperature of the metal and the nearby region increases. Despite its simplicity, this solution has many drawbacks. The most obvious one is that the amount of power dissipated to operate a single heater is in the order of some mW, up to tens of mW [51]. Hence, photonic circuits that rely on heaters cannot be the promised power-free solution for short link communication. Moreover, since the heat will diffuse in other portions of the circuit, the crosstalk between different devices has to be addressed [52]. Many foundries today offer the possibility of digging deep trenches in the SiO_2 in order to thermally isolate a portion of the circuit, increasing in this way the efficiency of the heater while also minimizing the cross-talk.

The bandwidth of these devices reaches some hundreds of kHz [53], that is sufficient to implement a closed-loop control actions to counteract the typical instabilities that affect photonic systems. The heaters can also be used to apply a slow modulation of the optical signal, but not an RF modulation for data-transmission links. Another issue of this class of devices is that its efficacy depends on the TOC of the material. Silicon has a fairly high TOC, around $1.8 \times 10^{-4}\,\mathrm{K}^{-1}$ at $300\,\mathrm{K}$ [54], and so few tens of mW are often enough to obtain 2π phase shift of the controlled device. Other popular materials, such has Silicon Nitride [55] or Titanium Dioxide [56], exhibit lower TOCs, meaning that more power has to be dissipated in order to obtain similar results.

Other types of thermal actuators are being deployed [51]. For instance, it is possible to pass current directly through the core of the waveguide [57], dissipating the power directly where it is necessary. While this solution would limit the quantity of power dissipated to achieve the same phase shift, these devices increase the free carriers in the core of the waveguide, increasing in this way the loss experienced by the optical field. Moreover, this solution is only feasible for waveguides with a semiconductor core. Finally, an active ion implantation step is necessary to electrically contact the waveguide, complicating the technological stack.

1.3.2 Electro-optical actuators

Electro-optical actuators exploits the dependency of the dielectric properties of certain materials from the electrical field. In particular, the refractive index experiences a variation that can be linear with respect to the electric field (Pockels effect) or quadratic (Kerr effect) [58]. The changes in the refractive index induced by the electrical field are however small, meaning that these devices cannot be used to fully tune the working condition of the photonic device. Vice versa, given the high bandwidth of this phenomena, they are suitable to realize RF modulators. Finally, typical materials that exhibit this effect are $LiNbO_3$, $BaTiO_3$, $LiTaO_3$ and others, while instead Silicon and other materials typically used in the CMOS industry do not show any Pockels or Kerr effect. This means that heterogeneous integration would be necessary to fabricate devices exploiting this effect [59, 60].

Another possibility for realizing electro-optical actuators is by exploiting the free-carrier plasma dispersion effect. This effect relates a variation in the free carrier concentration in a semiconductor (electrons and holes) to a change in the complex refractive index of the material. Indeed, as the free carrier concentration increases, the imaginary part of the refractive index increases, while the real part decreases [61]. This means that by embedding the silicon waveguide in a *pn* junction and by changing the applied voltage to its terminals, it is possible to realize an electro-optical modulator [62]. This class of devices features bandwidth above 50 GHz [63], and are then suitable to realize RF modulators. The quantity of phase shift obtained with normal operating voltages is however limited, meaning that they cannot be used to

fully tune the working condition of the controlled device. Finally, since the imaginary part of the refractive index depends on the free carrier concentration, the losses introduced by the modulator change during its operation.

1.3.3 Micro electro-mechanical systems actuators

A promising class of devices for solving the aforementioned issues are the optical micro electro-mechanical systems (MEMS) [64]. In integrated technologies, the phase of a traveling mode is tuned by changing the gap between the waveguide and a free silicon beam, connected to a shuttle displaced by the generated electric field. In [65], a sub-dB insertion loss MEMS achieving more than 2π phase shift was demonstrated. The very high non-linearity of the exponential dependence of the refracting index from the gap of the waveguide and the suspendeded beam is partially compensated by the quadratic relationship between the actuation voltage and the displacement obtained by the MEMS. The device shows a flat pass-band up to hundreds of kHz, comparable to thermal actuators, meaning that electro-optical actuators are still the best option to realize RF modulators.

One issue of these structures comes from the high voltages needed to actuate these devices, usually in the tens of volt. Despite being virtually power-free, such high voltages would require dedicated electronics to be handled. Moreover, the fabrication is not offered by any commercial foundry yet, meaning that the release of the MEMS structure is obtained by post-processing the die in third-party clean rooms [66]. Finally, this approach has only been studied for Silicon Photonics up to now, both because it is the most widespread technology for integrated photonics, but also because the mechanical properties of silicon are well-known and the fabrication can leverage the consolidated know-how of the MEMS industry.

1.4 Modular architecture of the electronic system

Finally, the way to physically implement the electronic system should be discussed. The photonic chip has to be placed on the optical bench in order to provide easy access for the optical signals. At the same time, a single chip could require multiple

electronic components to control all the sensors and the actuators present in the system. For these reasons, the electronic system is conceived in a modular fashion: the photonic chip is placed on a host PCB, shaped to best fit on the optical bench, while all the control electronics is placed on a different general-purpose motherboard, that can be designed for different applications.

1.4.1 Host PCB

Figure 1.3 shows a host board designed for housing a photonic chip. The board is shaped in order to have easy optical access on the east-west side of the chip, so that it is possible to couple light inside the chip, either through grating couplers or edge couplers. Vice versa, all the pads to electrically connect the devices are manufactured in the north-south direction. Furthermore, some applications might require the use of an ASIC, that is best placed next to the photonic chip. For instance, that is the case when using CLIPP sensors, whose electronic front-end is preferably integrated on an ASIC [48], but also for realizing control and calibration schemes directly on-chip [67, 68]. Hence, this chip should be placed on the host board as well, directly connected to the photonic system.

The presence of a host PCB allows also to glue a Peltier cell below the chip to control the global temperature of the system. A commercial thermistor is mounted right next to the chip in order to assess its temperature as precisely as possible. The board is then placed on a metal holder that provides mechanical stability and also acts as a heat sink.

The host PCB is electrically connected to the motherboard with two high-density shielded cables (TE Connectivity 2821385-1), each of them featuring 24 pairs of shielded signals.

1.4.2 Control motherboard

The control motherboard hosts the rest of the electronics necessary to control the optical devices. Since the same motherboard could be used for different photonic chips, the system should be designed to be as general-purpose as possible, even if not every single part of it will be used in every experiment. A photograph of the fabricated board is reported in Fig. 1.4, highlighting the

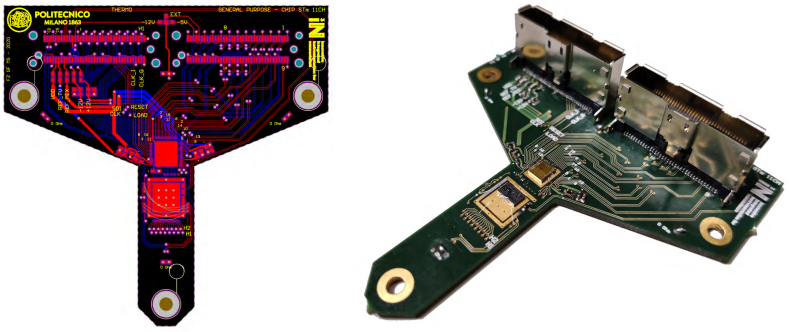

Figure 1.3: *Layout and photograph of the proposed host board.*

main tasks performed by each portion of the board. As also shown in the schematic view of the system reported in Fig. 1.5, the board is designed to perform the following tasks:

- **Acquisition:** the signals coming from the host PCB have to be further amplified, filtered and digitized, without introducing any degradation in the readout resolution. For the purpose, a low-noise front-end INA is mounted, followed by a 4th-order Butterworth anti-alias (AA) filter with a cut-off frequency of $f_{pole} = 150\,kHz$. After that, a programmable gain amplifier (PGA) further amplifies the signal to match the dynamic range of the ADC. Finally, a high-precision 16-bit ADC samples the signal at 625 kSps. Notice that this value is higher then twice the cut-off of the AA filter, so that it is possible to digitally filter out the quantization noise ("oversampling"). The board features 16 parallel acquisition chains.
- **CLIPP stimulation:** the CLIPP sensors need to be stimulated with a sinusoidal signal of programmable frequency (50 kHz - 10 MHz) and amplitude (1 V - 10 V) to adapt to different geometries and fabrication technologies. The board mounts only one CLIPP stimulation chain, since the preferred way of stimulating the sensor is to share the forcing pad between all the CLIPPs integrated on the chip [69].
- **Actuation:** the board can control up to 16 actuators on the photonic chip. A voltage accuracy of few mV is necessary, with a maximum voltage of 10 V. Since a maximum current

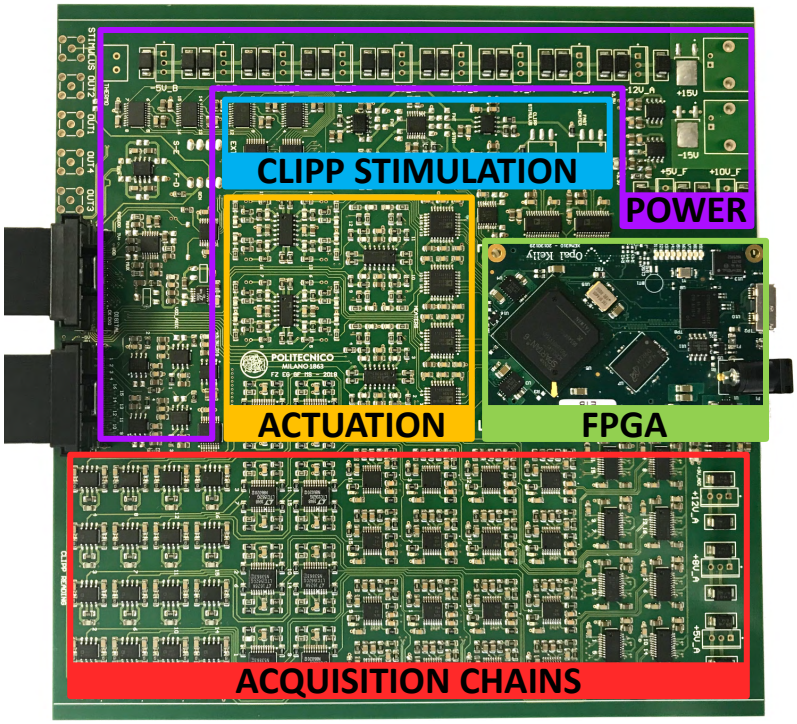

Figure 1.4: *Photograph of the motherboard, highlighting its main sections.*

of around 50 mA is required, a current driver is mounted for the purpose. The actuation chain is also able to produce sinusoidal signals up to tens of kHz, to generate the low-frequency thermal modulation necessary for the dithering technique.

- **Digital processing:** an FPGA performs the digital functionalities necessary to manage the components of the motherboard, process the acquired signals, properly drive the actuators and communicate with a personal computer. The controllers for stabilizing the photonic devices of interest are also implemented in the FPGA.

The choice of the FPGA and the details of the implementation of the digital system will be thoroughly discussed in Chapter 6.

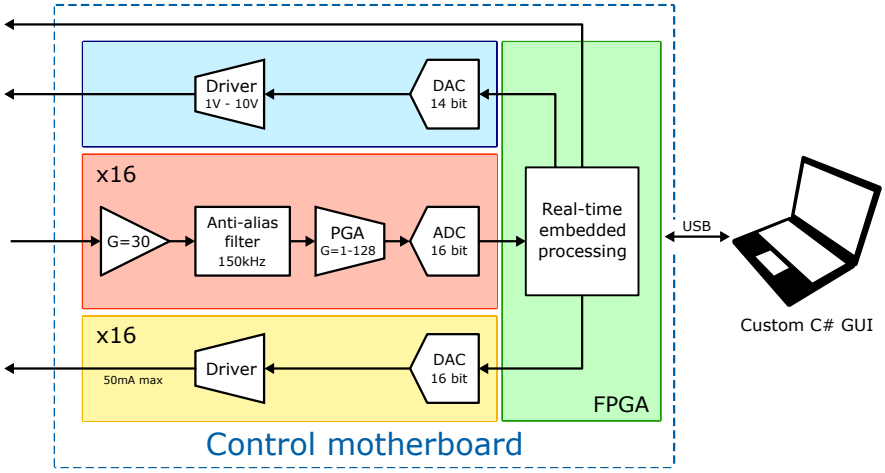

Figure 1.5: *Schematic view of the complete control platform.*

1.5 Organization of the thesis

This work falls within the scope described so far. A control system to stabilize the working condition of different photonic circuits is here presented. In particular, the control loop developed is both power-independent and calibration-free, and it is able to stabilize photonic devices both on the minimum (or maximum) of their transfer function, and in the maximum-slope condition. This means that the proposed assembly can be employed in complex heterogeneous photonic circuits operated in harsh conditions, such as modern datacenters. The strategies adopted were also extensively tested, and the main results are here reported. The digital implementation of the system is also discussed, with a particular care towards the digital signal processing blocks necessary for the purpose. Finally, a plasmonic-assisted sensor is designed, tested and validated in a simple control experiment. More in details, the thesis is organized as it follows:

- in chapter 2, the control strategy adopted, based on the dithering technique, is discussed in details, highlighting the main advantages of the proposed scheme with respect to other implementations. A control loop based on an integral controller is also described and developed, showing how the system is able to lock on the stationary points of the trans-

fer function of any photonic device. All the parameters of interest and their effect on the performance of the system are then fully described;

- in chapter 3, the control strategy previously described has been experimentally validated on a complex reprogrammable photonic circuit consisting of a mesh of Mach-Zender interferometers. To control the chip, only one external benchtop photodiode is used, exploiting the concept of orthogonality between different frequencies and frequency selectivity of the lock-in readout. The pilot tones technique is also introduced and its use in conjuncture with the dithering technique is examined. A mode-sorting experiment is reported and discussed;

- in chapter 4, the proposed control strategy is expanded in order to use the same control scheme for integrated modulators, that cannot be operated on the stationary points of their transfer function. By exploiting the device nonlinearities and extracting the second-derivative of the transfer function, it is possible to stabilize the system in the maximum-slope of the device transfer function, close to the optimum operation condition for a modulator. The system is experimentally validated on an integrated micro-ring modulator;

- in chapter 5, the whole assembly has been used to control the operation of a heterogeneous photonic chip integrating both resonant switches and ring modulators, that have to be stabilized on two different working conditions. As anticipated, the same control scheme can be used for both the classes of devices, and the system was correctly operated in long and slowly-varying scenarios;

- in chapter 6, the digital architecture implemented is discussed in details, highlighting the role of the FPGA in complex systems and presenting all the necessary implementations and optimizations to achieve the best results. In particular, the chapter includes the description of the main digital signal processing blocks necessary to implement the

lock-in for the extraction of the dithering signal and the integral controller used in the experiments;

- in chapter 7, a plasmonic-assisted bolometric sensor is discussed. With the aid of FEM simulations, the detector has been dimensioned to avoid undesired losses. The proposed geometry has been fabricated and extensively tested to fit the simulations results. At the end of the chapter, the sensor is finally used in a simple control experiment, using a minimum-chaser algorithm implemented on the external custom software interface.

CHAPTER 2

Control system architecture

In this chapter, the control system architecture is presented. In particular, the dithering technique, that allows to extract the first derivative of the transfer function of any photonic device, is thoroughly discussed. A control loop with an in integral controller is presented and all the main parameters of the system are discussed. Advanced use of the dithering technique to reduce the complexity of the assembly are finally introduced at the end of the chapter.
The results presented in this chapter were published in:

F. Zanetto, **V. Grimaldi**, F. Toso, E. Guglielmi, M. Milanizadeh, D. Aguiar, M. Moralis-Pegios, S. Pitris, T. Alexoudi, F. Morichetti, A. Melloni, G. Ferrari, M.Sampietro, "Dithering-based real-time control of cascaded silicon photonic devices by means of non-invasive detectors," IET Optoelectronics, vol. 15, no. 2, pp-111-120, 2021, doi: 10.1049/OTE2.12019.

2.1 The need for closed-loop control

As thoroughly explained in Chapter 1, once the sensor has transduced the information into the electrical domain, the electronic layer has to control the actuators in order to shift and tune the working condition of the device of interest. To do so, a common approach is the construction of look-up tables [70, 71] and the operation of the system in an open-loop fashion. This solution becomes quickly unfeasible as the complexity of the circuit grows, since it might render the necessary calibrations virtually impossible. In fact, a complete calibration of the system would require to take into account all the possible operating conditions of the assembly. Moreover, the crosstalk between different actuators should also be considered, further complicating the process [52]. Finally, the functional drifts and aging of the components might change the collective optical response of the circuit over time, invalidating the calibration performed in the first place.

A closed-loop control of the photonic devices based on a real-time monitoring of their working condition thus emerged as a better and more scalable alternative [21]. Several approaches have been proposed, with controllers implemented by means of FPGAs [72–74], microcontrollers [75], ASICs [67, 76], discrete components [77], or any combination of these, depending on the required level of parallelization, power consumption and flexibility required by the specific application.

2.1.1 Control techniques

Several locking techniques have been proposed in literature to stabilize the operation of photonic devices. The easiest way to do so would be to monitor the output power level [77] of the system and, whenever a change is detected, to shift the working condition of the device. However, a change of the output power level is not only due to a different working condition of the system, but it might be also due to an overall drift of the input power. This is particularly ineffective in cascaded photonic systems, where the change of optical power measured after a certain device can be the consequence of something happening upstream.

A possible solution would be to add a detector to monitor both the input and the output power of the controlled device [76].

This process would complicate the assembly and would require an extensive calibration of the system, nor it would be robust against the aging of the elements. It is thus necessary to develop a control scheme that is both power-independent and calibration-free to be applied to different complex and heterogeneous photonic systems.

The control scheme of choice in this work relies on the use of the dithering technique [78, 79], that will be now discussed in details. Since thermo-optic phase shifters are by far the most popular choice for actuating a photonic circuit, all the discussion in the following will be carried out implying their use. However, any actuator that has a sufficient bandwidth and dynamic range can be in principle chosen for the task.

2.2 The dithering technique

The dithering technique allows to extract in real-time the first-derivative of the transfer function of an optical device directly from the physical system and use it for control purposes. In particular, the power emerging from a photonic device can be maximized (or minimized) by driving to zero the first-derivative signal, since this condition indicates stationary points of the transfer function. This approach is suitable for cascaded architectures, because the set-point does not require any calibration and does not depend on the absolute amount of light that reaches each device. In addition, in case of variations of the optimal working point, it is possible to understand the direction of the shift from the sign of the derivative, leaving no ambiguity on the control action to be taken. Finally, variations of the dark readout baseline do not affect the measurement, allowing to effectively use the dithering readout as error signal of a robust closed-loop control system. The main drawback of this approach is that, to perform the extraction of the derivative information, only a small fraction of the optical signal reaching a sensor is exploited. Consequently, to obtain a proper sensitivity in the detection of light and efficiently use the dithering information in a control scheme, a narrow readout bandwidth is usually required. The speed of a dithering-based system might thus be slower than the other implementations mentioned before. A more complex lock-in based electronics is also required to extract the dithering information

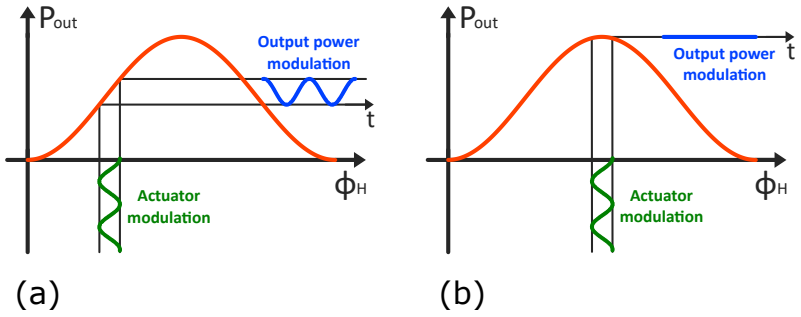

Figure 2.1: *Working principle of the dithering technique. (a) By driving an actuator with a modulated signal, an oscillation of the transfer function around its bias point is induced, causing an amplitude modulation of the output optical power. (b) When the device is in a stationary point, the amplitude modulation observed is null.*

from the sensor readout, as explained in the following.

2.2.1 Theory of operation

In order to extract the transfer function derivative, a small modulation signal is superimposed to the actuator voltage that controls the optical device of interest, causing an oscillation of the transfer function around its bias working point and a consequent modulation of the optical power at the output, as qualitatively depicted in Fig. 2.1(a). Notice that the wavelength of the laser is fixed, so it is the oscillation of the transfer function to cause a modulation of the output power. The modulation depth of the output power will depend on the slope of the transfer function, i.e. its first derivative, around its bias condition. This means that, by synchronously demodulating the output power signal, it is possible to extract the information about the first derivative of the transfer function around its working condition and use it for control purposes.

Formally, the idea is to linearize the output optical power $P_{OUT}(W)$, where W is the power dissipated by the heater, around the bias condition, here indicated as W_0. So it is

$$P_{OUT}(W) \approx P_{OUT}(W_0) + \frac{\partial P_{OUT}(W)}{\partial W}(W - W_0) \qquad (2.1)$$

A small modulation is being added around the bias condition of

the heater. This means that the power dissipated by heater can be expressed as

$$W = W_0 + W_{dith} \sin(2\pi f_{dith} t) \tag{2.2}$$

where W_{dith} is the amplitude of the modulation and f_{dith} is the frequency. Thus, the output optical power will be

$$P_{OUT}(W) \approx P_{OUT}(W_0) + \frac{\partial P_{OUT}(W)}{\partial W} [W_{dith} \sin(2\pi f_{dith} t)] \tag{2.3}$$

Equation (2.3) suggests that the output power can be expressed as the sum of a DC average value and an oscillation centered at f_{dith} whose amplitude is related to the first derivative of the optical power, as qualitatively expected.

Equation (2.3) also shows that the modulation of the output power P_{dith} is proportional to the dithering amplitude W_{dith}. To maximize the dithering signal, it is then possible to increase the amplitude of the modulation. However, depending on the application, a large modulation may disturb the operation of the optical device controlled. Therefore, the amplitude of the dithering sinusoid must be chosen carefully, evaluating the trade-off between a higher signal-to-noise ratio (SNR) and the perturbation of the optical functionality [78]. Moreover, the linearization remains a reasonable approximation only for small oscillation of the heater power, i.e. only until the oscillation is comparable with the optical features of the device of interest. For instance, a device with multiple maxima and minima requires an amplitude small enough to observe a response related to the single peaks.

It should be also noticed that, when a stationary point is reached, the transfer function is flat in that region, and so the residual perturbation is minimized, as qualitatively depicted in Fig. 2.1(b). Therefore, when locking to the maximum or minimum of a transfer function, a larger dithering signal is usually allowed, as compared to other working points. Although it might increase the complexity of the system, adaptive dithering amplitude could be envisioned for optimum results.

The considerations made so far are valid for a perfectly sinusoidal dithering power oscillation. However, resistive heaters are often driven by a voltage source. The approximation of perfectly

sinusoidal dithering oscillation is valid by using small dithering signals ($v_{dith} \ll V_0$) or by predistorting the actuator driving signal [68] to compensate for the quadratic relation between command signal and dissipated power.

2.2.2 Lock-in readout

The extraction of the dithering information can be obtained by placing a detector directly at the output of the photonic device. The signal measured by the sensor will contain both the DC information due to the average power and a modulated contribution at f_{dith}. This dithering fraction can be conveniently isolated and extracted with a lock-in processing. With respect to other implementations, like peak-stretchers or envelope detectors, the lock-in technique has the advantage of being frequency and phase selective, thus allowing to isolate and extract only the useful information and reject spurious signals and disturbances. In addition, the mixing action of the lock-in has the beneficial effect of upconverting the 1/f noise of the readout chain to high frequency, allowing to extract the small dithering signal with the highest resolution possible. In the system used, the lock-in scheme is implemented in the digital domain, allowing to simplify the electronic design of the motherboard and ensuring the maximum possible flexibility.

In principle, any periodic waveform can be used to modulate the actuator voltage. For instance, a square-wave modulation can be adopted, since it requires minimum electronics for the synthesis [67]. However, the high-frequency components of a square-wave would couple more easily with the nearby electrical lines, complicating the design for densely-integrated photonic and electronic system. Despite being more difficult to synthesize, a sinusoidal signal centered f_{dith} would result in a better spectral purity of the output power. In the following, a sinusoidal dithering oscillation will be assumed.

Regarding the choice of the dithering frequency f_{dith}, the upper limit is determined by the thermal relaxation time of the heater T_h, usually in the order of few µs [51]. The lower limit is instead due to the $1/f$ noise corner frequency of the readout electronics, typically around few hundreds of Hz. To maximize the system performance, the dithering frequency chosen should be between these two values, in the range of some kHz to some

tens of kHz.

2.3 Dithering-based automated control architecture

Locking the working point of a device to the maximum/minimum of its transfer function requires driving the dithering readout to zero. Several approaches can be used to minimize a signal, like stepper or gradient descent algorithms [80, 81]. The use of an integral controller is here presented and discussed. With respect to other techniques, an integral controller can be easily implemented both in digital and analog systems. In the following, its closed loop behavior is thoroughly analyzed in the continuous-time domain.

2.3.1 Control scheme

As an example of general validity, let us consider the locking of a microring resonator (MRR) obtained by means of a thermal actuator and a CLIPP sensor on its through port, as shown in Fig. 2.2(a). Figure 2.3 shows the transfer function of the MRR and its first-derivative signal, obtained by sweeping the value of the heater power and synchronously monitoring the CLIPP mean value and dithering signal. As expected, the zero of the first derivative corresponds to the minimum of the MRR transfer function.

The schematic block diagram of the proposed control system is reported in Fig. 2.2(b). To bring the MRR at resonance, the heater voltage V_H should be driven to a certain value $V_{H,RES}$. This value $V_{H,RES}$ is a function of the wavelength λ, of the instantaneous temperature of the system T, of the quantity of optical power inside the system [82], and so on. This means that the resonance condition changes during the normal operation of the device. The role of the control loop is then to track the variations of $V_{H,RES}$ in real-time and apply to the heater the correct voltage V_H to maintain the chosen set-point. To do so, the system exploits the extracted dithering amplitude A_{dith} to monitor the MRR working condition, and drive the heater so that the desired derivative value corresponds to the set-point $\bar{d}$ at equilibrium.

To understand the qualitative behavior of the control system, let us assume to target the resonance condition of the MRR. This

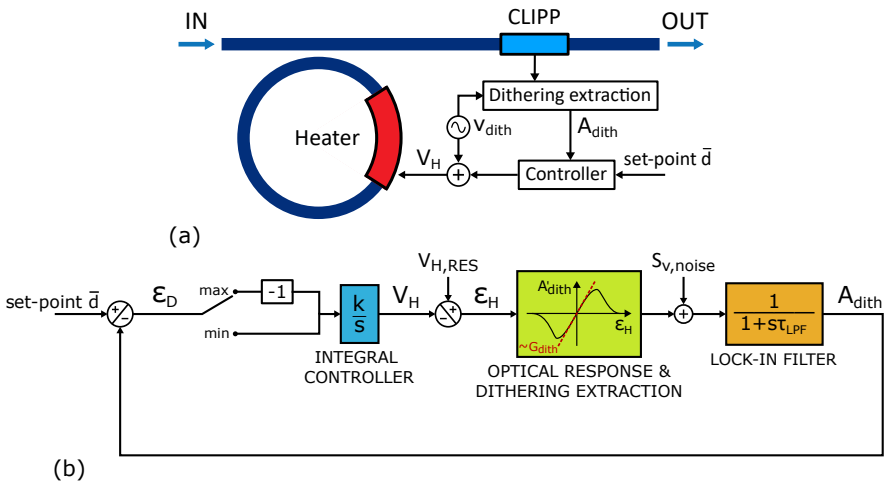

Figure 2.2: *(a) Example of the dithering-based control scheme applied to a microring resonator with a CLIPP sensor on its through port. (b) Block diagram of the closed-loop architecture that can be implemented to lock the working point of photonic devices, based on the dithering technique and an integral controller.*

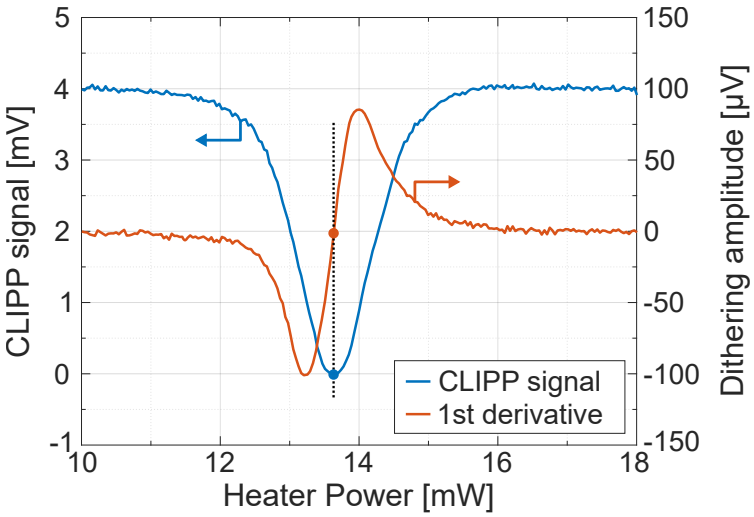

Figure 2.3: *Transfer function (in blue) and dithering signal (in orange) of a microring resonator extracted with a CLIPP. For the dithering signal, a* 2 mV *voltage oscillation, corresponding to around* 50 µW *of power oscillation, at* 6 kHz *has been used. The lock-in bandwidth was set at* 10 Hz*.*

means, as already stated, that the first-derivative has to be driven to zero, i.e. that $\bar{d} = 0$. If the controlled heater voltage does not match the resonance condition, the extracted dithering signal will have a non-zero value that depends on the instantaneous difference $\mathcal{E}_H$ between the target value $V_{H,RES}$ and the actual heater voltage V_H. This signal A_{dith} is then fed to an integral controller, that will update its output until its input is zeroed again. By connecting the output of the integral controller to the heater voltage, the system will keep moving until the measured dithering signal is zero.

Notice that this operation is in principle valid for any target point $\bar{d}$. For $\bar{d} \neq 0$, the integral controller updates V_H until the difference between V_H and $V_{H,RES}$ keeps the dithering function at the desired value. While technically feasible, this approach does not result in a control system that is power-independent, since any change in the input optical power will also change the amplitude of the dithering signal. Vice versa, the stationary points of a certain optical system do not change during its operation, meaning that the position of the zeroes do not depend on the quantity of power present in the system.

Finally, an inverting stage can be used to change the sign of the error signal $\mathcal{E}_H$, so that the system is able to lock on both maxima and minima of the transfer function of interest, depending on the requirement of the application.

2.3.2 Computation of the loop gain

As depicted in Fig. 2.2, the sensor readout can be described with a single *optical response & dithering extraction* block, that relates the heater detuning $\mathcal{E}_H$ to the measured dithering amplitude A_{dith}, and a single-pole low-pass filter having a transfer function $1/(1 + s\tau_{LPF})$, where τ_{LPF} is the time constant of the lock-in readout stage. The filter is needed to remove the harmonics resulting from the dithering demodulation, that would otherwise impair the loop functionality, and to reduce the noise of the system. Although the relation between heater voltage and measured dithering amplitude is generally described by a non-linear function, the case of small perturbations around resonance is here assumed for simplicity, meaning that a linearized model can be used. In this case, the *optical response & dithering ex-*

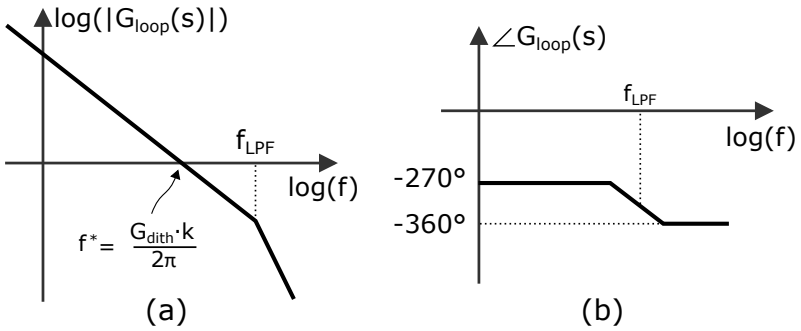

Figure 2.4: *Qualitative Bode plots of (a) magnitude and (b) phase of the loop gain of the proposed control scheme*

traction block can be simplified and considered as a constant gain G_{dith}, representing the slope of the curve around resonance. An integral controller, with transfer function k/s, completes the feedback loop. The parameter k defines the gain of the controller and can be tuned to ensure the stability of the system.

Based on this discussion, the loop gain G_{loop} of the system can be computed as:

$$G_{loop}(s) = -\frac{G_{dith}}{1 + s\tau_{_{LPF}}} \cdot \frac{k}{s} \tag{2.4}$$

The qualitative Bode plots of the loop gain are shown in Fig. 2.4. The Bode criterion states that a system is stable if the phase margin

$$\Phi_m = 360^\circ - \angle G_{loop}(s)|_{f=f^*} \tag{2.5}$$

is larger than 0°, where f^* is the unity gain frequency, i.e. the frequency at which the Bode plot of the magnitude of the loop gain crosses the 0-dB axis. For real electronic systems, the condition $\Phi_m > 45^\circ$ is usually considered the actual limit for having acceptable performance of the control loop. This condition is met as long as the sensor readout pole $f_{_{LPF}} = 1/2\pi\tau_{_{LPF}}$ is placed after the zero-crossing point of the loop gain. In this case, the closed-loop bandwidth of the system can be computed by solving for $|G_{loop}| = 1$:

$$BW_{CL} = f^* = \frac{G_{dith} \cdot k}{2\pi} \tag{2.6}$$

As expected, the bandwidth depends on G_{dith}, meaning that

a higher dithering signal would lead to a faster control loop, but also on k. Indeed, since the dithering signal cannot be increased indefinitely without disrupting the optical functionalities of the controlled device, the bandwidth can be tuned by choosing the value of k, that hence defines the maximum frequency at which the feedback loop can react to variations of the ring resonance voltage.

The verge of the stability of the system is the value of k for which $f^* = f_{LPF}$, that would lead to a phase margin $\Phi_m = 45°$. However, given the non-linearities of the device and the dependency of G_{dith} from the quantity of power inside the circuit, this condition should be avoided, since it may lead to an oscillatory behavior of the system.

Error signal

The effect of the loop on the behavior of the system can also be observed by computing the expression of the error signal $\mathcal{E}_D$, that is

$$\mathcal{E}_D(s) = \frac{\bar{d}(s)}{1 - G_{loop}} - \frac{V_{H,\,RES}(s) \cdot G_{dith}}{1 + s\tau_{_{LPF}}} \cdot \frac{1}{1 - G_{loop}} \tag{2.7}$$

that, at frequencies below the lock-in readout pole, when $\bar{d}$ is set to zero, can be approximated as

$$\mathcal{E}_D(s) \approx -\frac{V_{H,\,RES}(s)}{k} \cdot \frac{s}{1 + \frac{s}{G_{dith} \cdot k}} \tag{2.8}$$

The qualitative Bode plot of this signal is reported in Fig. 2.5(a). A high-pass action is clearly visible, meaning that, within the feedback bandwidth, the loop works to minimize the error signal, thus keeping the extracted derivative A_{dith} equal to zero and consequently keeping the ring at resonance. As the frequency increases, the loop cannot react any more to variations of the ring working point and the error signal follows the fluctuations imposed by $V_{H,RES}$. This behavior confirms that the proposed control strategy can be effectively exploited to lock the resonance of photonic devices against fluctuations within the feedback bandwidth.

Noise vs. bandwidth trade-off

As already stated, by changing the value of the integral controller gain k it is possible to tune the bandwidth of the control loop to satisfy the requirements of the application. The upper limit of the bandwidth is not only set by the stability of the feedback loop, but also by the maximum value of the actuation noise acceptable in the application, according to the well-known trade-off between noise and bandwidth.

In order to quantify the relationship between the bandwidth of the control loop and the residual noise, the transfer function from the readout noise $S_{v,noise}$ to the actuation voltage V_H should be computed. From Fig. 2.2, it is:

$$G_{noise}(s) = \frac{k}{s} \cdot \frac{1}{1 - G_{loop}} \approx \frac{1}{G_{dith}} \cdot \frac{1}{1 + \frac{s}{G_{dith} \cdot k}} \tag{2.9}$$

The readout noise is here considered as the main noise source and it is modeled with a voltage generator placed after the dithering readout block, since it represents the noise superimposed to the extracted dithering signal in the real system. A qualitative Bode plot of the noise transfer function is reported in Fig. 2.5(b). It is possible to observe that the noise of the readout is transferred to the actuation voltage for frequencies up to BW_{CL}, and then low-pass filtered. Consequently, the RMS noise that affects the actuation voltage for a given bandwidth can be computed as:

$$\begin{aligned} V_{H,noise} &= \sqrt{S_{v,noise}} \cdot \frac{1}{G_{dith}} \cdot \sqrt{\frac{\pi}{2} \cdot BW_{CL}} = \\ &= \sqrt{S_{v,noise}} \cdot \sqrt{\frac{k}{4 \cdot G_{dith}}} \end{aligned} \tag{2.10}$$

where $S_{v,noise}$ is the voltage noise power spectral density of the readout. This is true only if the contribution from $S_{v,noise}$ can be considered white noise. This is indeed the case, since the 1/f noise is removed thanks to the use of the lock-in technique. Equation (2.10) clearly states that the RMS value of the actuation noise increases by increasing the bandwidth of the system. Hence, a lower k provides lower actuation noise and consequently a more accurate locking action, with the drawback of a lower control

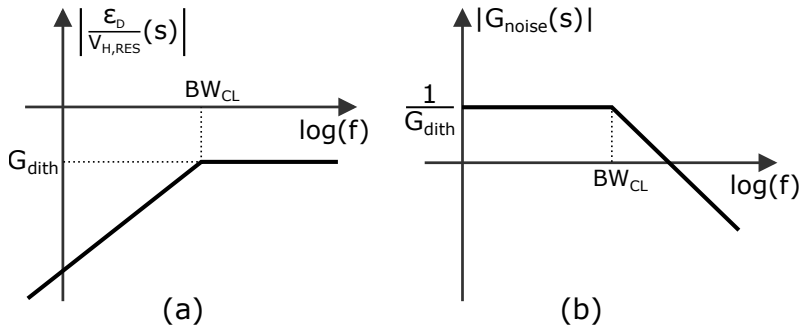

Figure 2.5: *Qualitative Bode plots of the (a) error signal and (b) the actuation noise*

bandwidth. The choice of k is therefore the result of a careful trade-off between noise and speed of response and must be tuned depending on the requirements of the application. Notice that a noise reduction of a factor n is obtained at the price of a bandwidth narrowing of a factor n^2, since the RMS noise is proportional to the square root of the bandwidth.

2.4 Advanced uses of the dithering technique

The system explained so far implements a power-independent and calibration-free control scheme, but without the right set of precautions it can easily become too complex and congestioned for densely-integrated photonic circuits. To scale-up the system, it is important to understand the limiting factor of each application. For instance, one architecture might benefit from a reduction in the number of sensors used in the circuit, since they might require too many electrical lines to be properly contacted and acquired. Alternatively, a system may benefit from a reduction in the number of dithering frequencies necessary to implement the control. In the following, some advanced uses of the dithering technique are discussed, showing how there is of a substantial trade-off between the number of sensors adopted and the number of frequencies necessary for the system.

