We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Merging Plasmonics and Silicon Photonics Towards Greener and Faster “Network-on-Chip” Solutions for Data Centers and High-Performance Computing Systems

Sotirios Papaioannou, Konstantinos Vyrsoinos, Dimitrios Kalavrouziotis, Giannis Giannoulis, Dimitrios Apostolopoulos, Hercules Avramopoulos, Filimon Zacharatos, Karim Hassan, Jean-Claude Weeber, Laurent Markey, Alain Dereux, Ashwani Kumar, Sergey I. Bozhevolnyi, Alpaslan Suna, Oriol Gili de Villasante, Tolga Tekin, Michael Waldow, Odysseas Tsilipakos, Alexandros Pitilakis, Emmanouil E. Kriezis and Nikos Pleros

Additional information is available at the end of the chapter

<http://dx.doi.org/10.5772/51853>

1. Introduction

In recent years it has become evident that the increasing need for huge bandwidth and throughput capabilities in Data Center and High Performance Computing (HPC) environments can no longer be met by bandwidth-limited electrical interconnects. Besides their limited capacity, electrical wiring technologies impose great energy and size limitations originating from the requirements to handle the vast amount of information that needs to be exchanged across all hierarchical communication levels within Data Centers and HPCs, i.e. rack-to-rack, backplane, chip-to-chip and on-chip interconnections. This reality has inevitably led to a clearly shaped roadmap for bringing optics into the spotlight and replacing electrical with optical interconnects, thereby overtaking the bottleneck in traffic exchange imposed by electrical wires [1]. Hence, the optical technology should penetrate into intra-rack and board-to-board transmission links, considering that the optical fiber as a large-bandwidth transmission medium [2] is used only in commercial systems for rack-to-rack communication [3].

Nevertheless, data communication and power consumption are still daunting issues in Data Centers and HPCs. According to recent predictions made in [4], the barrier of 10PFlops

computing performance should have been overcome in 2012 by a supercomputer that consumes 5MW of power [5]. In addition, [4] predicted that exascale supercomputing machines would consume 20MW having a power efficiency of 1mW/Gb/s [5]. Nonetheless, power consumption in such environments has been proven to be even higher than expected: Today's top-ranked supercomputer, the "K computer", has already reached the 10PFlops performance benchmark but at the expense of excessive consumed power that is more than twice [6] the value that was predicted in 2008. All the above imply that the use of optics at inter-rack communication level is not enough for delivering the necessary performance enhancements. Therefore, the optical technology should now be exploited at shranked networking environments: The penetration of low-energy photonic solutions at board-to-board, chip-to-chip and eventually intra-chip interconnects would yield remarkable savings in energy consumption [7]. The current mainstream photonic route with high integration and low-cost perspectives relies on the Silicon-on-Insulator (SOI) photonics platform, whose growing maturity is soon expected to release Tb/s-scale data transmission and switching capabilities in datacom and computercom units ensuring low latency, low power consumption and chip-scale integration credentials [8].

Even so, photonic devices cannot reach the compactness of their electronic counterparts: the dimensions of traditional optical structures are limited by the fundamental law of diffraction, preventing the way towards high density integration for interfacing with electronics at the nanoscale. This gap in size between photonic and electronic components is called to be bridged by a promising disruptive technology named plasmonics [9]-[11]. The emerging discipline of plasmonics has started to gain ground as the "beyond photonics" chip-scale platform that can enter the interconnect area [12]-[14], holding a great promise for additional reductions in circuit size and increase in energy efficiency. Plasmonics relies on the excitation of surface plasmon polaritons (SPPs) that are electromagnetic waves coupled to oscillations of free electrons in a metal and propagate along a metal-dielectric interface at near the speed of light. These "hybrid" surface waves have transverse magnetic (TM) polarization in nature and their exhibited electromagnetic field intensity reaches its maximum value at the metal surface whereas it decays exponentially while moving away from the metal-dielectric interface [15]. In this way strong intrinsic confinement is feasible even at sub-wavelength scale [16], breaking the size barriers of diffraction-limited optics and enabling the development of compact integrated nanophotonic circuits [17]. Plasmonic technology does not only succeed in providing light manipulation at sub-wavelength dimensions but at the same time allows for the injection of electrical pulses via the metallic layer, offering thereby a seamless energy-efficient platform for merging light beams with electrical control signals towards "active" operations [18]-[25].

Among the various plasmonic waveguide structures proposed so far (i.e. band-gap structures [26], metallic nanowires [27], V-groove waveguides [28]), the low-energy credentials of plasmonics has been mainly highlighted in the case of dielectric-loaded SPP (DLSPP) waveguides, where a dielectric (e.g. polymer) ridge is deposited on top of a smooth metallic film. As a result, strong sub-wavelength (DLSPP) mode confinement is achieved at the metal-dielectric interface. Nonetheless, this performance comes at the expense of

excessively increased propagation losses compared to conventional dielectric waveguides due to radiation absorption in the metallic stripe, yielding quite short propagation lengths lying in the order of few tens of micrometers [29],[30]. Apart from the strong mode confinement, the main advantage of the DLSPP waveguide technology stems from the nature of the chosen dielectric material which, depending on its thermo-optic (TO) [25],[31]-[35] or electro-optic [24],[36] properties, allows for the exploitation of the corresponding effects, enabling in this way the deployment of highly functional active plasmonic elements. Regarding the TO effect, the underlying metallic layer, which is in direct contact with the polymer and therefore with the largest part of the DLSPP mode field, serves as an energy-efficient electrode that yields immediate change of the effective index of the propagating mode at the presence of electric current, leading to fast TO responses. These advantages render the DLSPP waveguides ideal for TO manipulation of the plasmonic waves in various functional circuitry implementations like modulation [32],[33], ON/OFF gating [25] and switching [34],[35]. The progress made so far on active DLSPP-based circuits renders the TO effect as the most mature mechanism for bringing low-energy active plasmonics in true data traffic environments [25],[31].

To this end, the roadmap towards practical employment of active plasmonic circuits in the development of "greener" and faster "Network-on-Chip" (NoC) solutions for Data Centers and HPCs seems to be the following: the synergy between plasmonic, electronic and photonic components for the realization of NoC deployments with minimized power consumption, size and enhanced throughput capabilities, simply by taking advantage of the virtues of each technology. In this perspective, silicon photonics can be used for low-loss optical transmission in passive interconnection components, plasmonics can provide small footprint and low power consumption in active switching modules and electronics can be employed for intelligent decision-making mechanisms. However, considering that plasmonic technology is a premature technology with only a few years of development compared to silicon photonics and with limited functionality so far, the way towards the implementation of such hybrid NoC environments requires priorly: a) the interconnection between silicon and plasmonic waveguide structures in a SOI platform in order to avoid the employment of high-loss plasmonic waveguides for passive circuitry [30],[37], b) the proof of the credentials of plasmonics in wavelength division multiplexing (WDM) applications so as to support the complete portfolio of optics [38], c) the demonstration of functional active plasmonic circuitries in realistic data traffic environments [25],[31].

This chapter aims to cover all these issues starting from the 4×4 silicon-plasmonic router architecture, where its main building blocks and their principle of operation are described briefly. The chapter continues with the geometrical specifications of the silicon and plasmonic waveguides of the hybrid routing platform as well as with the silicon components and sub-systems, namely the photodiodes (PDs) and the multiplexers (MUXs), respectively. The silicon-plasmonic asymmetric Mach-Zehnder interferometer (A-MZI) is then presented and analyzed thoroughly since it is the constituent component of the 4×4 switching matrix of the router. Subsequently, the WDM credentials of plasmonics are experimentally confirmed by demonstrating WDM data transmission over a DLSPP

waveguide as well as WDM data switching operation via a hybrid Si-DLSPP A-MZI, along with its switching performance metrics. Various optimization procedures are also listed. The chapter concludes with an estimation of the router's total power consumption and losses.

