Electronics-photonics co-design for robust control of optical devices in dense integrated photonic circuits

F. Toso, F. Zanetto, V. Grimaldi, A. Perino, E. Guglielmi, F. Morichetti, A. Melloni, G. Ferrari, M. Sampietro

Politecnico di Milano, Department of Electronics, Information and Bioengineering, Piazza Leonardo da Vinci 32, 20133 Milano, Italy

The increasing requirements in terms of bandwidth and power efficiency demanded by telecommunication, automotive and datacenter applications are pushing copper-based interconnections close to their intrinsic limits [1]. On the same line, the emerging trends on new rack-scale multi-socket computational units are putting pressure on chip-to-chip interconnections, calling for high-speed and low-latency performance at a reduced energy and cost [2]. These technological requirements are making copper-based interconnects the bottleneck of high-performance systems, suggesting the use of point-to-point optical connectivity not only for long range communications but also at short or ultra-short distances [3]. With respect to other solid-state solutions, Silicon Photonics (SiP) seems to be the ideal candidate to answer to these needs, as it shares the same fabrication technologies of the microelectronic industry and is therefore expected to play a significant role in the field.

Similarly to MEMS technologies, which leveraged the mechanical properties of the silicon crystal to develop a new realm of devices and applications that are nowadays supporting the electronic industry with a large fraction of its revenues [4-6], Silicon Photonics is called to exploit the optical properties of silicon to create a new class of miniaturized systems and devices that bring innovative functionalities. The silicon crystal is in fact transparent to near-infrared radiation, in particular to 1550 nm, which is the typical working wavelength of long haul optical fiber links, and to 1310 nm, which is of increasing interest for short link interconnects. In addition, light confinement into a lithographed waveguide (WG) can be obtained by using silicon dioxide (or silicon nitride) as cladding material, thanks to the large difference of refractive index with respect to silicon ($n_{Si} = 3.45$, $n_{SiO2} = 1.45$). These characteristics have already made possible the demonstration of high-complexity architectures integrating different kind of interferometric devices, like Mach Zehnder interferometers (MZI), ring resonators, arrayed waveguide grating router (AWGR) [7-9], enabling advanced on-chip manipulation of light beams.

Even though photonic platforms have already demonstrated maturity for integrating thousands of devices in a compact footprint [10], the widespread use of complex silicon photonic architectures in real applications is still limited. The reason for this delay can be found in the intrinsic nature of photonic devices, that usually rely on interferometry and are consequently very sensitive to fabrication tolerances, temperature variations and mutual crosstalk. As an example, the resonance frequency of a ring-based filter is observed to shift of 10 GHz due to a 1°C temperature variation [11], making it very difficult to reliably operate complex photonic architectures in harsh environments like datacenters without a real-time monitoring of their working conditions. Local light detection, possibly with CMOS compatible sensors, and active control of each photonic device thus emerged as strong requirements to implement feedback stabilization and fully exploit the potential of photonic integration.

This paper reviews the recent advances in the design of non-invasive light sensors, that locally probe the optical power without introducing any perturbation of the photonic functionality, and describes the corresponding readout electronics. The limited area occupation and the negligible optical losses allow to place many non-invasive detectors in a complex photonic architecture to monitor all the points of interest, so multichannel CMOS ASICs are needed to perform an effective low-noise readout in a compact form factor. In the following, the requirements of the application are presented and several circuital topologies are discussed to address them. The information about the status of the photonic devices acquired with the ASICs is then used by an analog or digital feedback controller to drive the actuators on the photonic chip (heaters in the most common case) and stabilize a given working point. By carefully co-designing the electronics and the photonics (Fig. 1), this interplay can be optimized to perform advanced operations like reconfigurable routing schemes [12], wavelength-division multiplexing (WDM) [13] or optical mode manipulation and unscrambling in mode division multiplexing (MDM) systems [14]. These achievements potentially allow a cost and size reduction of optical components and an improvement of the capacity of optical networks. The paper ends with a technological perspective on how monolithic integration of photonics and electronics on the same chip can be obtained and can be a viable option to further scale up the complexity of photonic systems.

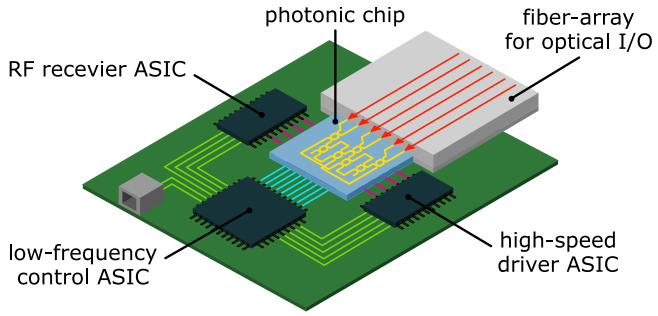

Fig. 1. Example of co-designed photonic and electronic chips, mounted together to obtain a compact high-performance electro-optical assembly.

I – THE CHALLENGE OF TRANSPARENT DETECTION

In order to reliably operate a complex photonic system, the status of each integrated device has to be inspected in real-time without altering its functionality. As long as the number of components in a circuit is limited to a few units, light monitoring can be achieved by tapping a small fraction of the optical power from the waveguide towards a photodetector. The use of integrated germanium photodiodes is the most common solution in this case [15,16], even though the addition of this material to the processing stack is not a trivial operation and it increases the overall cost of the die. When the number of devices integrated in the same chip increases to hundreds or thousands, this approach becomes quickly impossible, as the large quantity of probing points causes an unacceptable light attenuation and/or perturbation [17].

For these reasons, an increasing research effort is being carried out to develop and use in-line non-invasive photodetectors, that exploit the waveguide itself as a power monitor and promise to allow the control of large-scale architectures. Even though silicon is nominally transparent to near-infrared wavelengths, photocarrier generation has been demonstrated to happen in silicon waveguides because of two photon absorption (TPA) [18,19] and sub-bandgap mechanisms as surface-state absorption (SSA) [20,21], that are responsible of the intrinsic propagation losses observed in photonic waveguides. As the number of free carriers generated by these mechanisms is small, non-invasive detectors based on such effects generally have a lower sensitivity with respect to germanium photodiodes, but, as they measure the full optical power in the waveguide and not just a small fraction of it, their use is not so disadvantageous with respect to the standard approach. Different transparent sensors have been demonstrated, as in-line pin diodes [22] or photoconductors [23,24], yet using either extra ion-implantation or moderate waveguide doping to increase the sensitivity of the readout. This requires additional non-standard processing steps or causes non-negligible losses in the light propagation, limiting the possibility to increase the number of monitoring points.