2.4.1 Orthogonal dithering for discrimination of multiple actuators

The dithering technique can be effectively used also to discriminate the effect of multiple actuators, while still using only a single detector. This feature is very useful in case of devices that require more than one heater to be operated, like Mach-Zehnder interferometers (MZI), or in cascaded structures where many actuators affect the output optical power. In these situations, operating by only looking at the average output power requires complex techniques to take into account thermal crosstalk effects and perform an effective tuning action [71].

The orthogonality principle can be instead leveraged to easily separate the contribution of each actuator when using the dithering technique. Recalling that:

$$sin(2\pi f_1 t) \perp sin(2\pi f_2 t) \quad \text{if} \quad f_1 \neq f_2 \tag{2.11}$$

by using different frequencies for each actuator, the corresponding dithering signals can be independently recovered from the lock-in based sensor readout. Indeed, by simultaneously demodulating the CLIPP output at the different frequencies, each dithering can be extracted in parallel, having in this way the information of the effect of all the actuators at the same time. Notice that some residual crosstalk can arise due to the spurious harmonics produced either by the dithering generator or by the nonlinearities of the photonic device transfer function. This effect can be completely suppressed by choosing the dithering signals at non-harmonic frequencies. The phase selectivity of the lock-in technique can also be exploited to use couples of dithering signals at the same frequency, but in quadrature with each other. In this way, the number of frequencies to be generated in case of complex structures can be halved.

Experimental validation of orthogonality principle

An experimental demonstration of this principle was carried out on a MZI featuring two heaters and only one sensor on the output port, as shown in Fig. 2.6(a). A continuous-wave (CW) laser at 1550 nm was split and launched to both inputs of the MZI to perform the experiment. The voltage of the heaters was scanned

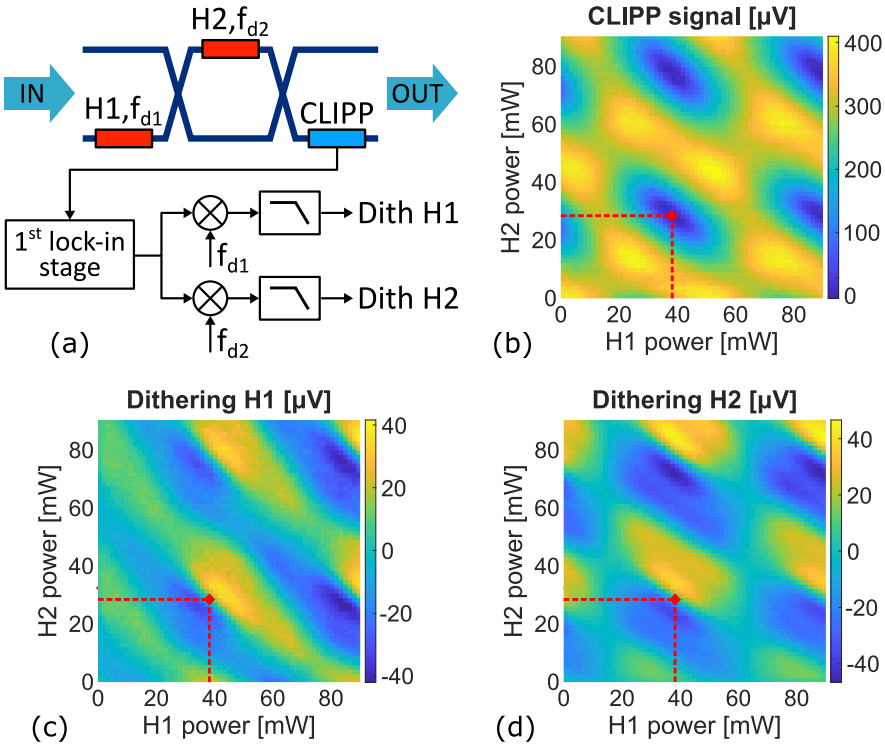

Figure 2.6: *(a) Mach-Zehnder interferometer control scheme when using two heaters with orthogonal dithering signals. (b) Color map of the MZI transfer function obtained with a CLIPP detector at the lower output branch. (c), (d) Partial derivatives of the MZI transfer function, extracted in parallel by using orthogonal dithering signals. The three maps, made of 60×60 points, were obtained simultaneously with an acquisition time of* 5 ms *per pixel.*

from 0 V to 6 V to obtain a map of the output light, measured with a CLIPP for each combination of heater powers, as shown in Fig. 2.6(b). Two dithering signals of 50 mV amplitude at 5 kHz and 7 kHz respectively were applied to the heaters and the CLIPP readout was demodulated in parallel at the two frequencies to extract the partial derivatives of the output power. Figure 2.6(c) and 2.6(d) show the obtained results. The system was able to correctly discriminate the effect of the two actuators and recover the respective dithering information. The optimal working point to completely steer the light on the top branch (red dot in Fig. 2.6(b), (c), (d)) can now be easily found by driving the two heaters with two completely independent control loops, as discussed in Section 2.3.

2.4.2 Dithering frequency re-use in cascaded systems

By leveraging the orthogonality principle, it is then possible to use a single sensor to discriminate the effect of different actuators, as long as the frequencies of their dithering modulations are different. Vice versa, it is possible to use the same dithering frequency for all the devices in a cascaded structure, as long as a dedicated sensor is used to control each of them. The peculiar nature of cascaded photonic systems helps in achieving this goal: their tuning procedure is in fact inherently sequential, as one stage can be completely configured only when the previous one has reached its final working point. This mode of operation makes it possible to recycle dithering signals.

Indeed, when a photonic device is configured to work in a stationary point, the residual dithering modulation that survives at its output is minimized, thanks to the flatness of the transfer function in that region. In addition, since photonic devices are normally symmetric around their maxima/minima, the frequency of the residual oscillation is doubled with respect to the dithering one. These two properties can be exploited to control cascaded architectures using the same dithering modulation for all the optical devices, as they guarantee that ideally no interaction is observed when the equilibrium condition is reached. The interaction between stages is indeed only happening during an initial transitory phase, but it does not prevent to reach the correct working point.

The configuration of a cascaded architecture, obtained with a single dithering frequency, thus works in a sequential way. The first stage, whose output contains only the information of its own dithering, can be configured immediately. While this tuning is in progress, the downstream stages receive an optical signal containing also the modulation of the first device, so they cannot correctly complete their configuration. When the first stage is set, the residual dithering oscillation at its output is minimized, so the second device can be tuned correctly. All the stages can be thus sequentially controlled until the full architecture is configured, making this procedure well suited for cascaded systems.

In real operations a small interaction between stages might survive due to offsets and non-idealities, but its impact is usually negligible. The residual modulation of a stage gets attenuated by each device it passes through, so crosstalk effects are observed almost only between consecutive stages. In order to minimize this issue, a couple of orthogonal dithering signals can be used, rotating them every two devices. In this way, optimal operation is still possible while keeping the complexity of the system reasonably low.

CHAPTER 3

Control of a reprogrammable photonic processor

This chapter demonstrates an application of the dithering technique on a large optical system. The circuit consists of a 9x2 mesh capable of separating and independently measuring two orthogonal beams. By leveraging the orthogonality principle of different dithering tones, a single commercial benchtop photodiode is used to extract up to 16 independent dithering signals, that are then fed to 16 integral controllers to stabilize the system. The pilot tones technique is also introduced and an experimental demonstration is carried out.

Part of the results presented in this chapter were published in:

M. Milanizadeh, S. SeyedinNavadeh, F. Zanetto, **V. Grimaldi**, C. De Vita, C. Kliti, M. Sorel, G. Ferrari, D.A.B. Miller, A, Melloni, F. Morichetti, "Separating arbitrary free-space beams with an integrated photonic processor", Light Sci Appl 11, 197 (2022), doi: 10.1038/s41377-022-00884-8.

3.1 The need for reprogrammable photonics

Most of the photonic circuits fabricated today are designed for a specific application, so that, in analogy to their electronic counterparts, are often called application-specific photonic integrated circuits (ASPICs). These systems can be tailored on the exact application desired and as such can be heavily optimized, resulting in very compact and power-efficient assemblies. However, the increase in complexity and in the number of different tasks required by each system has sparked an interest in a new generation of photonic circuits, that can be programmed and adapted to the needs of the application during the normal operation of the system [83].

These programmable PICs are based on the idea that the flow of light on the chip can be manipulated at run-time. To do so, electrically-controlled tunable beam couplers have to integrated and operated. Mach-Zender interferometers (MZIs), realized by cascading two directional couplers, are a 2x2 optical gate that can perform any linear operation between input and output, and as such are perfect candidates for the purpose [84, 85]. In order to change at run-time the operation performed, two adjustable parameters are necessary for each MZI, so to independently control the power splitting at the two outputs and the relative phase delay at the two inputs. Two independent actuators are hence necessary for each MZI.

The individual MZIs can then be arranged in different ways, in order to create waveguide meshes that implement various linear functions by interfering the light signal along different paths. These structures are usually divided into forward-only meshes, where the light flows from one side of the mesh to the other [85, 86], and recirculating meshes, where the light can be routed in loops and even back to the input ports [87]. A particular class of forwards-only meshes are the triangular meshes, that can implement any real-time matrix–vector multiplication [85], an essential operation in a variety of applications, such as information processing [88–90], neuromorphic computing and artificial intelligence [91–93], cryptocurrencies [94], and so on. In the following, a bi-diagonal mesh is introduced and controlled.

3.2 Description of the photonic circuit

The circuit used for this application is schematically depicted in Fig. 3.1(a), and was fabricated on a standard 220 nm SiP platform by AMF foundry. A microscope photo of the fabricated circuit is reported in Fig 3.1(b). It is a 9x2 diagonal mesh, comprising of 8+7 balanced MZIs designed to work at around the 1550 nm wavelength. The single-mode channel waveguides are designed with a width of 500 nm, and each MZI is realized with two 3 dB directional couplers, implemented by two 40 µm-long waveguides with a gap of 300 nm. Two TiN heaters (2 µm x 80 µm) are integrated in each MZI stage, one in the lower waveguide, one in the upper internal interferometer arm. Since each MZI needs 2 heaters to be fully configured, the circuit encompasses 30 heaters overall.

To couple light into the chip, a 2D 3x3 array of grating couplers (GCs) optimized for TE transmission is fabricated on the left of the chip. The nine waveguides connecting the GCs to the photonic processor share the same optical length in order to minimize the wavelength dependence of the multipath interferometer implemented by the mesh, so that the circuit can have the widest possible wavelength range of operation [85]. On the right side of the mesh, the top two waveguides outputs are used as output ports, while the remaining 7 waveguides outputs are available for monitoring and control purposes.

3.2.1 Theory of operation

The circuit was designed for free-space optics (FSO) operations. A certain beam of light with a random spatial phase distribution is shined onto the chip and gets sampled by the GCs array. If the MZIs are not programmed, the light is randomly distributed at the 9 output waveguides. Once each heater is tuned to maximize the optical power reaching the top arm of each MZI in the first diagonal, all the light can be rerouted towards the single top output. This means that, by measuring the output power, it is possible to assess the overall quantity of light that was launched in the circuit in the first place, since the rerouting operation is virtually lossless. Moreover, the voltage supplied to the heater to steer the input light into a common output is a function of

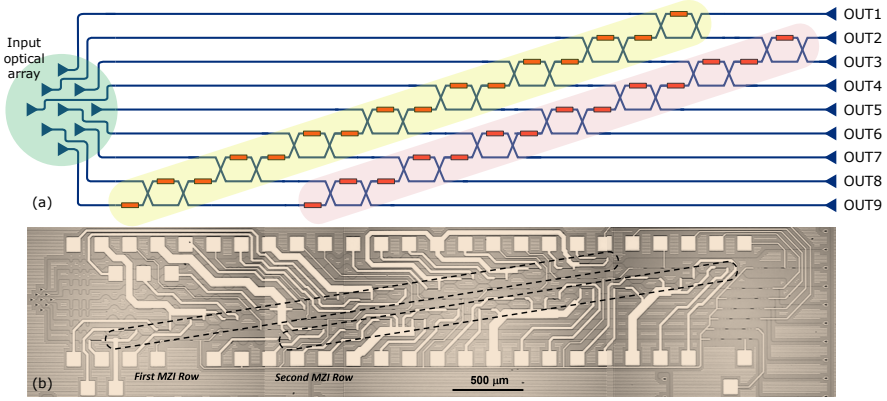

Figure 3.1: *(a) Schematic view and (b) microscope picture of the circuit used in this application. The two diagonals are highlighted.*

the spatial phase difference of the input light sampled by the GCs. This means that, with a prior calibration [95], it is possible to have an indirect measurement of this phase difference. The implication is that the circuit is able to measure both the intensity and the spatial phase distribution of the impinging beam.

Once the first diagonal of the mesh is fully configured, a second beam of light, orthogonal to the first one, is again shined onto the GCs array. Formally, two beams are defined as "orthogonal" if they lead to orthogonal complex vectors of amplitudes in the input waveguides inside the circuit [96]. For instance, two different modes in the free-space optical field, or two beams coming from different angles, can be orthogonal between each other. Due to the orthogonality of the two guided modes, the first diagonal of the mesh is transparent to this second mode, and all the light is hence forwarded to the second diagonal. Once again, the heaters in the second diagonal can be tuned to have the light redirected to the second output of the photonic processor. This means that the circuit is able to separate and individually measure the intensity and the spatial phase distribution of two sets of orthogonal input mode. This reasoning can be further expanded, meaning that an MxN mesh is able to sample, separate and individually measure N modes with a spatial resolution that depends on the number M of input optical antennas [85].

Thus, this system can be adopted to create high-capacity

transmission links based on multibeam space-division multiplexing (SDM) FSO communication [97]. Moreover, the adaptive nature of the photonic processor would also allow the possibility to compensate for dynamic changes in the FSO link, caused by, for instance, moving obstacles or atmospheric turbulence, maintaining in this way the optimum communication link [98]. Finally, many applications that require advanced processing of FSO beams using these kind of circuits can be envisioned, such as phase-front mapping and reconstruction, imaging through scattering media [99], chip-to-chip optical wireless communication [100] and many others [101].

3.3 Control scheme

In this application, each MZI has to steer all the input light towards the top output. The easiest way to do so is to integrate a sensor on the bottom output arm of each MZI and use it to minimize the measured quantity of light, as depicted in Fig. 3.2(a). The dithering technique, together with the control scheme explained in Section 2.3, can be exploited for the purpose, since by zeroing the first derivative extracted by the sensor the system gets locked on the correct stationary point. Since a single sensor would have to discriminate between the effect of two heaters, the orthogonality principle should be leveraged, as also explained in Section 2.4.1. In particular, the in-phase a quadrature component of a single frequency can be exploited, in order to reduce the number of frequencies that have to be generated by the system. Once the system is locked on a stationary point, the residual oscillation due to the dithering signal would be minimized and shifted at twice the oscillation frequency, meaning that the same frequency could also be reused in the upstream stages, as also explained in Section 2.4.2.

The presence of a dedicated sensor on each MZI is not always a viable possibility. For instance, the technology platform chosen might not offer active materials in the integration process (such as Ge for Germanium photodetectors [22], or ion implantation for doped-silicon sensors [44]). Moreover, each sensor would need its dedicated line of electronic components to be properly acquired and digitized, as depicted in Fig. 3.2(a). In the 9x2 mesh de-

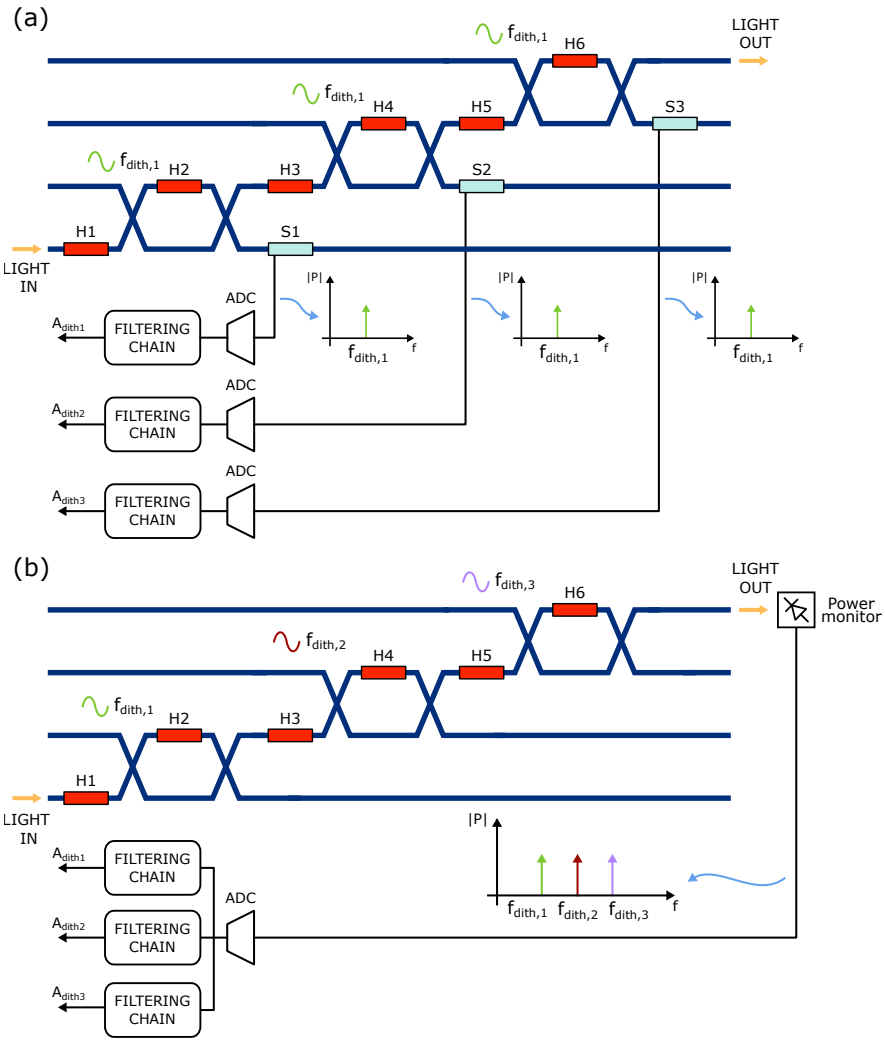

Figure 3.2: *Schematic view of a cascaded structure. In (a), each sensor is placed after every MZI, so each of them extracts the dithering information related to a single device. In (b), a single external photodetector collects the information related to three different MZIs. In order to discriminate the effect of each device, orthogonal tones have to be used.*

scribed above, in addition to the 30 electrical lines to control the heaters, other 15 lines for the sensors would be necessary. This might make the design too crowded and limit the scalability of the assembly [83].

A possible solution is the use of a single benchtop commercial power monitor to measure the emerging power from the output, as sketched in Fig. 3.2(b). In order to discriminate the effect of each heater, all the dithering tones would have to be orthogonal, meaning that the frequency-reuse technique cannot be used in this case. This means that both the in-phase and the quadrature component of 8 different dithering frequencies have to be generated and acquired to control the first diagonal of the mesh. This solution offers an easier electronic setup, but comes at the expense of the maximum speed of the system, since the low-pass filter bandwidth has to be narrow enough to reduce the effect of the unwanted harmonics in the output spectrum.

3.3.1 Guidelines for dimensioning the dithering frequencies

To better understand the limiting factor in the speed of the system, let us refer again to the situation sketched in Fig. 3.2(b). Each filtering chain acquires a signal with the spectrum depicted in the figure, and each of them has to separate the information of interest from the other ones. By extracting the i_{th} dithering tone, the spurious information of the $(i+1)_{th}$ tone will be down-converted by this operation at the frequency $f_{dith,i+1} - f_{dith,i}$, that will then be filtered out by the low-pass filtering stage of the lock-in. This means that the limiting factor of the speed of the system is the minimum distance between the two generated dithering tones. Thus, the best practice is to use equally-spaced dithering tones.

In reality, all the non-linearities of the system should be considered. Indeed, due to the non-perfect spectral purity of the generated sinusoidal signals, but also due to the non-linearities of the MZIs, the sensors and the actuators, the higher harmonics of each signal might mix with other dithering tones. This means that two dithering frequencies should never be close multiples with each other.

Given all these constraints, the practical way to implement a system of this kind is the following. First, a certain fundamental

frequency f_d should be chosen, that will be the spacing between two spectral lines, and as such it will be the limit for the bandwidth of the control loop. Then, all the dithering tones necessary will be multiples of this frequency, and they should all be prime between each other. For this system, the 8 dithering frequencies necessary for the control should be placed at $2f_d$, $3f_d$, $5f_d$, $7f_d$, $11f_d$, $13f_d$, $17f_d$, $19f_d$. If more frequencies are necessary, the same reasoning applies.

The upper limit in the chosen f_d depends on the maximum frequency that has to be fed to the actuators and their bandwidth. In case of thermal actuators, all the generated frequencies should not exceed 200 kHz. For this system, the maximum chosen frequency is 120 kHz, leading to a fundamental frequency of $f_d = 120\text{kHz}/19 \simeq 6\text{kHz}$.

3.4 Experimental result

Once the system has been correctly dimensioned, some preliminary control experiments were conducted, in order to assess the efficacy of the implemented scheme and the convergence time of the loop. The dependency of the convergence time vs. the number of MZIs in the mesh was also investigated. All the experiments reported in the following sections are performed by shining an FSO beam onto the input GCs array. The beam has a wavelength of 1550 nm and is TE-polarized in order to match the polarization sensitivity of the GCs.

3.4.1 Tuning of the mesh

A single FSO beam was shined while the control was inactive and the heaters were arbitrarily initialized, meaning that the power gets randomly distributed to all the output waveguides of the mesh. Then, after 50 ms, the control was turned simultaneously on the whole structure and the output power gets maximized by the system, as shown in Fig. 3.3. The experiment was then repeated starting from another random condition: to do so, the control was turned off, all the heaters voltages were randomized and then the control was again turned on. As expected, the power level measured at OUT1 before the control is activated changes from one experiment to the other, since it is a function of the

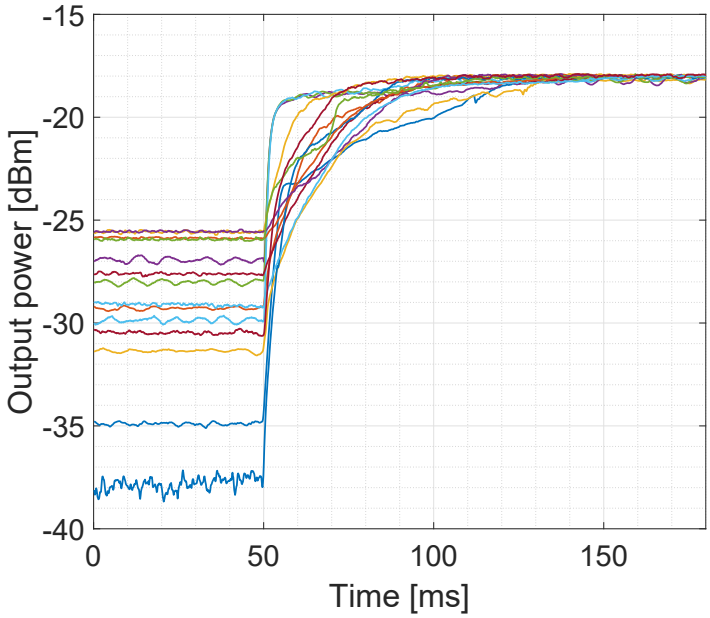

Figure 3.3: *Experimentally measured configuration transient of the full mesh, repeated for different starting conditions and overlapped for a convenient comparison.*

random value of the heaters. Vice versa, the output power restored by the control scheme is always the same, certifying how the system is able to maximize the output power regardless of the starting condition of the mesh. In Fig. 3.3, all the results of the different experiments are reported and overlapped for convenience. The rise time, defined as the time necessary to go from 10% to 90% of the final value, ranges from 10 ms to 80 ms in the worst-case condition.

3.4.2 Partial tuning

In order to assess how the rise time and the performance of the control system change with the number of stages, the same experiment was repeated on a smaller portion of the whole system. In practical terms, this means that, while the control was kept active for the first MZIs of a diagonal, the same routine described above (i.e. heaters randomly initialized and then control turned on simultaneously on all the remaining structures) was repeated

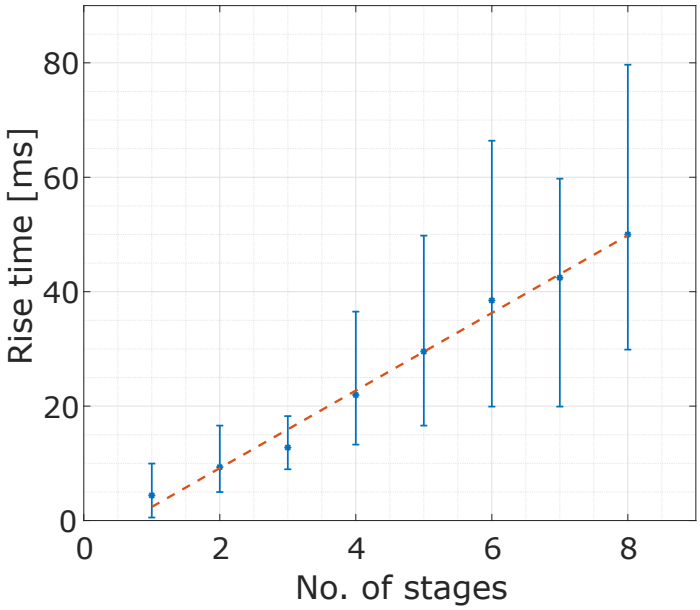

Figure 3.4: *Configuration time of the mesh as a function of the number of stages controlled. The error bars indicate the variance over the experiments performed for each configuration.*

only for the downstream sub-mesh. In this way, from a control standpoint, the same circuit could be used as it were a smaller mesh. For each possible sub-mesh, the same control experiment was then repeated 10 times, in order to gather a statistical distribution of the obtained performance. The results are shown in Fig. 3.4, where it is possible to see that the average convergence time is linear with respect to the number of stages controlled. This is expected, since the procedure is inherently sequential: each stage can be completely configured only when the previous one has reached its final working point. Figure 3.4 also shows how the variance of the convergence time significantly broadens as the dimension of the mesh increases. The qualitative reason is that, since the convergence time depends on the starting working condition of the single MZI, the statistic of each component sums up independently from each other.

3.5 Mode separation

As previously described, the system has to be able to separate two orthogonal beams and route them to the corresponding output port of the circuit. However, if the two beams are shined at the same time, there is no way to discriminate one from the other in order to calibrate and control the corresponding diagonal. Indeed, the signals passing through the first diagonal would be indistinguishably modulated by the dithering oscillation, and the sum of the two beams would be routed to the first output. A possible solution comes from the use of the pilot tones technique [102], that allows to individually label the incoming beams and separate their effect. Indeed, the pilot tones technique, used in conjuncture with the dithering technique, creates spectral components related to the individually-labeled tones, that can be independently extracted and used for control purposes.

3.5.1 Pilot tones

The pilot tones technique is used to discriminate between different optical signals regardless of the presence of other concurrent channels in the same waveguide. Indeed, pilot tones can be used to label specific signals by introducing a small modulation of their optical intensity. For instance, this might be useful in wavelength division multiplexing (WDM) systems, that rely on the use of multiple carrier wavelengths at the same time to expand the capacity of the network. In this case, it is useful to discriminate between two orthogonal beams, namely IN1 and IN2, shined on the mesh at the same time.

If the total power traveling in the waveguide due to the two beams is

$$P = P_{IN_1} + P_{IN_2} \tag{3.1}$$

by applying a small modulation to the first carrier, the power is

$$P = \left(P_{IN_1} + \tilde{P}_{IN_1} \sin\left(2\pi f_{tone} t\right)\right) + P_{IN_2} \tag{3.2}$$

where f_{tone} is the frequency of the pilot tone only applied to first beam and $\tilde{P}_{IN_1}$ is the quantity of optical power modulated. This means that, by synchronously demodulating the power at the

frequency f_{tone}, it is possible to extract the information regarding $\tilde{P}_{IN_1}$, that is only related to the first carrier. By properly labeling all the input carriers with different frequencies it is then possible to isolate the effect of each optical signal. This comes at the expense of the extracted power, that is now only a fraction of the overall power reaching the sensor. To increase the extracted signal, a bigger fraction of the signal $\tilde{P}_{IN_1}$ can be modulated. However, an excessive modulation of the input power may disrupt the optical functionality of the system, so the modulation depth should be decided on a case by case basis.

Notice that, while the implementation of the pilot tone technique is similar to the dithering technique (superimposing an oscillation to the optical power and synchronously demodulating the information of the sensor), the two are substantially different. Indeed, the pilot tone technique is used to label the individual carriers and discriminate one from the other, while the dithering technique is used to extract the first derivative of a photonic device around its working condition.

3.5.2 Pilot tones with the dithering technique

With the use of the dithering technique the information of the first derivative gets modulated around the frequency of the oscillation of the heater power. If a portion of the light is modulated with a pilot tone, two intermodulated spectral components appear at the sum and the difference of the pilot tone and the dithering frequency. To prove this result, the steps are very similar to the ones already shown in Section 2.2 for explaining the dithering technique. However, since the input power is not constant, it is now best to linearize the transfer function $H(W)$ of the device of interest around an average heater power W_0, rather than directly P_{OUT}.

Let us consider for the sake of simplicity to have only one carrier modulated by a pilot tone, so that

$$P_{IN} = P_{IN,0} + \tilde{P}_{IN} \sin(2\pi f_{tone} t) \tag{3.3}$$

where $P_{IN,0}$ is the average power of the beam and $\tilde{P}_{IN}$ is the quantity of power modulated by the pilot tone technique. The

output power will be

$$\begin{gathered} P_{OUT}\left(W\right) = P_{IN} \cdot H\left(W\right) = \\ \left[P_{IN,0} + \tilde{P}_{IN} \sin\left(2\pi f_{tone} t\right)\right] \cdot H\left(W\right) \end{gathered} \tag{3.4}$$

By linearizing $H(W)$ around its working condition and by considering a sinusoidal oscillation at f_{dith} of the heater power W, it is

$$\begin{aligned} P_{OUT}(W) = P_{IN,0} H\left(W_0\right) + \tilde{P}_{IN} \sin\left(2\pi f_{tone} t\right) H\left(W_0\right) + \\ + P_{IN,0} \frac{\partial H(W)}{\partial W} \sin\left(2\pi f_{dith} t\right) + \\ + \tilde{P}_{IN} \frac{1}{2} \frac{\partial H(W)}{\partial W} \cos\left[2\pi\left(f_{dith} \pm f_{tone}\right) t\right] \end{aligned} \tag{3.5}$$

For clarity, let us define the fraction of power modulated by the pilot tone with respect to the average value as

$$r = \frac{\tilde{P}_{IN}}{P_{IN,0}} \tag{3.6}$$

so that Equation (3.5) can be rewritten as

$$\begin{aligned} P_{OUT}(W) = P_{OUT,0} + r P_{OUT,0} \sin\left(2\pi f_{tone} t\right) + \\ + \frac{\partial P_{OUT}}{\partial W} \sin\left(2\pi f_{dith} t\right) + r \frac{1}{2} \frac{\partial P_{OUT}}{\partial W} \cos\left[2\pi\left(f_{dith} \pm f_{tone}\right) t\right] \end{aligned} \tag{3.7}$$

where $P_{OUT,0} = P_{IN,0} H(W_0)$ is the average power emerging from the device considered. The spectrum of the output power is depicted in Fig. 3.5. The average power $P_{OUT,0}$ is still correctly at DC, the contribution due to the pilot tone $r P_{OUT,0}$ is still correctly modulated around f_{tone}, the first derivative of the transfer function remains modulated around f_{dith}, but two more contributions, centered around $f_{tone} \pm f_{dith}$ appear. These components are proportional to the first derivative of the transfer function as well, but with respect to the contribution modulated around f_{dith} they are scaled down of a factor $r/2$. This means that, by extracting the signal around $f_{tone} \pm f_{dith}$, it is possible to extract the first derivative of the output power only related to the labeled tone considered.

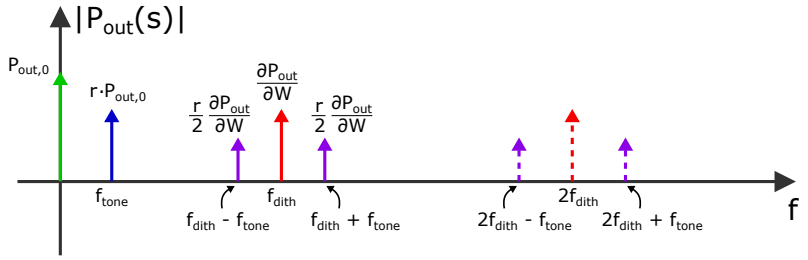

Figure 3.5: *Spectrum of the output power when both the pilot tone technique and the dithering are used. The direct contribution due to the pilot tones is in blue, the dithering and higher-order harmonics in red (solid and dashed, respectively), and the intermodulated tones in purple.*

This signal can be simply extracted by performing a double demodulation in cascade: the first one at f_{dith}, so that the intermodulated tones gets downconverted at f_{tone}, and then a second one at f_{tone}, to shift them to baseband. Another possibility is a single demodulation directly at $f_{tone} \pm f_{dith}$. Both the solutions achieve the same SNR, as reported in the appendix A, so the most convenient one should be chosen depending on the details of the system adopted.

3.5.3 Dimensioning of the pilot tones frequencies

Let us assume that the two input beams have been individually labeled, each with its own $f_{tone,1}$ and $f_{tone,2}$. As described in 3.3.1, by choosing a fundamental frequency f_d and by placing the dithering tones at prime multiples of this frequency, in the spectrum there will be a spectral component every f_d, either a dithering tone or a higher-order harmonic (but never both). The intermodulated spectral components generated by the pilot tones technique have to be located in between a certain $N \cdot f_d$ and the following $(N+1) \cdot f_d$. In between these two, the number of intermodulated spectral components that has to be accommodated is twice the number of the pilot tones used in the system: indeed, in this frequency range, there would be the high-frequency components of the first dithering harmonic ($Nf_d + f_{tone,1}$ and $Nf_d + f_{tone,2}$) and the low-frequency of the second one ($(N+1)f_d - f_{tone,1}$ and $(N+1)f_d - f_{tone,2}$). The result can be generalized, meaning that, to find the fundamental

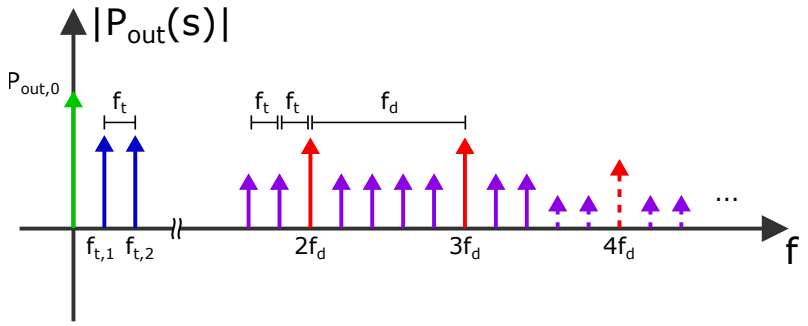

Figure 3.6: *Spectrum of the signal acquired by the photodiode at the end of the mesh. All the spectral components are spaced of a factor f_t.*

distancing of the pilot tones frequency, it is

$$f_t = \frac{(N+1)\cdot f_d - N\cdot f_d}{2\cdot \#\text{ of pilot tones} + 1} = \frac{f_d}{2\cdot \#\text{ of pilot tones} + 1} \tag{3.8}$$

In this case, since two tones have to be allotted, for $f_d \simeq 6\,\text{kHz}$, it is $f_t \simeq 1.2\,\text{kHz}$. Hence, the two pilot tones necessary for the experiments are $f_{tone,1} \simeq 1.2\,\text{kHz}$ and $f_{tone,2} \simeq 2.4\,\text{kHz}$. The possible non-linearities introduced by the pilot tones have not been considered, since they only depend on the performance of the external modulator used to label the input signal.

The resulting spectrum, depicted in Fig. 3.6, presents now a spectral component every f_t. This means that the use of the pilot tones technique comes also at the expense of the maximum bandwidth achievable by the control, since a narrower low-pass filtering action is now necessary in order to remove the downconverted unwanted harmonics.

3.5.4 Experimental results

Once the system was correctly dimensioned, the capacity of the mesh to separate orthogonal modes was finally assessed. For the purpose, a second photodiode on OUT2 was used to control the second diagonal of the mesh. The same dithering frequencies used for the first diagonal can be used for the second diagonal as well, since the information of interest is intermodulated with the specific pilot tone frequency.

As a first example of an application, it is shown how the

integrated photonic processor can be operated as a direction-diversity receiver, i.e. as a multibeam receiver capable of individually detecting beams that simultaneously arrive from different directions. This concept is demonstrated here by considering 2 beams, but this functionality can be generalized to N beams arriving from N directions utilizing a photonic processor with N diagonals of MZIs. As schematically shown in Fig. 3.7(a), two free-space beams, namely IN1 and IN2, with identical Gaussian shape, wavelength (1550 nm) and polarization status (TE polarization, to match the polarization sensitivity of the grating couplers) are shined from two different directions onto the 2D optical antenna array with a relative angle of 1.25°. This implies that, in this case, the orthogonality is given only by the direction of arrival of the beams. The control is then turned on, and the power at the respective output ports is assessed.

To quantify the mode-sorting capability, the level of crosstalk is measured. To this aim, the quantity of light coming from IN1 emerging from OUT2 is assessed. This can be easily done since, as shown in Fig. 3.6, in the output power spectrum there is a spectral line around $f_{tone,1}$ that is only associated to the beam IN1. The same operation is performed to quantify the light coming from IN2 emerging from OUT1. The results are reported in the bar chart in Fig. 3.7(b), showing how more that 25 dB optical crosstalk suppression is observed. If the mode sorting status is inverted, meaning that IN2 is sent to OUT1 and IN1 to OUT2, the same level of optical isolation is observed, proving how the converging of the mesh is independent from the initial status of the system and from the frequency chosen for each pilot tone.