2. Tb/s-scale hybrid Router-on-Chip platform

Data Center and HPC communication systems require the development of NoC environments with compact size, high-throughput and enhanced power saving capabilities for efficient traffic management. In this perspective, the benefits of silicon photonics and plasmonics can be combined towards realizing a 4×4 “Router-on-Chip” platform [13] with Tb/s-scale data capacity and ultra-low energy requirements for chip-to-chip and on-chip optical interconnects. This hybrid routing platform employs a silicon motherboard hosting all various optical and electrical circuitry, where the “heart” of this router chip is a hybrid 4×4 switching matrix comprising 2×2 TO Si-DLSPP switches, and exploits the best out of all technology worlds: the low optical loss of Si-waveguide motherboard for its passive interconnection parts, the high energy efficiency of plasmonics in the router's “active” parts, and the intelligence provided by electronics for its decision-making mechanisms. The following sections describe in detail the Tb/s-scale Si-DLSPP router architecture and its basic building blocks, the experimental evidences of the WDM data transmission and switching capabilities of the DLSPP waveguides, various optimization approaches and finally the performance of the 4×4 router with estimations in power consumption and optical insertion losses.

2.1. 4×4 silicon-plasmonic router architecture

The 4×4 silicon-plasmonic router architecture aims at achieving 1.12Tb/s aggregate throughput with small footprint and low power consumption. Figure 1 illustrates a schematic layout of this hybrid router architecture. The router consists of SOI MUXs, Si PDs, a hybrid Si-DLSPP 4×4 switching matrix and an integrated circuit (IC) microcontroller. Size minimization and power savings are promoted by the employment of four SOI MUXs composed of cascaded Si ring resonators (RRs) placed in parallel at the frontend of the router and four identical thermo-optically addressed Si-DLSPP A-MZIs interconnected in a Beneš topology. The plasmonic components are used only where switching functionality is required in order to achieve both low optical losses and low power consumption. The Si technology is also responsible for coupling light in and out of the platform, the interconnection of all subsystems and the optoelectronic (OE) signal conversion of the incoming packets' headers. Four 7×1 MUX circuits based on SOI RRs are used to multiplex 28 time-overlapped wavelengths, spaced by 100GHz and modulated with 40Gb/s non-return-to-zero (NRZ) packet traffic, into four optical data streams. Consequently, each one of these multi-wavelength data sequences carries an aggregate traffic of 280Gb/s that will enter one of the four inputs of the plasmonic switching device and will then follow the same route through the entire platform. Apart from the 28 input ports that are utilized for the insertion of the data carrying packets, the router also incorporates 4 additional input ports

where they are inserted four discrete low-rate wavelength beams containing the header information of the four incoming data streams. The header for each stream is a two pulse combination within $1\mu\text{s}$ time interval and specifies the desired output port, while the duration of the payload of the data packets can be variable. These header pulses are optoelectronically converted via low-rate all-silicon PDs into respective electrical signals that subsequently feed the IC control unit. Here the low speed electronics operating at the data packets rate carry out the intelligent decision-making process that is essential for the routing functionality: The IC microcontroller performs header processing operations and generates differential control electronic signals for driving appropriately the switching matrix. In this way, the incoming optical data streams are routed to the desired output ports as a result of fruitful cooperation between the silicon photonic, plasmonic and electronic technologies.

Figure 1. Schematic layout of the 4×4 silicon-plasmonic router architecture.

2.2. Silicon and DLSP waveguides

The different technologies and building blocks involved in the hybrid Tb/s-scale router require the employment of the appropriate waveguide technology in order to merge all of them successfully on the same motherboard. The existence of a slab layer in rib waveguides compared to the strip counterparts facilitates the fabrication process for the hetero-integration of all processing components, namely the PDs, the SOI MUXs and the DLSP waveguide structures, rendering this waveguide technology suitable for the low-loss Si-based communication links needed in the SOI routing platform. Taking into consideration the intrinsic TM nature of plasmonic waves, the geometrical dimensions of the Si rib waveguides need to be carefully chosen for compatibility with plasmonic waveguides. In this perspective, rib waveguides featuring a cross section of 400nm width by 340nm height with a 50nm-thick remaining slab are the Si waveguide technology of choice for the routing platform in order to support low-loss TM light propagation [37].

Nevertheless, the coexistence of silicon photonics and plasmonics on the same platform imposes proper formation of the SOI motherboard for efficient hetero-integration of the

DLSP waveguides [31]: By etching the SOI motherboard down to 200nm (Figure 2(a)), the formed recess in the 2 μ m-thick buried oxide (BOX) of the SOI substrate serves as the hosting region of a DLSP waveguide. Towards selecting the dimensions of the DLSP waveguides, it should be taken into account that the DLSP waveguide characteristics, i.e., the mode field confinement, effective index and propagation length, are strongly influenced by the width and height of the employed dielectric ridge [29] as a result of the SPP field confinement that occurs when this ridge is placed on top of a metal film. Considering methyl-methacrylate (PMMA) ridges placed on gold stripes to be operated at telecom wavelengths ($\sim$ 1550nm), the optimum ridge dimensions, ensuring tight mode confinement ($<1\mu$ m) and relatively long propagation ($\sim$ 50 μ m) of DLSP modes, are about 500nm and 600nm for the ridge width and height, respectively [29]. Moreover, the underlying gold stripes are about 3 μ m wide and 60nm thick. A microscope image taken from a Si-DLSP waveguide is shown in Figure 2(b) where the recessed area of the BOX and the underlying gold film are evident.

Figure 2. (a) Scanning electron microscope (SEM) image in bird's eye view of the silicon cavity after etching, including as inset the Si waveguide tip after the cavity patterning process, (b) microscope image of a straight DLSP waveguide coupled to Si waveguides. (c) Bird's eye and (d) side schematic views of a Si-to-DLSP waveguide transition.

Apart from the geometrical specifications of these two different waveguide technologies, it is necessary to implement an efficient transition from one type of waveguide to the other. In our case, the interface between the silicon and plasmonic waveguides is implemented by a butt-coupling approach, as depicted in bird's eye and side views in Figure 2(c)-(d), respectively. The design of this approach relies on numerical modeling conducted by means of full vectorial three-dimensional finite element method (3D-FEM) simulations. Towards finalizing the design specifications, various parameters should be examined, i.e., the Si waveguide's width, the vertical offset between the two types of waveguides, the existence of a longitudinal gap in the metallic stripe as usually exists in fabricated structures, so as to optimize the matching of the optical modes during transition [37]. The validity of the simulation outcomes is achieved by their comparison with cut-back measurements for a

variety of fabricated all Si and hybrid Si-DLSPP waveguide samples [31],[37]. As a result, a coupling loss per interface variance with a mean value of -2.5dB is feasible for 175nm-wide Si waveguides with 200nm vertical offset and DLSPP waveguides with 0.5 μ m metal gap. The same cut-back evaluation procedure leads to silicon and plasmonic propagation losses of approximately 4.5dB/cm and 0.1dB/ μ m, respectively, being in good agreement with typical results reported in the literature [30].

2.3. Silicon photodiodes

Since the silicon-plasmonic routing platform comprises also microelectronics, it is necessary to develop integrated C-band (1530-1565 nm) photodiodes that enable the OE conversion of the optical header pulses into electrical signals. However, a complementary metal-oxide-semiconductor (CMOS) -compatible SOI-based integration of a material that absorbs inside this wavelength window is really challenging. One detection scheme within the C-band on a SOI platform is the hybridization of integrated Si structures with other materials (e.g. Ge [39], InGaAs [40]) for direct (linear) absorption. This hybrid approach can lead to high-performance but complex devices, requiring complicated and cost-intensive fabrication procedures and modifications compared to the standard CMOS process. For this purpose, monolithical and fully CMOS-compatible detector concepts have been proposed for the development of comparably simple devices with less fabrication costs. In this perspective, two photon absorption detection can be exploited inside a non-resonant or resonant Si device [41], revealing however non-linear effect and requiring high light intensity. Alternatively, the generation of direct (linear) absorption inside the telecom window can be achieved by the introduction of midgap energy states via Si⁺-ion implantation [42]. Despite the lower (direct) absorption compared to the hybrid detectors, this type of monolithic Si detector approach, the defect state Si⁺-ion implanted detector, offers the best compromise for simple integration and low-rate detection.

Figure 3. Integrated Si-implanted PDs: (a) Schematic cross sectional view, (b) schematic top view, (c) test chip, (d) I-V curves at various input optical power levels, (e) sensitivity at different reverse bias.