A possible fully transparent in-line alternative to monitor the state of complex photonic circuits is the ContactLess Integrated Photonic Probe (CLIPP) [25]. Like a photoconductor, this detector measures the changes of the core conductivity induced by free carriers photogeneration. However, to perform a truly non-invasive measurement, the waveguide is not doped or processed and the access to its electrical proprieties is obtained in a capacitive way, by placing a couple of metal electrodes on top of the cladding at a

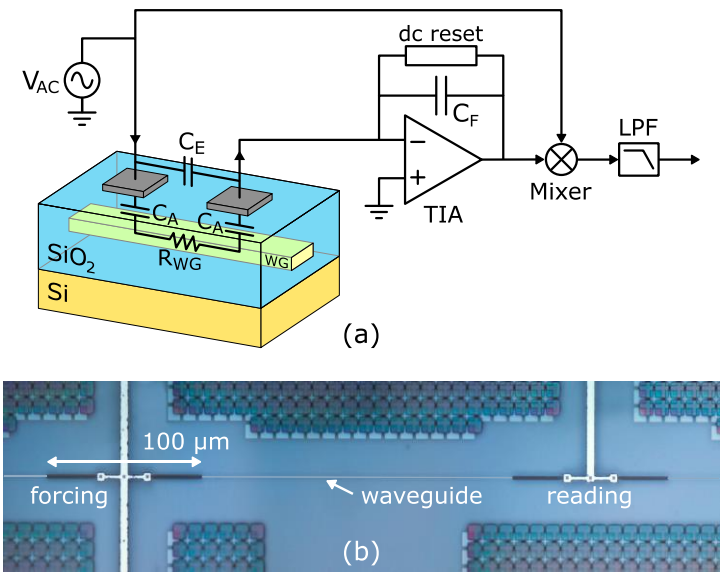

Fig. 2. (a) 3D representation of a CLIPP detector and its readout scheme. (b) Microphotograph of a 400 µm long device.

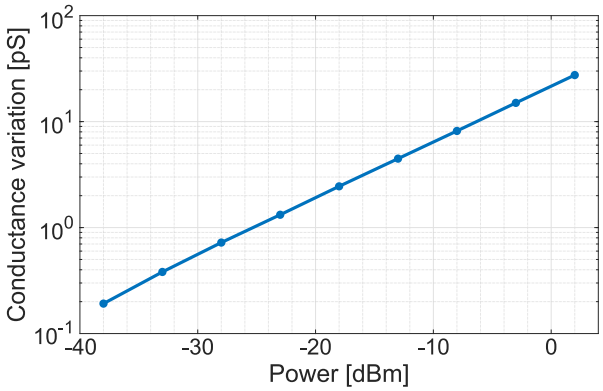

Fig. 3. Sensitivity curve of a 400 µm long CLIPP detector. A very low-noise electronic readout is required to detect light variations below -30 dBm, that cause changes of conductance smaller than 1 pS.

distance of around 700nm from the core. Fig. 2 shows a 3D representation of the CLIPP sensor and its simplified equivalent electrical model, neglecting the effect of the substrate [26]. In order to bypass the access capacitance C_A and probe the electrical resistance of the waveguide R_{WG}, one pad of the sensor is stimulated at a frequency $f_{CLIPP} > 1/\pi C_A R_{WG}$, usually between 100 kHz and 10MHz depending on the electrodes geometry, while the other is connected to the virtual ground of a TransImpedance Amplifier (TIA) to collect the current coming from the sensor and convert it into a voltage. A lock-in processing, obtained by demodulating the output of the TIA at the same frequency of the sensor stimulation, is thus a convenient readout scheme, as it allows to simultaneously reconstruct the complex admittance of the sensor while minimizing the electronic noise. The readout noise is in fact of particular relevance in this application. The conductance of the waveguide is in the order of 1 nS [25], so a resolution better than 1 pS is targeted to monitor the optical power in photonic devices with an accuracy below −30 dBm (Fig. 3). In order to sense such small admittance variations, CMOS integration of the sensor electronic readout is required. Indeed, an integrated solution allows to directly wire-bond the photonic chip to the electronic one and reduce the stray capacitances, with a beneficial effect on the readout signal-to-noise ratio (SNR). In addition, the small size and multichannel operation offered by the CMOS integration enables the control of many photonic devices in a compact footprint.

II – INTEGRATED LOCK-IN READOUT SYSTEM

2.1 Readout front-end architecture

Several requirements should be satisfied to read CLIPP sensors in an effective way. In addition to the very low-noise operation needed to target optical variations below -30 dBm, a sufficiently wide input bandwidth, roughly between 100 kHz and 10 MHz, should be guaranteed to adapt to different geometries of the sensors and fabrication technologies. Finally, a highly linear response over a rail-to-rail output is required. In fact, the sensor feed-through capacitance C_E usually injects in the TIA a wide spurious input signal superimposed to the much smaller waveguide current. Although not carrying any useful information, this signal would cause a significant error in the measurement if the linear range of the amplifier were not wide enough.

Two different multichannel ASICs, both based on the channel structure shown in Fig.4, were specifically designed for the optimized readout of CLIPP sensors [27,28]. Both designs take advantage of a capacitive-feedback architecture for the front-end TIA, which allows to satisfy all of the above requirements. The overall gain of the readout, which depends on the value of the feedback capacitance C_F and on the amplitude of the signal used to stimulate the CLIPP sensor, is ultimately limited by the value of the feedthrough capacitance C_E. In [28], a capacitance compensation system was proposed as a solution. A programmable capacitor, driven in counter-phase with respect to the stimulus of the sensor, is connected in parallel to the sensor, with the objective of sinking the spurious current injected by C_E (Fig. 4). The value of the compensation capacitor can be tuned by a 4-bit digital word to match the value of the C_E capacitance with a resolution of 10 fF, allowing to