Such a low crosstalk level allows to use the integrated photonic processor as a direction-diversity receiver in a FSO communication system, where the optical beams are employed to transmit two independent data channels. To this end, the two beams IN1 and IN2 were used for the transmission of independently externally-modulated 10 Gbit/s on-off keying (OOK) signals. At first, the heater of the MZIs are arbitrary initialized, thus the two data channels are randomly overlapped at both output ports OUT1 and OUT2 and the measured eye diagrams are completely closed (Fig. 3.7(c)). Vice versa, once the control is activated,

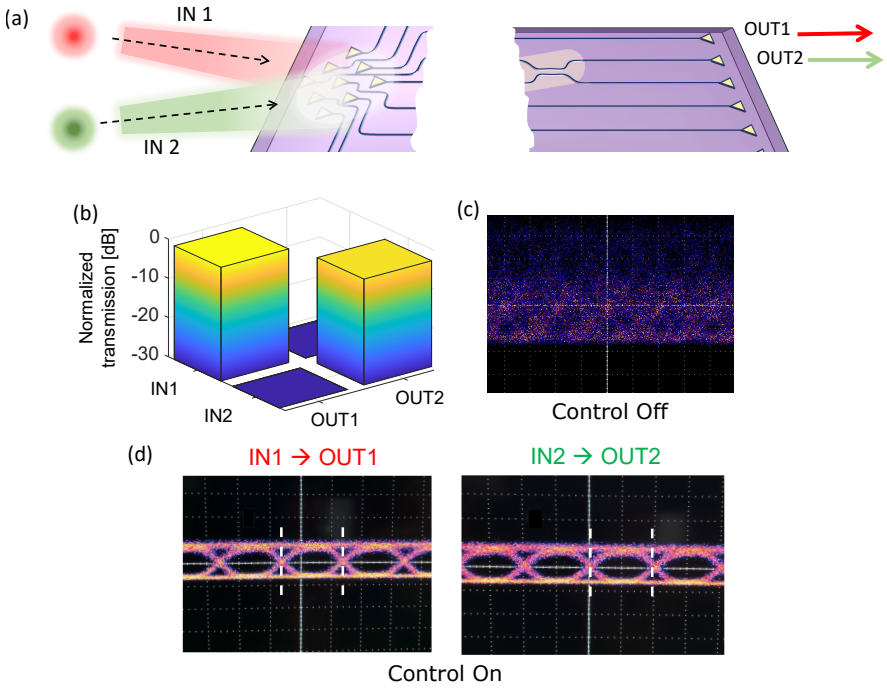

Figure 3.7: *(a) Schematic representation of two free-space beams (IN1 and IN2), sharing the same wavelength and state of polarization, and arriving at the receiver from different directions, with a relative angle of 1.25°. (b) Bar chart showing the normalized insertion loss of the beams IN1 and IN2 at the output waveguides OUT1 and OUT2. (c) The two beams were independently modulated by a* 10 Gbit/s *OOK signal. The eye diagram measured at one of the two ports when the photonic processor is not configured is completely closed. (d) After the control is activated, both eyes at the output ports are clearly open with neither evident distortion nor inter-symbol interference.*

open eye diagrams are recorded at both outputs (Fig. 3.7(d)), each corresponding to the correct transmitting channel. Both the eye diagrams show no degradation with respect to the reference eye diagram of the individual channels.

As already stated, in principle, any set of orthogonal beams can be separated and measured by the mesh. Indeed, the same experiment has been successfully conducted with very similar results between two free-space beams sharing the same wavelength (1550 nm) and state of polarization (TE), coming from the same direction, but being shaped according to different orthogonal spatial modes [103]. In this case, the circuit operates as a multibeam mode-diversity receiver.

CHAPTER 4

Second-derivative dithering-based control system for integrated modulators

This chapter presents a novel technique to stabilize the working condition of a photonic modulator. The technique is an extension of the dithering technique and relies on the extraction of the second-derivative signal by exploiting the intrinsic non-linearities of photonic devices, resulting in a scheme that is both power-independent and calibration-free. The approach is experimentally validated on a Silicon Photonics micro-ring modulator (MRM).

The results presented in this chapter were published in:

V. Grimaldi, F. Zanetto, F. Toso, I. Roumpos, T. Chrysostomidis et al., "Self-Stabilized 50 Gb/s Silicon Photonic Microring Modulator Using a Power-Independent and Calibration-Free Control Loop", in Journal of Lightwave Technology, vol. 41, no. 1, pp. 218-225, 1 Jan.1, 2023, doi: 10.1109/JLT.2022.3210756.

4.1 Extraction of the second derivative

To discuss the features of the proposed control scheme, let us refer to an MRM with an integrated heater to shift its resonance frequency and a sensor on its through port to monitor the output power, such as the one sketched in Fig. 4.1. The control strategy relies on the use of the dithering technique. As fully explained in Section 2.2, by adding a small oscillation to the heater dissipated power, the device output light is modulated around its average value. The modulation depth is proportional to the first derivative of the device transfer function around its working condition. By synchronously demodulating the signal measured by the detector at the output with a lock-in amplifier, it is possible to recover such information and use it for control purposes. However, a locking strategy exploiting this information is power-independent only when targeting the zero of the measured dithering signal, since the amplitude of the modulation of the optical output power depends also on the quantity of light inside the circuit. The optimum working condition for a modulator instead is located on the slope of the transfer function of an MRM.

For this reason, a new approach has been investigated for this class of devices. The proposed scheme relies on the fact that, given the non-linearity of the ring resonator transfer function, a perfectly sinusoidal modulation of the heater power causes a non-perfectly sinusoidal oscillation of the output power. Indeed, by expressing the heater power as

$$W = W_0 + W_{dith} \sin\left(2\pi f_{dith} t\right) \tag{4.1}$$

where W_0 is the average power, W_{dith} the amplitude and f_{dith} the frequency of the dithering oscillation, the transfer function $H(W)$ of the MRM can be Taylor-expanded around its bias working condition, so that the optical output power is

$$P_{OUT}(W) = P_{IN} \cdot \Bigg\{ H\left(W_0\right) + H'\left(W_0\right) \left[W_{dith} \sin\left(2\pi f_{dith} t\right)\right] + \\ + \frac{H''\left(W_0\right)}{2} \left[W_{dith} \sin\left(2\pi f_{dith} t\right)\right]^2 ... \Bigg\} \tag{4.2}$$

By limiting the Taylor expansion to the second-order term, Equation (4.2) can be rewritten as

$$P_{OUT}(W) = P_{IN} \cdot \Bigg\{ H(W_0) + H'(W_0) W_{dith} \sin(2\pi f_{dith} t) + \\ + \frac{H''(W_0)}{2} W_{dith}^2 \left[\frac{1}{2} - \frac{1}{2} \cos(2\pi 2 f_{dith} t) \right] \Bigg\} \tag{4.3}$$

suggesting that, by demodulating the sensor signal at twice the heater oscillation frequency, it is possible to extract the information regarding the second derivative of the MRM transfer function in a specific working point. By tuning the actuator DC bias to drive the second-derivative signal to zero, it is thus possible to lock the MRM in the point of maximum slope regardless of the quantity of optical power circulating inside the chip, making the control calibration-free and robust against power fluctuations. A feedback scheme based on an integral controller, such as the one previously discussed in Section 2.3, can be exploited for the purpose.

The considerations made so far are valid for a perfectly sinusoidal dithering power oscillation, so that the second derivative extracted from the optical output can be attributed only to the MRM transfer function non-linearity. In the case of resistive heaters driven by a current or a voltage source, this requirement is satisfied by using small dithering signals (amplitude of the dithering oscillation much smaller than the actuator bias value) or by pre-distorting the actuator driving signal [68] to compensate for the quadratic relation between command signal and dissipated power.

4.1.1 Optimum working condition of a MRM

To have a quantitative estimation of the functionalities of the proposed control loop, it is necessary to introduce the figures of merit commonly used to describe the performance of a MRM, that are i) the insertion loss (IL), defined as

$$IL = \frac{P_{in}}{P_{avg}} \tag{4.4}$$

where P_{in} is the input optical power and P_{avg} is the average power at the MRM output during its operation, determined by the device working point; and ii) the extinction ratio (ER), defined as

$$ER = \frac{P_1}{P_0} \tag{4.5}$$

where P_1 and P_0 are the transmitted power at the MRM output for a logical '1' and for a logical '0', respectively. In order to identify the theoretical optimum point for transmission of a MRM, the transmitter penalty (TP) is also used [104], defined as

$$TP[dB] = -10\log\left(\frac{P_1 - P_0}{2P_{in}}\right) = -10\log\left(\frac{1}{IL}\cdot\frac{ER-1}{ER+1}\right) \tag{4.6}$$

The theoretical optimal working condition for a modulator is found when the TP is minimized. In such operating point, a good compromise between modulation efficiency and insertion loss is observed. Indeed, the minimum TP does not coincide with the point of maximum ER, because by moving closer to the MRM resonance the insertion loss increases significantly, resulting in a degradation of the transmission quality.

4.2 System description

4.2.1 Fabricated photonic chip

To assess the performance of the control strategy based on the second derivative of the MRM transfer function, the circuit depicted in Fig. 4.1 was designed and fabricated in IMEC's silicon photonic ISIPP50G platform [63]. The microring resonator is an O-band all-pass ring with a radius of 7.5 µm, an FSR of 9.5 nm (1.7 THz), a bandwidth of 190 pm (33 GHz) and a Q-factor of 5000. To modulate the light, a highly-doped lateral *pn* junction is embedded into the ring, as shown in the cross-section in Fig. 4.1. The 70-nm Si slab is 10^{20} cm^{-3} p+ and n+ doped in the periphery to achieve low series resistance and improve the device bandwidth [63], while along the ring circumference a 10^{17} cm^{-3} doping is chosen to reduce the device insertion loss. To control the resonance frequency of the modulator, a tungsten integrated

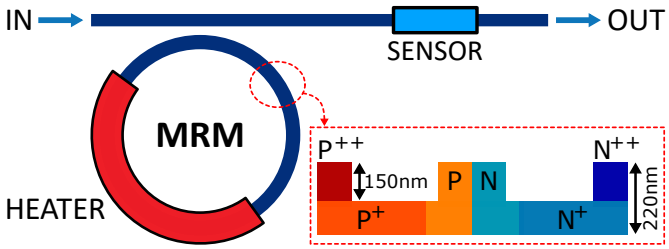

Figure 4.1: *Schematic view of an O-band silicon microring modulator. The heater is placed on top of the ring. In the inset, a cross-section of the waveguide is shown. The Si waveguide is partially etched to contact the pn junction with low series resistance while achieving sufficient guiding capability. For this reason, the slabs are p+ and n+ doped and the electrical contacts p++ and n++ doped.*

heater is fabricated on top of the ring at a distance of 700 nm from the core, showing a resistance of about 100 Ω. At the ring resonator output, a CLIPP sensor [47] is used to monitor the average optical power. Two TE-polarization grating couplers (GCs), each introducing 7 dB of insertion loss, allow optical access with the input and output fibers.

4.2.2 Electronic system

The photonic chip was mounted on a custom interface board designed to provide easy optical coupling with the input and output fibers and to allow electrical access with the RF tip needed to drive the modulator, as shown in Fig. 4.2. The photonic chip was wire-bonded to a custom ASIC for CLIPP readout [105], so to minimize the stray capacitance of the connection and maximize the sensor sensitivity. The interface board was connected to the FPGA-based multichannel motherboard with the rest of the electronics already described in Section 1.4.2.

4.2.3 Experimental setup

The performance of the modulator and the control scheme were tested with the experimental setup shown in Fig. 4.3. A tunable laser (TUNICS-T100S-HP) was used to generate a 0 dBm CW signal at $\lambda = 1307.1$ nm. An arbitrary waveform generator (AWG, MICRAM-DAC10002) was employed to create a 25 Gbit/s NRZ PRBS9 sequence that was sent to a high-speed

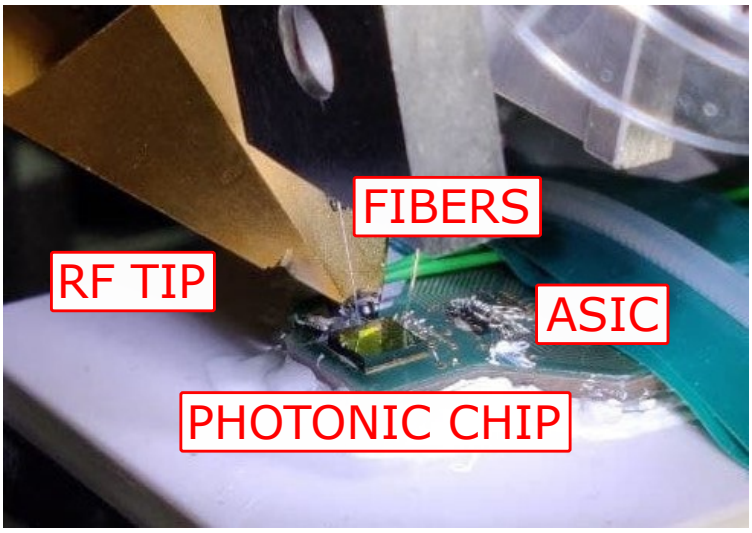

Figure 4.2: *Custom printed circuit board designed to house the photonic chip, the readout ASIC and allow easy optical coupling and RF access. The board is connected to the FPGA-based motherboard with all the rest of the electronics.*

amplifier (SHF-M804B) to drive the MRM with a $3\,V_{PP}$ signal through a GS RF tip. The bias voltage of the modulator *pn* junction was set to $-2.5\,V$ with a bias-tee. The optical signal at the output of the modulator was amplified with a PDFA, filtered with a 10 nm tunable bandpass filter and then measured with a photoreceiver. The electrical signal was sent to a samplescope (Keysight-N1000A/N1046A) to assess the quality of the transmission. Polarization controllers (PC) were used in the setup to provide the right polarization for each component. The temperature of the chip was measured with a $10\,k\Omega$ thermistor placed on the interface board next to the photonic chip, and the whole setup was kept at $28\,°C$ with a thermo-electric cooler (TEC).

Figure 4.3 also shows the implemented control scheme. The dithering oscillation is a sine wave generated by the FPGA with the direct digital synthesizer (DDS) DDS1. The generated dithering signal is summed with a digital adder to the DC value and fed to the heater with a DAC. The optical power in the waveguide is measured with a CLIPP placed at the MRM through port. The second derivative of the transfer function is extracted with the

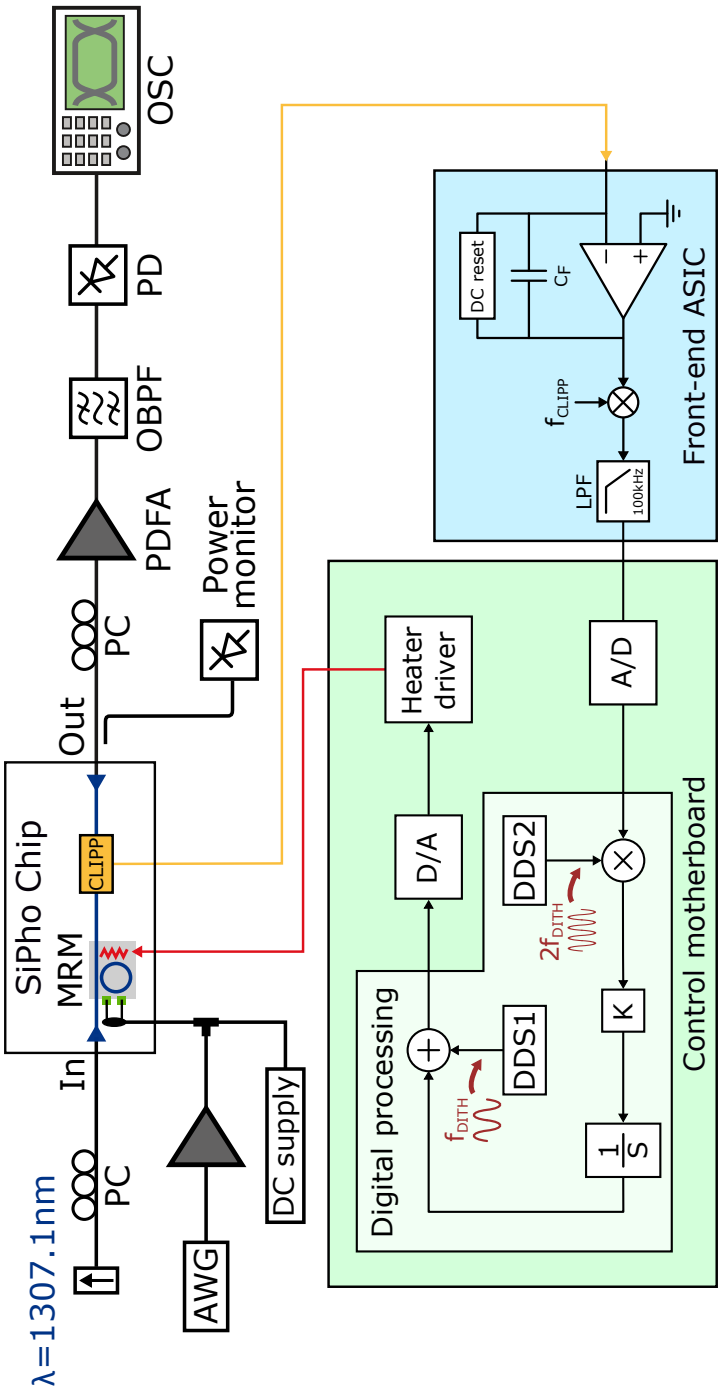

Figure 4.3: *Schematic view of the experimental setup used to assess the modulator and the control strategy performance.*

use of a digital mixer, fed by another DDS (DDS2). The phase of the sine wave generated by DDS2 is tuned in order to compensate for the phase shift introduced by the acquisition chain before the digital mixer. The extracted signal is finally low-pass filtered and fed to a digital integrator, that moves the heater DC power in order to drive to zero the amplitude of the measured second-derivative signal.

4.3 Validation of the control strategy

4.3.1 MRM characterization

The described system was first used to measure the first and second derivative of the MRM transfer function. This operation was performed by ramping the heater power while summing a 6 kHz dithering signal of around 50 µW (practically obtained with a voltage amplitude of 2 mV) to the average value and simultaneously acquiring the sensor signal. The electronic system was set to synchronously demodulate the CLIPP readout at 6 kHz and 12 kHz. The results are shown in Fig. 4.4. As expected, the zeroes of second derivative correspond to the stationary points of the first derivative, that in turns correspond to the points where the slope of the ring transfer function is maximum. The position of the zeroes does not depend on the quantity of light circulating inside the ring, that only affects the magnitude of the measured signals.

4.3.2 Second-Derivative MRM Locking

To assess the effectiveness of the proposed control strategy, the sole ring modulator performance was first investigated. To do so, the heater power was manually ramped while keeping the dithering oscillation disabled and the RF modulation active. The ER of the output light signal was continuously measured with the samplescope and the obtained results are shown in Fig. 4.5(a) in the trace "dith OFF". By measuring the IL with a low-speed power monitor and by using Equation (4.6), the TP was computed, as reported in Fig. 4.5(b). The working condition that achieves the best modulation performance is the point where the TP is minimized. In this situation, the measured output light power

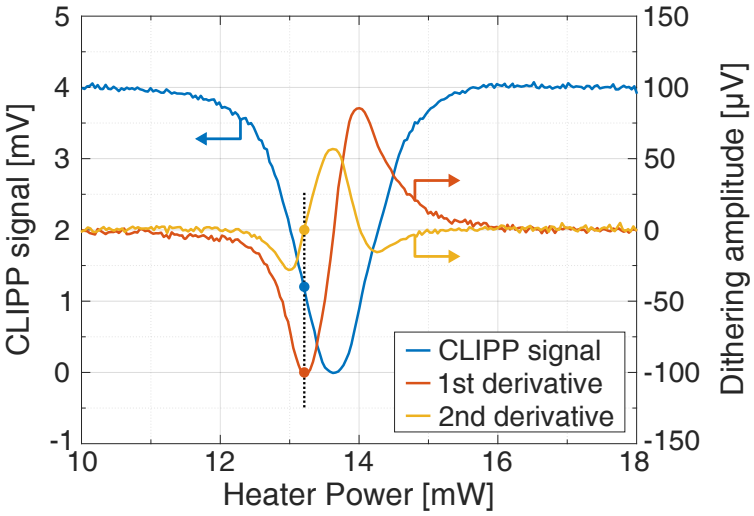

Figure 4.4: *Transfer function of the ring resonator (blue) and its first (red) and second derivatives (yellow) obtained by means of the dithering technique. The measurement was obtained by sweeping the heater power and by demodulating the CLIPP signal both at the dithering frequency and at twice that frequency. As expected, the point of maximum slope of the ring transfer function corresponds to the maximum of the first derivative signal and to the zero of the second derivative.*

was −18.9 dBm, corresponding to an estimated modulator insertion loss of 4.9 dB. Figure 4.5(c) depicts the respective captured eye-diagram, with a Q-factor of 8.45 and an ER of 11 dB.

To quantify the effect of the dithering signal on the transmission, the experiment was repeated by activating different oscillation amplitudes. Once again, the ER of the output signal was continuously measured with the samplescope and, together with the insertion losses, used to estimate the overall TP. The obtained results are shown in Fig. 4.5. For oscillations of the heater power of 50 µW amplitude, the ER is worsened of 1 dB at maximum, while negligible effects can be observed in the TP. A slightly higher degradation is instead observed in the 100 µW case, that for this reason was discarded in the following experiments.

Since the control scheme searches for the zeroes of the MRM transfer function second derivative, the device is locked in the point of its maximum slope. This point does not coincide with

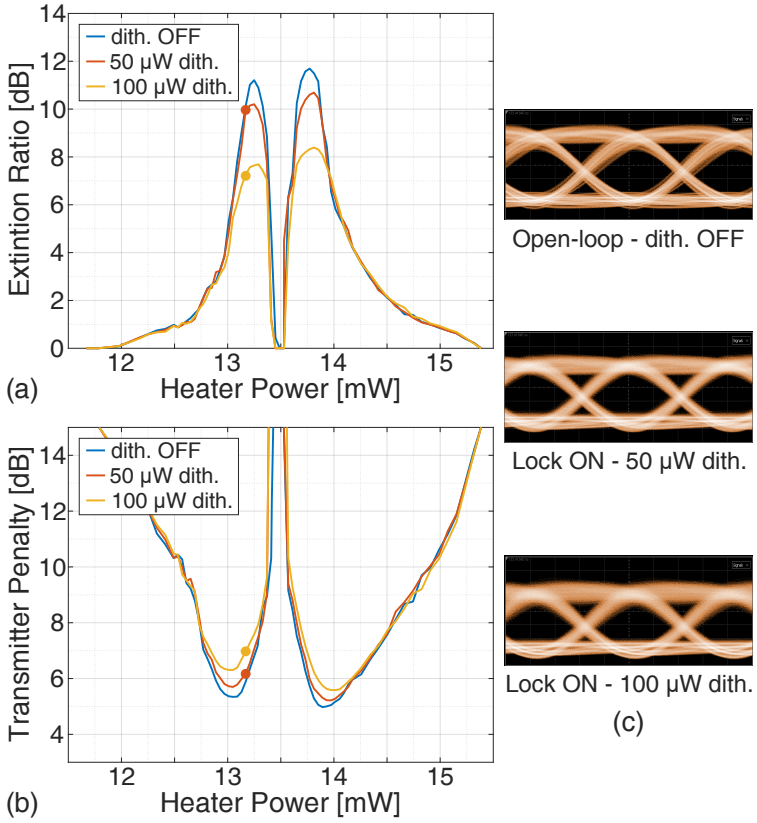

Figure 4.5: *(a) Extinction Ratio and (b) Transmitter Penalty measured at the MRM output. The curves are obtained for different dithering amplitudes (dithering off,* $50\,\mu W$ *and* $100\,\mu W$*). The circles indicate the working point targeted by the control scheme, showing that the system is able to converge to a condition very close to the TP minimum. (c) Eye diagrams measured when manually tuning the MRM in the optimum operating point, compared to those obtained when automatically locking the device with the proposed control strategy.*

the TP minimum. To assess how far these two working conditions are, the locking points of the system are marked in Fig. 4.5(a) and 4.5(b). It is possible to see how these values are very close to the minimum of the TP and well within a 1 dB tolerance region. This leads to no measurable deterioration of the modulator performance [63], as also confirmed by the eye diagrams in Fig. 4.5(c). The proposed technique allows to lock the MRM on both the blue and the red side of the resonance, simply by changing the sign of the acquired dithering signal to invert the behavior of the integrator. The blue side is usually preferred, in order to avoid feedback instabilities due to self-heating effects [82, 106].

4.3.3 Automatic recovery from wavelength variations

The performance of the control scheme was first assessed in terms of its capability to swiftly recover from abrupt MRM working point variations. While using a dithering power of 50 µW to have minimum degradation in the transmission performance, a wavelength variation was imposed to the system. Figure 4.6(a) reports the output power measured with the external power monitor while the control feedback was active. After 130 ms, the wavelength of the input laser was abruptly changed from 1307.6 nm to 1307.7 nm (+100 pm). The signal measured by the power monitor instantaneously changed from −18 dBm to −14 dBm, due to the shift of the MRM working condition. The reaction of the control loop is demonstrated by the evolution of the heater power shown in Fig. 4.6(b), that increases to compensate for the wavelength shift with a rise time of 30 ms. The control scheme correctly restores the initial value of the output power, proving that the system is able to recover the original working condition. The quality of the transmission is also restored, as indicated by the two eyes captured before and after the disturbance, that show no difference between each other. Being the wavelength and temperature fluctuations in real systems usually characterized by a slow evolution, often in the hundreds of milliseconds timescale, the achieved control speed is sufficient for most applications. If a faster reaction time is needed, the CLIPP sensor can be replaced with a more sensitive detector, since a higher readout signal-to-noise ratio (SNR) can be traded to obtain a larger control bandwidth, while keeping the same modulation accuracy.

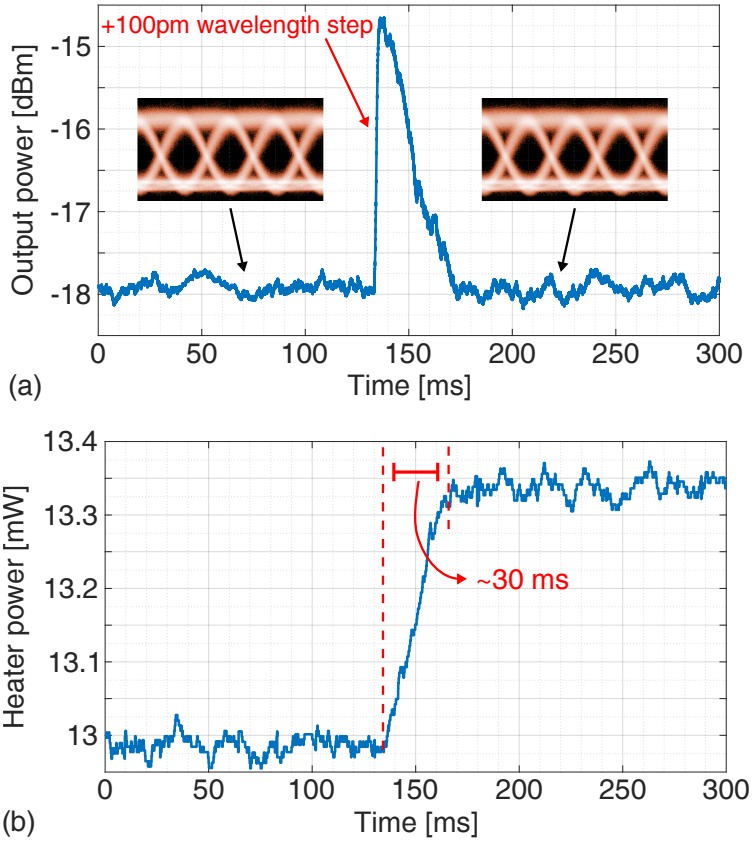

Figure 4.6: *Time evolution of (a) the optical power at the MRM output and (b) the heater power when abruptly changing the laser wavelength of* $+100\,\mathrm{pm}$ *(at* $t = 130\,\mathrm{ms}$*) with the feedback control loop activated. After the disturbance, the MRM working point shifts to the right, so the heater power has to increase to move the device transfer function in the same direction.*

4.4 50-G transmission experiment

The same experimental setup shown in Fig. 4.3 was also used to drive the MRM with a 50 Gbit/s on-off keying (OOK) modulation scheme. Notice that, since the control algorithm is not based on the direct assessment of the transmission performance, the modulation scheme or the frequency have no effect on the locking accuracy. In the following experiments, that had a duration of around 10 minutes, one external variable was manually swept in order to simulate a slowly-varying environment. The quality of the modulation was assessed during the whole experiments and reported to prove the efficiency of the control scheme in practical scenarios.

4.4.1 Wavelength drift compensation

During a 15-minute-long experiment, the wavelength of the input laser was slowly increased from 1306 nm to 1314 nm with the control loop constantly active. The results are reported in Fig. 4.7. During the whole experiment, the power measured at the input of the PDFA remains constant. As shown in Fig. 4.7(b), the heater power follows closely the laser shift, showing how the control action is able to always restore the correct working point of the modulator. The quality of the transmission, assessed by measuring the Q-factor of the output signal throughout the whole experiment, is thus not impaired. The average Q-factor value of 4.75 corresponds to a theoretical BER of 10^{-6} [107], well below the forward error correction threshold normally required for OOK transmission [108]. Figure 4.7 also reports three eyes taken at three different instants during the experiment, showing no measurable difference between each other.

4.4.2 Thermal drift compensation

The main challenge in employing Silicon Photonics circuits in datacenters comes from temperature fluctuations. Thus, the proposed control scheme was also tested in a temperature unstable scenario, to demonstrate that it can effectively compensate thermal drifts in real time. To this aim, temperature oscillations were intentionally generated in the setup using the TEC controller. As shown in Fig. 4.8, the temperature was increased from 28 °C to

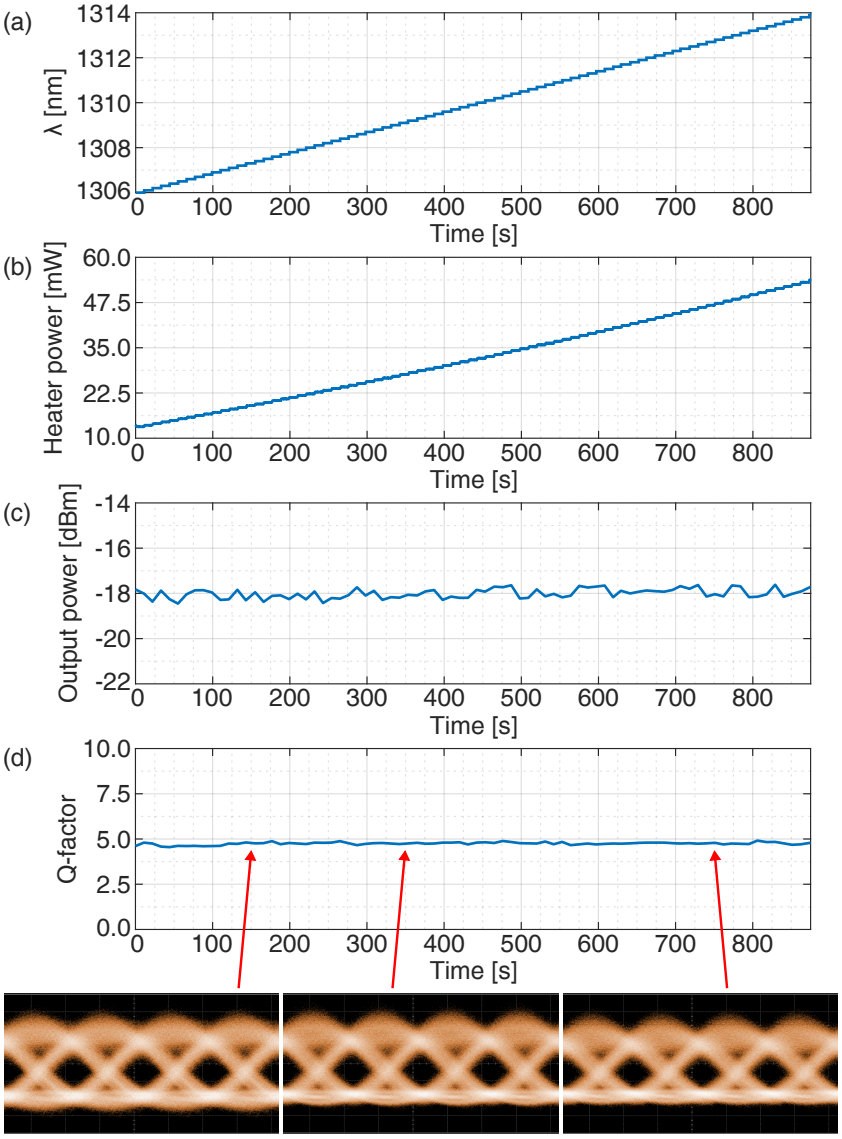

Figure 4.7: *Time evolution of (a) input laser wavelength, (b) heater power, (c) output optical power and (d) output eye diagram Q-factor, obtained when moving the laser wavelength from* 1306 nm *to* 1314 nm *while transmitting a* 50 Gbit/s *signal. The feedback loop updates the heater power in real-time to compensate the wavelength shift, preserving the transmission quality.*

38 °C during an 8-minute-long experiment. The controlled heater power mirrors perfectly the temperature profile, confirming that the control action is able to compensate any shift introduced by external perturbations. As a result, the output power measured at the input of the PDFA remains constant, apart from a negligible variation at high temperature due to a small light coupling drift in the setup. The Q-factor of the output signal, continuously assessed during the whole experiment, remains constant around 4.94 regardless of the global temperature of the system. In this condition, the estimated theoretical BER is $3.8 \cdot 10^{-7}$. The quality of the transmission link is therefore preserved, as also demonstrated by the eyes captured at different temperatures reported at the bottom of Fig.4.8.

4.5 Comparison with other control schemes

The control scheme proposed has been proven to be very effective in stabilizing the working condition of a microring modulator in the region of maximum slope of its transfer function, close to the theoretical optimum point of minimum transmitter penalty. A comparison with different state-of-the-art solutions is reported in Table I. In [77] and [109], the DC output power is monitored, resulting in a control scheme that is not power-independent. In [68] and [106], both the input and the output power of the ring modulator are assessed, in a ratiometric approach. While this strategy could lead to a power-independent control system, it is unclear how the fabrication mismatches and the non-linearities of the MRM and the sensor would affect the convergence point, hence requiring a prior calibration. In [110], the OMA is measured and used to control a ring, a valid approach as well that requires however expensive and complicated RF electronics.

Moreover, the dithering technique allows to extract the first derivative of any device transfer function as well, meaning that the proposed control scheme also allows to stabilize components, such as MZIs or add/drop ring resonators, in the point of minimum (or maximum) transmission. The same electronic system can thus be used to fully control and stabilize heterogeneous architectures performing multiple functionalities in the same chip, as it will be shown in the following chapter.

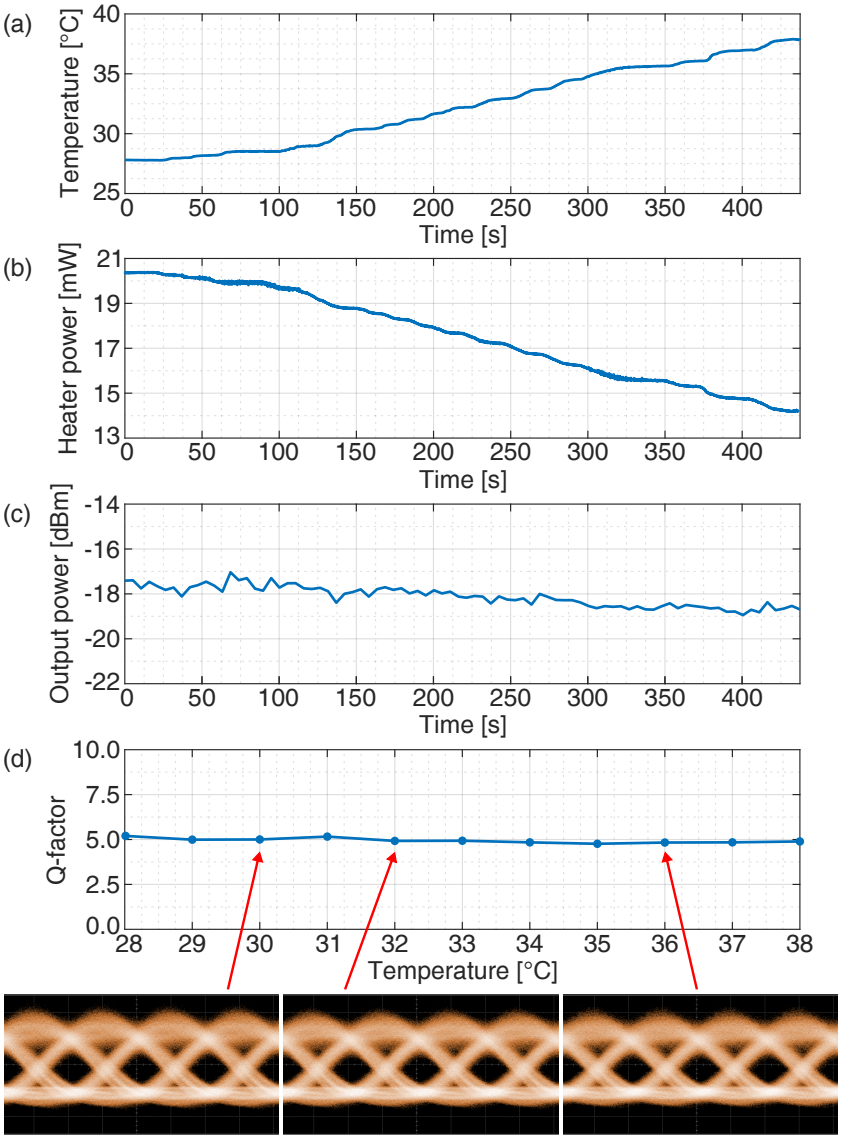

Figure 4.8: *(a) Chip temperature profile obtained by operating the TEC in the setup. (b) The heater power mirrors perfectly the temperature evolution to compensate for any shift in the MRM transfer function. As a result, the output power (c) remains constant, apart for a minor fiber-chip misalignment at high temperature, and the output eye diagram Q-factor (d) does not change with temperature.*

Table 4.1: *Comparative table of state-of-the-art control strategies for MRMs.*

	[77]	[68]	[110]	[109]	This work
Laser band	C-band	O-band	C-band	C-band	O-band
Control strategy	Output power monitoring	Input and through port power monitoring	OMA monitoring	MRM *pn* junction current monitoring	MRM second-derivative extraction
Number of sensors	1 (external)	2	1	Ring *pn* junction	1
Requires calibration	Yes	Yes	No	Yes	No
Power-independent	No	Yes	Yes	No	Yes
Modulation scheme	OOK	PAM4	OOK	OOK	OOK
Data-Rate	10Gbps	112Gbps	2Gbps	12.5Gbps	50Gbps
Requires RF electronics	No	No	Yes	No	No

CHAPTER 5

Dual-control system for multi-socket interconnect architecture

The dithering technique and its second-derivative extension can be used together to control more complex and heterogeneous systems. In this chapter, a WDM-based multi-socket interconnect architecture is presented and validated, where both resonant switches, locked on the maxima of their transfer function, and micro-ring modulators (MRM), locked on the maximum slope working condition, are simultaneously controlled.
The results presented in this chapter were published in:

F. Zanetto, **V. Grimaldi**, M. Moralis-Pegios, S. Pitris, K. Fotiadis et al., "WDM-based silicon photonic multi-socket interconnect architecture with automated wavelength and thermal drift compensation," Journal of Lightwave Technology, vol. 38, no. 21, pp. 6000-6006, 2020, doi: 10.1109/JLT.2020.3008001.

I. Roumpos, T. Chrysostomidis, **V. Grimaldi**, F. Zanetto, F. Toso et al., "Temperature and wavelength drift tolerant WDM transmission and routing in on-chip silicon photonic interconnects," Opt. Express 30, 26628-26638, 2022, doi: 10.1364/OE.455107.

5.1 System description

With the demand for data traffic capacity in intra datacenter applications roughly doubling every year [111], novel techniques must be developed to cope with traffic growth. A key problem is the power consumption and the latency in the communication between different processors, that do not scale with the rest of the electronics [2]. These problems have driven research interest on novel multi-socket-boards, that integrate several processor sockets on a single board interconnected with a low latency interface.