Taking into account that the 4×4 silicon-plasmonic router architecture employs low-rate optical header pulses, the defect state all-silicon implanted PDs are suitable for the procedure of the OE conversion, pursuing a trade-off between minimum signal distortion during header detection and integration complexity. Schematic illustrations of this type of PD interconnected with a Si waveguide are shown in cross sectional and top views in Figure 3(a)-(b), respectively, whereas the top view of a test chip including fabricated all-Si implanted PDs with different absorption lengths is depicted in Figure 3(c). Figure 3(d)-(e) illustrate the performance of a 0.75mm-long Si⁺-ion implanted PD, exhibiting linear absorption and sensitivity values of 0.09A/W and 0.11A/W for -6V and -20V bias voltage, respectively. By increasing the device length, the sensitivity is getting higher. However, even for a three times longer device, the dark current is well below 100nA.

2.4. Silicon multiplexers

In order to support Tb/s-scale bandwidth, the hybrid router takes advantage of SOI multiplexing devices. Till now various multiplexing/demultiplexing configurations based on SOI technology have been proposed in the scientific community, such as echelle diffraction gratings (EDGs) and arrayed waveguide gratings (AWGs) [43], first or higher order RRs [44]-[49] and A-MZIs [50]. In many cases, due to fabrication limits and discrepancies, thermal tuning is performed to attain ideal spectral positioning of the MUX resonances. This controlling technique is adopted for tunability of the second-order SOI RRs employed in the silicon-plasmonic router, aiming at optimized multiplexing performance. Considering that the router is planned to handle 40Gb/s line-rates, the 3-dB passband bandwidth for each data channel should be at least 40GHz and the adjacent channel crosstalk should be kept lower than -15dB. These issues can be met by the 2nd order RRs employed in the router, since such structures enable increased flexibility in the formation of transmission peaks' spectral shape and bandwidth within the same channel spacing constraints due to the additional coupling stage employed between the two RRs [51].

A fabricated 2nd order RR featuring the Si waveguide specifications of Section 2.2 is illustrated in Figure 4(a). Indicative experimental spectra obtained at the Drop port of 2nd order ring structures are presented with blue lines in Figure 4(b)-(c). In particular, Figure 4(b) refers to a fabricated 2nd order RR with 9μm radius, 190nm bus-to-ring gaps and 380nm ring-to-ring gap, exhibiting 42.5GHz 3-dB bandwidth and 9.2nm free spectral range (FSR). In the case of a ring structure with 12μm radius, 170nm bus-to-ring gaps and 300nm ring-to-ring gap, the corresponding 3-dB bandwidth and FSR values are 48.5GHz and 6.9nm (Figure 4(c)). Replicas of these measured transfer functions are also illustrated in Figure 4(b)-(c) with red dashed line by using the circuit-level modeling tool described extensively in [13].

The aforementioned 2nd order RRs are employed as the building blocks for the design of the 7:1 SOI MUX devices required in the router. Figure 5 depicts a more generic 8:1 multiplexing layout. The MUX design comprises eight cascaded thermo-optically tunable 2nd order RRs clustered in two groups with two different radii (R_1 and R_2). Based on this clustering approach, where the ring structures of each group have the same radius, it is feasible to

comply with the limits dictated by the fabrication resolution and the maximum possible thermo-optically induced wavelength tuning of the rings. The 100GHz spacing between the MUX stages as well as the wavelength resonance matching for each pair of RRs forming the MUX's stages are achieved by considering micro-heaters on top of each ring. However, the Si-DLSPP router requires also four different SOI MUXs operating efficiently at different spectral regions within the C telecom band. This requirement can be addressed by choosing appropriate ring radii combinations and values for the bus-to-ring (*gap1*) and ring-to-ring (*gap2*) gaps and for the lengths (*L1* to *L7*) of the inter-ring waveguide sections during the design procedure.

Figure 4. (a) SEM image of a fabricated 2nd order RR. Experimental spectra (blue solid line) at the Drop port for RRs with (b) 9 μm and (c) 12 μm radii. Red dashed lines illustrate the corresponding simulated spectral responses.

Figure 5. Generic layout of a SOI MUX based on 8 cascaded 2nd order thermo-optically tunable RRs that are grouped in two different radii clusters (R_1 and R_2). The sections with H illustrate the heaters controlling the straight inter-ring waveguide lengths.

Following this rationale, the desired spectral responses of the four MUXs are generated within the 1530-1565 nm wavelength window, revealing 100GHz channel spacing, at least -20dB optical crosstalk and 40GHz 3-dB bandwidth, as illustrated in Figure 6. The first three designs rely on rings with 12 μm and 11.7 μm radii for the first and second cluster, respectively, whereas the corresponding ring radii values for the fourth design are 9 μm and

9.2 μm . Moreover, the power coupling coefficients between ring-to-ring and bus-to-ring waveguides for all MUX designs are in the range of 0.33-0.007, corresponding to gap dimensions between 170-380 nm. It should be mentioned that the low crosstalk values in adjacent channels are the result of uneven straight waveguide sections between the successive MUX stages that are long enough ($\sim 100 \mu\text{m}$) also to prevent thermal crosstalk. Since the slightest variation in their length seems to greatly affect the optical crosstalk, micro-heaters are considered to be placed on the middle of every straight waveguide section so as to restore length deviations.

Figure 6. Spectral responses of the (a) first (1533-1540 nm band), (b) second (1540-1547 nm band), (c) third (1547-1554 nm band) and (d) fourth (1557-1563 nm band) 8:1 MUX designs.

2.5. Silicon-plasmonic A-MZIs

Targeting the best compromise between losses, footprint and power consumption, the A-MZI modules of the router rely on a hybrid implementation, where the coupling stages are based on low-loss Si waveguides and only the active phase branches exploit the DLSP structures. The silicon-plasmonic A-MZI structures play a key role in the accomplishment of routing operation. These are the elements of the router that are capable of performing optical switching by taking advantage of the TO effect at the plasmonic parts. The principle of operation regarding the TO switching is the following: In the $500 \times 600 \text{ nm}^2$ PMMA ridge, the DLSP mode fills practically the whole ridge, while the largest fraction of the mode is concentrated at the metal-dielectric interface. The strong confinement at the metal-dielectric interface renders the mode very sensitive to any temperature variation induced through current injection in the metal; in this way the underlying gold film acts as a heating electrode for the dielectric ridge, modifying in a very efficient way the DLSP mode effective index via changing the PMMA refractive index [32]-[34],[52] and requiring only a small amount of consumed power. At the same time, the inherent instantaneous heating of the metal is immediately transferred to the propagating SPP mode, again due to its strong confinement at the metal-dielectric interface, leading to small response times of the TO effect in the DLSP waveguide. Such low driving powers ($\sim 1 \text{ mW}$) and response times ($\sim 1 \mu\text{s}$) in

TO modulation and switching have already been shown in theoretical estimations [32] and calculations [52] that make use of these ridge dimensions in DLSP waveguide structures.

Considering operation at telecom wavelengths ($\sim 1550\text{nm}$) and an increase in temperature of about 61K in PMMA loadings (where the thermo-optic coefficient TOC is $-1.05 \cdot 10^{-4}\text{K}^{-1}$), a full π phase shift can be achieved for a waveguide length of $120\mu\text{m}$, according to the formula $L\pi = \lambda / (2 \cdot \Delta T \cdot TOC)$. This length value is almost twice the maximum propagation length of the DLSP mode and results in high propagation losses [32]. However, this length can be reduced by a factor of 2 by adopting the asymmetric MZI configuration shown in Figure 7(a). In this configuration, a permanent $\pi/2$ phase shift is induced in one of the two branches by slightly widening the PMMA ridge over a short distance, resulting in this way to a higher effective refractive index value in the wider ridge section and to a natural biasing of the MZI at its quadrature point of operation. To this end, the necessary thermo-optically imposed phase shift is reduced down to $\pi/2$. Following this rationale, the length of the phase arms in the employed hybrid A-MZIs has been set to $60\mu\text{m}$ ($L1$), while the $\pi/2$ phase asymmetry was achieved by widening the lower 600-nm-thick DLSP waveguide from 500nm ($W1$) to 700nm ($W2$) for a length of $6\mu\text{m}$ ($L2$), as a result of an increase of the DLSP mode effective index by ~ 0.06 [29],[52]. These dimensions ensure the existence of a single TM mode inside the polymer at telecom wavelengths, both for the nominal and the widened DLSP waveguides, as depicted in Figure 7(b)-(c).

