# Demonstration of a multiplane OAM-wavelength packet switch controlled by a two-step scheduler implemented in FPGA

M. N. Malik, J. C. Borromeo, M. Scaffardi, F. Scotti, N. Zhang, C. Klitis, M. Lavery, G. Preve, V. Toccafondo, R. Reyes, P. Castoldi, M. Sorel, A. Bogoni, N. Andriolli

***Abstract*— In order to meet the demand for highly scalable and energy efficient data center switches, multiplane architectures, where multiple optical switching domains are leveraged, are an attractive solution. In this scenario, the orbital angular momentum (OAM) of light can be exploited as an additional domain along with wavelength, space, and time. In this work we report the demonstration of an OAM-wavelength photonic integrated packet switch controlled by an FPGA implementing a two-step scheduler suited for multiplane architectures. Packet-based bit error rate (BER) measurements show transmission at 20 Gb/s with BER < $10^{-9}$ also for the most critical switch configurations, with penalty < 2 dB with respect to the back-to-back. The measurements demonstrate that the hardware-controlled switch can support up to 32 ports distributed in 4 cards with 8 ports each.**

***Index Terms*— Networks, packet-switched, Optical vortices**

## I. Introduction

THE data center switch scalability is highly affected by a vigorously growing data center traffic [1]. In this context, optical interconnection networks based on optical switching can help to realize high capacity, energy efficient, and reconfigurable interconnection networks [2]-[5]. To this aim, many optical interconnection network architectures have been studied, typically exploiting a single switching technology (or domain), e.g., space, wavelength, time, etc. [6]-[8]. Pioneering studies were performed in the past on this topic [9]-[12], pertaining to the field of optical packet switching and optical label switching, and involving typically bulk demonstrators. In these so-called single-plane architectures, each switch port can communicate with each other either on a designated link, or using a specific wavelength, or through an allocated time interval. Recently the use of orbital angular momentum (OAM) of light as an additional switching/multiplexing domain has attracted a lot of interest in order to increase the switch capacity [13]. However, single-plane architectures are inherently limited in the maximum number of ports and throughput by the physical or technological constraints typical of each domain (e.g., optical signal to noise ratio, optical bandwidth, time compression efficiency, OAM tuning range of the OAM emitters), effectively limiting their performance and scalability. Scalability of single-plane switch is also impaired by the superlinear (quadratic in [8]) growth of elementary switches necessary to make up a large switch.

Multiplane architectures were conceived to jointly leverage multiple optical switching domains. Multiplane architectures are organized in cards, each one with multiple ports: ports in a card are addressed using one domain, while cards are addressed using another domain. Each domain contributes to the switching through a reduced number of interconnections with respect to the single-plane architectures, thus overcoming the scalability limitations of each domain alone. Space-wavelength and space-time [14]-[16] are examples of multiplane architectures investigated in the past years.

We have recently proposed to exploit the OAM as a switching domain in multiplane switches together with wavelength [17]. An OAM beam of order $l$ has an azimuthal phase term $\exp(i \cdot l \cdot \phi)$, where $\phi$ is the azimuthal angle and $l$ is the topological charge (with an integer value) of the OAM state. OAM beams with different OAM order can propagate together ideally without crosstalk owing to the property of being orthogonal [18].

The use of photonic integrated devices for the generation, multiplexing and demultiplexing of OAM beams can help to realize compact, stable and energy efficient OAM switches. OAM beams can be generated using bulky passive devices like spiral phase plates (SPPs) [19] or active devices like spatial light modulators (SLMs) [20]. Compact solutions for OAM generation include integrated microrings [21], circular grating couplers cascaded to star coupler [22], and hybrid 3D integrated circuits [23].

Manuscript received December 13, 2018. This study has been supported by the Department of Science Technology - Engineering Research and Development Technology (DOST-ERDT) of the Philippines.