Optical architectures based on Arrayed Waveguide Grating Router (AWGR) have already been demonstrated to enable any-to-any, low-latency, low-energy communication between more than 8 nodes [3]. Figure 5.1 shows the AWGR-based NxN interconnection scheme proposed within the H2020 ICT-STREAMS European project [112]. Each processor is connected to a transmission engine (Tx), featuring N-1 lasers at different wavelengths, each equipped with an integrated modulator (MOD) to encode the data that has to be transmitted. These are then combined in a WDM-encoded data stream on a single bus by an optical (N-1):1 multiplexer (MUX), connected to the corresponding input of the AWGR. The carriers of each data stream get routed by the AWGR to different receivers depending on their wavelength, with collision-less transmission achieved thanks to the cyclic-frequency routing properties of the AWGR. At the receiving side (Rx) of each socket, the WDM-encoded data stream gets demultiplexed with a 1:(N-1) optical demultiplexer (DEMUX) so that each wavelength is acquired by a separate receiver, allowing to discriminate the transmission from each sender. In this way, every processor can simultaneously communicate with all the others by simply encoding its electrical information on the wavelength of the target receiver, allowing any-to-any direct interconnection without the need of switching mechanisms.

AWGRs have been demonstrated to be effective routing engines in any large integration scale platforms [113–116]. In particular, Silicon Photonics enables the largest integration scale and includes all the functional building blocks required by the proposed interconnection scheme [117]. However, as the num-

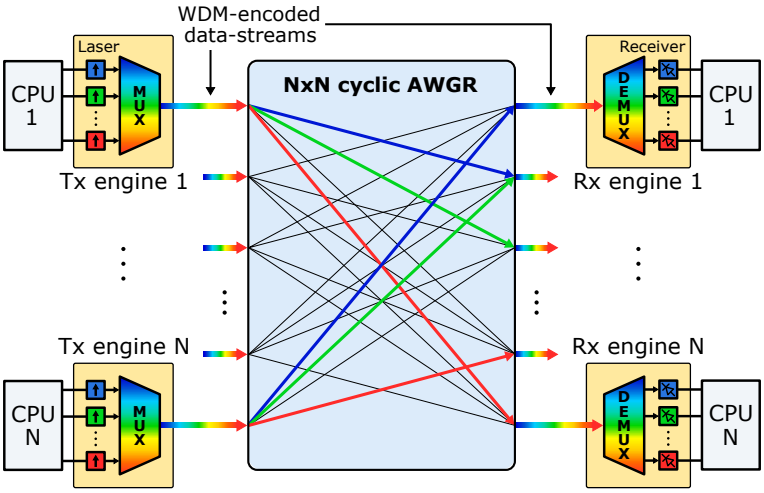

Figure 5.1: *CPU interconnection scheme proposed within the H2020 ICT-STREAMS European project. Thanks to WDM data encoding and to the cyclic-frequency routing properties of the AWGR, any-to-any direct communication is possible between the CPUs without the need of switching mechanisms.*

ber of components on a single chip grows, the requirements on their control become critical. Thus, algorithms and strategies to tune, reconfigure, calibrate and operate these devices are needed to maintain the performance of the system at the required levels [16, 21, 71, 118].

5.1.1 Photonic chip

To prove the advantages of a feedback scheme to keep the performance of the system at the required levels, a subsection of the proposed interconnection scheme was designed and integrated in IMEC's ISIPP50G platform in a single photonic chip. As depicted in Fig. 5.2, the system comprises of two Tx sockets, each equipped with a micro-ring modulator (MOD) such as the one described in Chapter 4: it is an O-band all-pass ring with a radius of 7.5 µm, an FSR of 9.5 nm (1.7 THz), a bandwidth of 190 pm (33 GHz) and a Q-factor of 5000, with a highly-doped lateral *pn* junction embedded into the ring to modulate the light. The two sockets are each equipped with a single laser to simplify the complexity of the assembly, but the approach and the obtained

results can be extended also to a multi-laser situation.

Each socket is interconnected to the AWGR optical engine through a 8x1 MUX with an FSR of 9.6 nm, based on cascaded double-ring resonators to provide sufficient selectivity and out-of-band attenuation in a very compact design [119]. The two Tx sockets are connected to a 16x16 O-band AWGR [9], with a channel spacing of 1.1 nm and a free spectral range (FSR) of 17.6 nm, allowing non-blocking communication of the two through wavelength-based routing.

Thus, the AWGR combines the two input signals coming from the transmitters and routes them towards a common receiver socket, that routes them back in two separated output through a 1x8 DEMUX, based on cascaded double-ring resonators similar to the ones employed for the MUX. The access to the optical inputs and outputs, denoted in figure as In_1, In_2, Out_1 and Out_2, is achieved via TE-polarization grating couplers (GC), all at the left side of the chip. Two auxiliary outputs, denoted as Aux_1 and Aux_2, are also present, in order to monitor the through port of the DEMUX on the Rx socket and the second output of the AWGR, respectively. Integrated heaters allow fine thermal tuning of the MODs, MUXes and DEMUX resonances, while CLIPP sensors are employed to monitor the optical power all around the system. Both the CLIPP sensors and the integrated heaters are accessible through electrical DC pads at the right side of the chip and are wire-bonded to the electronic system. A microscope photograph of the fabricated chip is reported in Fig. 5.3.

5.2 Experimental setup

As depicted in Fig. 5.4, two CW laser beams ($\lambda_1 = 1311.7$ nm, $\lambda_2 = 1310.6$ nm) were launched into the chip Tx1 and Tx2, respectively. The two employed wavelengths correspond to the channels of the AWGR optical spectrum. An arbitrary waveform generator (AWG, MICRAM-DAC10002) was employed to create a 25 Gbit/s NRZ PRBS9 sequence that was sent to a high-speed amplifier (SHF-M804B) to drive MOD1 with a $3\,V_{PP}$ signal through a GS RF tip. The bias voltage of the modulator *pn* junction was set to -2.5 V with a bias-tee. The optical signal at the output of interest (either OUT1, OUT2 or AUX1) was sampled

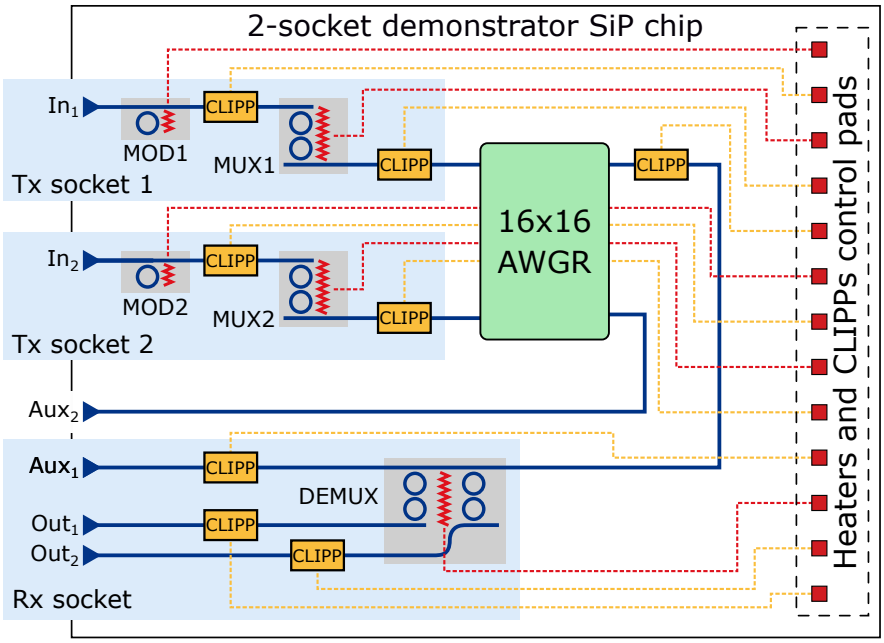

Figure 5.2: *Schematic view of the WDM-based Silicon Photonic two-socket interconnect architecture integrated in a single chip.*

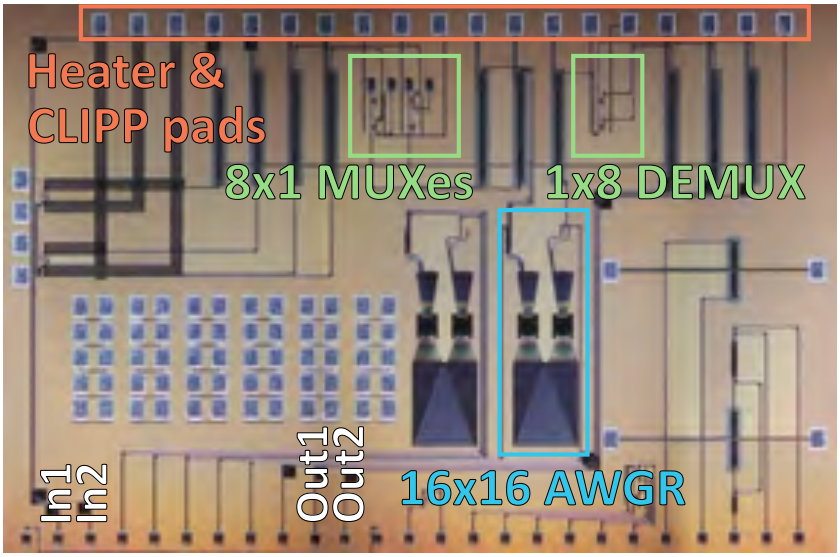

Figure 5.3: *Microscope photo of the chip fabricated in IMEC's ISIPP50G platform.*

with a 90/10 splitter in order to be measured by a power monitor and simultaneously amplified with a PDFA, filtered with a 10 nm tunable bandpass filter and then measured with a photoreceiver, that was then connected to a samplescope (Keysight-N1000A/N1046A) in order to assess the quality of the transmission. Polarization controllers (PC) were used in the setup to provide the right polarization for each component. The temperature of the chip was measured with a 10 kΩ thermistor placed on the interface board next to the photonic chip, and the whole setup was kept at 28 °C with a thermo-electric cooler (TEC).

The electronic system architecture implemented to independently lock MODs, MUXes and DEMUX is also reported in Fig. 5.4. The CLIPP readout is carried out through the use of a custom 32-channel ASIC that encompasses the electronic front-end and a mixer to downconvert the information of interest [120]. This chip is directly wire-bonded to the photonic chip in order to minimize the effect of the stray capacitances, so to maximize the CLIPP sensitivity [49, 69]. Both the photonic chip and the ASIC are mounted on a host PCB, that is then connected to the FPGA-based motherboard where the signal gets further conditioned and digitized. The lock-in is completed in the digital domain: by choosing the value of f_{demod}, either the first or the second derivative of the transfer function can be extracted and used as the error signal of each control loop, depending on the device that has to be stabilized. The loops are then completed by integral controllers that drive the correct heaters in order to zero the error signal. All the loops are operated independently from each other by the FPGA.

5.3 Characterization of the interconnect subsystem

In order to fully understand and exploit the system capabilities, the assembly had to be fully characterized. While the modulator characteristics were fully discussed in Chapter 4, the performance of the control system applied to the only resonant switches have to be assessed. Since the target working point of these devices is the resonance condition, the control strategy relies on the use of the dithering technique to extract the first-derivative signal, as explained in 2.2. Figure 5.5 shows the MUX transfer function and

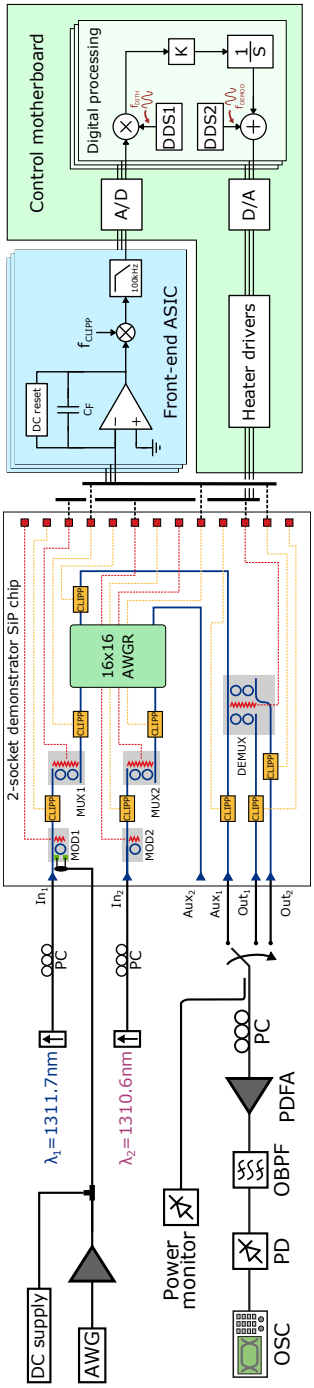

Figure 5.4: *Schematic view of the experimental setup used to assess the modulator and the control strategy performance.*

its first-derivative signal as a function of the heater power, measured with a CLIPP at the drop port of the filter. The application of a dithering modulation of 10 mV amplitude (corresponding to around 600 µW in terms of power) allows to obtain a sufficiently high signal-to-noise ratio in the extraction of the derivative information, yet without affecting the quality of the transmitted signal [78]. The double peaked transfer function of the filter is due to a slight overcoupling of the rings, that can be solved by using two separated heaters to control each of them independently. The use of a single collective heater was here preferred to simplify the control strategy while still allowing to match the MUX grid to the respective laser wavelength.

Figure 5.5 shows the spectral response of the MUX aligned to the AWGR one, that is instead operated in open-loop. Hence, the maximum transmission is obtained when the peak of the MUX transfer function is aligned to the AWGR maximum. This condition is met if the correct wavelength is launched in the system and the control of the system is activated.

5.3.1 Automatic locking and recovery vs. wavelength shifts

The performance of the implemented closed-loop control was evaluated in terms of capability to react to wavelength variations. To this aim, a single input laser, the MUX, AWGR and the DEMUX were first tuned and locked to maximize the output power. The optimal operating point of the system was then intentionally perturbed by shifting the laser wavelength with a step of 200 pm. Figure 5.7 shows the time evolution of the optical power at the output of the chip, measured with an external photodiode in four different scenarios.

The first experiment was conducted in open loop, i.e. with no control action, by applying a positive wavelength shift. As shown in the curve "NOT-locked (+200 pm)", a loss of about 6 dB was measured at the output of the system, since after the step the laser is outside the MUX and DEMUX passband. A second experiment was then conducted with the control of MUX and DEMUX activated in real time. The curve "Locked (+200 pm)" shows how the system is able to swiftly recover its working point after the wavelength detuning. The final power level is slightly different from the initial one because of the not perfectly flat

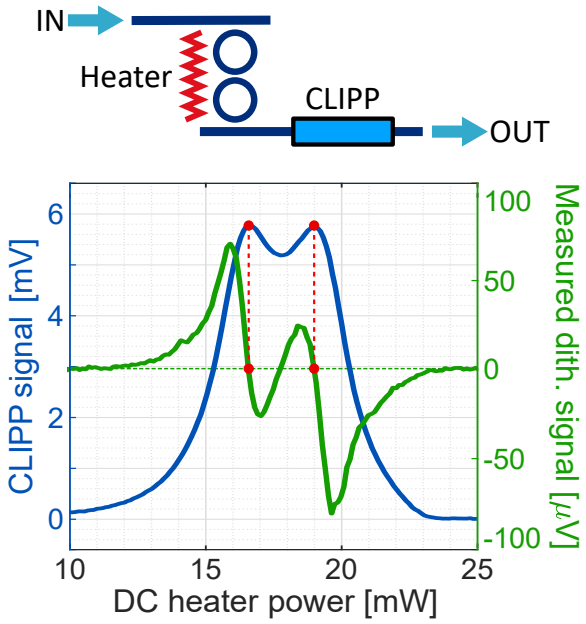

Figure 5.5: *Transfer function (blue) and output dithering signal (green) of a double-ring resonator filter, experimentally measured with a CLIPP. As expected, the stationary points of the transfer function correspond to the zeroes of the dithering signal.*

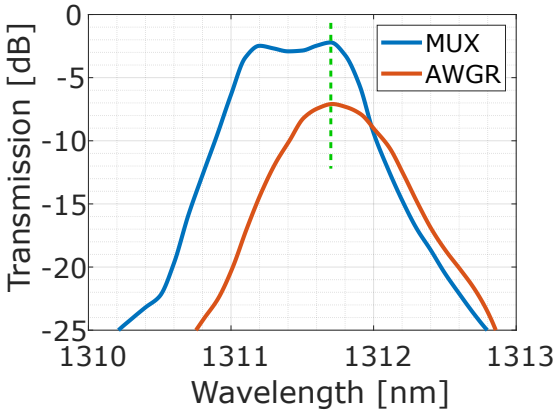

Figure 5.6: *Optical response of the MUX (blue) tuned to the optical response of the AWGR (orange), that is instead operated in open loop.*

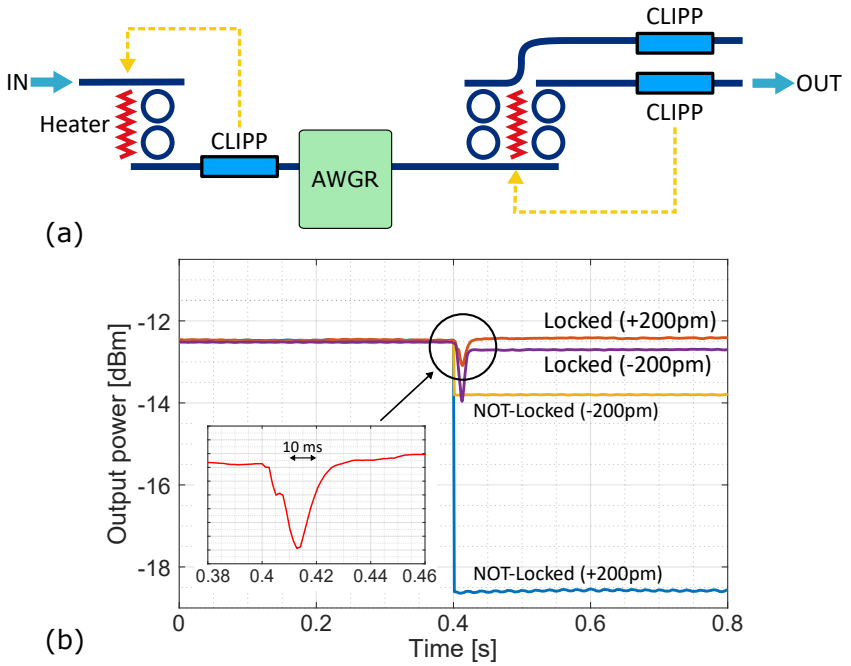

Figure 5.7: *(a) Schematic view of the portion of the circuit controlled in the experiment. The CLIPPs and the heaters used are also highlighted. (b) Time evolution of the power at the output of the chip when perturbing the laser wavelength with* ± 200 pm *steps. The comparison between locked and not locked situations is reported.*

response of the AWGR in its passband (Fig. 5.6). The inset of Fig. 5.7 expands the transient of the experiment, showing an estimated recovery time of the system of around 30 ms.

A similar experiment was conducted from the same starting point by shifting the input laser wavelength with a -200 pm step. Figure 5.7 also shows these results, in which a steady loss of about 3 dB is measured when the system is operated in open-loop. The loss is different than before because MUX, DEMUX and AWGR are now operating in different working points with respect to the previous experiment. Instead, if the control is active, also in this case the system realigns both MUX and DEMUX in real-time to the new signal wavelength within the 30 ms of the closed loop response time.

5.4 Dual-control system evaluation

The losses of each component in the fabricated chip are quite high: the input and output GCs cause an overall 10 dB of losses, the MOD around 4 dB when operated, the MUXes and the DEMUX around 2.5 dB each when fully tuned, and the AWGR around 7 dB, for a total attenuation of around 26 dB. Even by launching the highest possible power that could be generated by the laser, that is 5 dBm, the emerging power from OUT1 and OUT2 would be around −21 dBm. The PDFA used was not able to correctly amplify this signal without destroying the quality of the transmission. This means that, because of the high losses of the link, the full interconnect system cannot be validated. In particular, the biggest contribution that can be easily limited in future fabrications comes from the 10 dB from the GCs, that can be reduced by choosing edge couplers, which are reported to have insertion loss lower than 0.5 dB [121].

However, if the DEMUX is bypassed, the power emerging from AUX1 is around −18.5 dBm, that instead can be correctly amplified by the PDFA without impairing the transmission. For this reason, in the following experiments, only the MOD and the MUXes are controlled, allowing in this way to validate the dual-control system, since the two classes of devices are stabilized on two different working conditions. A 2-channel full-link demonstration with externally-modulated input signals is reported in [74].

Two experiments in the 5-minutes range with an external slowly-varying variable were conducted. The signal injected in the Tx1 socket was modulated through MOD1 with a 25 Gbit/s OOK modulation scheme. The modulator in the second transmitting socket Tx2 was instead not operated. The MUXes in Tx1 and Tx2 were both controlled.

5.4.1 Wavelength drift test

As a first test, an instability in the to Tx1 laser emitted wavelength was simulated. To this end, the CW light beam injected in In1 was swept from 1309.90 nm to 1310.85 nm, with 50 pm steps, during a 4.5 min-long experiment. The wavelength of the CW laser injected at In2 signal was kept constant at 1311.3 nm, and

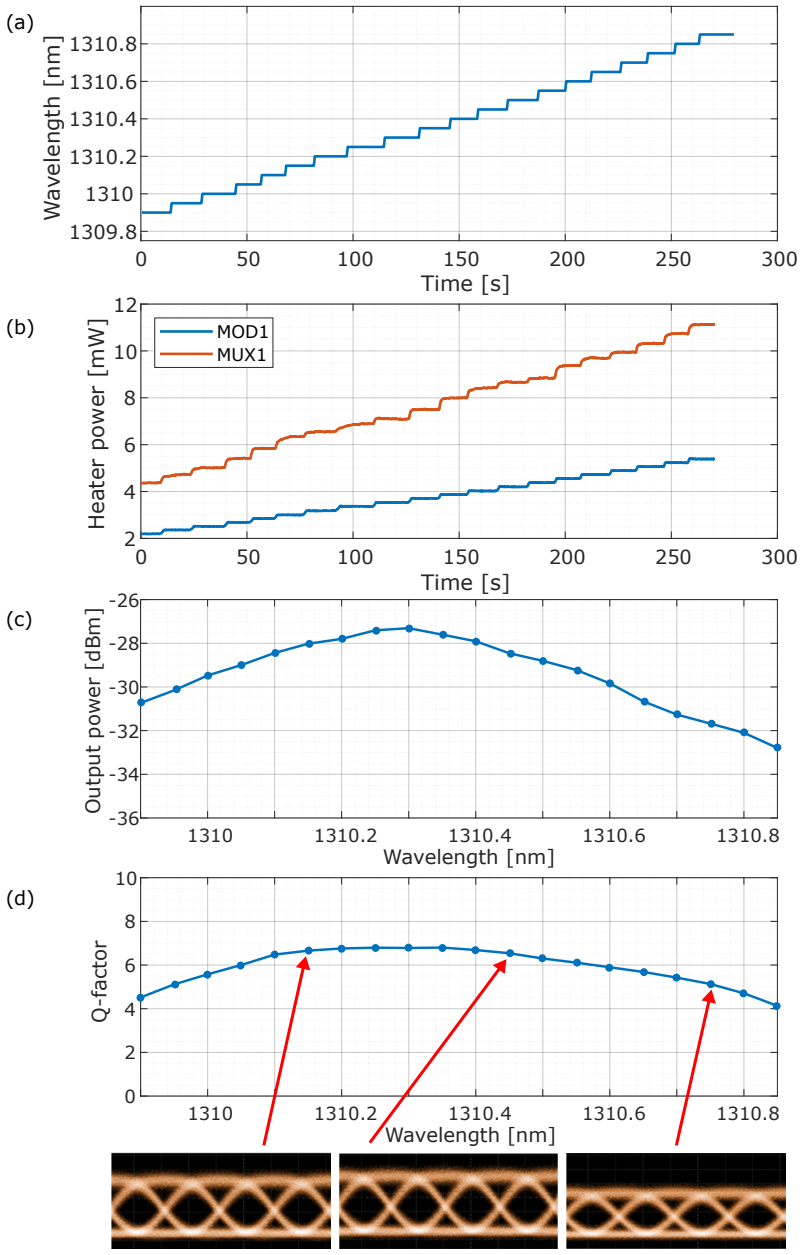

Figure 5.8: *(a) Time evolution of the input laser wavelength profile, obtained by sweeping the laser wavelength from* 1309.90 nm *to* 1310.85 nm. *(b) The controlled heaters power (MOD1 and MUX1) perfectly mirrors the wavelength evolution. (c) The output optical power and (d) the output eye diagram Q-factor, obtained while transmitting a* 25 Gbit/s *signal, were also recorded, showing how the feedback loop is able to preserve the transmission quality.*

the MUX2 was controlled to work at resonance and route the In2 signal inside the AWGR. In this sense, the second socket is only operated as a CW disturb in the AWGR interconnect. The results of the control are reported in Fig. 5.8. In particular, it is possible to see how the time evolution of MOD1 and MUX1 heater power perfectly mirrors the wavelength profile of the input laser. The operating voltage of MUX2 remains constant throughout the experimental procedure, corresponding to the constant value of the incoming light beam wavelength at 1311.3 nm. The signal from In1 was modulated through MOD1 and, after emerging from Aux1, was evaluated in terms of average optical power and Q-factor values extracted from the eye diagrams captured during the experiments. Some eye diagrams at specific wavelength values are also reported in Fig. 5.8. Both the received average optical power and the Q-factor follow the non-perfectly-flat AWGR channel response. Finally, it should be noted that the recorded Q-factor values ranged from 4.1 to 6.8 with an average value of 5.9, that corresponds to a theoretical bit error rate value of $1.82 \cdot 10^{-9}$, well below the forward error correction threshold required for OOK transmission [108].

5.4.2 Thermal stress test

In real datacenter environments, temperature instabilities are one of the main issues to be faced when employing Silicon Photonics devices. Consequently, the proposed optical interconnection scheme was also tested in a temperature unstable scenario, to demonstrate that a closed-loop control can effectively compensate thermal drifts in real time. For the purpose, external temperature fluctuations were applied to the chip through the on-board TEC. The CW signal is injected in Tx1 and modulated by MOD1, passes through MUX1, gets routed through the AWGR and emerges at Aux1, where it is evaluated with a power monitor and a sampling oscilloscope.

Figure 5.9(a) shows the time evolution of the on-chip temperature during the 6-minute-long thermal stress test, with the temperature varying from 27 °C to 36 °C. Figure 5.9(b) illustrates the time evolution of MOD1, MUX1 and MUX2 heater voltages. The curves mirror the time evolution of the on-chip temperature, certifying the correct response of the locking sys-

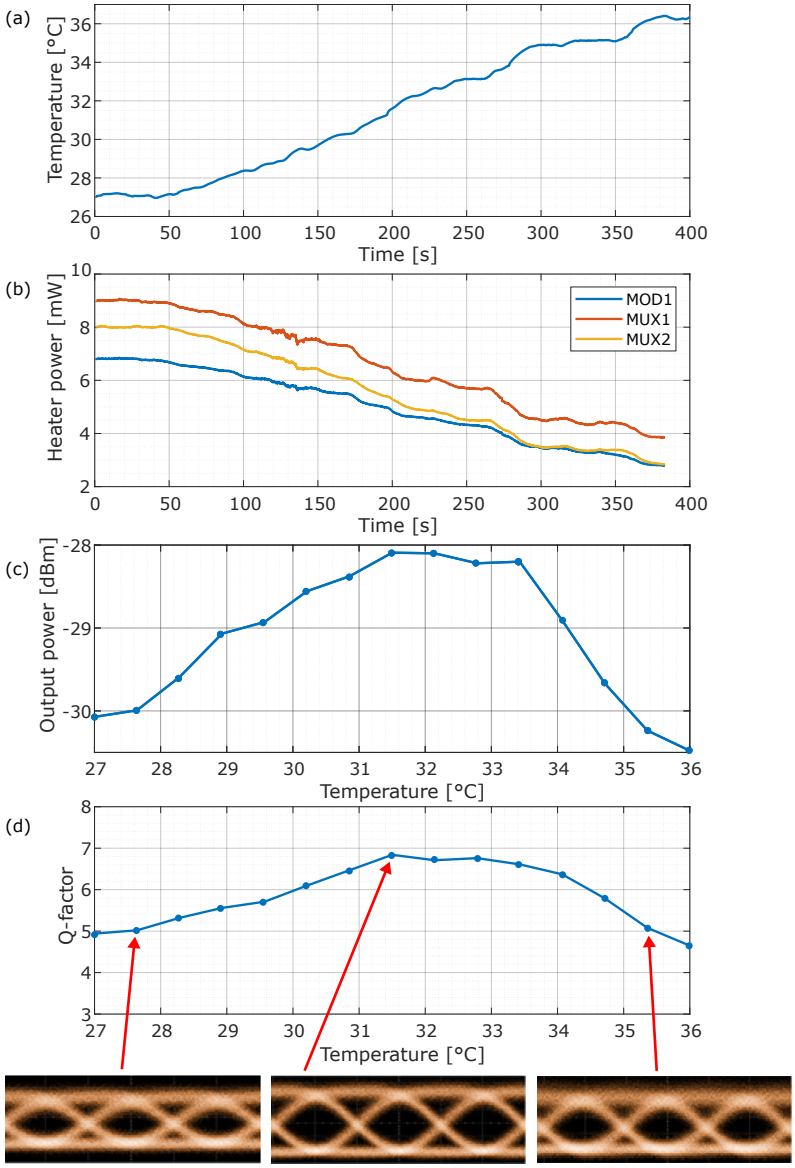

Figure 5.9: *(a) Chip temperature profile obtained by operating the TEC in the setup. (b) The heater power mirrors perfectly the temperature evolution to compensate any shift in the MRM transfer function. As a result, the output power (c) and the output diagram Q-factor follow the non-perfectly-flat response of the AWGR.*

tem to the applied temperature fluctuations. As expected, the optical power emerging from Aux1 vs. the on-chip temperature, depicted in Fig. 5.9(c), follows the non-flat transfer function of the AWGR. Similarly, Fig. 5.9(d) shows the extracted Q-factor of the modulation, ranging from 4.6 to 6.8. Once again, the dual-control system is able to maintain the quality of the transmission link during the whole stress test.

CHAPTER 6

Guidelines for the digital design of the system

This chapter discusses the practical implementation of the digital system previously described, with a focus on the main strategies employed in the design of the firmware for the FPGA. In particular, the overview on how to implement an any-to-any input-output control loop that can be dynamically reconfigured and used for different photonic chips is discussed. The digital signal processing (DSP) chain is also presented, showing how to efficiently filter out the spurious harmonics in the acquired spectrum while still keeping the size of the single stages limited. The main strategies employed to implement a design that could fit into a commercially-available FPGA are discussed. Finally, a quick overview on how to efficiently program a software to control the assembly is introduced.

6.1 The role of the FPGA

The electronic system described at the end of Chapter 2 requires a digital core in order to individually communicate with each electronic component. Moreover, a digital core is also necessary to complete the lock-in scheme to extract the dithering information, further process the readout signal, and implement the integral controller to close the loop. Given the high complexity of the tasks and the high number of independent parallel chains present on the board, an FPGA was preferred over a simple microcontroller. Indeed, FPGA are classified as "spatial-computing" devices, since they do not execute the instructions one after the other as a normal processor would do, but rather translate the code into a dedicated piece of hardware. Each piece of code would then be mapped individually into a certain region of the FPGA and work independently from the other ones. This means that, if properly implemented, the control loops are independent one from the other, facilitating future expansions of the dimension of the overall system.

Rather than directly integrating the FPGA on the motherboard, a commercial module was preferred, in order to simplify the design of the PCB and the interface with an external computer. The board chosen is the XEM7310, by OpalKelly, shown in 6.1, that mounts:

- an Artix-7 FPGA, manifactured by Xilinx;
- 1 Gibit DDR3 SDRAM;
- a 16 Mibit-flash memory to store the FPGA bit file when powered off;
- a SuperSpeed USB 3.0 interface for configuration and data transfer handled by a dedicated peripheral controller (Cypress FX3);
- eight LEDs, especially useful for debugging purposes;

. This board is connected to the motherboard through two 80-pin expansion connectors, that provide a convenient access to the FPGA I/O pins. The USB connection with the PC allows to change the parameters of the control system at run time and

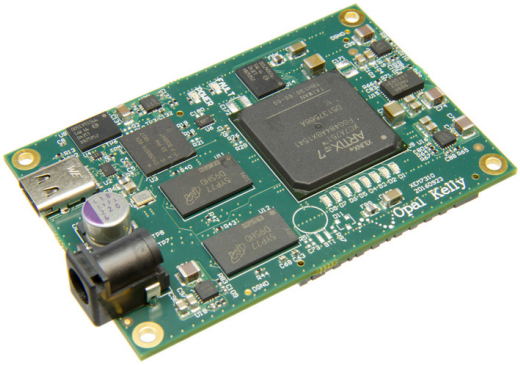

Figure 6.1: *Board XEM7310, sold by Opal Kelly. An Artix-7 FPGA is mounted on it.*

to read, display and possibly save the result of the acquisition chain. OpalKelly provides the HDL modules and the related software APIs to handle the communication between the FPGA and the USB interface. The huge flexibility offered by the FPGA comes at the expense of cost and power consumption. For specific commercial applications, a less power-hungry and cheaper microcontroller could be envisioned, that can be then tailored it on the exact needs of the system.

6.2 Firmware high-level architecture

The high-level view of the implemented firmware is shown in Fig. 6.2. The 16 acquisition channels present on the board are mapped to 16 independent DSP chains, each of them extracting the in-phase and quadrature component of the acquired signal. The digital chains encompass two sinusoidal mixers to downconvert the dithering information and separate the in-phase and the quadrature component of each acquired signal. A low-pass filter is also present to filter out the noise of the system and the unwanted harmonics of the dithering tones of other optical devices. The practical implementation of the digital chain will be discussed in details in Section 6.4. The 32 outputs of these 16 DSP chains are each mapped to all the heater control chains, each connected to an actuation chain on the motherboard. This means that the input of each controller is dynamically selected

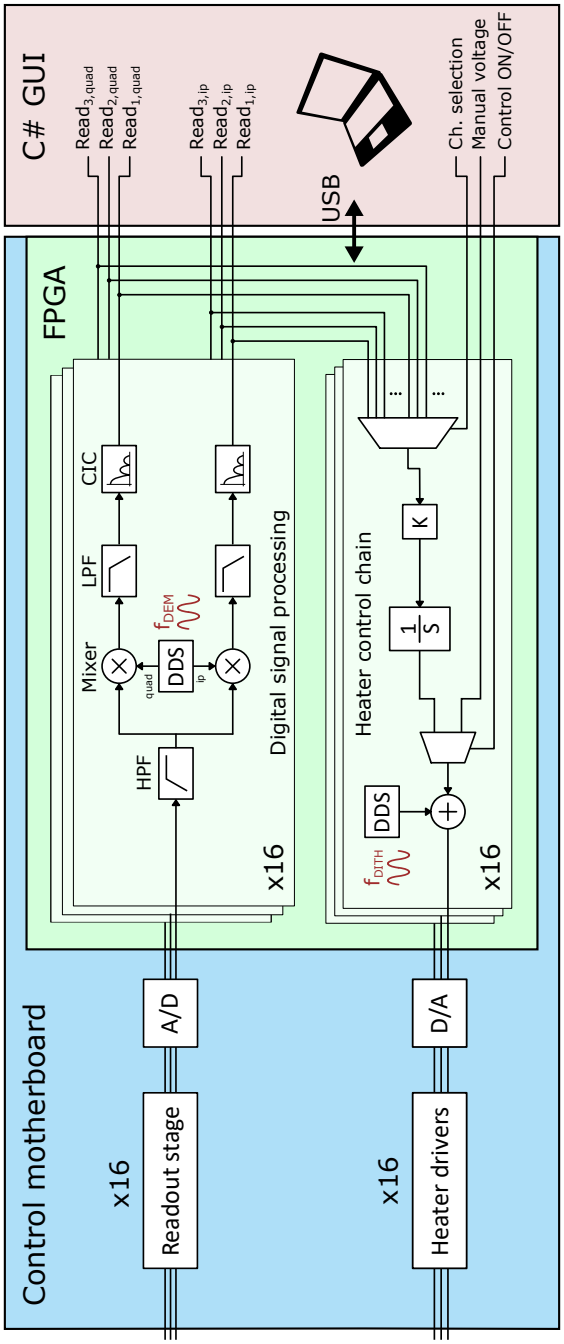

Figure 6.2: *High-level view of the implemented firmware. Each integral controller is connected through a 32:1 multiplexer to all the signals from the acquisition chains, allowing a reconfigurable any-to-any control system.*

with a 32:1 multiplexer connected to all the dithering signals, allowing to realize an any-to-any input-output control action for every different photonic circuit, without having to reprogram and recompile the firmware.

Finally, the DC value of the heater is connected either to a value set by the external user in the UI, or, if the control is activated, to the output of the integral controller. The DC value of the heater is then summed to a sinusoidal signal necessary to extract the dithering information.

6.3 Implementation of the DDSs

An efficient way to synthesize a sinusoidal oscillation in the digital domain is through the use of a direct digital synthesizer (DDS). The DDS is basically a look-up table (LUT) that stores the amplitude value of a single full sine cycle [122]. The faster the system goes through the LUT, the faster the frequency of the generated oscillation.

The easiest way to implement a DDS-based digital lock-in would be to use the same DDS to generate both the stimulus and the demodulating signal fed to the mixer. While appealing for its simplicity, this solution has many drawbacks. The most obvious issue is that the acquired signal will undergo an unwanted phase shift due to the electronic components in the acquisition chain and the digital registers inside the FPGA before reaching the input of the mixer. Notice that this phase shift depends on the frequency of the dithering signal. This means that the phase of the demodulating signal has to be independently tuned in order to compensate for these effects. Moreover, it is best to set the frequency of the demodulating signal independently from the dithering tone, since this would allow to perform more complex control action, such as the one to extract the second-derivative signal described in Chapter 4. Hence, the best approach is to implement two separated DDSs for each channel, one to generate the dithering oscillation, one for the demodulating signal. It is important, however, that the two DDSs are connected to the same digital clock, so that the frequency generated by two different and independently-programmed DDSs are perfectly related.

For the sake of efficiency, the DDS is not programmed and im-

plemented from scratch, but an IPCore from the Xilinx libraries is preferred [123]. Moreover, rather than designing 16 independent single-channel DDSs (one for each channel of the system), it is advised to implement a single multichannel DDS. In this way, a single time-multiplexed LUT gets implemented, saving many resources in the firmware that would otherwise not fit into a commercially-available FPGA. The drawback of this solution is that the clock connected to these multichannel DDSs has to be 16 times faster than the one normally necessary, complicating the design and the timing closure of the digital design [124].

6.3.1 Practical implementation of dithering and pilot tones frequencies

Due the digital nature of the DDS, the exact frequency generated depends on its frequency resolution. For a 5 MHz clocked DDS with a resolution of 16 bit, it is

$$f_{min} = \frac{5\text{MHz}}{2^{16}} = 76.29\text{Hz} \tag{6.1}$$

As stated in Section 3.3.1, to fit all the necessary dithering contributions in the spectrum without impairing the loop bandwidth, the dithering tones were dimensioned to be prime multiples of a fundamental frequency f_d, that was designed to be around 6 kHz. The exact frequency will depend on the frequency resolution of the DDS. Moreover, the dimensioning of the pilot tones frequencies explained in Section 3.5.1 followed from there: the best use of the spectrum is to place the intermodulated spectral components equally spaced in the interval between two dithering spectral lines. This means that the fundamental dithering frequency f_d has to be the a multiple of f_{min} that is as close as possible to 6 kHz but also divisible by twice the number of pilot tones + 1, i.e. 5 for the system used in Chapter 3.