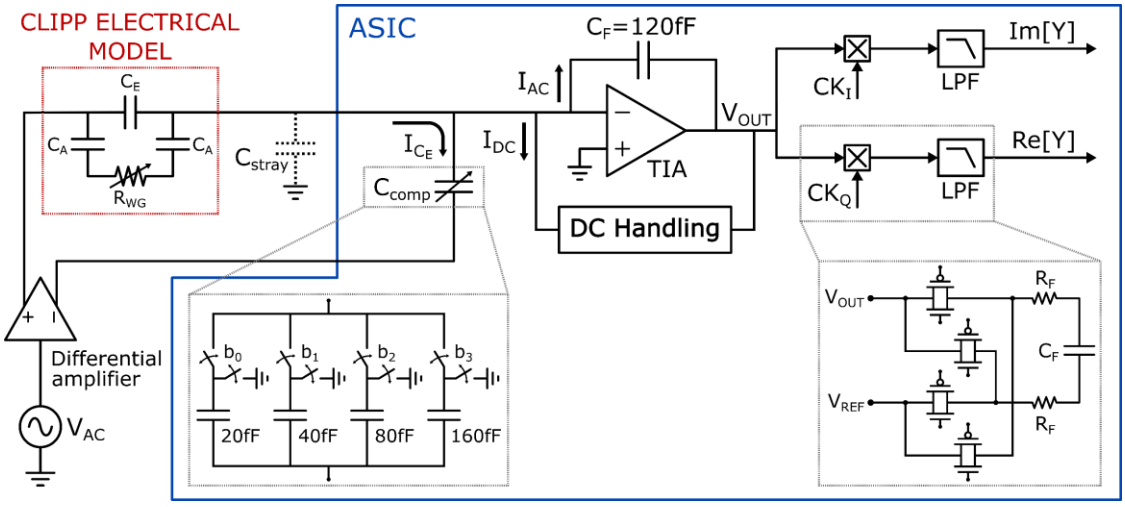

Fig. 4. Schematic structure of a CLIPP readout channel.

compensate feedthrough capacitances up to 300 fF and to use stimulation signals as high as 10 V with a C_F of 120 fF.

The readout chain is completed by a couple of double-balanced passive mixers based on transmission gates, that provide in-phase and quadrature on-chip demodulation of the TIA output and enable the reconstruction of the complex admittance of the sensor. Passive mixers have a significantly lower $1/f$ noise compared to active analog multipliers [29] and the degradation of the SNR caused by the folding of high-frequency noise is mitigated and limited to only 10% by the presence of the capacitive-feedback TIA, which produces an almost white output noise spectrum. The double-balanced architecture helps instead in improving linearity and clock feedthrough rejection.

The ASIC proposed in [27] features 4 independent channels, each equipped with an 8-to-1 analog input multiplexer, enabling the read-out of up to 32 different CLIPP sensors. This solution is optimized for applications where a high number of optical devices has to be monitored and stabilized against slow drifts. The other solution [28] is instead equipped with 11 independent channels, without input multiplexing, and it is therefore more suitable for applications where a smaller number of devices need to be constantly controlled and tuned in parallel.

2.2 Leakage current discharge subcircuits

Capacitive-feedback TIAs require an additional DC feedback network to set the bias of the amplifier and to discharge any currents coming from the sensor or from the electrostatic discharge protection circuits, that would otherwise cause the saturation of the stage. This problem is of particular relevance in photonic applications, where the electronic front-end is operated in close proximity to the photonic device and light leaking from the photonic circuit may produce photo-generated leakage currents at the input node of the TIA up to the nA range. The DC-handling network of choice should therefore be able to handle currents ranging from few pA to up to several nA, without impacting the response of the circuit in the frequency range of the AC signal of the sensor.

In [27], a solution based on a passive continuous-time pseudo-resistor topology was implemented. The advantages of this kind of architecture are its extreme design simplicity and its reduced footprint and power consumption. A major limitation of these implementations is the extremely high value of the equivalent resistance of the pseudo-resistor network, that is in the order of several GΩ to several TΩ. Here, a leakage current of few hundreds of pA is sufficient to cause the saturation of the output of the TIA. A solution to this problem was recently proposed in [30], where the passive pseudo-resistor network is replaced by a tunable pseudo-resistor topology (Fig. 5a). The proposed bias network is built around a symmetrical pseudo-resistor cell made of two pMOS transistor connected in parallel. The linearity of the cell is extended by mean of a resistive voltage divider and the tunability of the resistance value is obtained thanks to a pair of matched floating voltage generators used to regulate the gate-source voltage of the two transistors. The resulting structure has an equivalent resistance that can be tuned between 20MΩ and 20GΩ and that can be externally regulated depending on the amount of leakage current that needs to be handled (Fig. 6).

A disadvantage common to both passive and active pseudo-resistor networks is the trade-off between maximum leakage current that can be handled (i.e. causing a reasonable output offset) and resulting bandwidth of the circuit (Fig. 7). In fact, the frequency of the pole of the feedback network is inversely proportional to the equivalent resistance of the pseudo-resistor ($1/2\pi C_F R_{EQ}$), while the output offset for a given leakage current is directly proportional to it

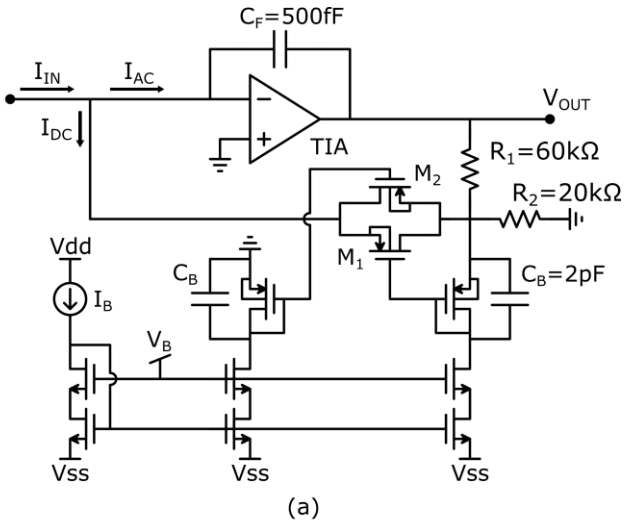

Fig. 5. DC Handling networks: a) tunable pseudo-resistor [30] and b) active servo loop [28].

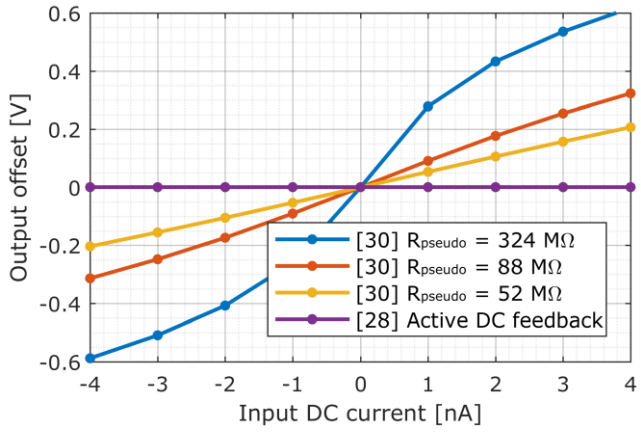

Fig. 6. DC offset at the output of the TIAs proposed in [28,30] as a function of the input leakage current.