M.N. Malik and Antonella Bogoni are with Scuola Superiore Sant'Anna, 56124 Pisa, Italy and also with CNIT, 56124 Pisa, Italy (e-mail: 08bicsenmalik@seecs.edu.pk, antonella.bogoni@santannapisa.it).

J. C. Borromeo and R. Reyes are with Ateneo de Manila University, Katipunan Ave., Quezon City, 1108 Metro Manila, Philippines (e-mail: justine.borromeo@obf.ateneo.edu, rsjreyes@ateneo.edu).

M. Scaffardi, F. Scotti, V. Toccafondo and G. Preve are with CNIT, 56124 Pisa, Italy (e-mail: {mirco.scaffardi, filippo.scotti, veronica.toccafondo, gbpreve}@cnit.it).

N. Zhang, C. Klitis, M. Lavery, and M. Sorel are with the University of Glasgow, Glasgow G12 8LT, U.K. (e-mail: {ning.zhang, charalambos.klitis, martin.lavery, marc.sorel}@glasgow.ac.uk).

P. Castoldi and N. Andriolli are with Scuola Superiore Sant'Anna, 56124 Pisa, Italy (e-mail: {piero.castoldi, nicola.andriolli}@santannapisa.it).

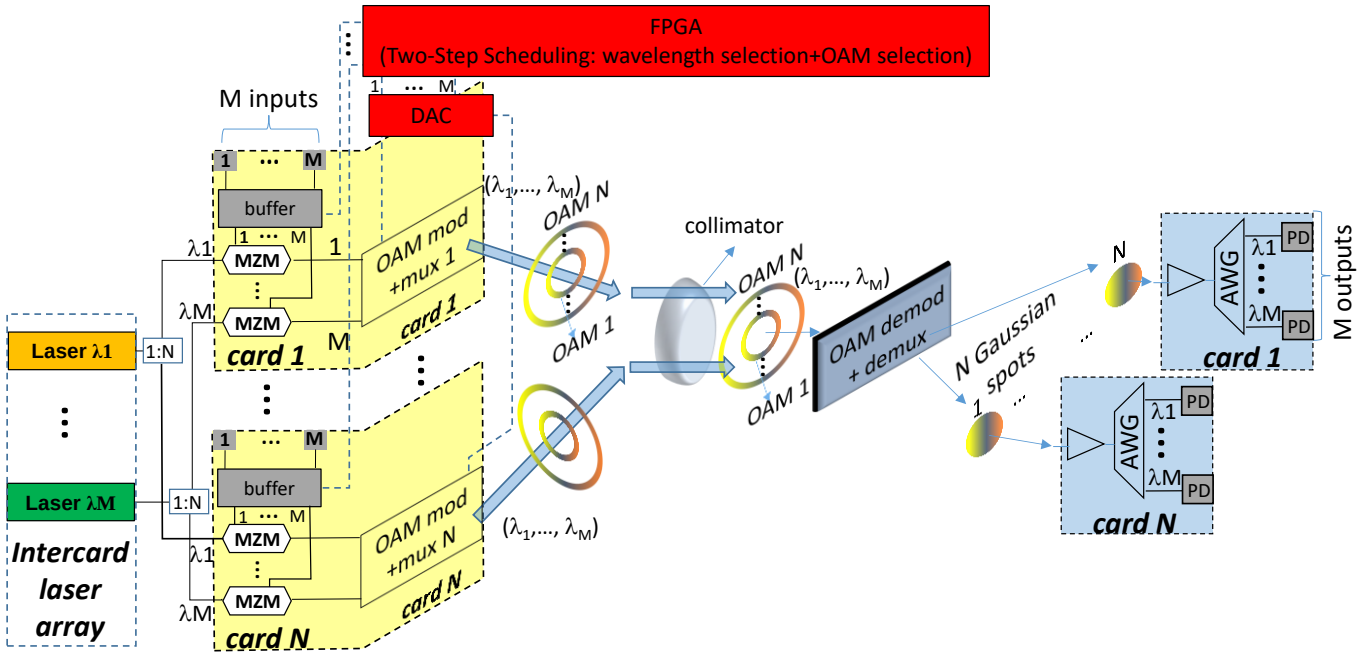


Fig. 1: Multiplane switch architecture exploiting OAM and wavelength as switching domains. MZM: Mach-Zehnder Modulator; DAC: digital-to-analog conversion; AWG: arrayed waveguide grating; PD: photodetector.