For this assembly, the fundamental dithering frequency f_d is rounded to 6.1 kHz (corresponding to 90 times the frequency resolution of the DDS), while consequently the pilot tone fundamental frequency f_t is 1.22 kHz (corresponding to $90/5 = 16$ times the frequency resolution of the DDS).

6.4 Digital-signal processing chain

In order to implement the lock-in scheme for extracting the dithering signal, the only essential building blocks are a mixer to bring back the signal to baseband and a low-pass filter to remove the high-frequencies harmonics due to the demodulation and reduce the readout noise of the system. However, in this system the digital processing chain is expanded in order to obtain better performance, as detailed shown in Fig. 6.3. In particular, this chain features:

- a **high-pass filter**, used to remove the offset and attenuate the 1/f noise of the system;
- two **mixers**, used to downconvert the signal to baseband and separate between the in-phase and quadrature component;
- two **low-pass filters**, used to filter-out the readout noise of the system and attenuate the spurious harmonics created by the demodulation;
- two **cascaded integrator-comb filters**, used to completely suppress the residual harmonics that are not attenuated enough by the low-pass filter.

Each one of these components will be now discussed in details.

6.4.1 High-pass filter

As already stated, the role of the high-pass filter is to cancel out the offset of the system. Notice that the lock-in scheme is usually enough to cancel it out: indeed, the mixer upconverts the offset around f_{demod} and the low-pass filter reduces its contribution. However, the acquired offset of the system can easily be in the order of mV. Even by placing the cut-off frequency at few Hz, a first-order LPF only reduces that contribution by a factor $\simeq 10^3$, so the up-converted offset will be visible as a sinusoidal signal with an amplitude of some µV, well above the noise of the system, that is around few hundreds of nV. This means that a high-pass filter would be beneficial in order to cancel out this contribution before getting upconverted. If the pole of the high-pass filter is

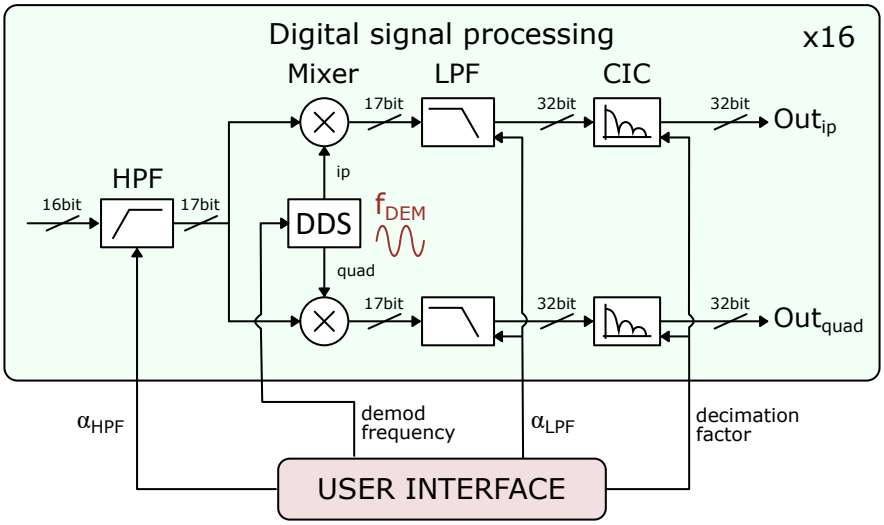

Figure 6.3: *Detailed view of the implemented digital processing chain.*

well before f_{demod}, the signal of interest does not get affected by the presence of the high-pass.

Practical implementation

To find the discrete-domain representation of this filter, it is necessary to find its Z-domain transfer function and anti-transform it. The transfer function of a HPF in the Laplace domain is

$$H_{HPF}(s) = \mu \frac{s\tau}{1+s\tau} \tag{6.2}$$

where μ is the gain at infinite frequency. Applying the bilinear transform $s = 2f_s \frac{1-z^{-1}}{1+z^{-1}}$, the Z-transform of the transfer function is

$$H_{HPF}(Z) = \mu \frac{2f_s \frac{z-1}{z+1}}{1+\tau 2f_s \frac{z-1}{z+1}} = ... = \mu \frac{2f_s\tau}{1+2f_s\tau} \frac{1-z^{-1}}{1-\frac{2f_s\tau-1}{2f_s\tau+1}z^{-1}} \tag{6.3}$$

Setting

$$\mu' = \mu \frac{2f_s\tau}{1+2f_s\tau} \quad \text{and} \quad \alpha = \frac{2f_s\tau-1}{2f_s\tau+1} \tag{6.4}$$

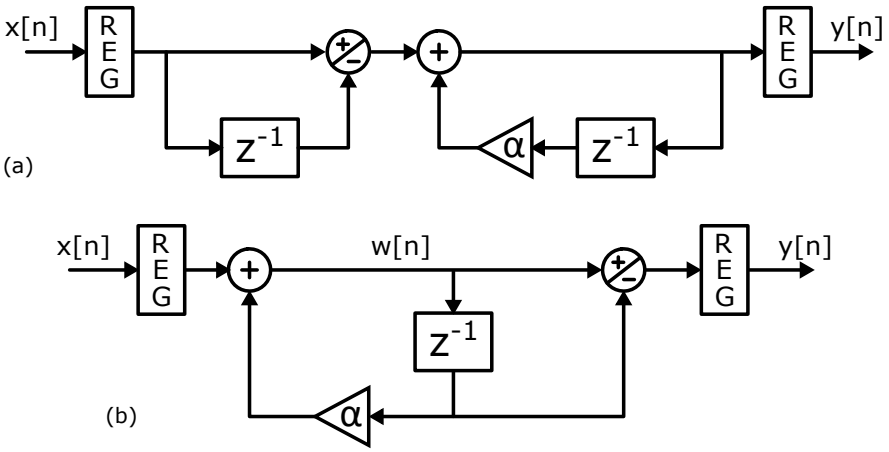

Figure 6.4: *Implementation of a first-order HPF in its (a) direct form I and (b) direct form II. The input and output signals are registered in order to avoid metastability issues.*

it is

$$H_{HPF}(Z) = \mu' \frac{1 - z^{-1}}{1 - \alpha z^{-1}} \tag{6.5}$$

From the Z-transform of the transfer function, it is possible to anti-transform it into its representation in the discrete time domain. In fact, recalling that $H(Z) = Y(Z)/X(Z)$ and that the multiplication with z^{-1} represent a one-step delay in the discrete time domain, the relationship between input and output is

$$y[n] = \mu'(x[n] - x[n-1]) + \alpha y[n-1] \tag{6.6}$$

According to Equation (6.3) and (6.4), the choice of μ' only affects the high-frequency gain of the filter. So, by setting

$$\mu' = 1 \Rightarrow \mu = \frac{1 + 2f_s\tau}{2f_s\tau} \tag{6.7}$$

it is possible to avoid the use of a multiplier, saving resources in the practical implementation of the filter. This comes at the expense of a high-frequency gain that is dependent from the position of the pole, that has to be properly normalized in the software. Therefore, Equation (6.6) is directly translated into the structure shown in Fig. 6.4(a), known in literature as "direct

form I" [125].

The feedforward block and the feedback one both perform linear time-invariant operation. This means that they can be swapped, as shown in Fig. 6.4(b), allowing the possibility of sharing the inner register. This solution is known in literature as "direct form II" [125].

According to Equation (6.4), it is possible to move the position of the pole by choosing different values for α. Being

$$\alpha = \frac{2f_s\tau - 1}{2f_s\tau + 1} = \frac{f_s - \pi f_p}{f_s + \pi f_p} \tag{6.8}$$

the possible values of α range from +1 to -1. The closer to +1 the value of α chosen, the closer to zero the frequency of the pole obtained. Therefore, α is a fractional binary number, so it has to be represented in 2's complement fixed-point binary format. For this design, it was chosen to use a 16-bit wide α, so that the largest fractional number that can be represented is $\frac{2^{15}-1}{2^{15}}$. This leads to a minimum pole of

$$f_p = \frac{1-\alpha}{1+\alpha}\frac{f_s}{\pi} = 3 \text{ Hz} \tag{6.9}$$

Dimensioning the filter

To dimension the filter, it is necessary to compute its high-frequency gain. This can be done directly from the Z-domain transfer function. For $Z^{-1} \to -1$, it is

$$H(-1) = \frac{2}{1+\alpha} \tag{6.10}$$

Hence, the maximum gain is, for $\alpha \to 0$, 2 at most. This means that, at the end of the chain, only one extra bit is necessary to avoid any overflow errors. This gain will be then normalized in the user interface, as already stated. The digital signal processing blocks following the HPF will have a 17-bit input word, as depicted in 6.3.

A problem arises due to the nature of binary multiplication. In fact, since the output of the multiplier is fed back to its input, the number of bits of the result would increase indefinitely if no action is taken. For this reason it is necessary to discard some

bits of the result. The number of bits necessary to avoid overflow has to be computed. To do so, it is necessary to compute the transfer function from the input signal x[n] to w[n], shown in Fig. 6.4, that is

$$W(Z) = W(Z) + \alpha W(Z)Z^{-1} \Rightarrow \frac{W(Z)}{X(Z)} = \frac{1}{1 - \alpha z^{-1}} \qquad (6.11)$$

The highest value of this transfer function is obtained for $Z^{-1} \rightarrow +1$, i.e. at low frequencies. Since α is a 16-bit signed fractional number, it is, at most

$$\frac{W[-1]}{X[-1]} = \frac{1}{1 - \frac{2^{15}-1}{2^{15}}} = 2^{15} \qquad (6.12)$$

meaning that it is necessary to keep 15 bits more after the multiplication. These bits are not useful for the final result, but they are necessary to avoid overflow in the inner stages. For this reason, they are often referred to as "guard bits".

Transposed form

The proposed solution needs 15 additional internal registers to store the guard bits, even though at the end of the chain they will almost all be discarded. This scheme can be further optimized by applying the transposition theorem [126], that allows to change the block diagram implementation of a linear time-invariant digital network leaving the network transfer function unchanged.

To apply the transposition theorem, the following steps must be followed

1. reverse the direction of all signal-flow arrows;
2. convert all adders to signal nodes;
3. convert all signal nodes to adders;
4. swap the x[n] input and y[n] output labels.

By applying this to the HPF described above, the result is the so-called "transposed direct form II", represented in Fig. 6.5. The transfer function from x[n] to y[n] is the same as before, and

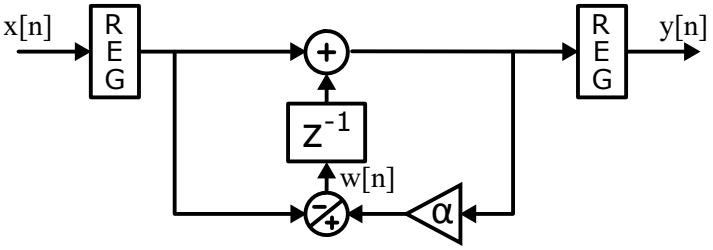

Figure 6.5: *Implementation of a HPF in the transposed direct form II. The input and output signals are registered in order to avoid metastability issues.*

the maximum gain is again 2. However, computing the transfer function from x[n] to w[n], it is

$$W(Z) = \alpha Y(Z) - X(Z) = \alpha(X(Z) + W(Z)Z^{-1}) - X(Z) \Rightarrow$$
$$\Rightarrow \frac{W(Z)}{X(Z)} = \frac{\alpha - 1}{1 - \alpha z^{-1}} \tag{6.13}$$

As long as $\alpha > 0$, the magnitude of this transfer function is limited between 0 and 1: this means that no extra guard bits are necessary for this implementation, allowing to save 15 registers for each channel and to use a 32-bit rather than a 47-bit multiplier.

6.4.2 Mixer

Once the DC offset is correctly filtered out by the HPF, the dithering signal has to be brought back to baseband with the use of a mixer. In order to maximize the SNR of the system, the spectrum of the demodulating signal should match the one of the signal of interest. However, as already stated in Chapter 2, the most convenient way to modulate a heater is to use a sinusoidal signal, despite being more resource-demanding. For this reason, a perfectly-sinusoidal oscillation of the heater power is assumed in the following.

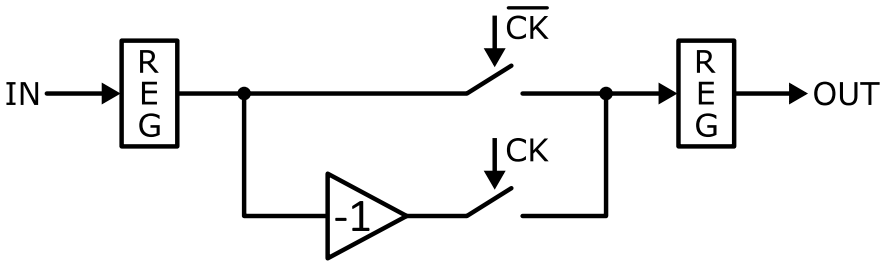

Figure 6.6: *Implementation of a square-wave digital mixer. No multipliers are necessary to perform the operation. Two mixers in parallel perform the required I/Q demodulation.*

Square-wave mixer vs. sinusoidal

For a perfectly-sinusoidal modulated signal, the SNR obtained with a sinusoidal demodulating signal would be

$$SNR_{sin} = \frac{A}{\sqrt{2n_y^2}} \tag{6.14}$$

where A is amplitude of the modulated signal and n_y^2 is the variance of the noise computed with the system bandwidth.

While it does not grant the same SNR, it is worth discussing a demodulation performed with a square-wave reference. Indeed, a square-wave mixer is far less resource-demanding, since no actual multiplier is necessary for its implementation, and as such it might be adopted to save resources whenever necessary. As depicted in Fig. 6.6, a square-wave multiplier can be implemented as a simple process that forwards the input word when the demodulating signal is '0', and inverts it when it is '1'. For a perfectly-sinusoidal modulated signal, the SNR obtained with a square-wave reference would be

$$SNR_{sqr} = \frac{A}{\frac{\pi}{2}\sqrt{n_y^2}} \tag{6.15}$$

which is $\simeq 10\%$ worse.

6.4.3 Low-pass filter

After the mixer, the dithering signal is brought back to baseband, and the 1/f noise of the electronic readout chain gets upconverted

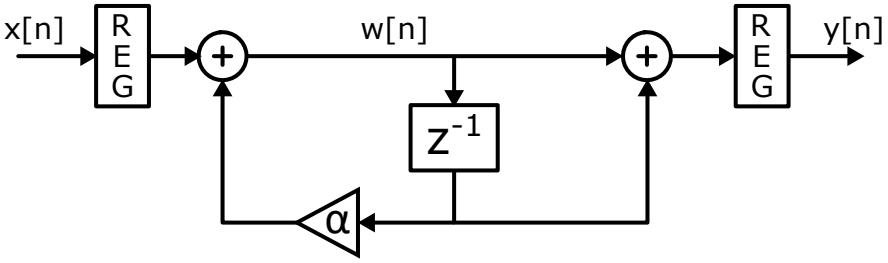

Figure 6.7: *Implementation of a first-order LPF in direct form II.*

around f_{demod}. The low-pass filter is hence useful to filter out the upconverted 1/f noise, the white noise, the quantization noise of the ADC, but also the high-frequency harmonics due to the demodulations and the other dithering tones present in the spectrum of the signal. The pole frequency of the low-pass filter should be set according to the well-known noise-bandwidth trade off.

Practical implementation

The steps followed to realize the LPF are very similar to the ones already described for the HPF. The goal is again to get a Z-domain transfer function starting from a Laplace-domain one in order to anti-transform it and get the discrete-domain representation. Only the final result is here reported, that is

$$y[n] = x[n] + x[n-1] + \alpha y[n-1] \tag{6.16}$$

with

$$\alpha = \frac{2f_s\tau - 1}{2f_s\tau + 1} = \frac{f_s - \pi f_p}{f_s + \pi f_p} \tag{6.17}$$

directly translated in the structure depicted in Fig. 6.7. Similarly to what discussed for HPF, the position of the pole is set by by the choice of α. By using a 16-bit wide 2's complement fixed-point binary representation, the minimum pole is, for $\alpha = \frac{2^{15}-1}{2^{15}}$, $f_p = 3$ Hz.

In order to proper dimension the filter, it is necessary to compute the maximum gain possible. So, for $Z^{-1} \to +1$, it is

$$H(1) = \frac{2}{1-\alpha} \tag{6.18}$$

that is, at most

$$H(1)_{max} = \frac{2}{1 - \frac{2^{15}-1}{2^{15}}} = 2^{16} \tag{6.19}$$

Hence, 16 extra bits should be used. Recalling that, after the HPF, the input word was 17-bit long, the output word will be 33-bit long. Since the minimum gain of the filter is 2, it is possible to compensate for this factor by doing a 1-bit right shift at the end of the stage. In this way, the output word is only 32-bit long, allowing an easier communication with the software through the USB 3.0 interface [127]. These bits are also sufficient to avoid overflow in the intermediate steps, as it can be easily seen by computing once again the transfer function from x[n] to w[n].

Detrimental effect of residual offset

Similarly to what happened with the HPF, the output of the LPF multiplier is fed back to its input, and the number of bits of the result would increase indefinitely if no action is taken. For this reason it is necessary to discard some bits. In particular, since the inner signal is 33-bit long word multiplied with 16-bit alpha, 16 bits have to be discarded to have a 33-bit long output word. The easiest way to do so is to simply discard the 16 least significant bits. This operation is called "truncation", and it means that the result of the multiplication would always been rounded down, as graphically depicted in Fig. 6.8(a). By assuming randomly-distributed results of the multiplication, the error caused by the truncation has the statistical distribution shown in Fig. 6.8(b). The variance of the distribution is [128]

$$\sigma^2_{trunc} = \frac{LSB^2}{12} \tag{6.20}$$

where LSB is the minimum value represented by the digital word after the truncation. The distribution is not centered around zero, since the mean value is:

$$\mu_{trunc} = -\frac{LSB}{2} \tag{6.21}$$

This non-zero average value of the error means that the sig-

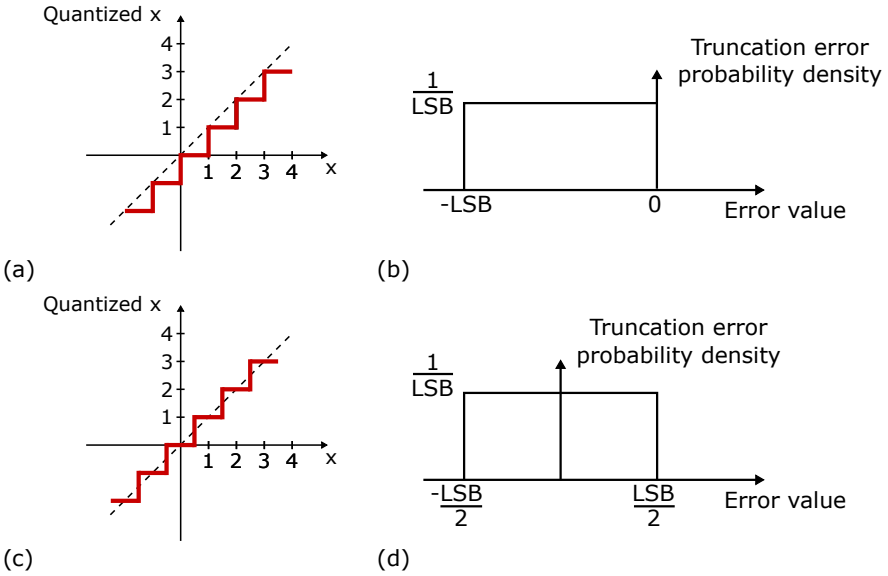

Figure 6.8: *Visual representation of the truncation (a) and the rounding (c) strategy. These approaches lead to an error statistical distribution reported in (b) and (d), respectively.*

nal, at the output of the LPF, is centered around μ_{trunc}. This offset has a detrimental effect on the control action since it is directly summed to the extracted dithering signal, effectively shifting down the whole transfer function by a factor μ_{trunc}. The effect of this error on the output optical power will depend on the slope of the transfer function around the locking condition, so it will be more accentuated for high-Q ring resonators.

A solution to solve this problem is to perform a rounding operation: rather then discarding the bits no matter what, the resulting word is rounded to the closest integer, as depicted in Fig. 6.8(c). In this way, the variance of the error stays the same, but the statistical distribution ranges from -LSB/2 to +LSB/2, meaning that it is now centered around zero (Fig. 6.8(d)).

Notice that this offset was not a concern in the implementation of the HPF: indeed, a non-zero average error of -LSB/2 gets summed as well at the output of the HPF, but it gets upconverted by the mixer and then removed by the following filtering stages.

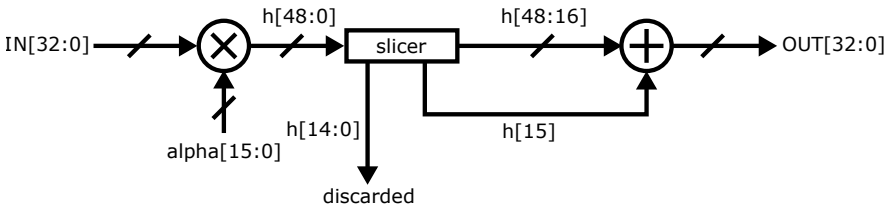

Figure 6.9: *Implementation of the custom rounding multiplier.*

Implementation of rounding multiplier

FPGAs usually offer dedicated logic slices designed for the efficient implementation of multiplication-based operations [129]. These often feature embedded rounding capability, and as such they should be exploited for the implementation of a rounding multiplier. However, this was not possible for this system, since the necessary multiplier was too wide to be implemented with one of these dedicated resources.

For this reason, a custom rounding multiplier was designed. To do so, a normal multiplication is performed and the 16 least significant bits are discarded. Then, the MSB of the discarded bits is assessed, and if it was a '1', a +1 is summed to the result of the operation, otherwise it is forwarded as it is. Practically speaking, this is implemented by simply summing the MSB of the discarded bits to the result of the truncation, as depicted in Fig. 6.9.

For the sake of completeness, it should be pointed out that this implementation does not yield to a mean error value that is exactly zero, since it is always rounding up the case where the discarded bits are exactly at half dynamics. Stated differently, over time with this implementation the result of the multiplication is rounded up slightly more often than it is rounded down. While more sophisticated versions of a rounding multiplier that avoid this problem are possible, the reduction of the mean value of the error was deemed sufficient for this application, since the resulting DC offset is now far lower than the noise of the system.

6.4.4 Cascaded-Integrator Comb filter

After the mixer, half of the power of the dithering signal gets downconverted to baseband, while the other half gets upcon-

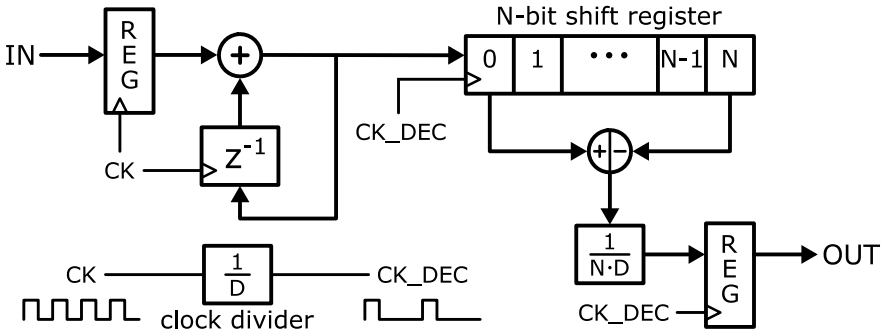

Figure 6.10: *Structure of the cascaded integrator-comb filter.*

verted at twice the dithering frequency, usually in the few tens kHz range. With the lowest possible bandwidth of the LPF, that is 3 Hz, the noise of the system has a RMS value around some hundreds of nV. With a typical dithering signal around 10 mV, the double-frequency harmonic of the extracted signal due to demodulation would be attenuated by factor $\simeq 10^4$, and hence would still be above the noise level, and would impair the control action. However, this harmonic is placed at a well-known and precise frequency, so a notch filter can be designed to remove it. A moving average filter has a periodic notching action that could be exploited for the purpose, but a large amount of memory elements is necessary to implement it. Moreover, after the mixing and low-pass filtering action, there is a lot of useless information in the data stream, since the bandwidth of the signal is greatly reduced. This is due to the high sampling frequency used, that was necessary since the signal was originally modulated around tens of kHz. A decimation should then be employed. Both these issues are efficiently tackled by the use of a cascaded integrator-comb (CIC) filter [130].

Practical implementation

The practical implementation of the filter is reported in Fig. 6.10. The stage features the cascade of an accumulator and a recursive moving averager operated by a decimated clock, resulting in a cardinal-sine transfer function. The filter should then be dimensioned so that the zeroes of the transfer function remove all the dithering tones that might be present in the spectrum of the

acquired signal. As explained in Section 3.5.1, all the spectrum components are multiple of a fundamental frequency f_t. The notches of the cardinal-sine transfer function of the filter are all multiples of

$$f_{notch} = \frac{f_s}{N \cdot D} \tag{6.22}$$

where f_s is the sampling frequency, N the moving average width and D the decimation factor. By setting $f_{notch} = f_t$, all the unwanted spectral components will be canceled out. Since $f_s = 625\,\text{kHz}$, by choosing $\text{D} = 8$ and $\text{N} = 64$, $f_{notch} = 1.22\,\text{kHz}$. Notice that this is the exact same frequency synthesized by the DDS, thanks to the fact that the clock of the DDS (from which the DDS frequency resolution depends on) is chosen to be an integer multiple of the sampling frequency (from which the CIC notches depend on).

Notice that the CIC filter has a low-passing behavior itself, setting in this way the ultimate limit of the lock-in bandwidth to $f_{notch}/2.258 \simeq 540\,\text{Hz}$, thus limiting to the maximum bandwidth of the control loop.

6.5 Integral Controller

The integral controller is implemented starting from its continuous transfer function and then using the bilinear transform. Only the final result is here reported, and the obtained discrete-domain representation is

$$y\,[n] = \frac{k}{2f_s}\,(x\,[n] + x\,[n-1]) + y\,[n-1] \tag{6.23}$$

that can be directly translated into the structure depicted in Fig. 6.11. In literature, this implementation is often called "trapezoidal integrator" [131].

The integral controller gain $\mu = \frac{k}{2f_s}$ can be tuned at runtime to set the bandwidth of the control system, as explained in Section 2.3. For the sake of completeness, in Fig. 6.11 the operating working condition of the loop $\bar{d}$ is also represented, that has to be subtracted to the input of the controller. As already stated, in order to have a power-independent calibration loop, that value should be set to zero, so that the loop will stabilize on

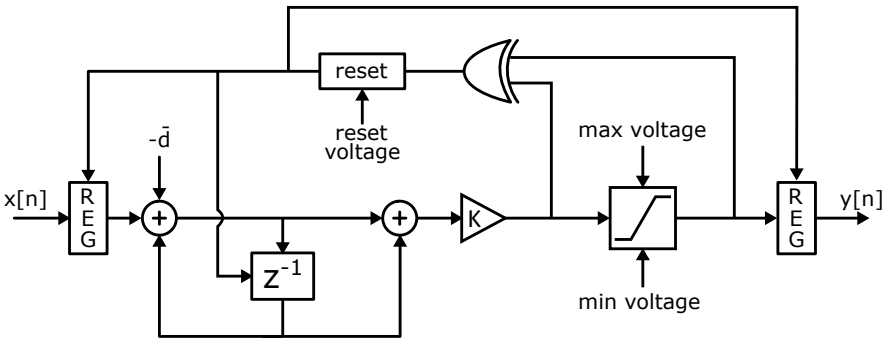

Figure 6.11: *Implementation of the integral controller. The gain K is placed at the end to avoid resolution problems. The saturation controller prevents the system from damaging the heater by limiting the maximum values of voltages allowed. When a saturation occurs, a reset circuit programs the correct value into all the internal registers of the stage so that the system is restarted from a certain value programmed by the user. The thresholds for the saturation controller can also be programmed by the user at run-time.*

a maximum (or minimum) of the device transfer function.

The thermal actuators driven by the integral controller have a limited operating range. In particular, the actuators are damaged by high currents, so the maximum voltage fed to the heaters should be limited. For this reason, a saturation-detection circuit is necessary, which should monitor if the output of the controller exceed the limits of the operating range (*max voltage* and *min voltage* in the figure). When a saturation is detected, the output has to be limited and the controller needs to be reset to a valid working point. The reset condition and the lower and upper threshold can all be chosen from the external user, since they depend on the application and on the technology platform used in the experiment. The reset voltage should be chosen in order to ensure that the system is able to find and lock on a stable working condition.

6.6 Resources sharing through time-division multiplexing

Another way to save more resources in the implementation of the described architecture comes from the observation that the

sampling frequency of the system (625 kHz) is much lower than the working frequency that can be handled by an FPGA. It is then possible to implement a solution to share the multiplier and the adders of each stage, in a time-division multiplexed fashion. This solution is feasible as long as the logic is able to process all the channels before the next samples are available. Notice that the new approach doesn't introduce any phase delay in the digital demodulation, since all the signals necessary for the purpose are still updated at 625 kHz.

This solution effectively reduces the number of multipliers and adders, that are the most critical resources in the FPGA. Notice that the number of memory elements is instead fixed by the number of channels, and hence does not change with the introduction of this technique. This solution complicates the HDL design of the filters, whose logic, for this reason, is implemented as finite-state machines (FSMs). All the digital processing stages presented in this chapter were thus updated to be operated in this way.

As an example, Fig. 6.12 shows the time-multiplexed implementation of the LPF. When a new data is available, the implemented logic cycles through all the memory elements of the structure and process each of them sequentially through a single multiplier and two adders, dramatically decreasing the use of these resources. The operation of the FSM is regulated by proper digital logic clocked at 40 MHz, that is enough to process up to 40 MHz / 625 kHz = 64 channels.

6.7 Timing closure

The digital system implemented on the FPGA must perform an increasingly number of operations. The complexity of the firmware increases as well, meaning that the tool for compiling the design (Vivado, from Xilinx) has to place and route a huge quantity of components into the design. In complex multirate systems such as this one, it is necessary to take particular care on the timing analysis, since even a single path which does not respect the timing requirements of the internal registers can compromise the entire functionality of the design. Solving all the timing issues makes the implementation of the design more

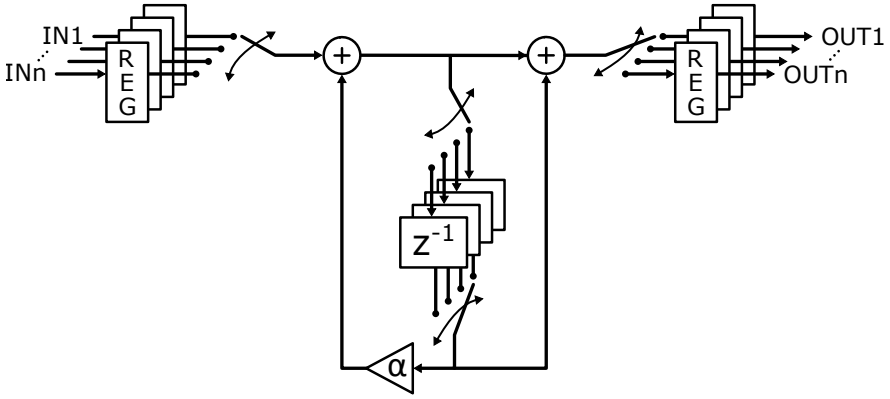

Figure 6.12: *Time-multiplexed implementation of the IIR low-pass filter, allowing to use a single multiplier to process all the 16 channels of the system*

deterministic and stable, reducing the differences between one run and the following, and may also reduce the time required to compile. This process is called "timing closure". However, over-constraining the design may lead to the opposite result, slowing down the implementation time required by the tool.

This system presents multiple clock domains that can interact with each other. These interactions can cause timing problem if the clocks are asynchronous. For this reason their frequencies were chosen so that each one is a multiple of the other one of a factor two. In particular, the clocks chosen for the design are:

- 160 MHz to handle the logic of the CLIPP stimulus signal;
- 80 MHz to be fed to the 16-channels DDSs, so that each channel gets updated with an effective frequency of 5 MHz, enough to ensure a good spectral purity even for the fastest dithering tones required in the assembly;
- 40 MHz to handle the acquisition logic of the ADC and the time-division multiplexed logic of the DSP stages;
- 20 MHz to handle the logic of the DACs connected to the heaters.

Another issue that might compromise the stability of the design is the high capacitive load of the internal signals. These

signals can seriously degrade the performance of the entire design: indeed, when a register output is connected to many nodes, it is difficult to find a location where all the fanouts paths are short and fast. To overcome this issue, the registers of the interested signals are duplicated, allowing the tool to better organize the location of the components of the whole system.

Finally, an accurate timing analysis must be performed not only to all the FPGA internal paths, which is automatically done by the tool, but also to all the external components. To do so, the user must provide explicitly the trace delays and the timing specifications of each external device to the tool.

6.8 User Interface

The possibility of interfacing the FPGA to a personal computer allows to implement a custom software to configure and monitor the behavior of the full system in real time. A C# user interface was thus implemented, using the environment "Visual Studio", provided by Microsoft. The software runs on any Windows 64-bit machine and communicates with the board through an USB 3.0 interface thanks to the OpalKelly APIs.

The main functionalities that the software performs are:

- load the configuration file onto the FPGA and check the connection with the board;
- set amplitude, frequency and phase of the dithering and demodulation signals;
- acquire the data from the USB interface and provide a graphical interface to properly display and possibly save them;
- set the voltage of the heaters or activate the control action by individually connecting each heater to their respective integral controller;
- perform some simple and low-speed control algorithm, as a preliminary step before implementing them directly in the FPGA. An application will be shown in Chapter 7.

For most of these tasks, the system stays idle waiting for an input from the user, and the operations do not take a considerable amount of time to be executed. The notable exception is the transfer of the acquired data from the FPGA, that is instead very resource-demanding, and should work in parallel to the other tasks. This means that the code handling the data transfer should be executed in a separated thread, so that it does not stop the main application from working. The code is then "non-blocking", since it is implemented as an asynchronous task.

In general, asynchronous programming can be defined as the way of executing programming code in a dedicated thread without having to wait for an I/O-bound or CPU-bound task to finish, such as file-system accesses or API calls. Working with asynchronous tasks has many advantages, such as keeping UI of the application responsive, improving the performance of the application and avoiding thread pool starvation.

CHAPTER 7

Control on a TiO_2 platform with a plasmonic-assisted sensor

In this chapter, the design of a plasmonic-assisted bolometric sensor is presented. Aided by simulations results, different chips have been fabricated in order to assess the best design parameters. The sensor was then integrated at the through and the drop port a micro-ring resonator, and together with a thermal actuator has been successfully used for a control experiment.

Part of the results presented in this chapter were published in:

A. Martinez*, **V. Grimaldi***, C. De Vita, D. Kumar Sharma, I. Gallet, A. Andrieux, J. Arocas, K. Hammani, L. MArkey, J. Weber, A. Bouhelier, M.Sampietro, A. Melloni, F. Morichetti, "Inline photothermal surface plasmon detector integrated in titanium dioxide waveguides", in IEEE Journal of Quantum Electronics (JQE), doi: 10.1109/JQE.2023.3244281.

**These authors contributed equally to this work*

7.1 Plasmonics

Photonics does not profit from the progressive technology scaling achieved in the recent years, since the minimum dimensions of the devices are defined by the diffraction limit of the light. Stated otherwise, photonic devices have dimensions that are similar to the wavelength of the light, usually in the order of some µm. This limit can be broken with the adoption of plasmonics [132]. Surface plasmon polaritons (SPPs) are waves that propagate along the surface of a conductor, typically a metal, thanks to the interaction of the light wave with the surface electrons. In this interaction, the free electrons respond collectively by oscillating in resonance with the light wave at the interface with a dielectric medium, hence propagating the electromagnetic energy along the surface of the conductor. In this sense, these structures are called "plasmonic waveguides" [133].

To create these miniaturized circuits, the light would first have to be converted into SPPs, which would then propagate and be processed by the different plasmonic devices, before being finally converted back into light. Many plasmonic devices have been fabricated, such as SPP logic gates or plasmonic MZIs [133]. In particular, high-speed plasmonic phase modulators have been demonstrated [134], achieving the energy efficiency of a ring modulator while promising to overcome the barrier of the 100 GHz modulation data rate [135].

If electrically contacted, the metal used to propagate the SPPs can be also used to carry electrical signals. This allows to envision the use of the same component both as a plasmonic device, but also as a detector to assess the quantity of power circulating in the waveguide. In this sense, a surface plasmonic detector (SPD) should be studied and developed, that could be offered as a building block for the next-generation plasmonic circuits, but also used as an ultra-compact detector as it is.

7.2 Design of a plasmonic waveguide

The waveguide used for this work is realized with an air-cladded 1.2 µm-wide and 375 nm-high titanium dioxide (TiO_2) core on top of a 2 µm-thick silicon dioxide (SiO_2) on a Si substrate, as shown

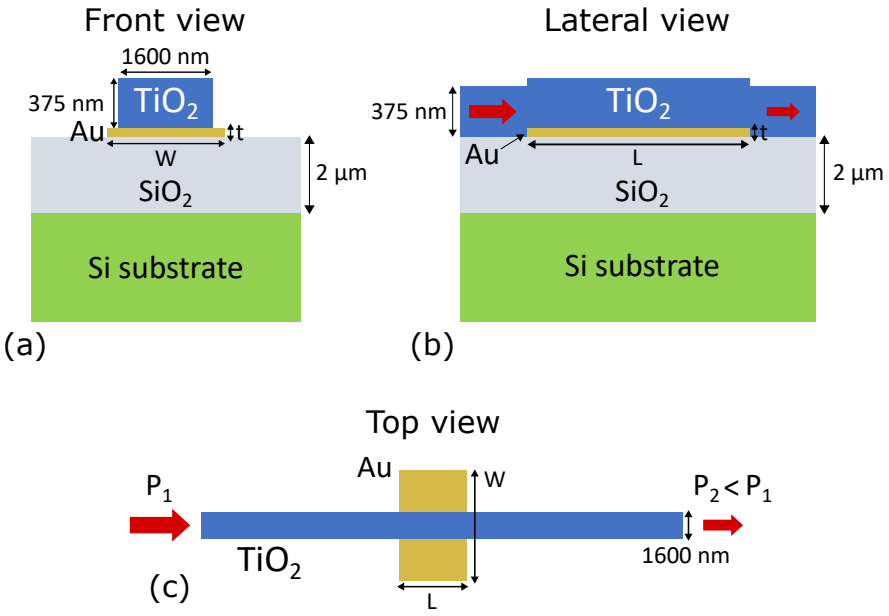

Figure 7.1: *Schematic view of the proposed plasmonic waveguide. The gold is placed below the TiO_2.*

in Fig. 7.1(a). To realize the plasmonic waveguide, a thin strip of gold is fabricated underneath the waveguide, in contact with the TiO_2 core. At the transition between the pure dielectric waveguide (no gold strip beneath) and the plasmonic waveguide (TiO_2 waveguide and gold strip), the TM fundamental mode of the light propagating in the dielectric waveguide may excite different SPP modes [136, 137], meaning that the light will propagate along the gold surface. Because of the losses associated with the plasmonic mode, part of the optical power will be absorbed by the metal and will turn into heat. At the other end of the plasmonic waveguide, the remaining power will excite again the mode of the pure-dielectric waveguide. Notice that the TiO_2 core is not interrupted where the metal is fabricated, since the dielectric medium is still necessary to support the propagation of the mode.