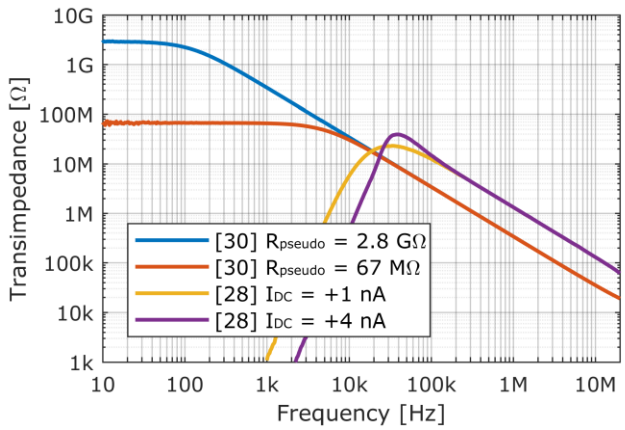

Fig. 7. Transfer function of the TIAs proposed in [28,30] for different bias conditions and different DC leakage currents.

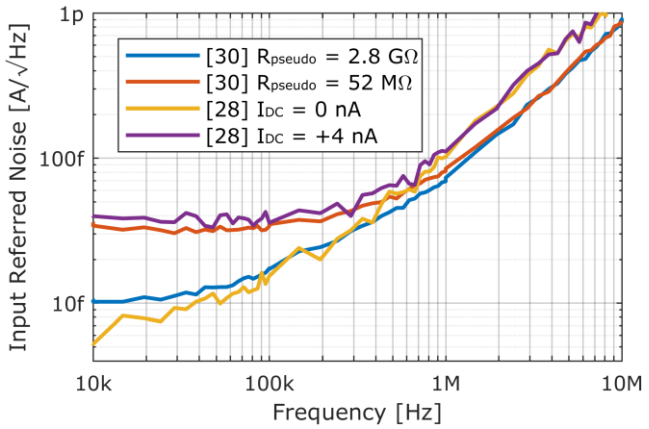

Fig. 8. Equivalent input current noise of the circuits proposed in [28,30] for different values of DC leakage current, for an overall 6 pF capacitive load on the virtual ground.

($\approx I_{LEAK}R_{EQ}$). To solve this trade-off, in [28] an active DC feedback network has been implemented (Fig. 5b). Here, the output voltage of the TIA is read by an ancillary active circuit, which filters out the AC signal and feeds the DC and low-frequency components back to the input node, effectively creating a DC feedback path. The structure is made by an active integrator (OA_2), where the input resistance is amplified by a current-reducer circuit (OA_1, M_a, M_b) [31] and the output voltage is converted into a signal current by a passive pseudo-resistor device (R_{pr}). By using an integrator in the feedback, this topology operates with a zero residual offset (Fig. 6), apart from mismatch effects and non-idealities. The use of a current reducer was necessary to set the DC feedback bandwidth at around 35 kHz, before the required operative bandwidth of the main circuit. In fact, in this topology the bandwidth of the TIA stage does not depend on the value of the leakage current (Fig. 7), but it is set by the bandwidth of the DC handling loop, since the signal can be amplified by the main feedback capacitance only at frequencies where the ancillary loop is disabled. The downsides of this solution are the increased area occupation, power dissipation and design complexity, mainly related to stability issues.

2.3 Sensitivity and noise performance

The noise performances of the two presented systems are very similar (Fig. 8). In both architectures, the input-referred low frequency noise is dependent on the level of the DC leakage current, while at high-frequency the noise of the transimpedance amplifier becomes dominant, shaped by the total capacitance connected to the virtual ground node and proportional to the frequency. In [27,30] the low frequency noise is dominated by the thermal noise of the equivalent resistance of the feedback pseudo-resistor, which has to be tuned based on the amount of leakage current present at the input. In [28] the dominant contribution is the shot noise of the pseudo-resistor R_{pr}, operating in sub-threshold regime. Both systems introduce a noise lower than 40 $fA/\sqrt{Hz}$ in the worst case at low frequency and lower than 100 $fA/\sqrt{Hz}$ around 1MHz.

III – MULTICHANNEL ELECTRONIC PLATFORM FOR AUTOMATED CONTROL OF PHOTONIC CIRCUITS

The analog preamplification performed by the front-end ASIC must be supplemented by a complete electronic system to effectively control the behaviour of photonic architectures. The system is in our case conceived in a modular way (Fig. 9) [32]: a small PCB, shaped to best fit in the optical bench, houses the photonic chip and the front-end ASIC, while a cable-connected motherboard hosts the rest of the electronics. This approach allows to maximize the readout sensitivity by placing the two chips close to each other, while allowing easy optical access and good flexibility in the design of the electronic system.

The custom mixed-signal motherboard (Fig. 10) is designed to: i) generate the sinusoidal signal necessary to stimulate the CLIPP, with programmable frequency (50kHz - 10 MHz) and amplitude (1V - 10V) to adapt to different sensor geometries; ii) amplify, filter and digitize the signal coming from 8 CLIPP sensors, without introducing any degradation in the readout SNR; iii) drive up to 16 actuators on the photonic chip, with a voltage accuracy of few mV and maximum current of 50 mA each; iv) perform the control algorithms for tuning or locking the working points of photonic devices.

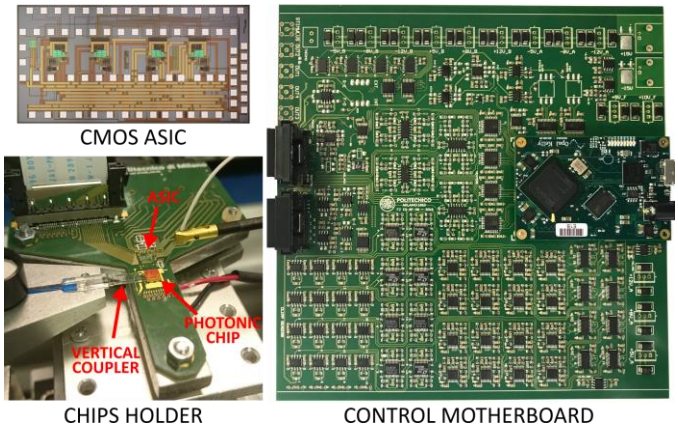

Fig. 9. Photograph of the modular control platform composed of a holder for the chips and a custom motherboard.