Figure 7. (a) Schematic layout of a Si-DLSP A-MZI. The lower plasmonic branch is widened in order to introduce a default asymmetry, (b) fundamental quasi-TM mode of the $500 \times 600 \text{ nm}^2$ PMMA-loaded SPP waveguide, (c) fundamental quasi-TM mode of the $700 \times 600 \text{ nm}^2$ PMMA-loaded SPP waveguide.

Taking advantage of the default $\pi/2$ phase asymmetry introduced at the lower plasmonic branch, this asymmetric interferometric configuration is capable of performing switching with only a $\pi/2$ phase shift compared to all-plasmonic symmetric MZI switches [52]–[54]. The injection of electric current to a MZI arm yields a temperature rise that alters the effective index of the mode propagating on the heated arm, leading to a phase shift via this TO effect [32],[34],[52]. By electrically controlling the upper A-MZI arm, a negative phase shift is experienced by the propagating DLSP waveguide mode as a result of the PMMA's TOC. When the induced phase shift equals $\pi/2$, the phase difference between the modes travelling through the two MZI branches equals π and therefore the whole mode power is exported to the BAR output of the device. On the contrary, when the same current level applies only to the lower MZI plasmonic branch, the default MZI phase asymmetry is cancelled out due to the $-\pi/2$ thermo-optically induced phase shift and, thus, the whole mode power emerges at the CROSS port of the MZI. To this end, the ON switching

operation is achieved by electrically driving the upper MZI plasmonic arm, whereas the lower branch has to be driven by the same amount of electric current for the OFF switching operation. Nevertheless, due to the default $\pi/2$ phase asymmetry of the MZI, high-performance switching can be reached even if one of the two plasmonic branches is thermo-optically addressed, since the MZI is initially biased at the quadrature point that lies in the linear domain of its output transfer function. Consequently, this asymmetric formation constitutes a simple and passive mechanism for reducing the required energy level and the active plasmonic arm length for a given maximum service temperature.

2.6. WDM data transmission through a DLSP waveguide

The evaluation of the WDM data transmission capabilities of plasmonics is initially carried out via a $60\mu\text{m}$ -long PMMA-loaded straight waveguide that is included in a Si-plasmonic chip [31],[38]. This chip comprises Si rib waveguides, DLSP waveguides and Si-to-DLSP interfaces that follow the specifications presented in Section 2.2 as well as Si TM grating couplers so as to enable optical communication with the outside world. At 1545nm TM-polarized light the cut-back measurements reveal $0.1\text{dB}/\mu\text{m}$ and $4.6\text{dB}/\text{cm}$ plasmonic and silicon propagation losses as well as 2.5dB and 12dB coupling losses for each Si-to-DLSP and TM grating coupling interface, respectively.

Figure 8. Experimental setup consisting of the WDM transmitter, the Si-DLSP waveguide and the receiver.

The experimental setup used to transmit a 480Gb/s WDM optical stream through the DLSP waveguide is presented in Figure 8. The transmitter involves twelve continuous-wave (CW) optical signals stemming from distributed feedback (DFB) lasers and spaced by 200GHz within the $1542\text{--}1560\text{ nm}$ spectral range. The multiplexing of these individual light beams into a single optical fiber is realized by using AWG and 3-dB couplers. The multiplexed signal is then encoded by a $2^{31}-1$ pseudo-random bit sequence (PRBS) at 40Gb/s NRZ line rate in a Ti:LiNbO_3 Mach-Zehnder modulator (MZM). Subsequently, the 480Gb/s WDM signal is amplified by a high-power erbium-doped fiber amplifier (EDFA), providing 24dBm output power towards tackling the high losses induced by the hybrid chip. After ensuring TM polarization conditions for compliance with the inherent TM nature of plasmonics, via a polarization controller, the incoming optical data signal is inserted into the DLSP waveguide. The multi-wavelength data stream that transmitted through the whole Si-DLSP waveguide is amplified at the chip's output by a low-noise EDFA and demultiplexed into individual data channels by an optical bandpass tunable filter (OBPF). Each optical data

channel is then converted into an electrical signal by a 40GHz 3-dB bandwidth photoreceiver that is connected to a 1:4 electrical demultiplexer, where each electrical data signal is received by a 10Gb/s error detector for bit error rate (BER) measurements.

Figure 9. Spectra of the 12×40Gb/s WDM signal (a) at chip’s input, (b) at chip’s output, in comparison with chip’s spectral response and B2B flat losses, (c) after post-chip amplification, (d) BER curves for all 40Gb/s B2B and transmitted channels, (e) BER curves for B2B and transmitted channel 1, (f) BER curves for B2B and transmitted channel 8.

The 12-channel data signal is examined, in terms of its spectrum, in three crucial places along the transmission line: just before the Si-DLSPP chip’s input, after the chip’s output and after the receiver’s EDFA, as depicted in Figure 9(a)-(c). The TM grating couplers of the chip reveal a non-flattened spectral response (red dashed line in Figure 9(b)), dictating thereby the chip’s spectral response and resulting in wavelength-dependent fiber-to-fiber transmission losses which range between 40dB and 48dB within the wavelength window of interest. As a consequence, the spectrum of the WDM signal constantly alters, even after the final amplification stage due to the absence of gain flattened filters (GFFs) at the EDFAs. To this end, a wavelength-dependent performance is observed in the receiver with respect to the optical noise-to-signal ratio (OSNR). The performance of the 480Gb/s WDM transmission through the 60μm-long DLSPP waveguide is evaluated via BER measurements that are shown in Figure 9(d)-(f). In particular, Figure 9(d) illustrates an overview of the BER curves obtained for all 40Gb/s discrete channels. Six out of the twelve channels (channels #1-5, 12) perform error-free, exhibiting a power penalty in a range between 0.2dB and 1dB for a 10^{-9} BER value against back-to-back (B2B) measurements. However, the remaining six channels (channels #6-11) reveal an error-floor at $\sim 10^{-7}$ BER. Figure 9(e) corresponds to the best performing channel #1 that yields 0.2dB power penalty, whilst Figure 9(f) refers to the worst performing channel #8 that reveals an error-floor at 10^{-7} BER value. The B2B measurements are generated by replacing the hybrid chip by a variable attenuator that induces constant, wavelength-independent losses equal to the losses experienced by channels 1-4 that reside in the 1542-1548 nm spectral region (flat reference loss with green line in Figure 9(b)). The

different performances between the channels transmitted via the hybrid chip originate from the unequal OSNR distribution due to the wavelength-dependent spectral responses of the grating couplers and the EDFAs. It should be noted that the error-floor in channels #1-5, 12 can be eliminated by selecting their spectral position to reside within the almost flat low-loss chip's response and by the employment of EDFAs with GFFs.

2.7. WDM data traffic switching with a hybrid Si-plasmonic A-MZI

After verifying the WDM data transmission capabilities of plasmonics through the DLSP waveguide, the next step towards enriching the WDM portfolio of plasmonics is the demonstration of data switching [25]. In this perspective, the TO electrically controlled PMMA-based A-MZI that analyzed in Section 2.5 is employed towards performing switching of a 4×10Gb/s WDM traffic according to the rationale described in the same section. This A-MZI, along with other structures, is included in a Si-plasmonic chip [25] that is equipped with TM grating couplers, similar to the ones presented in the previous section, in order to couple light in and out of the hybrid chip. At 1542nm TM-polarized light the cut-back measurements reveal 0.1dB/μm and 4.4dB/cm plasmonic and silicon propagation losses as well as 2.5dB and 13dB coupling losses for each Si-to-DLSP and TM grating coupling interface, respectively.