The simulated cross-sectional shape of the fundamental mode of the SPP waveguide (TM-polarized) is shown in Fig. 7.2(a). The plasmonic waveguide supports the propagation of higher order modes, as the one shown in Fig. 7.2(b), whose effects will be studied in the following.

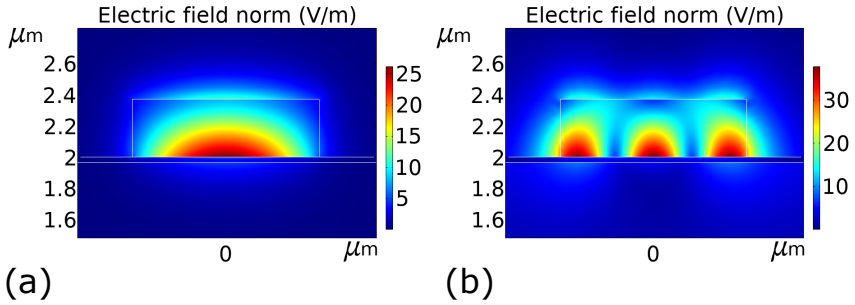

Figure 7.2: *Electric field profile of the guided modes of the plasmonic waveguide with* $T = 35\,\text{nm}$ *and* $W = 3\,\mu\text{m}$*: (a) first order and (b) third order SPP mode. The second order mode (not shown) is not excited by the fundamental TM mode of the dielectric waveguide because its electric field has an odd symmetry in the horizontal direction.*

7.3 Operating principle of the detector

The variation of the resistance of a metal due to temperature changes can be linearized around a certain temperature as

$$R = R_0 \cdot (1 + \alpha \Delta T) \tag{7.1}$$

where R_0 is the resistance of the metal at the reference temperature, α is the thermal coefficient of the metal and ΔT is the temperature difference. Since ΔT is proportional to the absorbed optical power, Equation (7.1) suggests that it is possible to assess the quantity of optical power in the waveguide just by measuring the variation of the resistance of the gold strip. To this aim, a certain amount of stimulus current I_{stim} should be forced through the metal and the voltage across its terminals should be measured. The information of interest is the voltage difference caused by the presence of the optical power, i.e.

$$\Delta V = I_{stim} R_0 \alpha \Delta T \tag{7.2}$$

Equation (7.2) suggests that, to maximize the quantity of signal available, the stimulus current should be as large as possible. However, the cross-section of the metal is in the μm^2 range, and so the capability of the structure to withstand large currents is limited.

7.3.1 4-probes lock-in detection scheme

The easiest and most-compact way to access the electrical information of the device is by using two pads, as sketched in Fig. 7.3(a). The same pad used for the stimulus current would be connected to the measurement apparatus to assess the voltage drop, while the other pad is grounded. The problem with this approach is that all the parasitic resistance due to the connections to the sensor (PCB tracks, bondings, on-chip connections, ...) would sum to the resistance of the sensor, hindering a correct electrical characterization. Moreover, any voltage drop between the ground of the sensor and the ground of the amplifier could worsen the readout performance.

For these reasons, a 4-probes approach is preferred, as sketched in Fig. 7.3(b). The current is injected through one of the terminals, while another one is grounded. The readout is then performed through the other 2 pads with an instrumentation amplifier (INA). In this way, the spurious voltage drop across the current path are not measured, while no current flows through the sensing tracks connecting to the differential amplifier, since they show a high-impedance path.

Another important aspect is to minimize the noise of the readout in order to maximize the minimum signal detectable and, in turn, the sensitivity of the system. To do so, it is better to rely on a lock-in detection scheme: by modulating the current I_{stim} it is possible to upconvert the information of interest around f_{stim}, that has to be reported in baseband with the use of a mixer. The mixing action upconverts the 1/f noise of the electronics, that can be then simply filtered out with a LPF.

7.3.2 Electronic implementation

The practical implementation of the sensor readout scheme is implement as sketched in Fig. 7.4. Given the limited current handling capability of the sensor, the best way to implement the system is to use a voltage DAC with a series resistance correctly dimensioned. Indeed, since the load resistance of the detector is very small, in the order of few Ohms up to few hundreds of Ohms, the voltage drop due to the resistive shunt is negligible. In this setup, a R_{stim} of $10\,k\Omega$ is usually used.

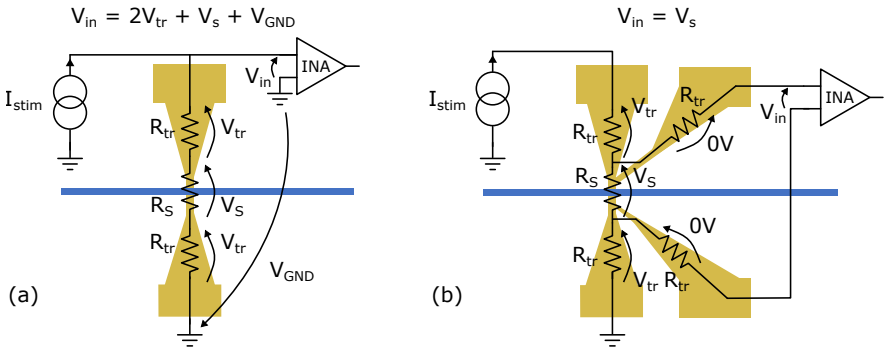

Figure 7.3: *Schematic view of a 2-probes measurements vs. a 4-probes measurements.*

The differential readout is carried out with an instrumentation amplifier (INA), that is then connected to an ADC in order to digitize the signal. The lock-in is completed in the digital domain, as previously described in Chapter 6. With this setup, the limiting contribution in terms of noise comes from the input voltage noise of the INA.

7.4 Technology

The thickness of the gold in the plasmonic waveguide has to be carefully assessed, since it effects both the real and the imaginary part of the effective refractive index of the plasmonic mode. The metal of the tracks and the pads instead should be as thick as possible, in order to provide a reliable electrical connection to the detector. For this reason, the proposed technology stack relies on two separate gold depositions, so that the two thicknesses can be controlled independently one from the other.

7.4.1 Fabrication process

The author thanks the CNRS for the details on the fabrication process.

All devices are fabricated on a 22 mm by 22 mm silicon chip covered by a 2 µm thick thermally grown SiO_2 layer. The devices are produced by three consecutive steps of electron-beam lithography. In a first step, the plasmonic gold detectors are patterned in a double-layer of PMMA (poly(methyl methacrylate)). This

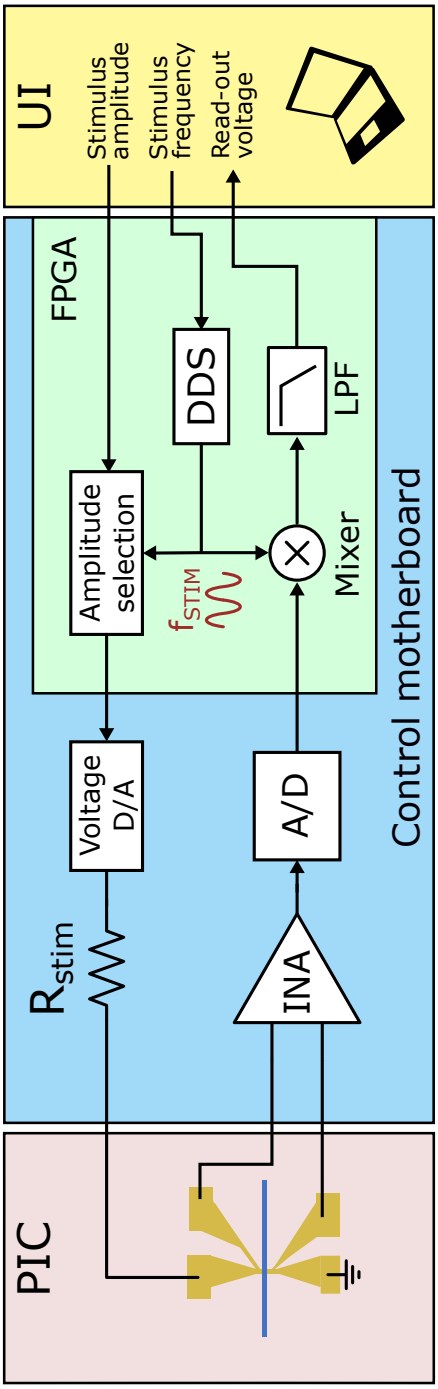

Figure 7.4: *Schematic view of the proposed electronics system. A series resistance R_{stim} of 10 kΩ is enough to minimize the effect of the resistor shunt on the stimulus current.*

initial step is finalized by a subsequent evaporation of a 3 nm Cr adhesive layer and a sub-40 nm Au film followed by a lift-off process. A 370 nm amorphous TiO_2 layer is then deposited on the entire chip by electron-beam evaporation of a suitable target material on a sample holder maintained at 470 °C. Ellipsometry of this layer measured a refractive index of 2.2 at 1550 nm. The TiO_2 PIC [138] including waveguides and grating couplers is patterned in a second electron-beam lithography step in a negative tone resist which is then used as a hard mask during the following inductively-coupled plasma reactive ion etching (ICP-RIE) of the TiO_2 layer. Here, to etch the layer a CF4 plasma (20sccm) with a radiofrequency power of 50 W, an ICP power of 200 W, under a pressure of 4.5 mTorr at 0 °C is used. The TiO_2 waveguides are 1200 nm-wide and 370 nm-high. The last electron-beam lithography step is to fabricate the pads and the tracks to electrically connect the metal of the plasmonic waveguide. The pads are made of 10 nm thick Cr and 200 nm thick Au layers to insure mechanical robustness against wire-bonding.

7.5 Dimensioning of the sensor

The design of the in-line SPD has to be carefully studied in order to control the insertion loss of the device. Particular attention should be paid to all the sources of loss that do not stem from the absorption of the propagating mode in the metal strip, since they would only increase the insertion loss of the device without providing any useful signal. To this end, electromagnetic simulations were carried out to evaluate the dependence of the complex refractive index of the SPP modes from the geometry of the metal strip. The refractive indices used for the SiO_2 bottom oxide and for the TiO_2 layer are 1.445 and 2.2, respectively.

7.5.1 Thickness

Figure 7.5 shows the simulated effective index of the two SPP modes as a function of the thickness t of the metal strip. The simulated curve shows how the real part of the index is barely affected by the thickness of the gold. Vice versa, the propagation

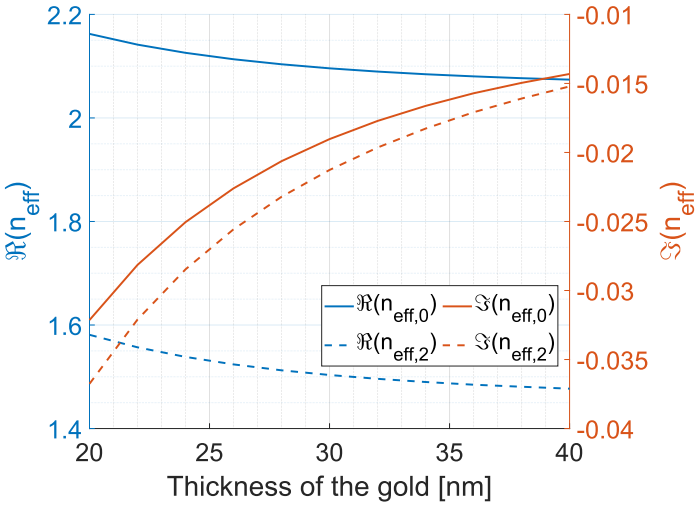

Figure 7.5: *Real (blue) and imaginary (orange) parts of the SPP modes refractive index vs. the gold thickness.*

loss, that depends on the imaginary part as

$$\alpha = -\frac{4\pi\Im\{n_{eff}\}}{\ln 10\lambda_0} 10^{-5} \text{ [dB/}\mu\text{m]} \tag{7.3}$$

substantially increases for thinner strips. In order to minimize the insertion losses of the detector, the gold should then be as thick as possible. However, if the metal is too thick, the significant discontinuity at the edges of the metal strip would severely increase the back reflections and reduce the light coupling to the detector. Thus, the metal thickness should be designed in order to reduce the gold absorptance while also minimizing back reflection loss. In the following, a thickness T = 35 nm is assumed.

7.5.2 Width

Regarding the width W of the gold, the lateral confinement of the SPP modes should be assessed. The width of the detectors should be designed accordingly, so to maximize the interaction of the light with the free electrons of the metal. At the same time, the gold strip should be as narrow as possible in order to increase the thermal resistance of the detector: in this way,

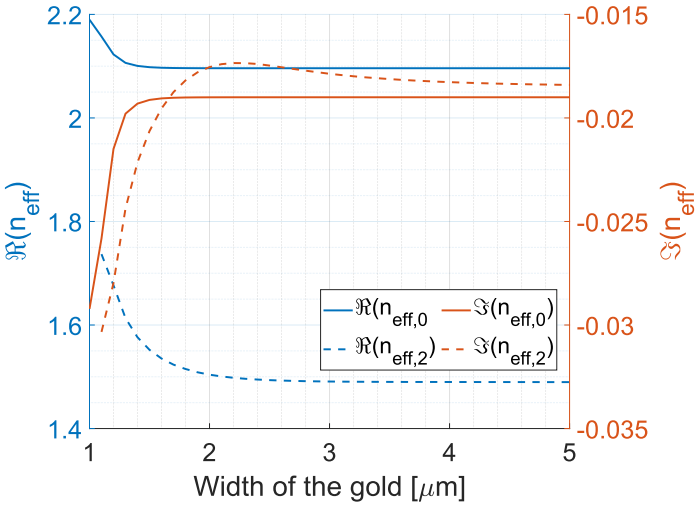

Figure 7.6: *Real (blue) and imaginary (orange) parts of the SPP modes refractive index vs. gold width.*

the same amount of absorbed optical power would cause a larger variation of the temperature of the detector and, consequently, of the read-out voltage.

In Fig. 7.6, the simulated refractive index as a function of the width is reported. It is possible to notice how the refractive index does not change after 2.5 µm, meaning that the guided light does not interact with any portion of the metal that is further than that. Thus, the width of the detector was set to 3 µm, in order to account for process variability and make it more robust to fabrication mismatches.

7.5.3 Length

The length of the metal strip L also determines the intensity of the transmitted optical wave and the amount of optical power that is absorbed by the metal. Figure 7.7 shows the simulated normalized transmittance across the SPP detector versus the length L (where W = 3 µm and t = 35 nm). Results show that the TM transmission has a decaying sinusoidal behavior with a period of 2.4 µm. The exponential decay is due to the fact that the SPD has a fundamental mode with a complex-refractive index $n_{eff} = 2.084 - i0.016$ which corresponds to a propagation

loss of 0.56 dB/µm. The periodic oscillation is instead due to the presence of the third order SPP mode (Fig. 7.2), which propagates independently along the surface of the metal at a different speed. Indeed, the real part of the refractive index of the first and third mode is substantially different, as shown in Fig. 7.5 and Fig. 7.6, meaning that the two modes travel with a different phase velocity and accumulate a relative phase shift along the SPP waveguide. When the relative phase difference between the two SPP modes is a multiple of 2π, they interfere constructively and P_{out} is maximum. Vice versa, when the relative phase difference is a multiple of π, the two modes interfere destructively and P_{out} is minimum. Since transmission is almost zero in all the minima (thus corresponding to complete destructive interference), it is possible to conclude that the two modes are excited with a similar efficiency at the transition between the dielectric and the SPP waveguide and that they have a similar propagation loss along the SPP waveguide, as confirmed by the fact that the values of the imaginary part of the respective refractive indexes are similar, as shown in Fig. 7.5 and Fig. 7.6. The second order mode instead is not excited, and thus does not contribute to the transmittance of the device. These considerations suggest that, in order to minimize the sensor insertion losses, it is fundamental to make the two modes interact constructively, targeting the first maximum of the transmission curve. The length chosen in the following is then $L = 2.2\,\mu m$.

Figure 7.7 shows also how the device does not exhibit an oscillating behavior for a TE dielectric mode, proving in this way that only TM-polarized light is able to excite the electrons of the metal and sustain a plasmonic mode. In this case the normalized transmittance simply shows an exponentially decay due to absorption in the metal strip. The shallow fluctuations on top of both TM and TE curves, with a period of about 500 nm, depends on Fabry-Perot (FP) cavity effects due to internal back reflections at the metal-dielectric interface; as the detector length L increases, this FP effect reduces because of the increased cavity loss due to metal absorption. As a result, the TE transmittance flattens after about 3 µm, while TM mode maintains the beating effects between the two SPP modes.

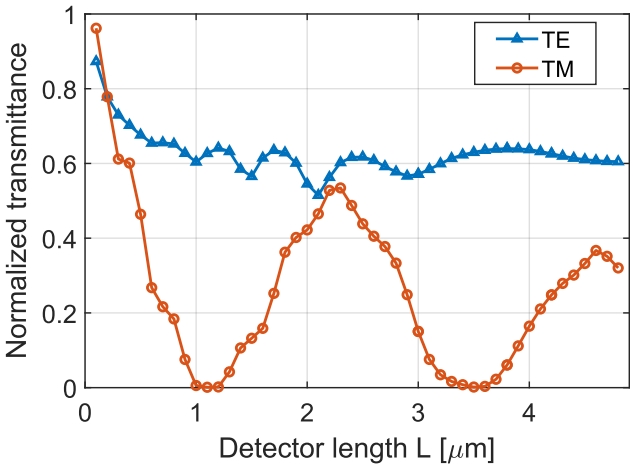

Figure 7.7: *Normalized transmittance of the SPD vs. the metal strip length L for TE (blue) and TM (orange) dielectric mode.*

Practical considerations

Given the importance of the beating effect between the two SPP modes on the transmission of the light, it is important to understand how this information is affected by the fabrication tolerances and the wavelength chosen.

The beat length ΔL, i.e. the distance between two maxima of the transmission curve in Fig. 7.7, is related to the difference between the real part effective index of the two SPP modes as

$$\Delta L = \frac{\lambda_0}{\Re\{n_{eff,0}\} - \Re\{n_{eff,2}\}} \tag{7.4}$$

In the considered SPP waveguide geometry (W = 3 µm and t = 35 µm), where $\Re\{n_{eff,0}\} = 2.084$ and $\Re\{n_{eff,0}\} = 1.489$, a beat length $\Delta L = 2.6\,\mu m$ is found, which is in very good agreements with the result shown in Fig. 7.7. Figures 7.5 and 7.6 also show that the difference between the real part of the effective index of the two SPP modes is almost independent from the thickness and the width W of the gold strip, meaning that the beat length ΔL is not critically affected by fabrication tolerances.

Moreover, the beat length does not change significantly with the wavelength of the light: Figure 7.8 shows how the simulated

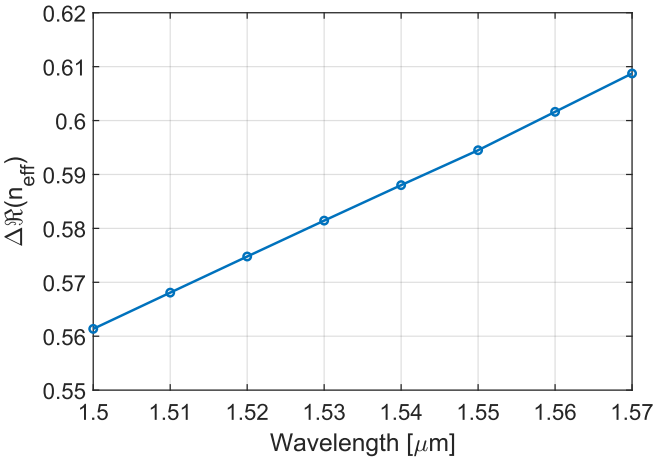

Figure 7.8: *Simulation of the difference between the real part of the effective refractive index of the first SPP mode vs. the input wavelength. The simulation suggests that the difference between the real part of the modes remains constant over a large bandwidth.*

refractive index difference $\Delta\Re\{n_{eff}\} = \Re\{n_{eff,0}\} - \Re\{n_{eff,2}\}$ between the two SPP modes changes by only $\pm 4\%$ across a wavelength range of more than 60 nm, thus enabling the SPD to work on a broad wavelength range.

7.6 Experimental results

7.6.1 Experimental validation of the beating effect

The curve shown in Fig. 7.7 was also experimentally validated. For the purpose, a chip containing many different sensors at different lengths was fabricated, as shown in Fig. 7.9. Light coupling into the chip was achieved by using vertically-emitting grating couplers designed to operate on TM polarization. Straight waveguides integrating SPDs with increasing length L were fabricated to assess the insertion loss (IL) and the back reflections of the SPD. Figure 7.10 shows the transmitted optical power P_{tx} versus L normalized to a reference all-dielectric waveguide (L = 0). Experimental data are averaged over a wavelength range from 1520 nm to 1580 nm, and the error bars indicate the standard deviation evaluated on several nominal identical samples. A very

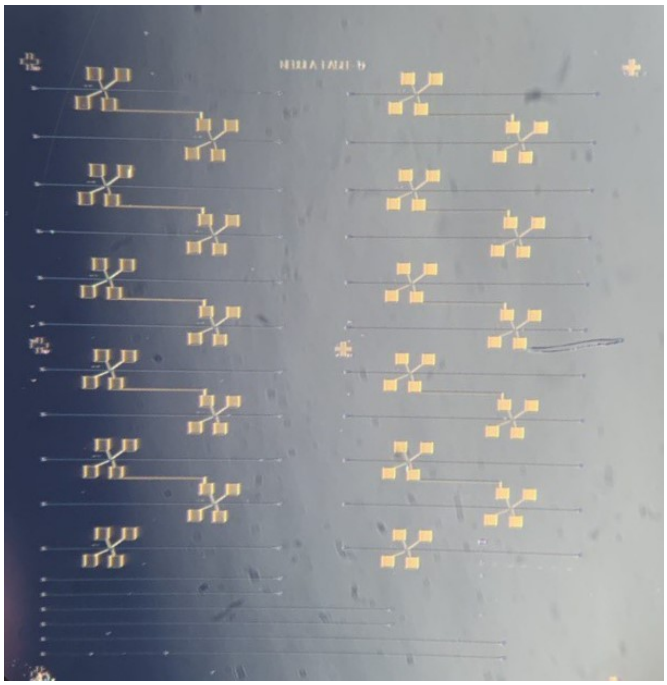

Figure 7.9: *Microscope picture of the fabricated sample. The chip integrates sensors with different lengths in order to experimentally demonstrate the oscillating dependency of the transmission vs. the length of the detector.*

good agreement is found with numerical simulations (data from Fig. 7.7 are reported in dB-scale for a direct comparison) on the position of the minimum transmission (around 1.1 µm) and of the maximum transmission (2.2 µm). The insertion loss in the transmission maxima (5.5 dB) is higher than the value expected from simulations (2.5 dB), probably due to the additional scattering losses in the plasmonic waveguide due to surface roughness at the metal-dielectric interface [139], not accounted for in the simulations.

The back reflection of each SPD was also measured by using coherent optical frequency domain reflectometry (OFDR) [140] and amounts to −17 dB, in agreement with electromagnetic simulations. This means that mode adapters from the dielectric to the plasmonic waveguide are not necessary.

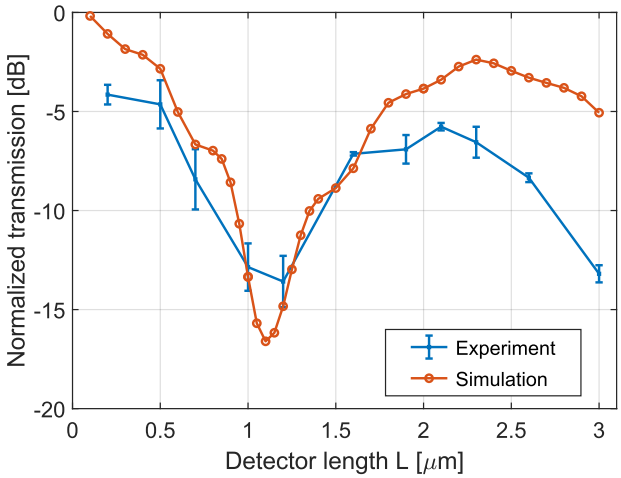

Figure 7.10: *Experimental measurements of the transmission of the SPD vs. its length (in blue). The simulated curve is also reported (in orange) for comparison.*

7.6.2 Passive characterization

The fabricated SPDs were also electrically characterized. As shown in Fig. 7.11(a), the chip was mounted on a host PCB that was glued to a Peltier cell. A commercial thermistor was mounted on the PCB next to the chip, in order to have a measurement of the temperature of the chip that is as precise as possible. With the use of a commercial TEC-controller, the temperature of the system was ramped from 10 °C up to 80 °C and then back to 10 °C, and the electrical resistance of the SPD was assessed every 10 °C. To do so, a stimulus current I_{stim} of 10 µA modulated around 20 kHz was forced through the sensor, and the voltage was measured through the lock-in scheme already described. The results are reported in Fig. 7.11(b), where it is possible to see how the resistance is linear with respect to the temperature, as expected. The thermal coefficient α of the gold for this technology, that is the slope of the reported curve, is estimated to be around 3.4 mΩ/°C, very similar to values reported in literature [141]. Finally, the sensor shows no hysteresis, since the values of resistance assessed stepping up the temperature (blue marks) are very similar to the ones assessed stepping it down

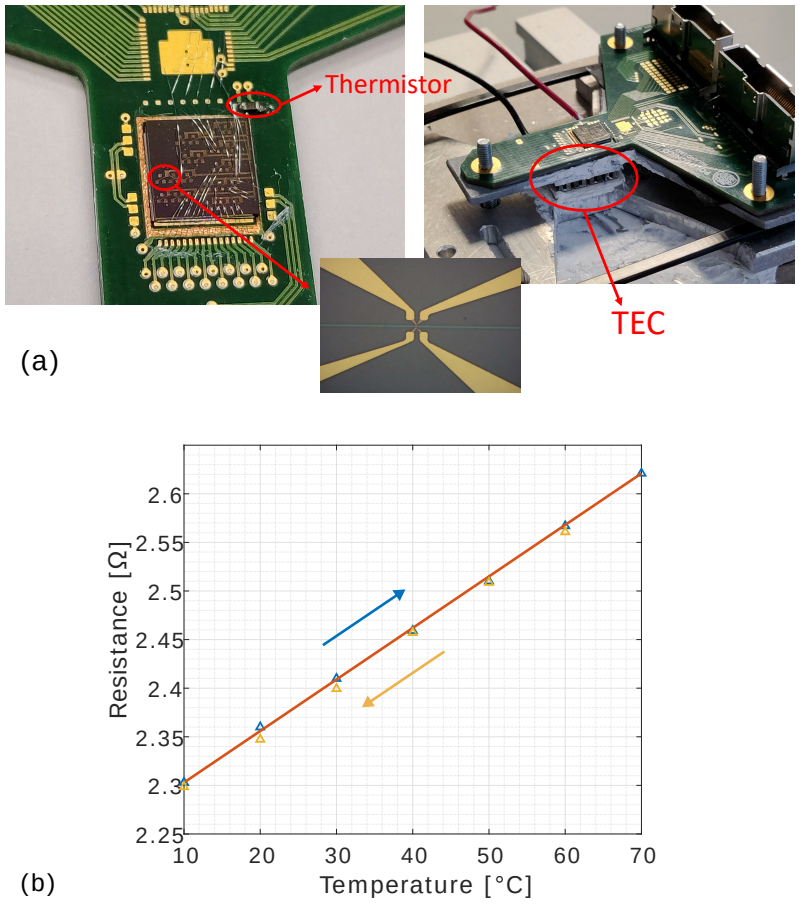

Figure 7.11: *(a) Electrical setup used to assess the thermal coefficient of the gold. The fabricated chip was mounted on a host board, glued to a Peltier cell used to change the global temperature of the chip. (b) Electrical resistance of the SPD versus temperature. The extracted a is* $3.4\,\mathrm{m\Omega/^{\circ}C}$*, and the sensor shows no hysteresis.*

(yellow marks).

7.6.3 Responsivity

The chip shown in Fig. 7.9 was also used to extract the sensitivity curve of the sensor. For the reasons described before, in the following only the sensor with $L = 2.2\,\mu m$ is considered. The sensor was stimulated with a current $I_{stim} = 25\,\mu A$ and the bandwidth of the lock-in was set to 3 Hz. Figure 7.12 shows the light-dependent change of the voltage ΔV across the SPD for increasing power P_{in} in the waveguide. As expected, ΔV is linear with respect to the input power P_{in}. The responsivity R_d, which is given by the slope of the linear fit (in red), is equal to $7.5\,\mu V/mW$. The minimum detectable power is around $-20\,dBm$, which corresponds to a voltage change $\Delta V = 75\,nV$, that is induced by a temperature variation of only $1.75 \times 10^{-3}\,°C$. The detector sensitivity can be extended to lower optical power by increasing I_{stim}, the limit being given by the current handling capability of the detector. The ultimate low limit in the sensitivity is not set by the noise of the electronic instrumentation, but rather by the oscillations of the global temperature of the system, that are inevitably sampled by the detector.

7.6.4 Time response of the sensor

The time response of the SPD is shown in Fig. 7.13. To measure this value, the light was modulated with a square wave at 20 kHz with the use of function generator. The same modulating signal was used to trigger an oscilloscope, connected to the output of an external INA for reading the detector voltage. The resulting curve is averaged over many measurements in order to improve the SNR of the system. The time constant of the detector response (red curve), which is defined as the time it takes for the system step response to reach $1 - e^{-1} \simeq 63.2\%$ of its final value, is about $1.2\,\mu s$, that corresponds to a bandwidth of about 130 kHz. The measured response is in very good agreement with simulations (yellow curve) performed. The limit of the SPD speed is set by the heat capacity of the dielectric materials and the gold thermal constant.

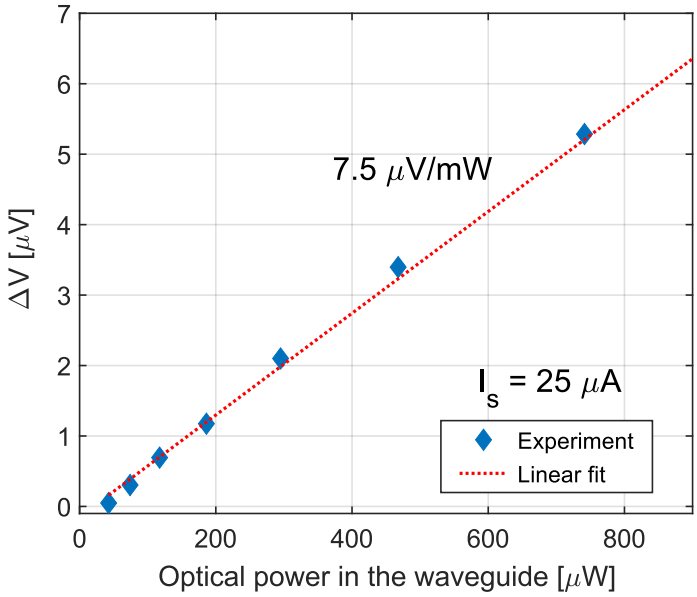

Figure 7.12: *Measured sensitivity curve of the sensor with* $I_{stim} = 25\,\mu A$ *and a lock-in bandwidth of* 3 Hz.

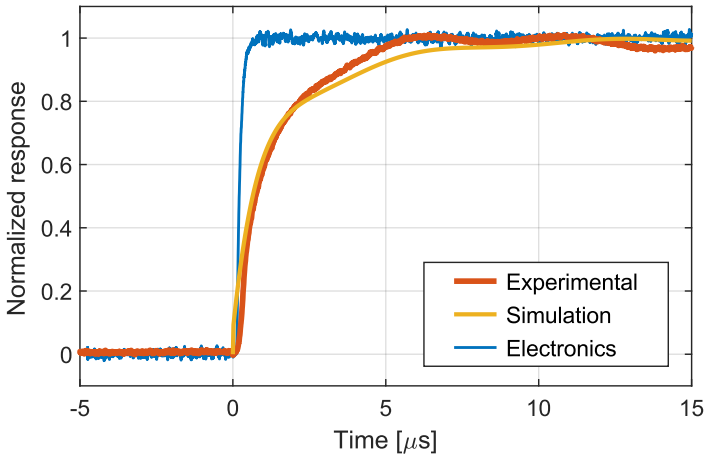

Figure 7.13: *Normalized time response of the SPD. In blue, the response time of the SPD when electronically excited, i.e. when directly connected to the function generator. In orange, the response when the modulated light is coupled into the system. In yellow, the simulated response time.*

7.7 SPD-based control of a ring resonator

The presented SPD was integrated on the through port of a ring resonator and used for a preliminary control experiment. This result is, to the best of the author's knowledge, the first example of controlling and stabilizing the working condition of an integrated photonic device through a plasmonic-assisted sensor.

A microscope photo of the chip realized for the control experiment is reported in Fig. 7.14(a). The chip features a micro-ring resonator (MRR) equipped with a through and a drop port. The gap between the ring and the waveguides is 130 nm, and the radius of the ring is 45 nm, leading to an FSR of 3.2 nm. The structure was experimentally characterized, as shown by the spectral response of the ring reported in Fig. 7.14(b), obtained by sweeping the wavelength of the input laser and measuring the emerging power. An SPD is integrated both on the through and on the drop port of the MRR. The SPD is designed with the optimum dimensions previously assessed, i.e. thickness t = 35 nm, width W = 3 µm, and length L = 2.2 µm.

7.7.1 Dimensioning of the actuator

To actively shift the working condition of the MRR, the integration of a thermo-optic actuator is necessary. Since the guides are air-cladded, it is not possible to fabricate a heater on top of the cladding of the structure of interest, as it usually done in commercial technologies. For this reason, the heater has been fabricated on the side of the ring, next to the TiO_2 waveguide. The expected TOC of TiO_2 is around $-2.3 \times 10^{-5}\,K^{-1}$ [56], an order of magnitude smaller than silicon. This means that it is pivotal to maximize the efficiency of the heater in order to have the maximum dynamic range possible.

The cross section and a microscope picture of the MRR and the heater are reported in Fig. 7.15. The drop port waveguide is fabricated orthogonally with respect to the through waveguide, and is then curved before reaching the output GC, as shown in the full picture in Fig. 7.14(a). In this way it was possible to maximize the length of the heater embracing the ring, increasing its electrical resistance and hence the power dissipated. For the fabrication, the metal layer devoted to the pads is preferred over

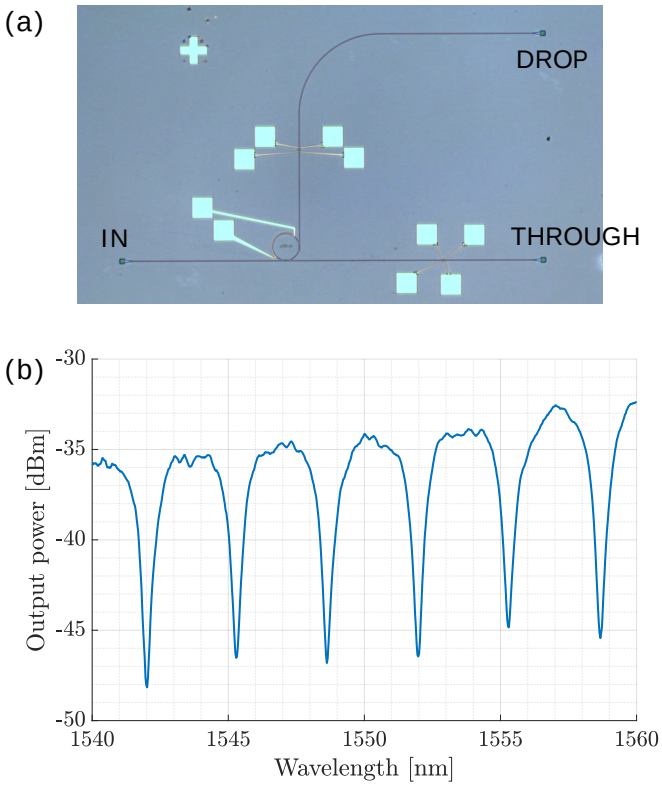

Figure 7.14: *(a) Microscope picture of the micro ring resonator. Both the through and the drop port are equipped with an SPD. (b) Spectral response of the MRR.*

the one used for the SPDs, since it is thicker and as such is able to carry more current before breaking down.

The distance between the heater and the ring should also be optimized. The closer the heater is fabricated to the ring, the more dissipated heat is directly transferred to the waveguide. However, the proximity of the metal next to a waveguide will increase the losses experienced by the optical field. For this reason, FEM simulations were carried out in order to find the minimum gap that would lead to acceptable performance. The results of the simulations suggested that a distance of 400 nm has very limited effect on the transfer function of the MRR, and it was then adopted for this design.

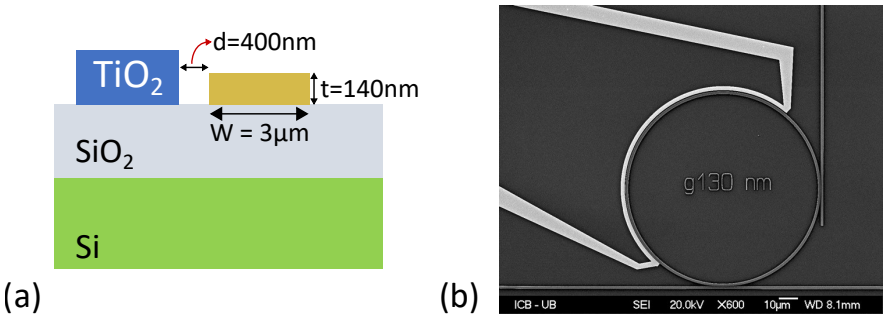

Figure 7.15: *(a) Cross section of the ring waveguide and the heater. All the dimensions of interest are indicated. (b) Microscope photo of the fabricated MRR and heater.*

The last parameter to be discussed is the width of the heater. Indeed, wider heaters can sustain more current before failing, but, for the same amount of current, they perform worse than thinner ones, since the same amount of heat would be dissipated in a wider area, so partly further from the ring. This means that the heater width should be tailored on the maximum current that the electronic driver is able to provide. To this aim, the current density limit for the deposited gold was experimentally assessed. A dummy heater was fabricated and tested by forcing increasing current until failure. The sustainable current density measured in this way was around $J = 10^7\,\mathrm{A}\cdot\mathrm{cm}^{-2}$. For a maximum output current of the electronic driver of 50 mA, a width of 3 µm was hence chosen.

7.7.2 Experimental setup

The circuit proposed cannot be operated without some special precautions. The main issue is that the heater changes the temperature of the chip, and hence, inevitably, the temperature of the SPD. This means that the dark resistance of the SPD changes during the normal operation of the chip, drowning the contribution of the light. A possible solution is to upconvert in frequency the light-induced temperature changes in the SPD. For the purpose, it is possible to externally AC-modulate the light with an OOK scheme. By biasing the sensor with a DC current, the spectrum of the voltage drop across its terminals will have a DC component related to the dark resistance and the heater cross-talk,

and an AC component only related to the effect of the light. By synchronously demodulating at the frequency of the OOK scheme with a lock-in readout, it is then possible to recover and isolate the effect of the light signal. The price to pay is a reduction of a factor 2 in the SNR, since only half of the signal of interest is upconverted, while half of it, corresponding to the mean value of the OOK-modulated light, is still at DC.

This solution is possible as long as the frequency of the modulation of the light is slower than the response time of the sensor, previously estimated to be around 130 kHz. At the same time, the upconverted light-induced signal should be placed above the corner frequency of the electronic instrumentation, in order to reduce the noise of the system and fully exploit the lock-in potential. For this reason, a modulation of few kHz is chosen.