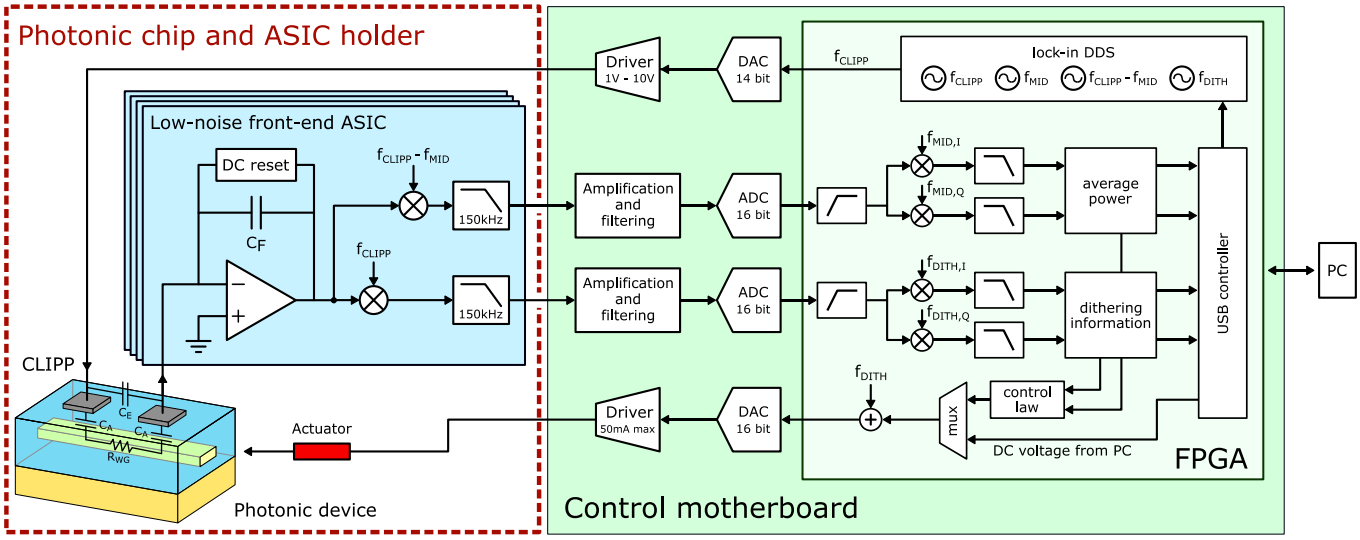

Fig. 10. Schematic view of the multichannel control platform for automated control of photonic circuits.

The digital core of the platform is an FPGA for versatile, real-time parallel processing. The FPGA is mounted on a commercial module, including the necessary components to interface it with a PC via USB. The FPGA manages the whole platform and generates the signals needed for the lock-in processing. In order to achieve the best possible resolution in the measurement, a heterodyne lock-in down-conversion is chosen: a first analog demodulation made by the ASIC brings the signal to an intermediate frequency (f_{MID}, around some kHz) above the 1/f noise corner frequency of the acquisition chain, then, after acquisition by the ADC, a second digital I/Q demodulation is performed in the FPGA to bring the signal to baseband as required by the lock-in processing. A tuneable digital filter defines the overall readout bandwidth, which is set depending on the resolution and speed requirements of the application.

The FPGA also implements the algorithms to control the photonic devices. The automatic operations can be grouped into tuning, i.e. scanning the heater voltage until the maximum or minimum of a desired cost function is reached, or locking, i.e. real-time feedback stabilization of the working point of a photonic device. Different control laws can be implemented: we successfully tested integral and PI controllers [33], as well as iterative gradient methods. The FPGA also generates modulation tones (indicated as f_{DITH} in Fig. 10), which can be used for two purposes: i) "labelling" different wavelengths by modulating the input light source, in order to identify different channels in the optical circuit; ii) using the dithering technique, by superimposing the tones to the heaters voltage, to obtain a signal proportional to the derivative of a photonic device transfer function, useful for amplitude-independent locking on the maximum (or minimum) [34]. A set of additional lock-in demodulators and filters is implemented in the digital domain for an efficient detection of these tones.

IV – EXPERIMENTAL DEMONSTRATIONS

Complex photonic circuits are usually made by combining elementary building blocks to obtain the desired optical functionality. One of the photonic components most commonly used to implement routers and filters [13,35] is the micro-ring resonator (Fig. 11a), a device able to steer the input light from one output port (THROUGH) to another (DROP) if the circumference length is an integer multiple of the wavelength, condition indicated as resonance. Being the transfer function of a micro-ring very narrow (Fig. 11b), active stabilization is required to reliably maintain the resonance condition. This can be obtained by placing a thermal actuator inside the device to actively modify its optical length, while monitoring the working point with a CLIPP detector at the output. The use of the dithering technique in combination with an integral controller is particularly convenient to implement an automated feedback loop in this case [33]. We have tested this configuration by tuning a ring resonator of 8 µm radius at a wavelength λ=1550 nm, then suddenly increased of 200 pm to introduce a perturbation after 150 ms. As it is possible to see from Fig. 11c, when the feedback control is off, the laser shift causes the ring to go out of resonance and a consequent drop of the output power is observed. On the other hand, when the active stabilization is on, the system is able to rapidly move the resonant wavelength of the ring by updating the heater voltage, restoring the full original output signal. In case of slower fluctuations, for example due to a thermally unstable environment, the feedback can constantly update the heater voltage to safeguard the resonance condition in real time, keeping the performance of the system at the required level.

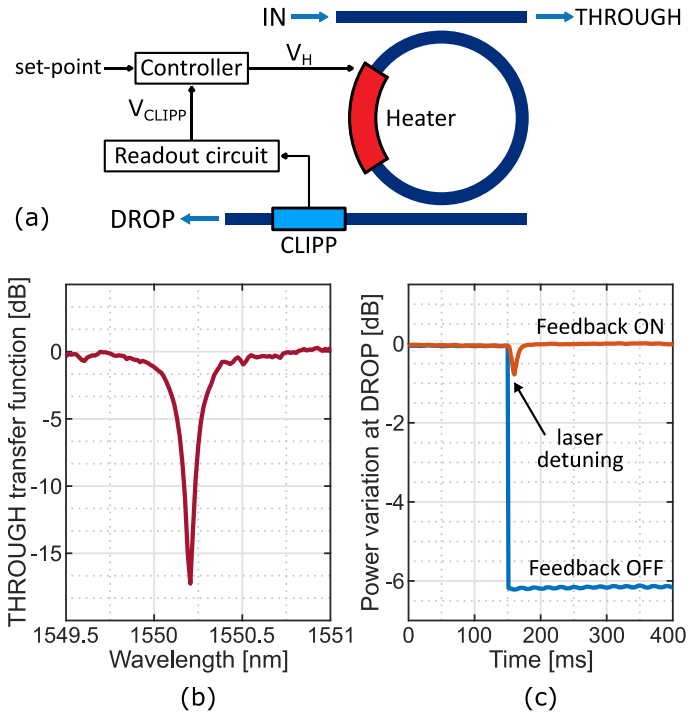

Fig. 11. (a) Ring resonator structure and CLIPP-based control scheme applied to it. (b) Transfer function of the device, measured at the THROUGH port. (c) Time evolution of the power at the DROP port when shifting the laser wavelength 200 pm out of resonance, comparing the cases of feedback control enabled/disabled.