A multiplane OAM-wavelength switch exploiting an integrated OAM modulator/multiplexer and a refractive element-based OAM demodulator/demultiplexer was demonstrated in [24]. Nevertheless, in that demonstration only circuit switching was implemented. In [25] the performance of a two-step scheduler (TSS) specifically designed for the multiplane architectures, e.g., the OAM-wavelength switch, was assessed through simulation. TSS operation is composed of two sequential steps, the first one devoted to the intra-card scheduling, and the second one performing the inter-card scheduling [26].

In this paper we demonstrate, for the first time to the best of our knowledge, a multiplane OAM-wavelength packet switch based on an integrated OAM modulator/multiplexer with packet-based slotted operation under FPGA control. Such a switch is modular, i.e. composed of cards, each of them hosting fixed number of ports. OAM modes are used to address the cards, while wavelengths are used to address the ports on each card. The TSS scheduler has been implemented in FPGA to schedule the packet transmissions and reconfigure the switch accordingly. We report the results in terms of packet-based bit error rate (BER) measurements on the OAM-wavelength switch supporting 4 OAM modes (i.e., cards) and 8 wavelengths (i.e., ports per card), showing packet transmission with BER < $10^{-9}$. Preliminary measurements on a smaller switching architecture with half the number of ports were described in [27]. The manuscript is organized as follows. Sec. II describes the OAM-wavelength based switch architecture. Sec. III introduces the TSS scheduler and its FPGA implementation. Sec. IV describes the experimental demonstration of the OAM-wavelength switch with FPGA control for different configurations. The conclusions are drawn in Sec. V.

## II. Multiplane Architectures and OAM-Wavelength Based Switch

The proposed OAM-wavelength switch is shown in Fig. 1. The architecture is composed of *N* cards, each of them equipped with M ports, hence, the whole interconnection network hosts MN ports. For ease of visualization, the transmitting and receiving sides are reported on the left and right hand sides of the figure, respectively. The destination cards are addressed by the OAM domain, while the destination ports on each card are addressed by the wavelength domain. The choice of the packets to be switched, the wavelength to use (i.e., the destination port), and the destination card is performed by a scheduler that operates synchronously, i.e., at each time slot a fixed-length packet can be transmitted from each input port. At the transmitter side of each card an electronic buffer stores the incoming packets.

The core element on the transmitting side of each card is the photonic integrated OAM modulator/multiplexer (OAM mod+mux) based on concentric vortex emitters [24], which converts multiple Gaussian beams, i.e., beams propagating in standard fibres, into OAM beams of selectable order and simultaneously multiplexes them. On each card, packets to be delivered to a specific card-port destination are optically modulated on one of the *M* wavelengths ($\lambda 1$, ...., $\lambda M$) using *M* Mach-Zehnder modulators (MZM). The MZM outputs are connected to *M* inputs the of OAM modulator/multiplexer, which converts each of the *M* signals from Gaussian spatial distribution into one of the *N* OAM modes (*OAM 1*, ..., *OAM N*). The OAM modulator/multiplexer has been implemented using accurately designed silicon photonic integrated waveguides, which emit the OAM beam of the desired order in an orthogonal direction with respect to the

waveguide plane. The multiplexed OAM beams from different cards are further multiplexed by means of a free space beam combiner, propagate in free-space to reach the OAM demodulator/demultiplexer which is used to spatially separate the OAM beams of different order (destined to different switch cards) and convert them back to Gaussian. The OAM demodulator/demultiplexer is implemented using two passive cascaded refractive elements [28][29]. On each destination card, the different wavelengths are then separated using an arrayed waveguide grating (AWG). Each AWG output is connected to a photodiode (PD) which converts the signal into the electrical domain.