Before proceeding to the WDM approach, the Si-DLSP A-MZI is tested in static and single-channel conditions regarding its electrical requirements and switching performance. Figure 10(a) shows the static TO transfer functions for the CROSS and BAR output ports of the A-MZI. The input signal is a 1542nm CW light beam with 6dBm optical power and the control medium is a direct current (DC) that drives the upper MZI arm and takes values up to 40mA. As it is obvious from this figure, the extinction ratio (ER) values for the CROSS and BAR ports are 14dB and 0.9dB when the DC reaches its maximum value. This unbalanced behavior between the two output ports stems from the 95:5, instead of 50:50, Si input/output couplers of the MZI as the result of unfortunate design error. Moreover, by comparing the obtained TO transfer functions with the theoretical transfer functions of a symmetric MZI, the default biasing point of the A-MZI is estimated at ~70°. Following the same procedure, the phase induced by the 40mA DC is ~-90° that corresponds to an increase in temperature by ~60K. These phase conditions in both MZI arms imply a total phase difference of ~160° between the two signal components that travel through the two branches. In terms of power requirements, the hybrid A-MZI consumes ~13.1mW when it is driven by 40mA DC, considering that the resistance of the 60μm-long DLSP phase arm is found to be 8.2Ω.

After the initial static and single-channel switching characterization, the Si-DLSP A-MZI is evaluated in terms of switching under dynamic and WDM data conditions. This evaluation is realized by using the experimental setup of Figure 10(b). Four CW optical signals emitted by DFB sources at 1545.1nm, 1546.7nm, 1547.7nm and 1549.1nm wavelengths are combined in pairs by 3-dB couplers in a way that channels 1 and 3 constitute the first pair and channels 2 and 4 form the other one. Subsequently, each channel pair is modulated by a 2³¹-1 PRBS at 10Gb/s NRZ line rate in a corresponding Ti:LiNbO₃ MZM. The final 4×10Gb/s WDM

signal is formed by using a 3-dB coupler to multiplex the two pairs of data sequences and is amplified by a high-power EDFA that provides 31dBm power at the input of the electrically controlled hybrid A-MZI. A polarization controller is used again towards establishing the required TM polarization conditions of the incoming WDM signal. Exploiting the TO effect, the switching state of the A-MZI is controlled dynamically by a pulse generator operating at 20KHz. After its transmission through the A-MZI, the multi-wavelength data signal is amplified by a two stage EDFA that comprises a 5nm midstage OBPF for out-of-band amplified spontaneous emission (ASE) noise rejection. The amplified WDM stream is then demultiplexed into its constituent channels by a 0.8nm OBPF. Each optical data channel is launched into a photoreceiver with 10GHz 3-dB bandwidth for OE conversion. Subsequently, the electrical signal exiting the photoreceiver is then received by a 10Gb/s error-detector for evaluation via BER measurements.

Figure 10. (a) Static TO transfer functions for the CROSS and BAR output ports of the A-MZI, (b) Experimental setup and the 4-channel spectrum at (c) MZI input, (d) directly at the MZI output before entering EDFA2, (e) after the receiver's pre-amplification stage. The spectral response of the chip including the A-MZI and the TM grating couplers is shown with the red dashed line in (d).

As depicted in Figure 10(c)-(e), the spectrum of the 4-channel data signal is observed at three places along the transmission line, that is just before the MZI's input, after the MZI's output and after the pre-amplifier (EDFA2) in the receiver. The high-power signal propagating in the fiber link between EDFA1 and the A-MZI causes the generation of four-wave mixing (FWM) terms, as it is easily noticeable in these figures. Nevertheless, by selecting unequal channel spacing ranging between 125GHz and 200GHz in the source signals, these terms are out-of-band compared to the data channels, giving the opportunity to be removed by subsequent filters. The energy transfer to the FWM components as well as the absence of GFFs after the transmitter's EDFA form the unequal power profile of the WDM signal at the

MZI's input, as shown in Figure 10(c). Due to the spectral response of the chip's TM grating couplers, the WDM spectral power profile at the MZI's output is now different (Figure 10(d)). The same figure also depicts with red dashed line the spectral response of the chip including the A-MZI and the TM grating couplers, revealing a 3dB loss variation within the 1545-1549 nm spectral band. After the receiver's pre-amplification stage, different power levels and also different OSNR values between the four channels are observed, as a result of the gain profile of EDFA2 (Figure 10(e)).

The dynamic control of the Si-DLSPP A-MZI is achieved by driving its upper arm with electrical rectangular pulses of $15\mu\text{s}$ duration, 20KHz repetition rate and 40mA peak value. The obtained data traces and eye diagrams for channel 1 (λ_1) and channel 2 (λ_2) signals at the CROSS and BAR output ports of the A-MZI are illustrated in Figure 11(a)-(h), where the electrical control signal is also depicted with red dashed line. According to these figures, inverted mode operation with an ER close to 14dB is attained at the CROSS port. In contrast, only 0.9dB ER performance is recorder at the BAR port as a result of the 95:5 Si couplers of the A-MZI. In terms of response time, the Si-DLSPP A-MZI exhibits fast rise and fall times lying in the 3-5 μs range, as it has been also demonstrated in a $90\mu\text{m}$ -long PMMA-based A-MZI [25]. It should be noted that similar results for both CROSS and BAR output ports are also obtained for data channels 3 and 4 at λ_3 and λ_4 wavelengths, respectively.

Figure 11. Modulation with $15\mu\text{s}$ electrical rectangular pulses at 20KHz repetition rate for 10Gb/s: (a) data trace at the CROSS port (channel 1), (b) data trace at the BAR port (channel 1), (c) eye diagram at the CROSS port (channel 1), (d) eye diagram at the BAR port (channel 1), (e) data trace at the CROSS port (channel 2), (f) data trace at the BAR port (channel 2), (g) eye diagram at the CROSS port (channel 2), (h) eye diagram at the BAR port (channel 2), (i) BER curves for ON/OFF operation of the A-MZI.

The performance of the $4\times 10\text{Gb/s}$ WDM switching via the $60\mu\text{m}$ -long PMMA-based A-MZI in ON and OFF switching states is evaluated via BER measurements that are depicted in Figure 11(i). In this experiment, the B2B measurements are obtained by using a straight Si waveguide, which is also included in the hybrid chip, as a reference point. In this way, only the signal degradation through the A-MZI is taken into consideration, excluding the lossy input/output TM grating couplers of the chip. Regarding the ON state, a DC value of 40mA is applied to the upper MZI branch, recording at the same time the BER measurements for the BAR output port. In case of no driving current, namely the OFF state, the BER curve is obtained at the CROSS output port of the MZI. As it is shown in Figure 11(i), the ON state reveals a negative power penalty close to the statistical error of 0.2dB for all four channels compared to the B2B curve. However, when operating at OFF state, the four channels

exhibit power penalties in a range between 1.7dB and 3.6dB at a 10^{-9} BER value. The enhanced power penalty values compared to the corresponding ON state performance originate mainly from the 8dB lower power level received at MZI's CROSS output (-26dBm) against its BAR port during ON state (-18dBm). Additionally, the power penalty shows a clear tendency to increase with the channel wavelength, due to the wavelength-dependent gain and OSNR profile experienced by the four channels during amplification in EDFA2. Finally, the different slopes observed between the BER graphs of channels 1, 3 and the BER curves of channels 2 and 4 are caused due to the different modulators per channel pair at the 4×10Gb/s transmitter end.

2.8. Switching performance metrics

A clear view about the switching credentials and advantages of plasmonics over the SOI and the polymer-based TO waveguide technology platforms can be obtained by comparing performance metrics regarding the power consumption, switching time and active phase arm length. Table 1 provides an overview of the respective performance metrics reported for different SOI and polymer-based TO MZI switches. The power consumption × switching time metric is also included in this table, since it has been commonly used for the comparison of TO switching elements in applications where both low energy and fast operation are a prerequisite for high performance [55]. This table verifies the advantages of plasmonic technology: The Si-DLSPP A-MZI switch attains the shorter active region and the smallest power-time product over every type of undoped thermo-optically addressed MZI switch. These findings come as a natural consequence of the strong confinement of the DLSPP mode at the metal-dielectric interface, the inherent instantaneous heating of the underlying metallic layer and the asymmetric configuration of the MZI. On the contrary, in SOI TO structures, where the electrode is located at certain distance from the Si waveguide, it is very difficult to achieve simultaneously ultra-low power consumption and response time values. The required power can be greatly reduced by using sophisticated waveguide

Active waveguide technology [reference]	Phase arm length (active region in μm)	Power consumption (P in mW)	Switching time (τ in μs)	Power-time product ($P \times \tau$ in $\text{mW} \cdot \mu\text{s}$)
SOI [58]	200	20	2.8	56
SOI [56]	1000	0.49	144	70.56
SOI [57]	100	0.54	141	76.14
SOI [61]	700	50	3.5	175
SOI [62]	6300	6.5	14	91
Polymer [55]	300	1.85	700	1295
Polymer [59]	100	4	200	800
PMMA-loaded SPP [Current work]	60	13.1	3.8	49.78

Table 1. Comparison with other SOI and polymer-based TO MZI switches.

engineering (e.g. suspended waveguides), thereby leading to sub-mW switching but at the expense of increased switching times and poor mechanical stability [56],[57]. A very good compromise between power consumption and response time is offered by using integrated NiSi waveguide heaters in single driving SOI-based TO switches [58]. Compared to the latter switches, lower power requirements but higher switching times are observed in polymer-based switches [55],[59]. It should be mentioned that TO switches that use the waveguide itself as the resistive heater and exploit differential driving schemes [60] have been excluded from this study to keep the same conditions for comparison.