The practical implementation of the setup is shown in Fig. 7.16. The MSB of the DDS output is connected to the laser and used as an external trigger to modulate the light. Notice that, since the signal is functionally a square-wave, the best SNR is achieved by adopting a square-wave mixer in the lock-in scheme.

The control algorithm is a simple minimum chaser implemented in the UI backend. During each iteration, the algorithm sums to the heater voltage a small step. Then, it computes the difference between the current sensor signal and the one assessed in the previous iteration, that is instead associated to the previous heater voltage. From the sign of result, the algorithm is able to understand whether the applied voltage step has shifted the device transfer function in the right direction or not. If the direction is wrong, the sign of the voltage step in the next iteration is simply inverted.

7.7.3 Experimental results

The sensor was biased with a current of 100 µA, the light was modulated with a 6 kHz signal, and the bandwidth of the lock-in was set to 3 Hz. The global temperature of the system was kept at 25 °C with the use of a Peltier cell glued below the host board, where the chip is housed, and a commercial thermistor. To prove how the proposed scheme is able to isolate and recover the information related to the luminous signal, the transfer function of the ring measured by the SPDs was first assessed. For the

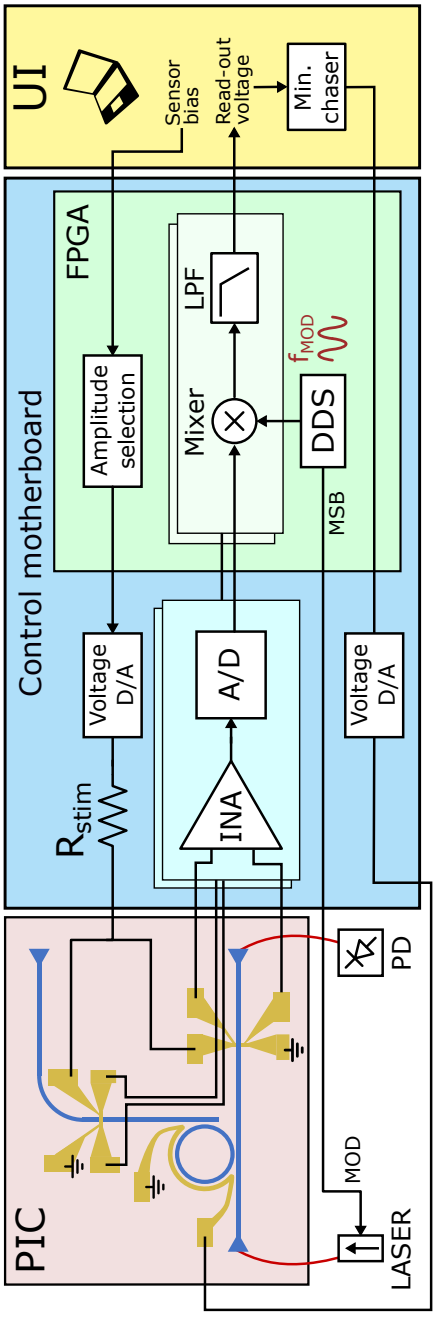

Figure 7.16: *Proposed experimental setup for separating the effect of the light from the dark-resistance variation induced by the heater.*

purpose, the wavelength of the input laser was manually swept, and the signal acquired by the lock-in scheme from the SPDs on the through and on the drop port was monitored, together with the power emerging from the chip. The results are reported in Fig. 7.17, showing how the transfer function extracted by the sensors are complementary and functionally related to the power assessed by the external PD.

After proving the validity of the approach, a control experiment was conducted. The sensor on the through port was used for the experiment, since it exhibits a higher signal than the one at the drop port, even at resonance. The results are shown in Fig. 7.18. The experiment starts with the ring in the resonance condition, thus the signal of the through sensor is correctly minimized and the heater is stable. Then, after 25 s, the wavelength of the input laser was manually perturbed with a step of −200 pm. The sensor signal and the power measured by the PD instantaneously increase, but within few seconds the control is able to restore the correct operating condition of the device, minimizing again the sensor and, in turn, the output power. Notice that, to compensate for a negative wavelength step, the heater power correctly increases, since TiO_2 exhibits a negative TOC. The system is perturbed again with a step-like disturb of −300 pm around 50 s, then again with −200 pm around 75 s, and finally with +500 pm around 95 s. Each time, the control is able to restore the correct working condition within a time scale of 10 s, proving that the SPD can be used to keep the MRR at resonance and counteract the effect of external wavelength fluctuations in the system.

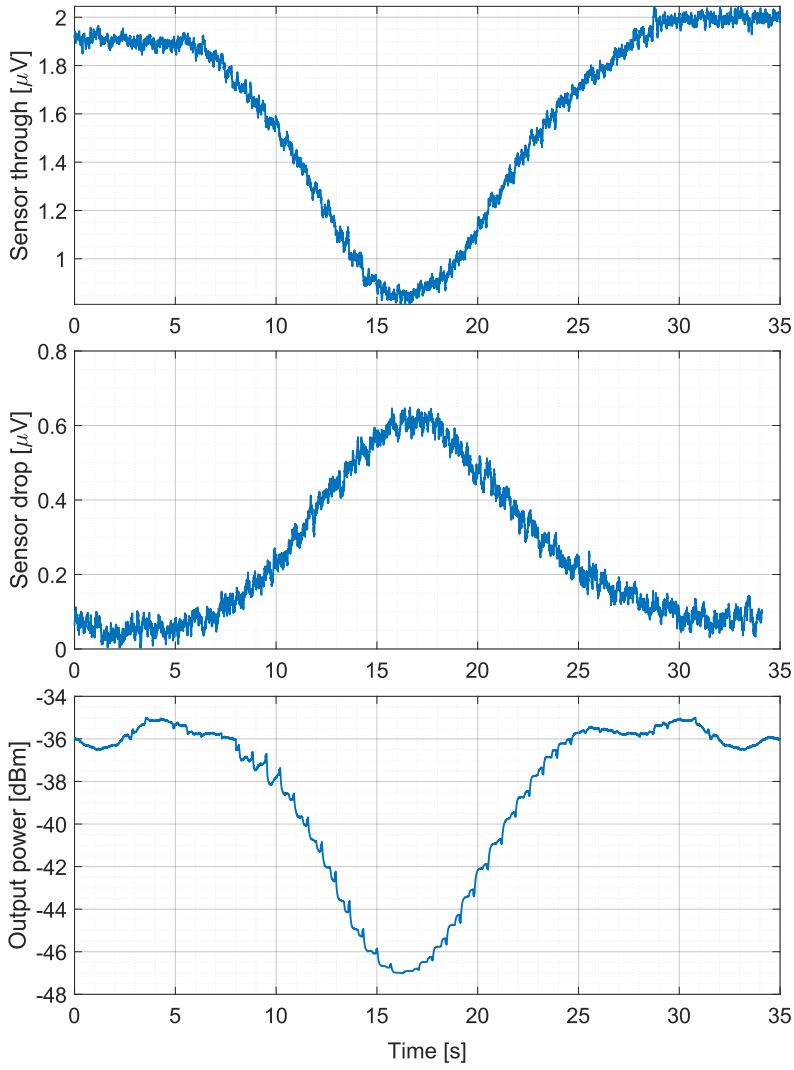

Figure 7.17: *Experimental validation of the proposed experimental setup to isolate the effect of the light. The transfer function of the MRR is measured with the sensor on the through and the drop port and compared with the power measured by the external PD.*

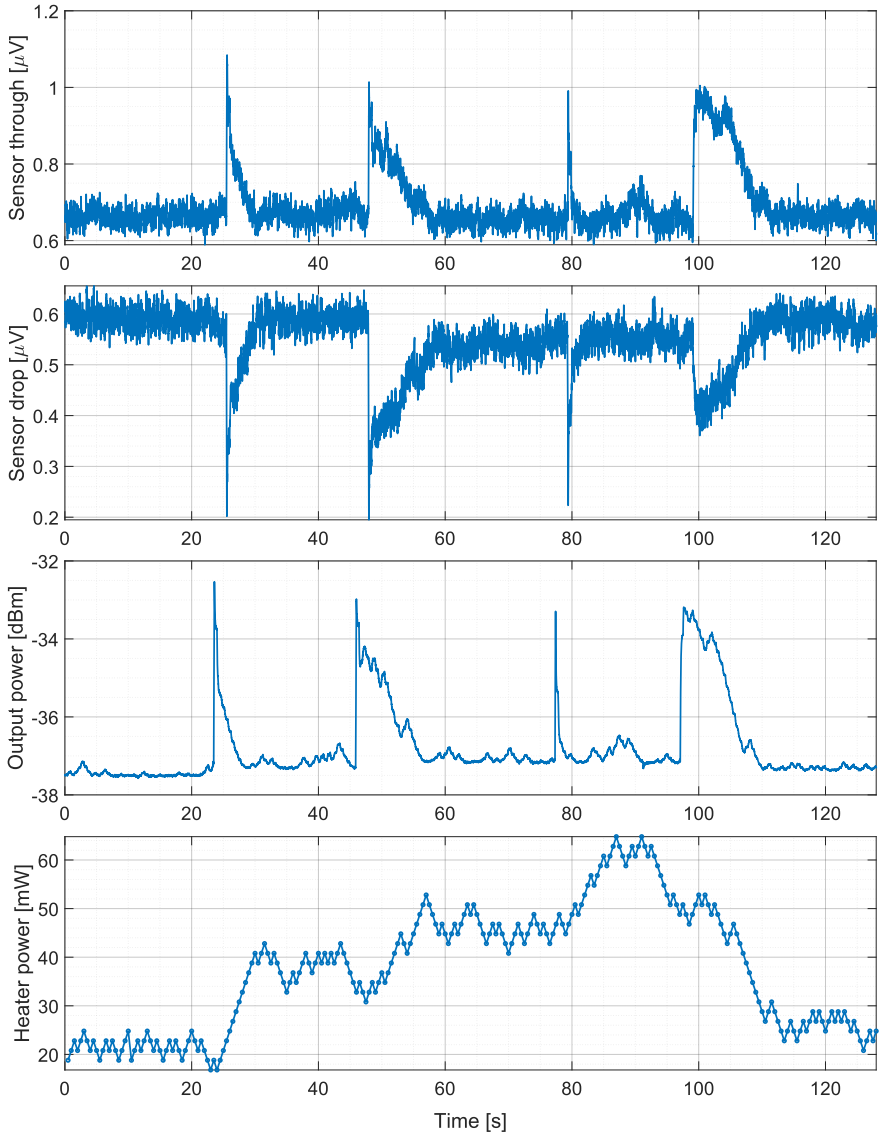

Figure 7.18: *Time evolution of the sensors signal, the power measured by the external PD, and the power of the controlled heater. Every time the system is perturbed, the algorithm is able to control the power dissipated by the heater to bring back the MRR at the resonance condition.*

CHAPTER 8

Conclusions

The market share of integrated photonics is expected to keep growing in the future years [142]. To fully unleash its potential and allow a large-scale diffusion of this technology in real scenarios, an electronic layer to operate and control the system has to be carefully designed and tailored for the application.

This thesis discussed efficient ways to implement the control electronics and the algorithm for stabilizing the behavior of different photonic devices. After having discussed the modular architecture of the system, the basic control loop was presented. The control scheme is based on the dithering technique that allows to extract the first derivative of the transfer function of the device of interest. The designed control loop was quantitatively discussed in order to understand the limits in terms of bandwidth and the effect of the system on the main parameters of interest. Advanced uses of the dithering techniques were also discussed, showing the possibility to discriminate between the effects of distinct actuators with the use of a single sensor by exploiting the orthogonality of different dithering frequencies and the frequency

selectivity of the lock-in readout.

The concept of orthogonality helps in simplifying the complexity of the assembly. In particular, the control of a programmable photonic mesh without integrated sensors was presented. The results show how it is possible to use a single external workbench power monitor to control 8 cascaded Mach-Zender interferometers, for which 16 different dithering signals were isolated from the same detector and hence the same electronic readout line. This approach comes at the expense of the maximum bandwidth of the loop, demonstrating how there is a substantial trade-off between the number of actuators controlled by a single sensor and the maximum speed of the control system.

The same mesh was also used to describe the pilot tones technique, particularly useful in circuits where multiple carriers are employed. By introducing a slow modulation of the optical intensity, it is possible to isolate the effect of one carrier from the others. The use of the pilot tone technique in conjuncture with the dithering technique causes the formation of intermodulated spectral components, each related to a single carrier but containing the information of the derivative of the transfer function, and as such they can be used with the same control loop previously described. However, the use of the pilot tones further reduces the maximum speed of the system. An experimental validation of this approach has been carried out, showing how a bi-diagonal mesh is able to separate and measure the effect of two orthogonal independently-labeled beams of lights shined at the same moment on the circuit.

The dithering technique is particularly useful to control the behavior of switching devices, that have to be stabilized on the maxima (or minima) of their transfer function. If a different working condition is targeted, the implemented loop can still be used with the first-derivative signal, but the resulting control scheme would not be power-independent nor calibration-free. Hence, an extension of the dithering technique was designed: by exploiting the non-linearities of photonic devices, it is possible to extract the second derivative of the transfer function and feed it to the control scheme. This is particularly useful for integrated modulators, especially high-Q micro-ring modulators (MRM), that have to be operated on the slope of their response. The perfor-

mance of the second-derivative control loop was experimentally assessed on an integrated Silicon Photonics MRM, showing how the locking point of the system is close to the theoretical optimum working condition of a modulator, i.e. the minimum of the transmitter penalty. The loop was also validated in long transmission experiments where an external variable was perturbed.

Since the same system can be used to extract both the first and the second derivative with the dithering technique, the same control loop can be used to stabilize both switching devices and modulators on the same chip. An experimental validation of this dual-control system was carried out. For the purpose, a subsection of a WDM-based interconnection scheme that allows different CPUs to communicate with each other was integrated in a single chip in a Silicon Photonics platform. This chip contains both resonant switching devices and MRMs, each stabilized and operated on the correct working condition, as proven by the long transmission and rerouting experiments carried out.

The enormous flexibility of the platform was possible thanks to the use of a digital core for operating the system. In order to ensure maximum versatility, an FPGA was used for the application. Despite the numerous resources offered in modern FPGAs, special care should be put in the design and in the implementation of the single structures, in order to fit and reliably operate the system with commercially-available FPGAs. The basic digital signal processing chain for implementing the lock-in amplifier was also expanded in order to minimize the effect of spurious spectral components on the performance of the loop.

Finally, a plasmonic-assisted bolometric sensor was designed. The sensor relies on the fact that the free electrons interact with the lattice of the metal during the plasmonic propagation, dissipating energy and hence heating up the device, changing its resistance. Electromagnetic simulations were carried out in order to understand the effect of the device dimensions on the plasmonic propagation. Because of the presence of a higher-order mode excited with a similar amplitude with respect to the first-order mode, a strong beating effect between the two modes was observed, and experimentally demonstrated. The responsivity and the time response of the detector were also assessed, showing a sensitivity down to $-20\,\mathrm{dBm}$ and a bandwidth of $130\,\mathrm{kHz}$. The

sensor was eventually used in a first simple control experiment on a micro-ring resonator, where a minimum chaser algorithm was programmed and operated in the custom UI. The experiment showed how the sensor could be exploited to maintain the device at resonance and counteract the effect of the externally superimposed wavelength disturbances.

8.1 Future developments

As fully explained in this thesis, the number of sensors used to control a complex photonic system can be reduced at the expense of the maximum bandwidth of the loop. Further studies have to be conducted to find the optimum condition in this trade-off, depending on the type of actuators adopted and the number of carriers used in the system. For instance, for triangular meshes such as the one used in this work, it could be envisioned to integrate one sensor after a certain number of MZIs, effectively splitting the control problem in smaller sub-meshes, allowing the re-use of the dithering tones in the system.

Another important aspect to address is the reduction of the non-linearities of the loop in order to increase the reliability of the implemented scheme. In particular, the dependence of the heater power from the square of typical electrical drive signals (either voltage or current) should be solved. Hence, a pre-distortion of the signal fed to the heaters should be performed in the feedback loop. This also solves the dependency in voltage-operated heaters of the oscillating power from the DC operating condition, that would further complicate the analysis and impair the linearity of the loop.

Furthermore, the reasoning behind the extraction of the second derivative through the dithering technique can be further expanded. The same system could then be exploited to extract higher-order derivatives that could be used to possibly stabilize the devices yet on different working conditions while preserving all the advantages of the discussed scheme. The higher-order derivatives will however feature decreasing power, and as such will require optimized sensors and readout scheme to be properly extracted.

Another promising research aspect is the use of sensors in

differentials and ratiometric readout schemes. These new approaches, used in conjecture with the dithering technique, can increment the flexibility in the operation of photonic circuits and pave the way to new unforeseen applications.

Regarding the plasmonic-assisted bolometric sensors, the biggest issue to solve is the thermal cross-talk between the sensor and the actuator. Relying on other forms of actuators can solve the problem entirely, but the technological maturity of these solutions is not sufficient for mass-integration. For this reason, a calibration might be employed to solve the issue and operate the sensor with the dithering technique. Another approach might come from the complete redesign on the plasmonic waveguide. Indeed, there is a growing research interest in hybrid-plasmonic waveguides, where a dielectric inter-layer is placed between the core of the waveguide and the metal to support the propagating mode. These redesigned waveguides allow to envision completely new sensing scheme, for instance by assessing the impedance of the dielectric inter-layer.

Whatever the future of photonics holds, it is hard to imagine one where electronics is not fully involved.

APPENDIX $\mathcal{A}$

Optimum filter for side lobes extraction

In this appendix, the problem of extracting two side lobes around a certain spectral component is discussed. The optimum filter is introduced and discussed, and the SNR obtained with this solution is computed. Afterwards, the solution where two subsequent demodulations are performed is also discussed, showing how it achieves the same SNR of the optimum filter.

A.1 Spectrum of the input signal

To understand the performance achievable in terms of signal to noise ratio, let us imagine the simplest spectrum of interest. This corresponds to the case where a small portion of the input light is modulated around f_{tone}. This signal will then pass through a certain photonic device whose transfer function is modulated at f_{dith} with the aid of some actuator. The output spectrum of the optical signal will have the shape shown in Fig. A.2, as also extensively explained in 3.5.1. Notice that, in order to be coherent with the convention typically adopted in the signal theory, the

bilateral spectrum is here used. Moreover, the absolute value of the signal extracted does not matter either, since the purpose of this appendix is to solely confront the SNR obtained with two (apparently) different solutions. Hence, a generic $\tilde{P}$ is employed for the signal of interest. The various components present in the spectrum are scaled accordingly.

A.1.1 Single demodulation

The theory of the optimum filter states that the best SNR is achieved when the weight function of the transfer function is exactly identical to the signal that has to be processed. In a lock-in stage, this means that the signal used for the demodulation must have the same spectral components of the portion of the signal of interest, as shown in Fig. A.1. In practical terms, this is easily constructed by summing two sinusoidal signals, one at $f_{dith} - f_{tone}$ and one at $f_{dith} + f_{tone}$. A unitary amplitude for each sinusoidal signal is here assumed for the sake of simplicity. After the mixer, a portion of the signal of signal of interest ($\frac{\tilde{P}}{2}$) is brought back to baseband, as expected. Notice that the overall power of the signal is doubled because the sum of two unitary sinusoidal signals is employed, so a gain of a factor two is expected. The noise spectral density at the output of the mixer will be

$$S_{n,out} = S_{n,in} \cdot 4 \left(\frac{1}{2}\right)^2 = S_{n,in} \tag{A.1}$$

where $S_{n,in}$ is the noise spectral density at the input of the mixer and $S_{n,out}$ the one at the output. The SNR achieved will be

$$\frac{S}{N} = \frac{\frac{\tilde{P}}{2}}{\sqrt{S_{n,in}\Delta f}} = \frac{\tilde{P}}{2\sqrt{S_{n,in}\Delta f}} \tag{A.2}$$

where Δf is the equivalent noise bandwidth used to compute the rms value of the noise.

A.1.2 Double demodulation

Nonetheless, purposely crafting the signal to realize the optimum filter is not the most straight-forward way to bring the spectral component of interest in baseband. Indeed, since the dithering

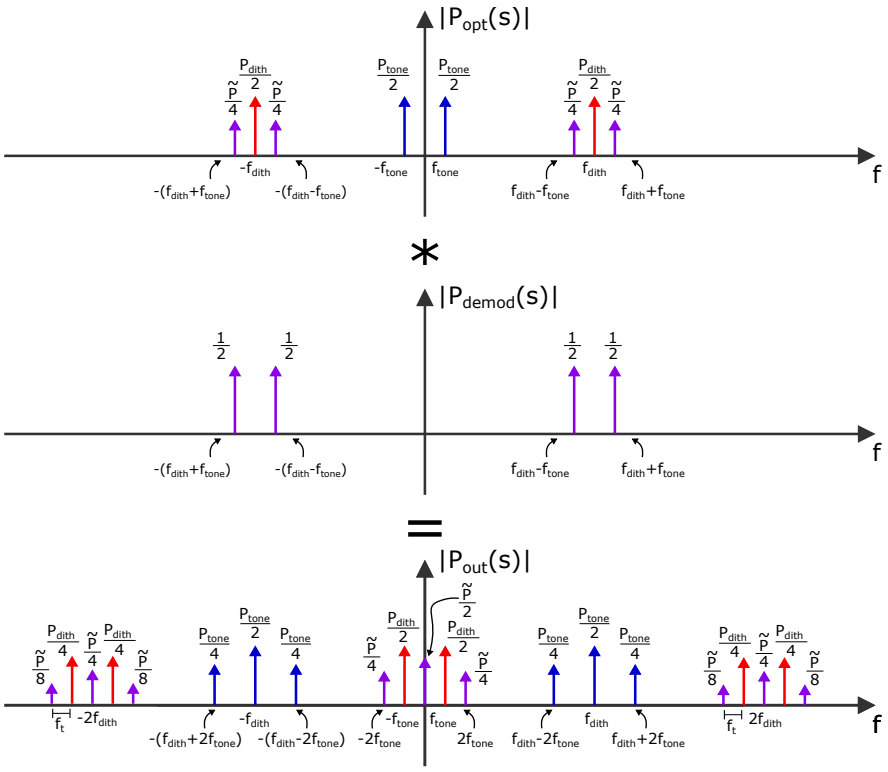

Figure A.1: *Spectrum obtained by performing a single demodulation with a signal matching the information of interest (optimum filter).*

oscillations and the pilot tone are both generated within the system in the first place, the two frequencies are already available to be used for the signal extraction. By performing two cascaded lock-in with those two signals, the side lobes will be downconverted to baseband, as desired. This solution, while it looks different from the one proposed above, grants the same SNR, since it overall implements the same transfer function. Indeed, by demodulating a first time with a signal at f_{dith}, the spectrum at the output of the signal would be the one reported in Fig. A.2. The noise spectral density will be

$$S'_{n,out} = S_{n,in} \cdot 2 \left(\frac{1}{2}\right)^2 = \frac{1}{2} S_{n,in} \tag{A.3}$$

where $S_{n,in}$ is the noise spectral density at the input of the mixer and $S'_{n,out}$ the one at the output of the first mixer. If a subsequent demodulation with a signal centered at f_{tone} is performed, the signal of interest will be brought back to baseband and will be worth $\frac{\tilde{P}}{4}$, as shown in Fig. A.3. For the noise, it is:

$$S''_{n,out} = S'_{n,out} \cdot 2 \left(\frac{1}{2}\right)^2 = \frac{1}{4} S_{n,in} \tag{A.4}$$

where $S''_{n,out}$ is the noise spectral density at the output of the second mixer. This means that the SNR of this implementation will be

$$\frac{S}{N} = \frac{\frac{\tilde{P}}{4}}{\sqrt{S''_{n,out}\Delta f}} = \frac{\tilde{P}}{2\sqrt{S_{n,in}\Delta f}} \tag{A.5}$$

where Δf is the equivalent noise bandwidth. As anticipated, this is the exact same SNR obtained before. This means that the two solution are identical in terms of signal-to-noise ratio, and the most convenient one should be chosen on a case-by-case basis.

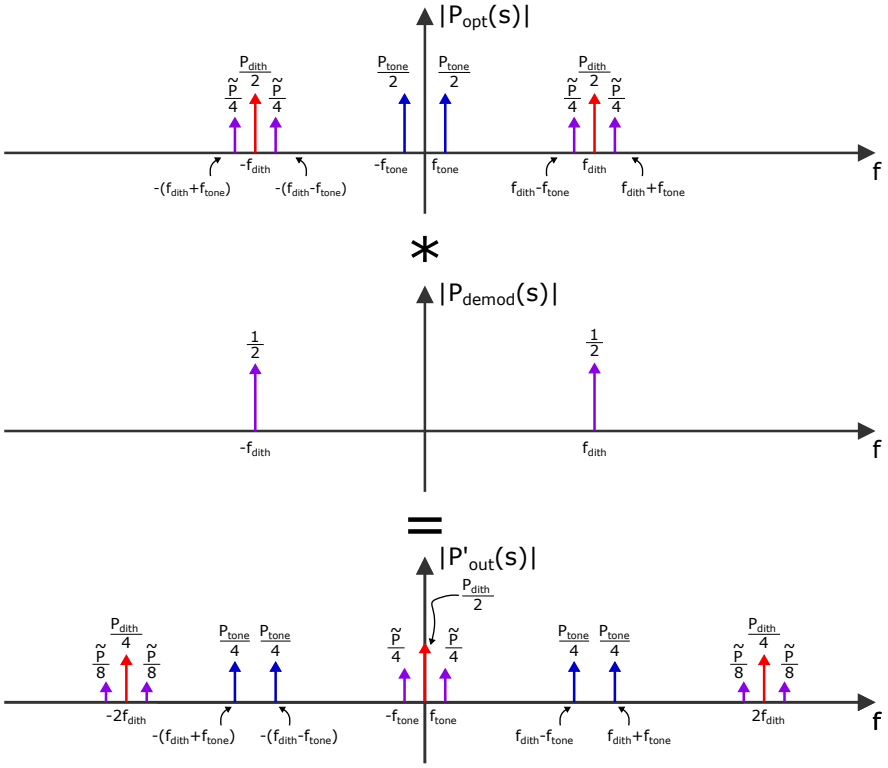

Figure A.2: *First demodulation performed with a first mixer and a signal centered around f_{dith}. The signal of interest is still modulated around f_{tone}.*

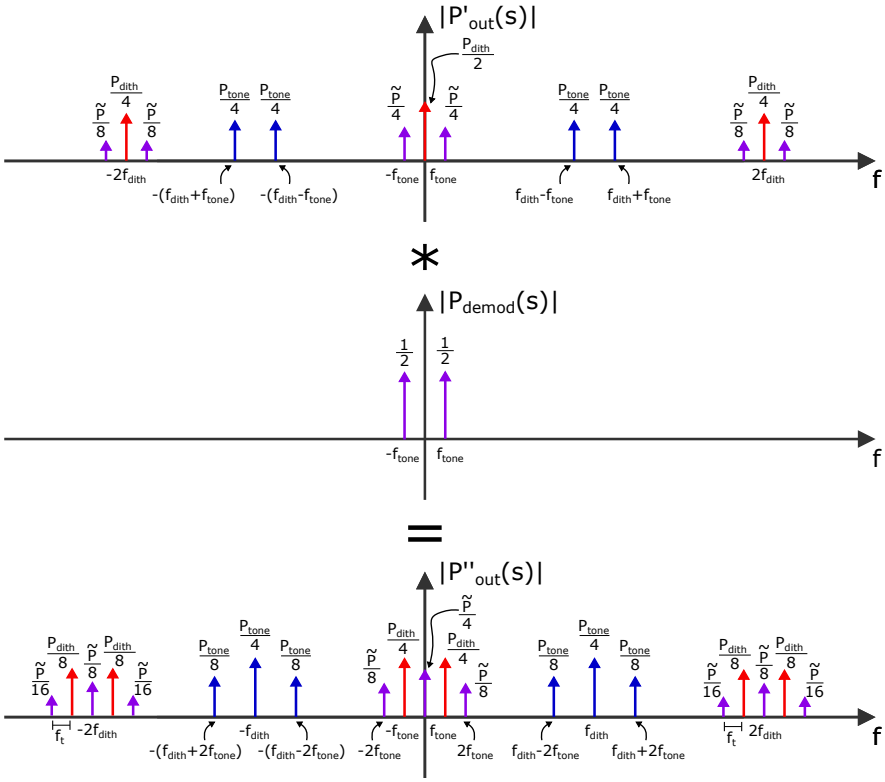

Figure A.3: *Second demodulation performed with a second mixer and a signal centered around f_tone. The signal of interest is downconverted to baseband. The output spectrum is the same obtained in the single demodulation case, except for a factor 2 in the gain.*

List of publications

Journal papers

F. Zanetto, F. Toso, **V. Grimaldi**, M. Petrini, M. Milanizadeh, A. Perino, F. Morichetti, A. Melloni, G. Ferrari, M. Sampietro, "Time-Multiplexed Control of Programmable Silicon Photonic Circuits Enabled by Monolithic CMOS Electronics", Laser & Photonic Reviews, *Under review.*

A. Martinez*, **V. Grimaldi***, C. De Vita, D. Kumar Sharma, I. Gallet, A. Andrieux, J. Arocas, K. Hammani, L. MArkey, J. Weber, A. Bouhelier, M.Sampietro, A. Melloni, F. Morichetti, "Inline photothermal surface plasmon detector integrated in titanium dioxide waveguides", in IEEE Journal of Quantum Electronics (JQE), doi: 10.1109/JQE.2023.3244281.
**These authors contributed equally to this work*

M. Petrini, S. Seyedinnavadeh, **V. Grimaldi**, M. Milanizadeh, F. Zanetto, G. Ferrari, F. Morichetti, A. Melloni, "Variable Optical True-Time Delay Breaking Bandwidth-Delay Constraints", Opt. Lett. 48, 460-463 (2023), doi: 10.1364/OL.478102.

V. Grimaldi, F. Zanetto, F. Toso, I. Roumpos, T. Chrysostomidis, A. Perino, M. Petrini, F. Morichetti, A. Melloni, N. Pleros, M. Moralis-Pegios, K. Vyrsokinos, G. Ferrari, M. Sampietro, "Self-Stabilized 50 Gb/s Silicon Photonic Microring Modulator Using a Power-Independent and Calibration-Free Control Loop", in Journal of Lightwave Technology, vol. 41, no. 1, pp. 218-225, 1 Jan.1, 2023, doi:

10.1109/JLT.2022.3210756.

I. Roumpos, T. Chrysostomidis, **V. Grimaldi**, F. Zanetto, F. Toso, P. De Heyn, Y. Ban, J. Van Campenhout, G. Ferrari, M. Sampietro, F. Morichetti, A. Melloni, T. Alexoudi, N. Pleros, M. Moralis-Pegios, K. Vyrsokinos, "Temperature and wavelength drift tolerant WDM transmission and routing in on-chip silicon photonic interconnects," Opt. Express 30, 26628-26638, 2022, doi: 10.1364/OE.455107.

M. Milanizadeh, S. SeyedinNavadeh, F. Zanetto, **V. Grimaldi**, C. De Vita, C. Kliti, M. Sorel, G. Ferrari, D.A.B. Miller, A, Melloni, F. Morichetti, "Separating arbitrary free-space beams with an integrated photonic processor", Light Sci Appl 11, 197 (2022), doi: 10.1038/s41377-022-00884-8.

F. Zanetto, **V. Grimaldi**, F. Toso, E. Guglielmi, M. Milanizadeh, D. Aguiar, M. Moralis-Pegios, S. Pitris, T. Alexoudi, F. Morichetti, A. Melloni, G. Ferrari, M. Sampietro, "Dithering-based real-time control of cascaded silicon photonic devices by means of non-invasive detectors," IET Optoelectronics, vol. 15, no. 2, pp-111-120, 2021, doi: 10.1049/OTE2.12019.

F. Zanetto, **V. Grimaldi**, M. Moralis-Pegios, S. Pitris, K. Fotiadis, T. Alexoudi, E. Guglielmi, D. Aguiar, P. De Heyn, Y. Ban, J. Van Campenhout, N. Pleros, G. Ferrari, M. Sampietro, A. Melloni, "WDM-based silicon photonic multi-socket interconnect architecture with automated wavelength and thermal drift compensation," IEEE Journal of Lightwave Technology, vol. 38, no. 21, pp. 6000-6006, 2020, doi: 10.1109/JLT.2020.3008001.

Conference proceedings

V. Grimaldi, F. Zanetto, F. Toso, C. De Vita, G. Ferrari, "Non-invasive light sensor with enhanced sensitivity for photonic integrated circuits", in PRIME conference, Villasimius, Italy, June 2022, doi: 10.1109/PRIME55000.2022.9816838. **Awarded as High Score Paper ("Gold Leaf Certificate", top 10% papers).**

V. Grimaldi, F. Zanetto, F. Toso, F. Morichetti, A. Melloni, G. Ferrari, M. Sampietro, "Enhancing Sensitivity and Reducing Temperature Dependence of Contactless Light Sensors", in European Conference on Integrated Optics (ECIO), Milan, Italy, May 2022, pp. 107-109.

F. Zanetto, F. Toso, **V. Grimaldi**, M. Petrini, M. Milanizadeh, A. Perino, P. Piedimonte, F. Morichetti, A. Melloni, G. Ferrari, M. Sampietro, "Monolithically Integrated Electronics in Zero-Change Silicon Photonics", in European Conference on Integrated Optics (ECIO), Milan, Italy, May 2022, pp. 15-17.
Awarded as High Score Paper (best paper in his category).

S. SeyedinNavadeh, M. Milanizadeh, F. Zanetto, **V. Grimaldi**, C. De Vita, G. Ferrari, D.A.B. Miller, A. Melloni, F. Morichetti, "Multi-channel free-space optical communication between self-configuring silicon photonics meshes", in European Conference on Integrated Optics (ECIO), Milan, Italy, May 2022, pp. 366-368.

A. Martinez, **V. Grimaldi**, D. K. Sharma, C. De Vita, F. Morichetti, A. Bouhelier, M. Sampietro, "In-line photo-thermal plasmonic detectors integrated in TiO_2 optical waveguides", in European Conference on Integrated Optics (ECIO), Milan, Italy, May 2022, pp. 248-250.

M. Milanizadeh, S. SeyedinNavadeh, F. Zanetto, **V. Grimaldi**, C. De Vita, C. Klitis, M. Sorel, G. Ferrari, D. A.B. Miller, A. Melloni, F. Morichetti, "Multi-Channel Free-Space Optics Receiver Enabled by a Programmable Photonic Processor", 2022 IEEE Photonics Society Summer Topicals Meeting Series (SUM), 2022, pp. 1-2, doi: 10.1109/SUM53465.2022.9858212.

S. SeyedinNavadeh, M. Milanizadeh, G. Benci, C. De Vita, C. Klitis, M. Sorel, F. Zanetto, **V. Grimaldi**, G. Ferrari, D.A.B. Miller, A. Melloni, F. Morichetti, "Self-Configuring Silicon-Photonic Receiver for Multimode Free Space Channels", 2021 IEEE 17th International Conference on Group IV Photonics (GFP), 2021, pp. 1-2, doi: 10.1109/GFP51802.2021.9674015.

I. Roumpos, T. Chrysostomidis, **V. Grimaldi**, F. Zanetto, F. Toso, P. De Heyn, Y. Ban, J. Van Campenhout, G. Ferrari, M. Sampietro, F. Morichetti, A. Melloni, K. Vyrsokinos, T. Alexoudi, N. Pleros, and M. Moralis-Pegios, "Temperature Tolerant On-Chip WDM Silicon Photonic Transmitter and AWGR-based Routing Interconnects", in Optical Fiber Communication Conference (OFC) 2022, S. Matsuo, D. Plant, J. Shan Wey, C. Fludger, R. Ryf, and D. Simeonidou, eds., Technical Digest Series (Optica Publishing Group, 2022), paper W4H.2, doi: 10.1364/OFC.2022.W4H.2.

F. Toso, M. Milanizadeh, F. Zanetto, **V. Grimaldi**, A. Melloni, M. Sampietro, F. Morichetti, G. Ferrari, “Self-Stabilized Silicon Mach-Zehnder Interferometers by Integrated CMOS Controller”, 2021 IEEE Group IV Photonics Conference (GFP), pp 1-2, doi: 10.1109/GFP51802.2021.9673835.

P. Piedimonte, F. Zanetto, F. Toso, **V. Grimaldi**, L. Sola, M. Cretich, A. Gori, M. Chiari, G. Ferrari, M. Sampietro, “Digital count of antibodies through differential impedance for high-resolution immunosensing”, 2021 Biomedical circuits and systems conference (BioCAS), pp. 1-4, doi: 10.1109/BioCAS49922.2021.9644983.

F. Toso, F. Zanetto, **V. Grimaldi**, A. Perino, E. Guglielmi, F. Morichetti, A. Melloni, G. Ferrari, M. Sampietro, “Electronics-photonics co-design for robust control of optical devices in dense integrated photonic circuits”, 2021 IEEE Custom Integrated Circuits Conference (CICC), 2021/4/25, doi: 10.1109/CICC51472.2021.9431512.

M. Moralis-Pegios, F. Zanetto, E. Guglielmi, **V. Grimaldi**, K. Fotiadis, S. Pitris, T. Alexoudi, P. De Heyn, Y. Ban, J. Van Campenhout, D. Aguiar, G. Ferrari, M. Sampietro, A. Melloni, N. Pleros, "Automated thermal drift compensation in WDM-based silicon photonic multi-socket interconnect systems," Optical Fiber Communications Conference and Exhibition (OFC), San Diego, 2020, doi: 10.1364/OFC.2020.W3G.2.

List of Figures

List of Tables

Bibliography

[1] "A way forward for optical fibers," *C&EN Global Enterprise*, vol. 98, no. 10, pp. 28–33, Mar 2020.

[2] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proceedings of the IEEE*, vol. 88, no. 6, pp. 728–749, 2000.

[3] T. Alexoudi, N. Terzenidis, S. Pitris, M. Moralis-Pegios, P. Maniotis, C. Vagionas, C. Mitsolidou, G. Mourgias-Alexandris, G. T. Kanellos, A. Miliou, K. Vyrsokinos, and N. Pleros, "Optics in computing: from photonic network-on-chip to chip-to-chip interconnects and disintegrated architectures," *Journal of Lightwave Technology*, vol. 37, no. 2, pp. 363–379, Jan 2019.

[4] R. Soref, "The past, present, and future of Silicon Photonics," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 12, no. 6, pp. 1678–1687, 2006.

[5] K. Worhoff, E. Klein, G. Hussein, and A. Driessen, "Silicon oxynitride based photonics," in *2008 10th Anniversary International Conference on Transparent Optical Networks*, vol. 3. IEEE, 2008, pp. 266–269.

[6] M. Smit, X. Leijtens, H. Ambrosius, E. Bente, J. Van der Tol, B. Smalbrugge, T. De Vries, E.-J. Geluk, J. Bolk, R. Van Veldhoven, L. Augustin, P. Thijs *et al.*, "An introduction to InP-based generic integration technology," *Semiconductor Science and Technology*, vol. 29, no. 8, p. 083001, 2014.

[7] Q. Cheng, S. Rumley, M. Bahadori, and K. Bergman, "Photonic switching in high performance datacenters," *Optics express*, vol. 26, no. 12, pp. 16 022–16 043, 2018.

[8] A. Ribeiro, A. Ruocco, L. Vanacker, and W. Bogaerts, "Demonstration of a 4× 4-port self-configuring universal linear optical component," in *2016 Progress in Electromagnetic Research Symposium (PIERS)*. IEEE, 2016, pp. 3372–3375.