## III. TSS SCHEDULER

The modular OAM-wavelength switch is controlled by the two-step scheduler (TSS), suited for multiplane switching architectures, as depicted in Fig. 2. Scheduling computations are performed at each time slot in parallel (independently in each card for the first step and in each switching layer in the second step, as sketched in Fig. 2) and most of the information is kept local on each card. On each card, the intra-card scheduler assigns a unique destination port $p_i$ to each input port, that is, decides the transmission wavelength for each input port on the card. Then, the inter-card scheduler (running in parallel on each wavelength layer) maps each port - assigned to a given wavelength - to a card, hence deciding the setting of all OAM mod+mux to avoid contention at the output ports.

For this demonstration an FPGA-based implementation of the TSS has been developed. Different scheduling algorithms (e.g., iSLIP or LQF [30][31]) can be implemented in each step, trading off scheduling complexity with switch throughput, as detailed next.

### A. *iSLIP Algorithm*

iSLIP algorithm operates based on the concept of parallel iterative matching. It aims at finding a maximal size matching based uniquely on the queue status (whether the queues are empty or not) [30]. Every iteration of iSLIP is composed of three phases that can be performed in parallel on the ports. In the first phase, each unmatched input port sends a request to each output port for which it has a queued packet. In the second phase, each unmatched output port receiving requests selects one of them. The selection is performed by scanning the port indices in a round-robin manner. Then, the corresponding input ports are notified. In the third phase, the input port accepts the granted request from the output port. If an input port receives one or more grants, it accepts one of them again in a round-robin way.

### B. *LQF Algorithm*

An iterative and parallel implementation of the longest-queue-first (LQF) maximal matching algorithm is exploited [31], where the scanning is performed in parallel among all input and output ports simultaneously. Each iteration consists of a request phase and an accept phase. In the request phase, each input (output) port selects the highest weighed output (input) port as a candidate for matching. The requests are then matched in the accept phase. In the subsequent iterations, the same phases are repeated solely on the unmatched ports. Both phases can be efficiently executed in hardware by resorting to matrices storing weights and matching requests.

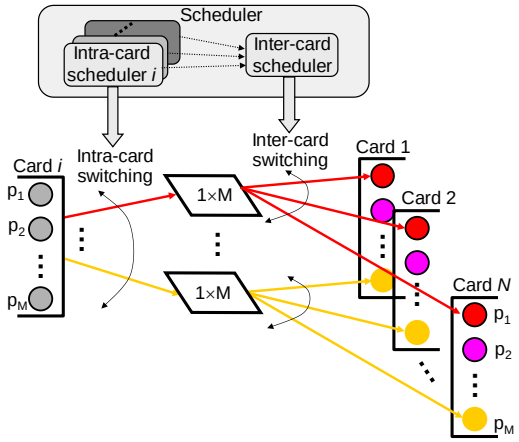


Fig. 2: Scheduler and switching operations in a multiplane architecture.

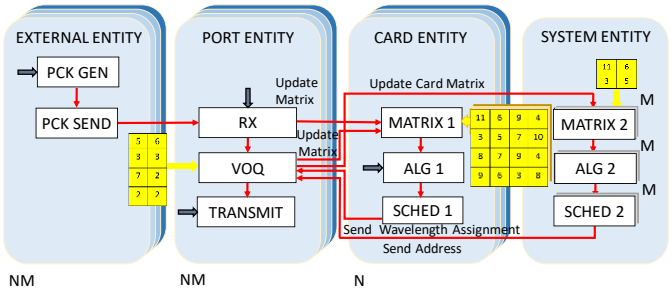


Fig. 3: Hardware design language entities and internal modules of the TSS FPGA implementation.