Another important photonic device is the Mach-Zehnder Interferometer (MZI) (Fig. 12a), that can be used as an optical switch. The splitting of the light that reaches the two outputs of this component can be controlled again with a thermal actuator, placed on one of the inner arms of the device. When the MZI is used to route all the light to a single output, a control strategy similar to what

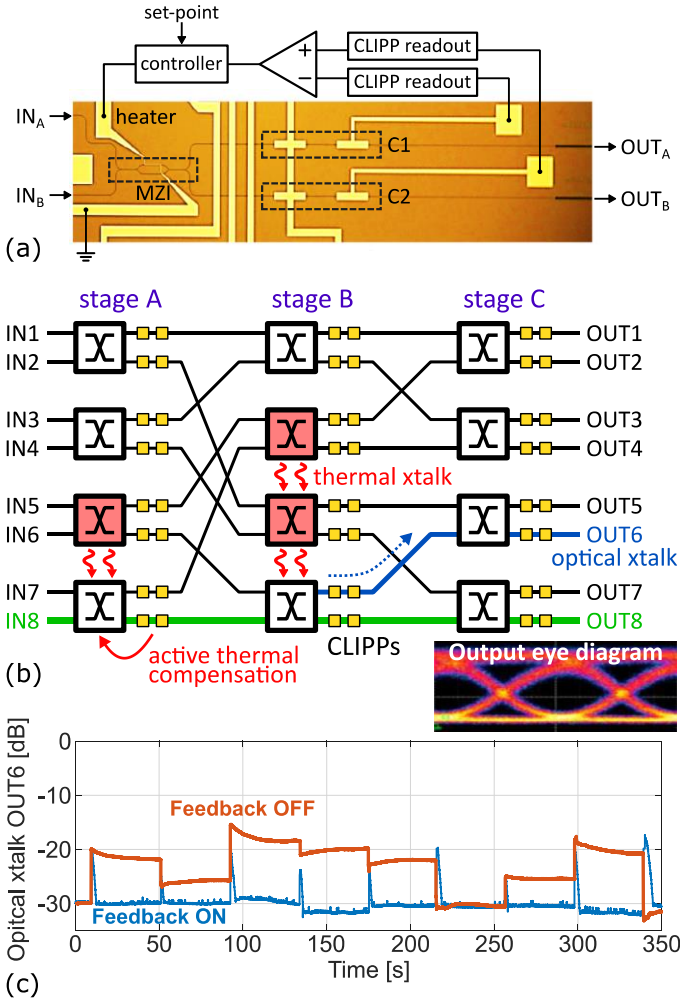

Fig. 12. (a) CLIPP assisted control of a MZI switch. (b) Schematic of the 8x8 switch fabric with active thermal compensation. (c) Optical crosstalk at port OUT6 during several random switching events of the surrounding MZIs, when active thermal compensation is OFF (green curve) and ON (blue curve).

discussed for ring resonators can be employed. A differential readout of the two output waveguides can be exploited in this case to improve the performance of the control action. We have experimentally tested this approach on an 8x8 silicon photonic switch fabric, made of 12 thermally actuated MZI switches arranged in 3 switching stages (Fig. 12b) [36]. This architecture allows reconfigurable routing of light arriving at any of its inputs to any of the outputs and, although very popular for its geometrical simplicity, poses several challenges in terms of control and stabilization. CLIPP detectors were integrated at both output ports of each MZI switch, allowing to monitor the optical power in each waveguide and track in real-time the status of the router. The active stabilization was tested by first configuring the matrix to route a 10 Gbit/s modulated input signal from IN8 to OUT8, then the equilibrium condition was perturbated by turning on and off the thermal actuators of the unused neighbouring paths. As it's possible to see, the feedback system is able to route the light to the correct output, leading to a clearly open output eye diagram, and to effectively counteract the external perturbations, avoiding optical crosstalk caused by fractions of the input signal that reaches the wrong output (OUT6 in our case). When the stabilization is turned off, the same perturbations induce an unacceptable worsening of the performance of the system (Fig. 12c), demonstrating how essential a closed-loop operation is.

V – TECHNOLOGY CHALLENGES TOWARDS MONOLITHIC INTEGRATION

We have so far discussed specific ASICs matching the requirements of a photonic integrated circuit (PIC) in terms of capability of reading the sensors with maximum sensitivity and of controlling the working point of photonic devices with the required precision and speed, through a real-time tuning of heaters properly placed over the waveguides of the PIC. Despite the mentioned successful results obtained with this approach, a problem of scalability certainly arises. Complex photonic circuits that might integrate hundreds of devices can hardly profit of this approach, simply considering the number of bonding wires that are necessary to connect them to the electronics. Flip-chip connection [37-40] has been experimented as an alternative approach and certainly successfully mitigated some issues, but it cannot be the solution. The answer seems to lead straight to a monolithic integration of photonics and electronics into a single technology.

Two directions are presently on discussion, whether to add the photonic devices to the list of manufacturable features within the highly mature microelectronic industry or to add some electronic capabilities into the highly refined processes available in the silicon photonics industry.