2.9. Optimization procedures

The performance of the Si-DLSPP A-MZI switch regarding its BAR and CROSS output ports is determined by the splitting ratio of its Si coupling stages. To this end, the incorporation of 50:50 couplers, instead of the considered 95:5, can greatly improve the switching performance, yielding more than 20dB ER for both output ports. Another aspect that can be optimized is the switching time of this hybrid interferometric element. As it is already mentioned, the adoption of differential driving schemes for actively controlling both of the DLSPP MZI branches has the potential for lower time responses, even in the order of sub- μ s [60]. Moreover, the power consumption in switching can be also reduced. One way is to form a kind of “suspended” plasmonic waveguides by etching the BOX layer under the gold film in order to achieve better confinement of the heat in the DLSPP waveguides and therefore enhanced power efficiency. This technique could be also exploited in the SOI MUX devices for energy-efficient thermal tuning of the RRs [63]. Another approach for bringing down the consumed power is the utilization of polymer materials with higher TOC value instead of PMMA loadings. For example, Cycloaliphatic acrylate (Cyclomer) exhibits almost three times higher TOC than PMMA and has been already successfully applied as the polymer loading in DLSPP switching structures [34],[38]. Though yet not optimized for low-energy switching operation, this route is expected to significantly decrease the required energy levels without compromising the switching performance. Alternatively, this scheme can be used to reduce the footprint and as such the losses of the A-MZI device. Considering the same amount of electric current to that used in the aforementioned WDM switching experiment, the length of the active plasmonic regions can be decreased from $60\mu\text{m}$ to $20\mu\text{m}$, resulting in 4dB lower plasmonic propagation losses. Besides, the Cyclomer loadings can boost the TO tuning performance of plasmonic resonant devices compared to respective PMMA-based structures [31],[33] when similar temperature changes occur [38].

Apart from the hybrid A-MZI switch, optimization acts should be carried out for the TM grating couplers that are responsible for coupling light in and out of the chips. The grating couplers employed in the WDM data transmission and switching experiments exhibit 24dB overall insertion losses at best that are too high for practical datacom applications where the employment of amplifiers is prohibited. Towards tackling this issue, new TM grating couplers are considered for a far better efficiency. These coupling structures, filled with Spin-on-Glass (SOG) of 800nm height, rely on a fully etched approach with 0.8 filling factor,

0.71 μm grating period, 0.13 μm groove width and 10 degrees incident angle of the light. Figure 12(a) presents a layout of the new grating coupler's schematic cross section and Figure 12(b) depicts the SEM image of the fabricated one. These optimized TM grating couplers reveal 3.25dB minimum coupling loss at 1557nm and about 32nm 3-dB bandwidth. The blue line in Figure 12(c) presents also the results obtained from simulation that indicate very good agreement between theory and experiment. These results ensure that the new TM grating coupler is a major achievement since it decreases the router's overall optical losses by 17.5dB.

Figure 12. TM grating coupler: (a) Schematic cross sectional view, (b) SEM image, (c) Spectral responses of the simulation and experimental results.

2.10. Router performance in terms of optical insertion losses and power consumption

As a NoC solution for Data Centers and HPC systems, the hybrid Router-on-Chip platform should consume low power towards "greener" computing environments and induce low losses towards intrachip point-to-point connection without using on-chip amplifiers.

Within this framework, an estimation of the total power consumption of the entire 4×4 silicon-plasmonic router is provided here based on the characterization and experimental results obtained from the router's components and subsystems. Considering 60 μm -long PMMA-phase arms in the branches of the hybrid A-MZIs that require 40mA driving current, the hybrid router is expected to consume at maximum ~2W. This power consumption stems from 1mW in the four PDs, 335mW in the four frontend SOI MUX devices due to the thermal tuning of their RRs and ~1.67W in the IC that is capable of providing at the same time two differential outputs with 48mA current. Therefore, the electronic IC is the router's component that consumes the largest part of power whereas the optical components reveal a power consumption of only ~400mW, considering the 13mW power requirement of each hybrid A-MZI. In comparison with current multicore central

processing units (CPUs) for servers and HPCs that consume 150W [64] and 55W [65] respectively, the hybrid router requires only 2W, rendering this technology ideal for such shrunk networking environments in terms of energy savings. Keeping in mind that the Si-DLSPP router has an aggregate throughput capability of 1.12Tb/s, the power efficiency of this hybrid routing platform is about 1.8mW/Gb/s that is a very low value compared to other hybrid OE routers [66].

The optimization procedures of the previous section could lead to significant reductions in the total power consumption of the silicon-plasmonic router. The usage of Cyclomer, instead of PMMA, loadings in the phase arms of the A-MZIs can bring the IC's power consumption down to ~600mW. Consequently, the total power consumption is limited to ~1W. Keeping the power consumption in the original level of 2W, the same switching performance in the 4×4 DLSPP switching matrix is expected by employing 20μm-long Cyclomer-based A-MZIs but with the profit of 8dB lower losses in the router. Towards minimizing the aggregate power consumption, the approach of suspended arms for the fabrication of the RRs composing the SOI MUXs can lead to increased tuning efficiency that results in ~30mW for the whole multiplexing circuitry and therefore in ~630mW for the entire routing platform. As a result, an outstanding performance with lower than 0.6mW/Gb/s power efficiency can be achieved.

The total insertion losses of the silicon-plasmonic router can be also estimated. The aggregate losses across an input-to-output route are expected to be ~31dB that are analyzed in 6.5dB for the input and output grating couplers, 2dB for the SOI MUX device, 0.5dB for a 90:10 monitor coupler, ~8dB for the silicon waveguides, 10dB for the four Si-to-DLSPP interfaces and 4dB for the two 20μm-long Cyclomer-loaded SPP branches at the A-MZIs. However, the losses concerning the Si waveguide interconnections can be quite lower after finalizing the router prototype, since the value of 8dB relies on first layouts, where the building blocks and their interconnections are not placed in an optimal -in terms of spatial density- way. Nevertheless, even the 31dB losses of the original design can be manageable with high-power vertical-cavity surface-emitting lasers (VCSELs) [67] and high sensitivity Si-PDs [68].

3. Conclusion

Power consumption and size appear as the main set of barriers in next-generation Data Center and HPC environments. Within this framework, the penetration of optics into shrunk networking environments for chip-scale interconnects is now more vital than ever. SOI technology has already demonstrated its low-cost and high-integration credentials towards supporting fundamental operations required for optical interconnect applications. At the same time, the emerging discipline of plasmonics appears as a promising candidate for further size minimization and power savings. To this end, the idea of merging plasmonics and silicon photonics into the same platform seems to be a great solution for the implementation of faster and “greener” NoC environments. Deriving the strengths of each technology, a high-throughput, energy efficient and compact hybrid “Router-on-chip” is

feasible. In this perspective, we present a 4×4 silicon-plasmonic router architecture for chip-scale applications with 1.12Tb/s aggregate bandwidth. The hybrid router relies on a SOI motherboard that incorporates Si waveguides for low-loss interconnection of SOI MUXs, PDs and the remaining subsystems, low-energy plasmonic waveguides in active A-MZI switches and an IC control unit for intelligent decision-making operations. Towards the demonstration of the router's functionality, we proceed to experimental evidences that turn the promises of plasmonics into real system-level application benefits: The transmission of a 12×40Gb/s WDM stream through a Si-DLSPP waveguide and the switching of a 4×10Gb/s WDM signal via a Si-DLSPP A-MZI. With the lowest reported power consumption × time response product among undoped SOI and polymer-based TO MZI switches, plasmonic technology becomes suitable for on-chip optical interconnects. According to experimental results, the hybrid router is estimated to consume ~2W with only ~400mW corresponding to the optical parts, yielding a power efficiency of 1.8mW/Gb/s. Moreover, the router is expected to induce about 31dB optical losses. With these power and optical loss characteristics, the Si-DLSPP routing platform seems to be appropriate for NoC environments where fast path establishment and route reconfiguration is necessary for efficient traffic management in Data Centers and HPCs.