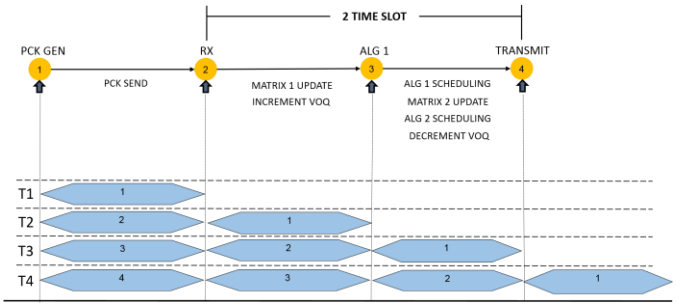


Fig. 4: Timing operation of the TSS.

### C. *FPGA Implementation of TSS*

The FPGA-based implementation of TSS, shown in Fig. 3, is composed of four main entities: external entity, port entity, card entity and system entity. The grey arrows in Fig. 3 represent the modules that are activated at the beginning of every time slot.

The external entity generates fixed-size packets and sends them to the corresponding port entity. Packets are created based on a Bernoulli process with uniform destination distribution. An 8-bit header containing the output port-card destination address is also generated.

After generating the packets, the RX module of the port entity receives and stores them in the virtual output queue (VOQ) buffers. Every VOQ module is composed of *NM* buffers used to store the packets that are waiting for transmission. Each VOQ buffer can accommodate up to 512

packets. The data packets inside these buffers will be transmitted in a first-in-first-out (FIFO) order after executing the two scheduling steps using the TRANSMIT module, responsible also for sending the switch control signals.

The card entity is responsible for the intra-card scheduling step. It is performed in parallel for each of the $N$ cards and returns the port assignment of each input port. The MATRIX 1 module arranges the queue lengths extracted from the VOQ modules in an $M \times M$ matrix. The matrix elements are incremented (decremented) every time a packet arrives to the VOQ (leaves the VOQ and is transmitted). The ALG 1 module executes the intra-card scheduling step on the matrix representation of the card VOQs, as detailed in [31]. After completing the first scheduling step, the SCHED 1 module translates the assignments into port addresses and sends them back to the VOQ modules, where the $N \times N$ card matrices, to be used for the inter-card scheduling step, are extracted and sent to the MATRIX 2 modules. The inter-card step is carried out by the system entity in parallel on the $M$ internal modules, one per wavelength. The ALG 2 modules compute the inter-card schedules and the SCHED 2 modules are used to collect, translate, and send the port and card addresses of the winning packets to the VOQ modules.

### *D. Timing Operation*

The scheduler is aimed to a synchronous switch, operating on fixed-duration time slots. For this reason, different scheduler operations are pipelined and synchronously triggered. More specifically, the beginning of a timeslot triggers the packet generation, matrix and VOQ update, scheduling, and transmission, as highlighted by the grey arrows in Figs. 3 and 4. The top of Fig. 4 shows the pipelined scheduling process, while the bottom details the time evolution in the first four timeslots (from T1 to T4). During the first timeslot, the external entities are active and generate packets while the other entities have nothing to process. In the second timeslot, the packets created in the previous timeslot are passed to the RX modules and queued in the VOQ modules. The MATRIX 1 modules are also updated in this timeslot. In the third time slot, besides the previously described actions, the matrices updated in the previous timeslot are used for the intra-card scheduling step performed by ALG 1 modules. Then, based on the first step scheduling decisions, MATRIX 2 modules are updated, ALG 2 modules perform the inter-card scheduling step, and SCHED 2 modules collect the results and trigger the update of the VOQ modules based on the final scheduling decisions. In the fourth timeslot, besides the previously detailed actions, the scheduled packets are transmitted and the process continues with the described pipeline.