5.1 Zero-change paradigm on standard microelectronic technology

Concerning the first approach, there are significant realizations adopting the 'zero-change' paradigm that relies on standard microelectronics foundry processes to accommodate photonics and all optical devices within the native process-manufacturing rules, thus reusing the same design tools and libraries [41-43]. An example of an electro–optic system on a single chip has been demonstrated by the group coordinated by Vladimir M. Stojanovic [44], integrating a microprocessor, memory and logic together with hundreds of photonic components (SiGe photodetectors, ring modulators etc.) using a commercial high-performance 45-nm CMOS silicon-on-insulator (SOI) process providing the same crystalline layer used for electronics for both the waveguides and the on-chip photonic transmitters and receivers. This system targets intra and inter data center communication, enabling direct interaction among chip components using light. The temperature control and the tuning circuitry essential for reliable operation of the system still rely on tape-out waveguides on each micro-ring modulator and standard SiGe photodetectors, given the fact that the ring modulator is the only optical element along the light path. Taking advantage of the densely integrated electronics, a digital controller monitors the photocurrent and controls the power of a microheater to keep the resonator locked to the laser wavelength under thermal variations, with a proved receiver sensitivity of about -5 dBm for a BER < 10^{-12}. Although a selective etching of the wafer substrate had to be added to reduce the optical leakage through the thin (200 nm) buried oxide layer, otherwise resulting in high WG losses, the standard microelectronic foundry process has proved to lodge silicon-strip waveguides of high quality, reporting optical propagation losses of about 4.3 dB/cm at the 1180 nm used laser wavelength, which well compares with the best practice of photonic foundries. This pioneering approach provides a good example of Electronic-Photonic co-design targeting monolithic integration of photonic functionalities into a high-end electronic platform.

5.2 Transistor integration on a Silicon Photonic technology platform

As an alternative approach, we might address the option to integrate basic electronic features into already existing photonic technology platforms not originally conceived for this role. Simple logic gates and multiplexing circuits would in fact be enormously advantageous in decreasing the number of required connections otherwise needed when managing large scale photonic architectures and would correspondingly solicit a truly intimate electronic-photonic co-design on the same photonic technology platform.

An interesting starting point is offered by active Silicon Photonics technologies, where inline thermal actuators and optical modulators or attenuators are implemented through specific process steps dedicated to doping implantation. Common doping levels are around $10^{17}\ cm^{-3}$ for devices and around $10^{20}\ cm^{-3}$ for ohmic contacts. The dopant species are implanted directly inside the waveguide core and both p-type and n-type dopings are available. Elementary electronic devices could be realized by exploiting these process steps to modulate the doping profile of dedicated portions of silicon. Since the doping implantation is available only on the waveguide layer, the conceived devices will be insulated from the main substrate of the chip and will have to be planar.

Given the low frequency of the required control feedback, on-chip multiplexing of signals, both in the sensing and actuation paths, would already downscale the interconnect complexity by a significant amount, allowing for a limited number of bonding wires with respect to the number of control/actuating points in the photonic fabric. As a first step in this direction, we designed a pnp bipolar transistor prototype (Fig. 13a), which could be used to locally drive each thermal actuator. The base of the device consists in a circular portion of waveguide layer with a $10^{17}\ cm^{-3}$ n-type doping, while the emitter is realized by a $10^{20}\ cm^{-3}$ p-type region in the middle of the base. The collector is made by a $10^{15}\ cm^{-3}$ p-type region (intrinsic doping of the waveguide) around the base and by a following $10^{17}\ cm^{-3}$ p-type ring. Both the base and the collector are contacted on dedicated $10^{20}\ cm^{-3}$ contact regions. The circular shape was chosen to maximize the collection of charges by the collector, while the intrinsic region between base and collector was inserted to improve the breakout voltage limit of the device. A BJT made of 16 of the described elementary cells was tested, with the promising results shown in Fig. 13b.

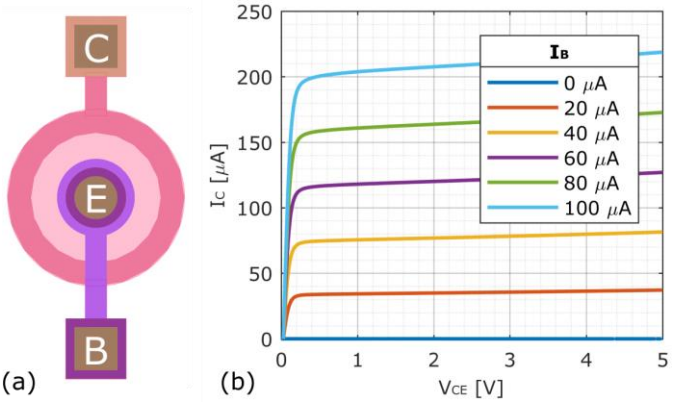

Fig. 13. pnp BJT prototype manufactured in a standard active Silicon Photonic technology: a) layout and b) measured characteristic curves.

VI - CONCLUSIONS

Electronics-Photonics co-design, either hybrid or monolithic, seems inevitable to bring Silicon Photonics to maturity and technological significance. The SiP weaknesses given by the high sensitivity to thermal fluctuations and process variations have been demonstrated to be efficiently compensated by proper feedback control. Furthermore, the availability of transparent detectors will boost the design of very complex photonic circuits, thanks to the possibility of distributed monitoring and control of the working point of many optical devices in parallel. While optical links can provide the additional bandwidth capacity that is required to scale the interconnect system with the increasing performances of today's processors, the electronic layer can provide the necessary feedback control and stabilization. The demonstrations summarized in this paper represent the turning point for chip-scale electronic-photonic systems to transform telecom architectures and equipment, enabling more energy-efficient and powerful computers for data centers, as well as reconfigurable network infrastructures.

ACKNOWLEDGEMENTS

Authors would like to acknowledge the contribution of European Commission through the Horizon 2020 research and innovation program, project NEBULA (H2020-ICT-2019-2) under contract n.871658 and project SUPER-PIXELS (FET Open) under contract n.829116. Authors also thank the staff of POLIFAB, the microelectronic facility of Politecnico di Milano, for support during fabrication and mounting of the systems.

References:

[1] 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.

[2] R. G. Beausoleil, “Large-scale integrated photonics for high-performance interconnects”, IEEE Photonic Society 24th Annual Meeting, Arlington, VA, pp. 326-327, 2011.

[3] T. Alexoudi *et al*., “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, 2019.

[4] R. Bogue, “Recent developments in MEMS sensors: a review of applications, markets and technologies”, Sensor Review, vol. 33, no. 4, pp. 300-304, 2013.

[5] M. Perlmutter *et al*., “The future of the MEMS inertial sensor performance, design and manufacturing”, 2016 DGON Inertial Sensors and Systems (ISS), Karlsruhe, pp. 1-12, 2016.

[6] S. Lloyd *et al*., "The age of sensors: how MEMS sensors will enable the next wave of new products”, 2016 IEEE Symposium on VLSI Technology, Honolulu, HI, pp. 1-4, 2016.

[7] Q. Cheng *et al*., “Photonic switching in high performance datacenters [Invited] ”, Optics Express 26, 16022-16043, 2018.

[8] A. Ribeiro *et al*., “Demonstration of a 4x4-port self-configuring universal linear optical component”, 2016 Progress in Electromagnetic Research Symposium (PIERS), Shanghai, pp. 3372-3375, 2016.