Author details

Sotirios Papaioannou

Corresponding Author

Informatics & Telematics Institute, Center for Research & Technology Hellas, Thessaloniki, Greece

Department of Informatics, Aristotle University of Thessaloniki, Thessaloniki, Greece

Konstantinos Vyrsoinos

Informatics & Telematics Institute, Center for Research & Technology Hellas, Thessaloniki, Greece

Dimitrios Kalavrouziotis, Giannis Giannoulis,

Dimitrios Apostolopoulos and Hercules Avramopoulos

School of Electrical & Computer Engineering, National Technical University of Athens, Athens, Greece

Filimon Zacharatos, Karim Hassan, Jean-Claude Weeber,

Laurent Markey and Alain Dereux

Institut Carnot de Bourgogne, University of Burgundy, Dijon, France

Ashwani Kumar and Sergey I. Bozhevolnyi

Institute of Sensors, Signals, and Electrotechnics, University of Southern Denmark, Odense M, Denmark

Alpaslan Suna, Oriol Gili de Villasante and Tolga Tekin

Fraunhofer IZM, Berlin, Germany

Michael Waldow

AMO Gesellschaft für Angewandte Mikro- und Optoelektronik GmbH, Aachen, Germany

Odysseas Tsilipakos, Alexandros Pitilakis and Emmanouil E. Kriezis
*Department of Electrical & Computer Engineering, Aristotle University of Thessaloniki,
 Thessaloniki, Greece*

Nikos Pleros
*Informatics & Telematics Institute, Center for Research & Technology Hellas, Thessaloniki, Greece
 Department of Informatics, Aristotle University of Thessaloniki, Thessaloniki, Greece*

Acknowledgement

This work was partially supported by the European FP7 ICT-PLATON (ICT- STREP no. 249135) project.

4. References

- [1] Miller D. A. B. Device Requirements for Optical Interconnects to Silicon Chips. Proceedings of the IEEE 2009; 97(7) 1166-1185.
- [2] Qian D. et al. 101.7-Tb/s (370×294-Gb/s) PDM-128QAM-OFDM Transmission over 3×55-km SSMF using Pilot-based Phase Noise Mitigation. Proceedings of OFC/NFOEC 2011, Paper PDPB5, 6-10 March 2011, Los Angeles, CA, USA.
- [3] Taira Y. et al. High Channel-Count Optical Interconnection for Servers. Proceedings of ECTC 2010; 5490959, 1-4 June 2010, Las Vegas, Nevada, USA.
- [4] Kash J. A. et al. Optical Interconnects in Exascale Supercomputers. Proceedings of IEEE Photonics Society 2010, 483-484, 7-11 November 2010, Denver, CO, USA.
- [5] http://www.igigroup.net/download/FFS09/0930_Benner_IBM.pdf.
- [6] <http://www.datacenterknowledge.com/archives/2011/06/20/new-top-500-champ-the-k-supercomputer/>.
- [7] Taubenblatt M. A. Optical Interconnects for High-Performance Computing. IEEE/OSA Journal of Lightwave Technology 2012; 30(4) 448-458.
- [8] Paniccia M. Integrating Silicon Photonics. Nature Photonics 2010; 4(8) 498-499.
- [9] Barnes W. L., Dereux A., Ebbesen T. W. Surface plasmon subwavelength optics. Nature 2003; 424(6950) 824-830.
- [10] Atwater H. A. The promise of plasmonics. Scientific American 2007; 296(4) 38-45.
- [11] Brongersma M. L., Shalaev V. M. The case for plasmonics. Science 2010; 328(5977) 440-441.
- [12] Zia R., Schuller J. A., Chandran A., Brongersma M. L. Plasmonics: the next chip-scale technology. Materials Today 2006; 9(7-8) 20-27.
- [13] Papaioannou S. et al. A 320Gb/s-throughput capable 2×2 silicon-plasmonic router architecture for optical interconnects. IEEE/OSA Journal of Lightwave Technology 2011; 29(21) 3185-3195.
- [14] Kim J. T. et al. Chip-to-chip optical interconnect using gold long-range surface plasmon polariton waveguides. Optics Express 2008; 16(17) 13133-13138.

- [15] Raether H. Surface Plasmons on Smooth and Rough Surfaces and on Gratings. Berlin: Springer-Verlag; 1988.
- [16] Gramotnev D. K., Bozhevolnyi S. I. Plasmonics beyond the diffraction limit. *Nature Photonics* 2010; 4(2) 83-91.
- [17] Grandidier J. et al. Surface plasmon routing in dielectric-loaded surface plasmon polariton waveguides. *Proceedings of SPIE* 2008, 7033, 70330S-70330S-8, 10-14 August 2008, San Diego, CA, USA.
- [18] Krasavin A. V., Zheludev N. I. Active plasmonics: Controlling signals in Au/Ga waveguide using nanoscale structural transformations. *Applied Physics Letters* 2004; 84(8) 1416-1418.
- [19] Nikolajsen T., Leosson K., Bozhevolnyi S. I. Surface plasmon polariton based modulators and switches operating at telecom wavelengths. *Applied Physics Letters* 2004; 85(24) 5833-5835.
- [20] Lereu A. L., Passian A., Goudonnet J.-P., Thundat T., Ferrell T. L. Optical modulation processes in thin films based on thermal effects of surface plasmons. *Applied Physics Letters* 2005; 86(15) 154101.
- [21] Pacifici D., Lezec H. J., Atwater H. A. All-optical modulation by plasmonic excitation of CdSe quantum dots. *Nature Photonics* 2007; 1(7) 402-406.
- [22] Pala R. A., Shimizu K. T., Melosh N. A., Brongersma M. L. A nonvolatile plasmonic switch employing photochromic molecules. *Nano Letters* 2008; 8(5) 1506-1510.
- [23] MacDonald K. F., Sámson Z. L., Stockman M. I., Zheludev N. I. Ultrafast active plasmonics. *Nature Photonics* 2009; 3(1) 55-58.
- [24] Krasavin A. V., Zayats A. V. Electro-optic switching element for dielectric-loaded surface plasmon polariton waveguides. *Applied Physics Letters* 2010; 97(4) 041107.
- [25] Kalavrouziotis D. et al. First demonstration of active plasmonic device in true data traffic conditions: ON/OFF thermo-optic modulation using a hybrid silicon-plasmonic asymmetric MZI. *Proceedings of OFC/NFOEC* 2012, OW3E.3, 4-8 March 2012, Los Angeles, CA, USA.
- [26] Bozhevolnyi S. I., Erland J., Leosson K., Skovgaard P. M. W., Hvam J. M. Waveguiding in Surface Plasmon Polariton Band Gap Structures. *Physical Review Letters* 2001; 86(14) 3008-3011.
- [27] Knight M. W., Grady N. K., Bardhan R., Hao F., Nordlander P., Halas N. J. Nanoparticle-mediated coupling of light into a nanowire. *Nano Letters* 2007; 7(8) 2346-2350.
- [28] Bozhevolnyi S. I., Volkov V. S., Devaux E., Ebbesen T. W. Channel plasmon-polariton guiding by subwavelength metal grooves. *Physical Review Letters* 2005; 95(4) 046802.
- [29] Holmgaard T., Bozhevolnyi S. I. Theoretical analysis of dielectric-loaded surface plasmon-polariton waveguides. *Physical Review B* 2007; 75(24) 245405.
- [30] Briggs R. M., Grandidier J., Burgos S. P., Feigenbaum E., Atwater H. A. Efficient coupling between dielectric loaded plasmonic and silicon photonic waveguides. *Nano Letters* 2010; 10(12) 4851-4857.