[9] K. Fotiadis, S. Pitris, M. Moralis-Pegios, C. Mitsolidou, P. De Heyn, J. Van Campenhout, T. Alexoudi, and N. Pleros, "16x16 silicon photonic AWGR

for dense wavelength division multiplexing (DWDM) O-band interconnects," in *Proceedings SPIE, Silicon Photonics XV*, vol. 11285, 2020, p. 112850F.

[10] K.-W. Ang, T.-Y. Liow, Q. Fang, M. B. Yu, F.-F. Ren, S. Y. Zhu, J. Zhang, J. W. Ng, J. F. Song, Y. Z. Xiong, G. Q. Lo, and D.-L. Kwong, "Silicon photonics technologies for monolithic electronic-photonic integrated circuit (EPIC) applications: Current progress and future outlook," in *2009 IEEE International Electron Devices Meeting (IEDM)*. IEEE, 2009, pp. 1–4.

[11] B. G. Lee, A. V. Rylyakov, W. M. J. Green, S. Assefa, C. W. Baks, R. Rimolo-Donadio, D. M. Kuchta, M. H. Khater, T. Barwicz, C. Reinholm, E. Kiewra, S. M. Shank *et al.*, "Monolithic silicon integration of scaled photonic switch fabrics, CMOS logic, and device driver circuits," *Journal of Lightwave Technology*, vol. 32, no. 4, pp. 743–751, 2014.

[12] C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, A. S. Waterman, J. M. Shainline, R. R. Avizienis, S. Lin, B. R. Moss, R. Kumar *et al.*, "Single-chip microprocessor that communicates directly using light," *Nature*, vol. 528, pp. 534–538, 2015.

[13] T.-Y. Liow, K.-W. Ang, Q. Fang, M. B. Yu, F.-F. Ren, S. Y. Zhu, J. Zhang, J. W. Ng, J. F. Song, Y. Z. Xiong, G. Q. Lo, and D.-L. Kwong, "Silicon photonics for monolithic electronic-photonic integrated circuit applications: Opportunities and challenges," in *16th Opto-Electronics and Communications Conference*. IEEE, 2011, pp. 837–839.

[14] L. C. Kimerling, D. Ahn, A. B. Apsel, M. Beals, D. Carothers, Y.-K. Chen, T. Conway, D. M. Gill, M. Grove, C.-Y. Hong, M. Lipson, J. Liu *et al.*, "Electronic-photonic integrated circuits on the CMOS platform," in *Silicon photonics*, vol. 6125. International Society for Optics and Photonics, 2006, p. 612502.

[15] K. Yu, C. Li, H. Li, A. Titriku, A. Shafik, B. Wang, Z. Wang, R. Bai, C. Chen, M. Fiorentino, P. Y. Chiang, and S. Palermo, "A 25 Gb/s hybrid-integrated silicon photonic source-synchronous receiver with microring wavelength stabilization," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 9, pp. 2129–2141, 2016.

[16] X. Zheng, E. Chang, P. Amberg, I. Shubin, J. Lexau, F. Liu, H. Thacker, S. S. Djordjevic, S. Lin, Y. Luo, J. Yao, J.-H. Lee *et al.*, "A high-speed, tunable silicon photonic ring modulator integrated with ultra-efficient active wavelength control," *Optics express*, vol. 22, no. 10, pp. 12 628–12 633, 2014.

[17] "Intel silicon photonics devices catalogue," https://www.intel.com/content/www/us/en/architecture-and-technology/silicon-photonics/silicon-photonics-overview.html.

[18] A. Melloni, R. Costa, G. Cusmai, and F. Morichetti, "The role of index contrast in dielectric optical waveguides," *International Journal of Materials and Product Technology*, vol. 34, no. 4, pp. 421–437, 2009.

[19] J. Komma, C. Schwarz, G. Hofmann, D. Heinert, and R. Nawrodt, "Thermo-optic coefficient of silicon at 1550 nm and cryogenic temperatures," *Applied Physics Letters*, vol. 101, no. 4, p. 041905, 2012.

[20] B. Guha, A. Gondarenko, and M. Lipson, "Minimizing temperature sensitivity of silicon mach-zehnder interferometers," *Optics express*, vol. 18, no. 3, pp. 1879–1887, 2010.

[21] K. Padmaraju and K. Bergman, "Resolving the thermal challenges for silicon microring resonator devices," *Nanophotonics*, vol. 3, no. 4-5, pp. 269–281, 2014.

[22] L. Vivien, J. Osmond, J.-M. Fédéli, D. Marris-Morini, P. Crozat, J.-F. Damlencourt, E. Cassan, Y. Lecunff, and S. Laval, "42 GHz pin germanium photodetector integrated in a silicon-on-insulator waveguide," *Optics express*, vol. 17, no. 8, pp. 6252–6257, 2009.

[23] M. Piels and J. E. Bowers, "Photodetectors for silicon photonic integrated circuits," in *Photodetectors*. Elsevier, 2016, pp. 3–20.

[24] D. Ahn, C.-Y. Hong, J. Liu, W. Giziewicz, M. Beals, L. C. Kimerling, J. Michel, J. Chen, and F. X. Kärtner, "High performance, waveguide integrated Ge photodetectors," *Optics express*, vol. 15, no. 7, pp. 3916–3921, 2007.

[25] L. Chen and M. Lipson, "Ultra-low capacitance and high speed germanium photodetectors on silicon," *Optics Express*, vol. 17, no. 10, pp. 7901–7906, 2009.

[26] S. Assefa, F. Xia, and Y. A. Vlasov, "Reinventing germanium avalanche photodetector for nanophotonic on-chip optical interconnects," *Nature*, vol. 464, no. 7285, pp. 80–84, 2010.

[27] R. Selim, R. Hoofman, M. Khoder, A. Masood, C. Littlejohns, D. Geuzebroek, R. Grootjans, T. Drischel, and K. Torki, "Silicon photonics open access foundry services review for emerging technology," in *Emerging Applications in Silicon Photonics II*, C. G. Littlejohns and M. Sorel, Eds., vol. 11880, International Society for Optics and Photonics. SPIE, 2021, p. 118800C.

[28] D. F. Logan, P. E. Jessop, and A. P. Knights, "Modeling defect enhanced detection at 1550 nm in integrated silicon waveguide photodetectors," *Journal of Lightwave Technology*, vol. 27, no. 7, pp. 930–937, 2009.

[29] H. K. Tsang, C. S. Wong, T. K. Liang, I. E. Day, S. W. Roberts, A. Harpin, J. Drake, and M. Asghari, "Optical dispersion, two-photon absorption and self-phase modulation in silicon waveguides at 1.5 μ m wavelength," *Applied Physics Letters*, vol. 80, no. 3, pp. 416–418, 2002.

[30] T. Baehr-Jones, M. Hochberg, and A. Scherer, "Photodetection in silicon beyond the band edge with surface states," *Optics Express*, vol. 16, no. 3, pp. 1659–1668, Feb 2008.

[31] Y. Li and A. W. Poon, "Active resonance wavelength stabilization for silicon microring resonators with an in-resonator defect-state-absorption-based photodetector," *Opt. Express*, vol. 23, no. 1, pp. 360–372, Jan 2015.

[32] M. W. Geis, S. J. Spector, M. E. Grein, R. T. Schulein, J. U. Yoon, D. M. Lennon, S. Deneault, F. Gan, F. X. Kaertner, and T. M. Lyszczarz, "CMOS-compatible all-Si high-speed waveguide photodiodes with high responsivity in near-infrared communication band," *IEEE Photonics Technology Letters*, vol. 19, no. 3, pp. 152–154, 2007.

[33] J. K. Doylend, P. E. Jessop, and A. P. Knights, "Silicon photonic resonator-enhanced defect-mediated photodiode for sub-bandgap detection," *Optics express*, vol. 18, no. 14, pp. 14 671–14 678, 2010.

[34] Y. Liu, C. W. Chow, W. Y. Cheung, and H. K. Tsang, "In-line channel power monitor based on helium ion implantation in silicon-on-insulator waveguides," *IEEE photonics technology letters*, vol. 18, no. 17, pp. 1882–1884, 2006.

[35] A. P. Knights, J. D. B. Bradley, S. H. Gou, and P. E. Jessop, "Silicon-on-insulator waveguide photodetector with self-ion-implantation-engineered-enhanced infrared response," *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, vol. 24, no. 3, pp. 783–786, 2006.

[36] T. Tanabe, H. Sumikura, H. Taniyama, A. Shinya, and M. Notomi, "All-silicon sub-Gb/s telecom detector with low dark current and high quantum efficiency on chip," *Applied Physics Letters*, vol. 96, no. 10, p. 101103, 2010.

[37] S. M. Sze and K. K. Ng, *Physics of semiconductor devices*. New York, USA: Wiley, 1981.

[38] C. R. Helms and E. H. Poindexter, "The silicon-silicon dioxide system: its microstructure and imperfections," *Reports on Progress in Physics*, vol. 57, no. 8, p. 791, 1994.

[39] V. Schmidt, S. Senz, and U. Gösele, "Influence of the Si/SiO_2 interface on the charge carrier density of Si nanowires," *Applied Physics A*, vol. 86, no. 2, pp. 187–191, 2007.

[40] P. Zhang, E. Tevaarwerk, B.-N. Park, D. E. Savage, G. K. Celler, I. Knezevic, P. G. Evans, M. A. Eriksson, and M. G. Lagally, "Electronic transport in nanometre-scale silicon-on-insulator membranes," *Nature*, vol. 439, no. 7077, pp. 703–706, 2006.

[41] E. H. Nicollian, J. R. Brews, and E. H. Nicollian, *MOS (metal oxide semiconductor) physics and technology*. Wiley New York, 1982, vol. 1987.

[42] H. Chen, X. Luo, and A. W. Poon, "Cavity-enhanced photocurrent generation by 1.55 μ m wavelengths linear absorption in a pin diode embedded silicon microring resonator," *Applied physics letters*, vol. 95, no. 17, p. 171111, 2009.

[43] L. Zhou, H. Zhu, H. Zhang, and J. Chen, "Photoconductive effect on pip micro-heaters integrated in silicon microring resonators," *Optics express*, vol. 22, no. 2, pp. 2141–2149, 2014.

[44] A. Perino, F. Zanetto, M. Petrini, F. Toso, F. Morichetti, A. Melloni, G. Ferrari, and M. Sampietro, "High-sensitivity transparent photoconductors in voltage-controlled silicon waveguides," *Opt. Lett.*, vol. 47, no. 6, pp. 1327–1330, Mar 2022.

[45] H. Jayatilleka, H. Shoman, L. Chrostowski, and S. Shekhar, "Photoconductive heaters enable control of large-scale silicon photonic ring resonator circuits," *Optica*, vol. 6, no. 1, pp. 84–91, 1 2019.

[46] A. Gazman, C. Browning, Z. Zhu, L. Barry, and K. Bergman, "Automated thermal stabilization of cascaded silicon photonic ring resonators for reconfigurable WDM applications," in *2017 European Conference on Optical Communication (ECOC)*, Gothenburg, Sweden, Sep 2017, pp. 1–3.

[47] F. Morichetti, S. Grillanda, M. Carminati, G. Ferrari, M. Sampietro, M. J. Strain, M. Sorel, and A. Melloni, "Non-invasive on-chip light observation by

contactless waveguide conductivity monitoring," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 20, no. 4, pp. 292–301, 2014.

[48] F. Zanetto, E. Guglielmi, F. Toso, R. Gaudiano, F. Caruso, M. Sampietro, and G. Ferrari, "Wide dynamic range multichannel lock-in amplifier for contactless optical sensors with sub-ps resolution," *IEEE Solid-State Circuits Letters*, vol. 3, pp. 246–249, 2020.

[49] V. Grimaldi, F. Zanetto, F. Toso, C. De Vita, and G. Ferrari, "Non-invasive light sensor with enhanced sensitivity for photonic integrated circuits," in *2022 17th Conference on Ph.D Research in Microelectronics and Electronics (PRIME)*, 2022, pp. 285–288.

[50] W. Bogaerts, L. Van Iseghem, P. Edinger, H. Sattari, A. Y. Takabayashi, X. Chen, H. Deng, P. Verheyen, A. Ribeiro, U. Khan, N. Quack, and K. B. Gylfason, "Low-power electro-optic actuators for large-scale programmable photonic circuits," in *2021 Conference on Lasers and Electro-Optics (CLEO)*, 2021, pp. 1–2.

[51] A. Masood, M. Pantouvaki, G. Lepage, P. Verheyen, J. Van Campenhout, P. Absil, D. Van Thourhout, and W. Bogaerts, "Comparison of heater architectures for thermal control of silicon photonic circuits," in *10th International Conference on Group IV Photonics*, 2013, pp. 83–84.

[52] M. Milanizadeh, D. Aguiar, A. Melloni, and F. Morichetti, "Canceling thermal cross-talk effects in photonic integrated circuits," *Journal of Lightwave Technology*, vol. 37, no. 4, pp. 1325–1332, 2019.

[53] H. Sun, Q. Qiao, Q. Guan, and G. Zhou, "Silicon photonic phase shifters and their applications: A review," *Micromachines*, vol. 13, no. 9, 2022.

[54] J. Komma, C. Schwarz, G. Hofmann, D. Heinert, and R. Nawrodt, "Thermo-optic coefficient of silicon at 1550 nm and cryogenic temperatures," *Applied Physics Letters*, vol. 101, no. 4, p. 041905, 2012.

[55] A. W. Elshaari, I. E. Zadeh, K. D. Jöns, and V. Zwiller, "Thermo-optic characterization of silicon nitride resonators for cryogenic photonic circuits," *IEEE Photonics Journal*, vol. 8, no. 3, pp. 1–9, 2016.

[56] S. Wiechmann and J. Müller, "Thermo-optic properties of tio2, ta2o5 and al2o3 thin films for integrated optics on silicon," *Thin Solid Films*, vol. 517, no. 24, pp. 6847–6849, 2009.

[57] F. Zanetto, A. Perino, M. Petrini, F. Toso, M. Milanizadeh, F. Morichetti, A. Melloni, G. Ferrari, and M. Sampietro, "Electrical conductance of silicon photonic waveguides," *Opt. Lett.*, vol. 46, no. 1, pp. 17–20, Jan 2021.

[58] E. Bahaa and A. Saleh, *Electro-Optics*. John Wiley and Sons, Ltd, 1991, ch. 18, pp. 696–736.

[59] P. O. Weigel, M. Savanier, C. T. DeRose, A. T. Pomerene, A. L. Starbuck, A. L. Lentine, V. Stenger, and S. Mookherjea, "Lightwave circuits in lithium niobate through hybrid waveguides with silicon photonics," *Scientific Reports*, vol. 6, no. 1, p. 22301, Mar 2016.

[60] S. Abel, F. Eltes, J. E. Ortmann, A. Messner, P. Castera, T. Wagner, D. Urbonas, A. Rosa, A. M. Gutierrez, D. Tulli, P. Ma, B. Baeuerle, A. Josten, W. Heni, D. Caimi, L. Czornomaz, A. A. Demkov, J. Leuthold, P. Sanchis, and J. Fompeyrine, "Large pockels effect in micro- and nanostructured

barium titanate integrated on silicon," *Nature Materials*, vol. 18, no. 1, pp. 42–47, Jan 2019.

[61] R. A. Soref and B. R. Bennett, "Electrooptical effects in silicon," *IEEE journal of quantum electronics*, vol. 23, no. 1, pp. 123–129, 1987.

[62] P. Dong, S. Liao, D. Feng, H. Liang, D. Zheng, R. Shafiiha, C.-C. Kung, W. Qian, G. Li, X. Zheng, A. V. Krishnamoorthy, and M. Asghari, "Low vpp, ultralow-energy, compact, high-speed silicon electro-optic modulator," *Opt. Express*, vol. 17, no. 25, pp. 22 484–22 490, Dec 2009.

[63] M. Pantouvaki, S. A. Srinivasan, Y. Ban, P. De Heyn, P. Verheyen, G. Lepage, H. Chen, J. De Coster, N. Golshani, S. Balakrishnan, P. Absil, and J. Van Campenhout, "Active components for 50 Gb/s NRZ-OOK optical interconnects in a Silicon Photonics platform," *Journal of Lightwave Technology*, vol. 35, no. 4, pp. 631–638, 2017.

[64] M. C. Wu, O. Solgaard, and J. E. Ford, "Optical mems for lightwave communication," *Journal of Lightwave Technology*, vol. 24, no. 12, pp. 4433–4454, 2006.

[65] P. Edinger, A. Y. Takabayashi, C. Errando-Herranz, U. Khan, H. Sattari, P. Verheyen, W. Bogaerts, N. Quack, and K. B. Gylfason, "Silicon photonic microelectromechanical phase shifters for scalable programmable photonics," *Opt. Lett.*, vol. 46, no. 22, pp. 5671–5674, Nov 2021.

[66] N. Quack, H. Sattari, A. Y. Takabayashi, Y. Zhang, P. Verheyen, W. Bogaerts, P. Edinger, C. Errando-Herranz, and K. B. Gylfason, "Mems-enabled silicon photonic integrated devices and circuits," *IEEE Journal of Quantum Electronics*, vol. 56, no. 1, pp. 1–10, 2020.

[67] F. Toso, M. Milanizadeh, F. Zanetto, V. Grimaldi, A. Melloni, M. Sampietro, F. Morichetti, and G. Ferrari, "Self-stabilized silicon mach-zehnder interferometers by integrated cmos controller," in *2021 IEEE 17th International Conference on Group IV Photonics (GFP)*, 2021, pp. 1–2.

[68] H. Li, G. Balamurugan, T. Kim, M. N. Sakib, R. Kumar, H. Rong, J. Jaussi, and B. Casper, "A 3D-integrated silicon photonic microring-based 112-Gb/s PAM-4 transmitter with nonlinear equalization and thermal control," *IEEE Journal of Solid-State Circuits*, vol. 56, no. 1, pp. 19–29, 2021.

[69] M. Carminati, A. Annoni, F. Morichetti, E. Guglielmi, G. Ferrari, D. Aguiar, A. Melloni, and M. Sampietro, "Design guidelines for contactless integrated photonic probes in dense photonic circuits," *Journal of Lightwave Technology*, vol. 35, no. 14, pp. 3042–3049, 2017.

[70] P. Dumais, D. J. Goodwill, D. Celo, J. Jiang, C. Zhang, F. Zhao, X. Tu, C. Zhang, S. Yan, J. He, M. Li, W. Liu, Y. Wei, D. Geng, H. Mehrvar, and E. Bernier, "Silicon photonic switch subsystem with 900 monolithically integrated calibration photodiodes and 64-fiber package," *Journal of Lightwave Technology*, vol. 36, no. 2, pp. 233–238, 2018.

[71] M. Milanizadeh, S. Ahmadi, M. Petrini, D. Aguiar, R. Mazzanti, F. Zanetto, E. Guglielmi, M. Sampietro, F. Morichetti, and A. Melloni, "Control and calibration recipes for photonic integrated circuits," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 26, no. 5, pp. 1–10, Sep. 2020.

[72] X. Zhu, M. Wang, K. Bergman, H. L. R. Lira, L.-W. Luo, and M. Lipson, "FPGA controlled microring based tunable add-drop filter," in *2013 Optical Interconnects Conference*, Santa Fe, USA, May 2013, pp. 102–103.

[73] W. A. Zortman, A. L. Lentine, D. C. Trotter, and M. R. Watts, "Bit-error-rate monitoring for active wavelength control of resonant modulators," *IEEE Micro*, vol. 33, no. 1, pp. 42–52, Jan 2013.

[74] F. Zanetto, V. Grimaldi, M. Moralis-Pegios, S. Pitris, K. Fotiadis, T. Alexoudi, E. Guglielmi, D. Aguiar, P. De Heyn, Y. Ban, J. Van Campenhout, N. Pleros, G. Ferrari, M. Sampietro, and A. Melloni, "WDM-based silicon photonic multi-socket interconnect architecture with automated wavelength and thermal drift compensation," *Journal of Lightwave Technology*, vol. 38, no. 21, pp. 6000–6006, 2020.

[75] A. Annoni, E. Guglielmi, M. Carminati *et al.*, "Automated routing and control of silicon photonic switch fabrics," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 22, no. 6, pp. 169–176, Nov 2016.

[76] H. Li, G. Balamurugan, T. Kim, M. N. Sakib, R. Kumar, H. Rong, J. Jaussi, and B. Casper, "A 3-d-integrated silicon photonic microring-based 112-gb/s pam-4 transmitter with nonlinear equalization and thermal control," *IEEE Journal of Solid-State Circuits*, vol. 56, no. 1, pp. 19–29, 2021.

[77] K. Padmaraju, J. Chan, L. Chen, M. Lipson, and K. Bergman, "Thermal stabilization of a microring modulator using feedback control," *Optics Express*, vol. 20, no. 27, pp. 27 999–28 008, 2012.

[78] K. Padmaraju, D. F. Logan, T. Shiraishi, J. J. Ackert, A. P. Knights, and K. Bergman, "Wavelength locking and thermally stabilizing microring resonators using dithering signals," *Journal of Lightwave Technology*, vol. 32, no. 3, pp. 505–512, Feb 2014.

[79] F. Zanetto, V. Grimaldi, F. Toso, E. Guglielmi, M. Milanizadeh, D. Aguiar, M. Moralis-Pegios, S. Pitris, T. Alexoudi, F. Morichetti, A. Melloni, G. Ferrari, and M. Sampietro, "Dithering-based real-time control of cascaded silicon photonic devices by means of non-invasive detectors," *IET Optoelectronics*, vol. 15, no. 2, pp. 111–120, 2021.

[80] Y. Zhang, Y. Li, S. Feng, and A. W. Poon, "Towards adaptively tuned silicon microring resonators for optical networks-on-chip applications," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 20, no. 4, pp. 136–149, July 2014.

[81] E. Timurdogan, A. Biberman, D. C. Trotter, C. Sun, M. Moresco, V. Stojanović, and M. R. Watts, "Automated wavelength recovery for microring resonators," in *2012 Conference on Lasers and Electro-Optics (CLEO)*, San Jose, USA, May 2012, pp. 1–2.

[82] C. Sun, M. Wade, M. Georgas, S. Lin, L. Alloatti, B. Moss, R. Kumar, A. H. Atabaki, F. Pavanello, J. M. Shainline, J. S. Orcutt, R. J. Ram *et al.*, "A 45 nm CMOS-SOI monolithic photonics platform with bit-statistics-based resonant microring thermal tuning," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 4, pp. 893–907, 2016.

[83] W. Bogaerts, D. Pérez, J. Capmany, D. A. B. Miller, J. Poon, D. Englund, F. Morichetti, and A. Melloni, "Programmable photonic circuits," *Nature*, no. 586, pp. 207–216, Oct 2020.

[84] J. Carolan, C. Harrold, C. Sparrow, E. Martín-López, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M. Oguma, M. Itoh, G. D. Marshall, M. G. Thompson, J. C. F. Matthews, T. Hashimoto, J. L. O'Brien, and A. Laing, "Universal linear optics," *Science*, vol. 349, no. 6249, pp. 711–716, 2015.

[85] D. A. B. Miller, "Self-configuring universal linear optical component," *Photon. Res.*, vol. 1, no. 1, pp. 1–15, Jun 2013.

[86] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walmsley, "Optimal design for universal multiport interferometers," *Optica*, vol. 3, no. 12, pp. 1460–1465, Dec 2016.

[87] D. Pérez, I. Gasulla, L. Crudgington, D. J. Thomson, A. Z. Khokhar, K. Li, W. Cao, G. Z. Mashanovich, and J. Capmany, "Multipurpose silicon photonics signal processor core," *Nature Communications*, vol. 8, no. 1, p. 636, Sep 2017.

[88] J. Carolan, C. Harrold, C. Sparrow, E. Martín-López, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M. Oguma, M. Itoh *et al.*, "Universal linear optics," *Science*, vol. 349, no. 6249, pp. 711–716, 2015.

[89] N. C. Harris, D. Bunandar, M. Pant, G. R. Steinbrecher, J. Mower, M. Prabhu, T. Baehr-Jones, M. Hochberg, and D. Englund, "Large-scale quantum photonic circuits in silicon," *Nanophotonics*, vol. 5, no. 3, pp. 456–468, 2016.

[90] J. Notaros, J. Mower, M. Heuck, C. Lupo, N. C. Harris, G. R. Steinbrecher, D. Bunandar, T. Baehr-Jones, M. Hochberg, S. Lloyd *et al.*, "Programmable dispersion on a photonic integrated circuit for classical and quantum applications," *Optics express*, vol. 25, no. 18, pp. 21 275–21 285, 2017.

[91] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund *et al.*, "Deep learning with coherent nanophotonic circuits," *Nature Photonics*, vol. 11, no. 7, pp. 441–446, 2017.

[92] N. C. Harris, J. Carolan, D. Bunandar, M. Prabhu, M. Hochberg, T. Baehr-Jones, M. L. Fanto, A. M. Smith, C. C. Tison, P. M. Alsing *et al.*, "Linear programmable nanophotonic processors," *Optica*, vol. 5, no. 12, pp. 1623–1631, 2018.

[93] S. Pai, Z. Sun, T. W. Hughes, T. Park, B. Bartlett, I. A. D. Williamson, M. Minkov, M. Milanizadeh, N. Abebe, F. Morichetti, A. Melloni, S. Fan, O. Solgaard, and D. A. B. Miller, "Experimentally realized in situ backpropagation for deep learning in nanophotonic neural networks," 2022.

[94] S. Pai, T. Park, M. Ball, B. Penkovsky, M. Milanizadeh, M. Dubrovsky, N. Abebe, F. Morichetti, A. Melloni, S. Fan, O. Solgaard, and D. A. B. Miller, "Experimental evaluation of digitally-verifiable photonic computing for blockchain and cryptocurrency," 2022.

[95] J. Bütow, J. S. Eismann, M. Milanizadeh, F. Morichetti, A. Melloni, D. A. B. Miller, and P. Banzer, "Spatially resolving amplitude and phase of light with a reconfigurable photonic integrated circuit," *Optica*, vol. 9, no. 8, p. 939, aug 2022.

[96] D. A. B. Miller, "Waves, modes, communications, and optics: a tutorial," *Adv. Opt. Photon.*, vol. 11, no. 3, pp. 679–825, Sep 2019.

[97] Y. Ren, Z. Wang, G. Xie, L. Li, Y. Cao, C. Liu, P. Liao, Y. Yan, N. Ahmed, Z. Zhao, A. Willner, N. Ashrafi, S. Ashrafi, R. D. Linquist, R. Bock, M. Tur, A. F. Molisch, and A. E. Willner, "Free-space optical communications using orbital-angular-momentum multiplexing combined with mimo-based spatial multiplexing," *Opt. Lett.*, vol. 40, no. 18, pp. 4210–4213, Sep 2015.

[98] Y. Ren, H. Huang, G. Xie, N. Ahmed, Y. Yan, B. I. Erkmen, N. Chandrasekaran, M. P. J. Lavery, N. K. Steinhoff, M. Tur, S. Dolinar, M. Neifeld, M. J. Padgett, R. W. Boyd, J. H. Shapiro, and A. E. Willner, "Atmospheric turbulence effects on the performance of a free space optical link employing orbital angular momentum multiplexing," *Opt. Lett.*, vol. 38, no. 20, pp. 4062–4065, Oct 2013.

[99] J. Bertolotti, E. G. van Putten, C. Blum, A. Lagendijk, W. L. Vos, and A. P. Mosk, "Non-invasive imaging through opaque scattering layers," *Nature*, vol. 491, no. 7423, pp. 232–234, Nov 2012.

[100] A. E. Willner, H. Huang, Y. Yan, Y. Ren, N. Ahmed, G. Xie, C. Bao, L. Li, Y. Cao, Z. Zhao, J. Wang, M. P. J. Lavery, M. Tur, S. Ramachandran, A. F. Molisch, N. Ashrafi, and S. Ashrafi, "Optical communications using orbital angular momentum beams," *Adv. Opt. Photon.*, vol. 7, no. 1, pp. 66–106, Mar 2015.

[101] M. P. J. Lavery, F. C. Speirits, S. M. Barnett, and M. J. Padgett, "Detection of a spinning object using light’s orbital angular momentum," *Science*, vol. 341, no. 6145, pp. 537–540, 2013.

[102] S. Grillanda, F. Morichetti, N. Peserico, P. Ciccarella, A. Annoni, M. Carminati, and A. Melloni, "Non-invasive monitoring of mode-division multiplexed channels on a silicon photonic chip," *Journal of Lightwave Technology*, vol. 33, no. 6, pp. 1197–1201, 2015.

[103] M. Milanizadeh, S. SeyedinNavadeh, F. Zanetto, V. Grimaldi, C. De Vita, C. Klitis, M. Sorel, G. Ferrari, D. A. B. Miller, A. Melloni, and F. Morichetti, "Separating arbitrary free-space beams with an integrated photonic processor," *Light: Science & Applications*, vol. 11, no. 1, p. 197, Jul 2022.

[104] M. Pantouvaki, P. Verheyen, G. Lepage, J. De Coster, H. Yu, P. De Heyn, P. Absil, and J. Van Campenhout, "20gb/s silicon ring modulator cointegrated with a ge monitor photodetector," in *39th European Conference and Exhibition on Optical Communication (ECOC 2013)*, 2013, pp. 1–3.

[105] E. Guglielmi, F. Toso, F. Zanetto, G. Sciortino, A. Mesri, M. Sampietro, and G. Ferrari, "High-value tunable pseudo-resistors design," *IEEE Journal of Solid-State Circuits*, vol. 55, no. 8, pp. 2094–2105, 2020.

[106] J. Sharma, Z. Xuan, H. Li, T. Kim, R. Kumar, M. N. Sakib, C.-M. Hsu, C. Ma, H. Rong, G. Balamurugan, and J. Jaussi, "Silicon photonic microring-based 4 x 112 Gb/s WDM transmitter with photocurrent-based thermal control in 28-nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 57, no. 4, pp. 1187–1198, 2021.

[107] W. Freude, R. Schmogrow, B. Nebendahl, M. Winter, A. Josten, D. Hillerkuss, S. Koenig, J. Meyer, M. Dreschmann, M. Huebner, C. Koos,

J. Becker, and J. Leuthold, "Quality metrics for optical signals: eye diagram, Q-factor, OSNR, EVM and BER," in *2012 14th International Conference on Transparent Optical Networks (ICTON)*, 2012, pp. 1–4.

[108] A. Udalcovs, T. Salgals, L. Zhang, X. Pang, A. Djupsjöbacka, S. Spolitis, V. Bobrovs, S. Popov, and O. Ozolins, "Optical power budget of 25+ Gbps IM/DD PON with digital signal post-equalization," *Applied Sciences*, vol. 10, no. 17, p. 6106, 2020.

[109] Z. Wang, D. Paez, A. I. A. El-Rahman, P. Wang, L. Dow, J. C. Cartledge, and A. P. Knights, "Resonance control of a silicon micro-ring resonator modulator under high-speed operation using the intrinsic defect-mediated photocurrent," *Opt. Express*, vol. 25, no. 20, pp. 24 827–24 836, Oct 2017.

[110] S. Agarwal, M. Ingels, M. Pantouvaki, M. Steyaert, P. Absil, and J. Van Campenhout, "Wavelength locking of a Si ring modulator using an integrated drop-port OMA monitoring circuit," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 10, pp. 2328–2344, 2016.

[111] A. Singh, J. Ong, A. Agarwal, G. Anderson, A. Armistead, R. Bannon, S. Boving, G. Desai, B. Felderman, P. Germano, A. Kanagala, J. Provost *et al.*, "Jupiter rising: a decade of clos topologies and centralized control in Google's datacenter network," *ACM SIGCOMM computer communication review*, vol. 45, no. 4, pp. 183–197, 2015.

[112] S. Pitris, M. Moralis-Pegios, T. Alexoudi, J. Lambrecht, X. Yin, J. Bauwelinck, Y. Ban, P. De Heyn, M. Pantouvaki, J. Van Campenhout, R. Broeke, and N. Pleros, "A 40 Gb/s chip-to-chip interconnect for 8-socket direct connectivity using integrated photonics," *IEEE Photonics Journal*, vol. 10, no. 5, pp. 1–8, Oct 2018.

[113] J. Wang, Z. Sheng, L. Li, A. Pang, A. Wu, W. Li, X. Wang, S. Zou, M. Qi, and F. Gan, "Low-loss and low-crosstalk 8x8 silicon nanowire AWG routers fabricated with CMOS technology," *Optics Express*, vol. 22, no. 8, pp. 9395–9403, Apr 2014.

[114] N. A. Idris and H. Tsuda, "6.4-THz-spacing, 10-channel cyclic arrayed waveguide grating for T- and O-band coarse WDM," *IEICE Electronics Express*, vol. 13, no. 7, Apr 2016.

[115] M. Zirngibl, C. Dragone, and C. H. Joyner, "Demonstration of a 15x15 arrayed waveguide multiplexer on InP," *IEEE Photonics Technology Letters*, vol. 4, no. 11, pp. 1250–1253, Nov 1992.

[116] S. Takenobu, Y. Kuwana, K. Takayama, Y. Sakane, M. Ono, H. Sato, N. Keil, W. Brinker, H. Yao, C. Zawadzki, Y. Morizawa, and N. Grote, "All-polymer 8x8 AWG wavelength router using ultra low loss polymer optical waveguide material (CYTOPTM)," in *Optical Fiber Communication Conference/National Fiber Optic Engineers Conference*, 2008, p. JWA32.

[117] M. Moralis-Pegios, S. Pitris, T. Alexoudi, N. Terzenidis, H. Ramon, J. Lambrecht, J. Bauwelinck, X. Yin, Y. Ban, P. De Heyn, J. Van Campenhout, T. Lamprecht, A. Lehnman, and N. Pleros, "4-channel 200 gb/s wdm o-band silicon photonic transceiver sub-assembly," *Optics Express*, vol. 28, no. 4, pp. 5706–5714, 2020.

[118] A. Melloni, A. Annoni, S. Grillanda, D. Melati, and F. Morichetti, "Control, configuration and stabilization of photonic integrated circuits," in *International Conference on Photonics in Switching (PS)*, July 2016, pp. 1–3.

[119] P. De Heyn, J. De Coster, P. Verheyen, G. Lepage, M. Pantouvaki, P. Absil, W. Bogaerts, J. Van Campenhout, and D. Van Thourhout, "Fabrication-tolerant four-channel wavelength-division-multiplexing filter based on collectively tuned Si microrings," *Journal of Lightwave Technology*, vol. 31, no. 16, pp. 2785–2792, Aug 2013.

[120] P. Ciccarella, M. Carminati, G. Ferrari, D. Bianchi, S. Grillanda, F. Morichetti, A. Melloni, and M. Sampietro, "Impedance-sensing CMOS chip for noninvasive light detection in integrated photonics," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 63, no. 10, pp. 929–933, 2016.

[121] X. Mu, S. Wu, L. Cheng, and F. H. Fu, "Edge couplers in silicon photonic integrated circuits: A review," *Applied Sciences*, vol. 10, no. 4, p. 1538, 2020.

[122] *A Technical Tutorial on Digital Signal Synthesis*, https://www.analog.com/media/cn/training-seminars/tutorials/450968421DDS_Tutorial_rev12-2-99.pdf, Analog Devices, 1999.

[123] *DDS Compiler v6.0*, https://docs.xilinx.com/v/u/en-US/pg141-dds-compiler, Xilinx, 2021.

[124] P. Simpson, *Timing Closure*. New York, NY: Springer New York, 2010, ch. 8, pp. 107–132.

[125] A. V. Oppenheim and R. W. Schafer, *Structures for discrete-time systems*. Prentice Hall, 1999, ch. 6, p. 354–356.

[126] R. G. Lyons, *Alternate IIR Filter Structures*. Pearson Education International, 2013, ch. 6.

[127] "Opalkelly api reference guide," https://library.opalkelly.com/library/FrontPanelAPI/classokCFrontPanel.html#a78d9da7ddea0a75c9ce408a64b14e51a.

[128] R. G. Lyons, *Digital data formats and their effects*. Pearson Education International, 2013, ch. 12.

[129] *User guide for DSP48E1*, https://docs.xilinx.com/v/u/en-US/ug479_7Series_DSP48E1, Xilinx, 2018.

[130] E. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, no. 2, pp. 155–162, 1981.

[131] R. G. Lyons, *Specialized Digital Networks and Filters*. Pearson Education International, 2013, ch. 7.

[132] W. L. Barnes, A. Dereux, and T. W. Ebbesen, "Surface plasmon subwavelength optics," *Nature*, vol. 424, no. 6950, pp. 824–830, Aug 2003.

[133] Y. Fang and M. Sun, "Nanoplasmonic waveguides: towards applications in integrated nanophotonic circuits," *Light: Science & Applications*, vol. 4, no. 6, pp. e294–e294, Jun 2015.

[134] A. Melikyan, L. Alloatti, A. Muslija, D. Hillerkuss, P. C. Schindler, J. Li, R. Palmer, D. Korn, S. Muehlbrandt, D. Van Thourhout, B. Chen, R. Dinu, M. Sommer, C. Koos, M. Kohl, W. Freude, and J. Leuthold, "High-speed plasmonic phase modulators," *Nature Photonics*, vol. 8, no. 3, pp. 229–233, Mar 2014.

[135] L. Alloatti, R. Palmer, S. Diebold, K. P. Pahl, B. Chen, R. Dinu, M. Fournier, J.-M. Fedeli, T. Zwick, W. Freude, C. Koos, and J. Leuthold, "100 ghz silicon-organic hybrid modulator," *Light: Science and Applications*, vol. 3, no. 5, pp. e173–e173, May 2014.

[136] P. Debackere, S. Scheerlinck, P. Bienstman, and R. Baets, "Surface plasmon interferometer in silicon-on-insulator: Novel concept for an integrated biosensor," *Optics express*, vol. 14, pp. 7063–72, 09 2006.

[137] S. Mirzanejhad, A. Ghadi, and M. E. Daraei, "Numerical study of nanoscale biosensor based on surface plasmon polariton propagation in mach-zehnder interferometer structure," *Physica B: Condensed Matter*, vol. 557, pp. 141–146, 2019.

[138] A. Andrieux, M.-M. Mennemanteuil, N. Geoffroy, M. Emo, L. Markey, and K. Hammani, "Optimized icpcvd-based tio2 for photonics," *Materials*, vol. 15, 2022.

[139] T. Stefaniuk, P. Wróbel, P. Trautman, and T. Szoplik, "Ultrasmooth metal nanolayers for plasmonic applications: surface roughness and specific resistivity," *Appl. Opt.*, vol. 53, no. 10, pp. B237–B241, Apr 2014.

[140] F. Morichetti, A. Canciamilla, C. Ferrari, M. Torregiani, A. Melloni, and M. Martinelli, "Roughness induced backscattering in optical silicon waveguides," *Phys. Rev. Lett.*, vol. 104, p. 033902, Jan 2010.

[141] A. Oliva, J. Lugo, R. Gurubel-Gonzalez, R. Centeno, J. Corona, and F. Avilés, "Temperature coefficient of resistance and thermal expansion coefficient of 10-nm thick gold films," *Thin Solid Films*, vol. 623, pp. 84–89, 2017.

[142] D. Thomson, A. Zilkie, J. E. Bowers, T. Komljenovic, G. T. Reed, L. Vivien, D. Marris-Morini, E. Cassan, L. Virot, J.-M. Fédéli, J.-M. Hartmann, J. H. Schmid *et al.*, "Roadmap on silicon photonics," *Journal of Optics*, vol. 18, no. 7, p. 073003, 2016.