[9] K. Fotiadis *et al*., “16×16 silicon photonic AWGR for dense wavelength division multiplexing (DWDM) O-band interconnects”, Proceedings SPIE 11285, Silicon Photonics XV, 112850F, 2020.

[10] J. Sun *et al*., “Large-scale nanophotonic phased array”, Nature 493, 195-199, 2013.

[11] S. Grillanda *et al*., “Post-fabrication trimming of athermal silicon waveguides”, Optics letters, vol. 38, no. 24, pp. 5450-5453, 2013.

[12] W. Bogaerts *et al*., “Programmable photonic circuits”, Nature 586, 207-216, 2020.

[13] F. Zanetto *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.

[14] A. Annoni *et al*., “Unscrambling light - automatically undoing strong mixing between modes”, Light: Science and Applications 6, e17110, 2017.

[15] J. Michel *et al*., “High-performance Ge-on-Si photodetectors”, Nature Photonics, vol. 4, no. 8, pp. 527-534, 2010.

[16] S. Assefa *et al*., “Reinventing germanium avalanche photodetector for nanophotonic on-chip optical interconnects”, Nature, vol. 464, no. 7285, pp. 80-84, 2010.

[17] J. K. Doylend *et al*., “The evolution of silicon photonics as an enabling technology for optical interconnection”, Laser & Photonics Reviews, vol. 6, no. 4, pp. 504-525, 2012.

[18] T. K. Liang *et al*., “Silicon waveguide two-photon absorption detector at 1.5 μm wavelength for autocorrelation measurements”, Applied Physics Letters, vol. 81, no. 7, pp. 1323-1325, 2002.

[19] T. Tanabe *et al*., “All-silicon sub-Gb/s telecom detector with low dark current and high quantum efficiency on chip”, Applied Physics Letters, vol. 96, no. 10, pp. 101103-101113, 2010.

[20] T. Baehr-Jones *et al*., “Photodetection in silicon beyond the band edge with surface states”, Optics Express, vol. 16, no. 3, pp. 1659-1668, 2008.

[21] H. Chen *et al*., “Cavity-enhanced photocurrent generation by 1.55 μm wavelengths linear absorption in a p-i-n diode embedded silicon microring resonator”, Applied Physics Letters, vol. 95, no. 17, pp. 171111-171113, 2009.

[22] Y. Li *et al*., “Active resonance wavelength stabilization for silicon microring resonators with an in-resonator defect-state-absorption-based photodetector”, Optics Express, vol. 23, no. 1, pp. 360-372, 2015.

[23] H. Jayatilleka *et al*., "Automatic configuration and wavelength locking of coupled silicon ring resonators”, Journal of Lightwave Technology, vol. 36, no. 2, pp. 210-218, 2018.

[24] A. Gazman *et al*., “Tapless and topology agnostic calibration solution for silicon photonic switches”, Optics Express, vol. 26, no. 25, pp. 32662-32674, 2018.

[25] F. Morichetti *et al*., “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.

[26] M. Carminati *et al*., “Design guidelines for contactless integrated photonic probes in dense photonic circuits”, Journal of Lightwave Technology, vol. 35, no. 14, pp. 3042-3049, 2017.

[27] P. Ciccarella *et al*., “Impedance-sensing CMOS chip for non-invasive light detection in integrated photonics”, IEEE Transactions on Circuits and Systems II: Express Briefs, vol. 63, no. 10, pp. 929-933, 2016.

[28] F. Zanetto *et al*., “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.

[29] S. Zho *et al*., “A CMOS passive mixer with low flicker noise for low-power direct-conversion receiver”, Journal of Solid-State Circuits, vol. 40, no. 5, pp. 1084-1093, 2005.

[30] E. Guglielmi *et al*., “High-value tunable pseudo-resistors design”, Journal of Solid-State Circuits, vol. 55, no. 8, pp. 2094-2105, 2020.

[31] F. Gozzini *et al*., “Linear transconductor with rail-to-rail input swing for very large time constant applications,” Electron. Lett., vol. 42, no. 19, p. 1069, 2006.

[32] E. Guglielmi *et al*., “16-Channel modular platform for automatic control and reconfiguration of complex photonic circuits”, 2017 IEEE International Symposium on Circuits and Systems (ISCAS), Baltimore, MD, pp. 1-4, 2017.

[33] S. Grillanda *et al*., “Non-invasive monitoring and control in silicon photonics using CMOS integrated electronics”, Optica, vol. 1, no. 3, pp. 129-136, 2014.

[34] K. Padmaraju *et al*., “Wavelength locking and thermally stabilizing microring resonators using dithering signals”, Journal of Lightwave Technology, vol. 32, no. 3, pp. 505-512, 2014.

[35] D. Aguiar *et al*., “Automatic Tuning of Silicon Photonics Microring Filter Array for Hitless Reconfigurable Add–Drop”, Journal of Lightwave Technology, vol. 37, no. 16, pp. 3939-3947, 2019.

[36] A. Annoni *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, 2016.

[37] K. Settaluri *et al*., “Demonstration of an optical chip-to-chip link in a 3D integrated electronic-photonic platform”, 41st European Solid-State Circuits Conference (ESSCIRC), Graz, AT, pp. 156-159, 2015.

[38] N. Pavarelli *et al*., "Optical and Electronic Packaging Processes for Silicon Photonic Systems", Journal of Lightwave Technology, vol. 33, no. 5, pp. 991-997, 2015.

[39] H. Nishi *et al*., “Monolithic integration of a silica AWG and Ge photodiodes on Si photonic platform for one-chip WDM receiver”, Optics Express 20, 9312-9321, 2012.

[40] K. Li *et al*., “Electronic–photonic convergence for silicon photonics transmitters beyond 100 Gbps on–off keying”, Optica 7.11 (2020): 1514-1516.

[41] L. Alloatti *et al*., “Waveguide-coupled detector in zero-change complementary metal–oxide–semiconductor”, Appl. Phys. Lett. 107, 041104 (2015).

[42] J. M. Shainline *et al*., “Depletion-mode carrier-plasma optical modulator in zero-change advanced CMOS”, Opt. Lett. 38, 2657–2659 (2013).

[43] M. T. Wade *et al*., “Energy-efficient active photonics in a zero-change, state-of-the-art CMOS process”, In Optical Fiber Communication Conference Tu2E.7 (Optical Society of America, 2014).

[44] C. Sun *et al*., “Single-chip microprocessor that communicates directly using light”, Nature, vol. 528, pp. 534-538, 2015.