- [31] Giannoulis G. et al. Data transmission and thermo-optic tuning performance of dielectric-loaded plasmonic structures hetero-integrated on a silicon chip. *IEEE Photonics Technology Letters* 2012; 24(5) 374-376.
- [32] Gosciniaik J. et al. Thermo-optic control of dielectric-loaded plasmonic waveguide components. *Optics Express* 2010; 18(2) 1207-1216.
- [33] Hassan K., Weeber J.-C., Markey L., Dereux A. Thermo-optical control of dielectric loaded plasmonic racetrack resonators. *Journal of Applied Physics* 2011; 110(2) 023106.
- [34] Hassan K. et al. Thermo-optic plasmo-photonic mode interference switches based on dielectric loaded waveguides. *Applied Physics Letters* 2011; 99(24) 241110.
- [35] Hassan K. et al. Characterization of thermo-optical 2×2 switch configurations made of dielectric loaded surface plasmon polariton waveguides for telecom routing architecture. *Proceedings of OFC/NFOEC 2012, OW3E.5, 4-8 March 2012, Los Angeles, CA, USA.*
- [36] Randhawa S. et al. Performance of electro-optical plasmonic ring resonators at telecom wavelengths. *Optics Express* 2012; 20(3) 2354-2362.
- [37] Tsilipakos O. et al. Interfacing Dielectric-Loaded Plasmonic and Silicon Photonic Waveguides: Theoretical Analysis and Experimental Demonstration. *IEEE Journal of Quantum Electronics* 2012; 48(5) 678-687.
- [38] Kalavrouziotis D. et al. 0.48Tb/s (12×40Gb/s) WDM transmission and high-quality thermo-optic switching in dielectric loaded plasmonics. *Optics Express* 2012; 20(7) 7655-7662.
- [39] Vivien L., Osmond J., Fédéli J., Marris-Morini D., Crozat P., Damlencourt J., Cassan E., Lecunff Y., Laval S. 42 GHz p.i.n Germanium photodetector integrated in a silicon-on-insulator waveguide. *Optics Express* 2009; 17(8) 6252-6257.
- [40] Sheng Z., Liu L., Brouckaert J., He S., Van Thourhout D. InGaAs PIN photodetectors integrated on silicon-on-insulator waveguides. *Optics Express* 2010; 18(2) 1756-1761.
- [41] Bravo-Abad J., Ippen E. P., Soljačić M. Ultrafast photodetection in an all-silicon chip enabled by two-photon absorption. *Applied Physics Letters* 2009; 94(24) 241103.
- [42] Geis M. W. et al. Silicon waveguide infrared photodiodes with >35 GHz bandwidth and phototransistors with 50 AW-1 response. *Optics Express* 2009; 17(7) 5193-5204.
- [43] Bogaerts W. et al. Silicon-on-Insulator Spectral Filters Fabricated With CMOS Technology. *IEEE Journal of Selected Topics in Quantum Electronics* 2010; 16(1) 33-44.
- [44] Geng M. et al. Compact four-channel reconfigurable optical add-drop multiplexer using silicon photonic wire. *Optics Communications* 2009; 282(17) 3477-3480.
- [45] Barwicz T. et al. Reconfigurable silicon photonic circuits for telecommunication applications. *Proceedings of SPIE* 2008; 6872, 68720Z-1-12.
- [46] Klein E. J. et al. Reconfigurable optical add-drop multiplexer using microring resonators. *IEEE Photonics Technology Letters* 2005; 17(11) 2358-2360.
- [47] Dong P. et al. Low power and compact reconfigurable multiplexing devices based on silicon microring resonators. *Optics Express* 2010; 18(10) 9852-9858.
- [48] Xiao S., Khan M. H., Shen H., Qi M. Multiple-channel silicon micro-resonator based filters for WDM applications. *Optics Express* 2007; 15(12) 7489-7498.

- [49] Dahlem M. S. et al. Reconfigurable multi-channel second-order silicon microring-resonator filterbanks for on-chip WDM systems. *Optics Express* 2011; 19(1) 306-316.
- [50] Liu A. et al. Wavelength Division Multiplexing Based Photonic Integrated Circuits on Silicon-on-Insulator Platform. *Journal of Selected Topics in Quantum Electronics* 2010; 16(1) 23-32.
- [51] Saeung P., Yupapin P. P. Generalized analysis of multiple ring resonator filters: Modeling by using graphical approach. *Optik-Int Journal for Light and Electron Optics* 2008; 119(10) 465-472.
- [52] Ptilakis A., Kriezis E. E. Longitudinal 2×2 switching configurations based on thermo-optically addressed dielectric-loaded plasmonic waveguides. *IEEE/OSA Journal of Lightwave Technology* 2011; 29(17) 2636-2646.
- [53] Krasavin A. V., Zayats A. V. Three-dimensional numerical modeling of photonic integration with dielectric-loaded SPP waveguides. *Physical Review B* 2008; 78(4) 045425.
- [54] Tsilipakos O., Ptilakis A., Tasolamprou A. C., Yioultsis T. V., Kriezis E. E. Computational techniques for the analysis and design of dielectric-loaded plasmonic circuitry. *Optical and Quantum Electronics* 2011; 42(8) 541-555.
- [55] Al-Hetar A. M., Mohammad A. B., Supa'at A. S. M., Shamsan Z. A. MMI-MZI polymer thermo-optic switch with a high refractive index contrast. *Journal of Lightwave Technology* 2011; 29(2) 171-178.
- [56] Fang Q. et al. Ultralow power silicon photonics thermo-optic switch with suspended phase arms. *IEEE Photonics Technology Letters* 2011; 23(8) 525-527.
- [57] Sun P., Reano R. M. Submilliwatt thermo-optic switches using free-standing silicon-on-insulator strip waveguides. *Optics Express* 2010; 18(8) 8406-8411.
- [58] Van Campenhout J., Green W. M., Assefa S., Vlasov Y. A. Integrated NiSi waveguide heaters for CMOS-compatible silicon thermo-optic devices. *Optics Letters* 2010; 35(7) 1013-1015.
- [59] Xie N., Hashimoto T., Uta K. Very low-power, polarization-independent, and high-speed polymer thermooptic switch. *IEEE Photonics Technology Letters* 2009; 21(24) 1861-1863.
- [60] Geis M. W., Spector S. J., Williamson R. C., Lyszczarz T. M. Submicrosecond Submilliwatt Silicon-on-Insulator Thermooptic Switch. *IEEE Photonics Technology Letters* 2004; 16(11) 2514-2516.
- [61] Espinola R. L., Tsai M.-C., Yardley J. T., Osgood R. M. Fast and low-power thermooptic switch on thin silicon-on-insulator. *IEEE Photonics Technology Letters* 2003; 15(10) 1366-1368.
- [62] Densmore A. et al. Compact and low power thermo-optic switch using folded silicon waveguides. *Optics Express* 2009; 17(13) 10457-10465.
- [63] Dong P. et al. Thermally tunable silicon racetrack resonators with ultralow tuning power. *Optics Express* 2010; 18(19) 20298-20304.
- [64] [http://ark.intel.com/products/64582/Intel-Xeon-Processor-E5-2687W-\(20M-Cache-3_10-GHz-8_00-GTs-Intel-QPI\)](http://ark.intel.com/products/64582/Intel-Xeon-Processor-E5-2687W-(20M-Cache-3_10-GHz-8_00-GTs-Intel-QPI)).
- [65] https://www.power.org/home/Blue_Gene_Q_Super_Efficient_Linley.pdf.

- [66] Takahashi R. et al. Hybrid Optoelectronic Router for Future Energy-efficient, Large-capacity, and Flexible OPS Networks. Proceedings of IEEE Photonics Society Summer Topical Meeting Series 2011; 153-154, 18-20 July 2011, Montreal, QC, Canada.
- [67] Hofmann W. et al. 44 Gb/s VCSEL for optical interconnects. Proceedings of OFC/NFOEC 2011, Paper PDPB5, 6-10 March 2011, Los Angeles, CA, USA.
- [68] Geis M. W. et al. CMOS-Compatible All-Si High-Speed Waveguide Photodiodes With High Responsivity in Near-Infrared Communication Band. IEEE Photonics Technology Letters 2007; 19(3) 152-154.