### *E. Latency Assessment through Emulation*

TSS has been implemented in VHDL using Quartus 13.0 and emulated with waveforms at 50 MHz clock frequency using ModelSim-Altera 10.1d. Since the present demonstrator targets a 32-port OAM-wavelength packet switch with 4 cards and 8 ports, the TSS latency performance is assessed on the same architecture. Both iSLIP and LQF algorithms are used in the first and second step (i.e., ALG 1 and ALG 2), leading to four TSS implementations. The average latency, i.e., the average number of time slots spent by packets in the queue before transmission, is reported in Fig. 5 as a function of the load, i.e., the probability that a packet reaches a switch input port in a given time slot. The figure shows that LQF-LQF achieves the lowest latency for all loads, due to the fact that LQF prioritizes the transmissions from the heavily loaded VOQs. The latency performance of LQF-iSLIP and iSLIP-LQF is similar, with the former one operating slightly better at lower loads. Finally, among the four TSS implementations, iSLIP-iSLIP suffers the highest latency because it considers only the VOQ occupation and not its length.

Therefore, in the following experimental demonstration, LQF is used as the scheduling algorithm in both intra-card and inter-card scheduling steps.

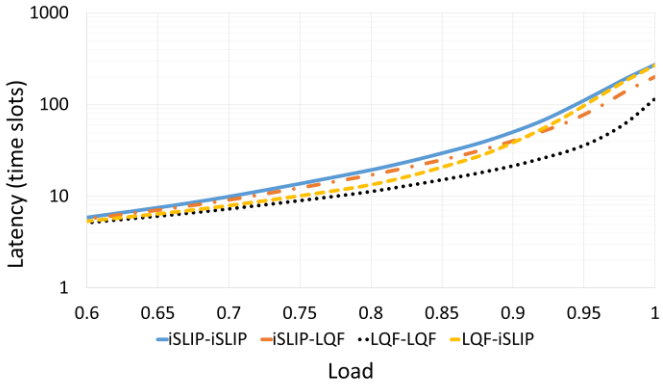


Fig. 5: Average latency vs. load of different combinations of intra-card and inter-card algorithm in TSS.

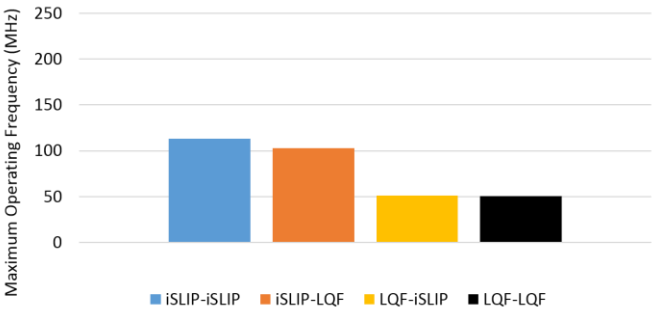


Fig. 6: Maximum operating frequency of different configurations of TSS.

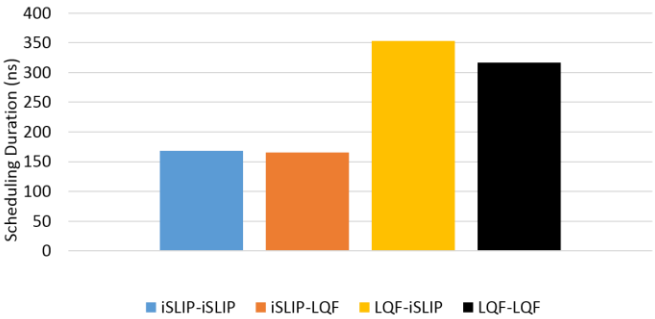


Fig. 7: Scheduling duration of different configurations of TSS.

### *F. Hardware Performance*

The four hardware implementations of 8Port-4Card TSS with different scheduling algorithms have been evaluated in

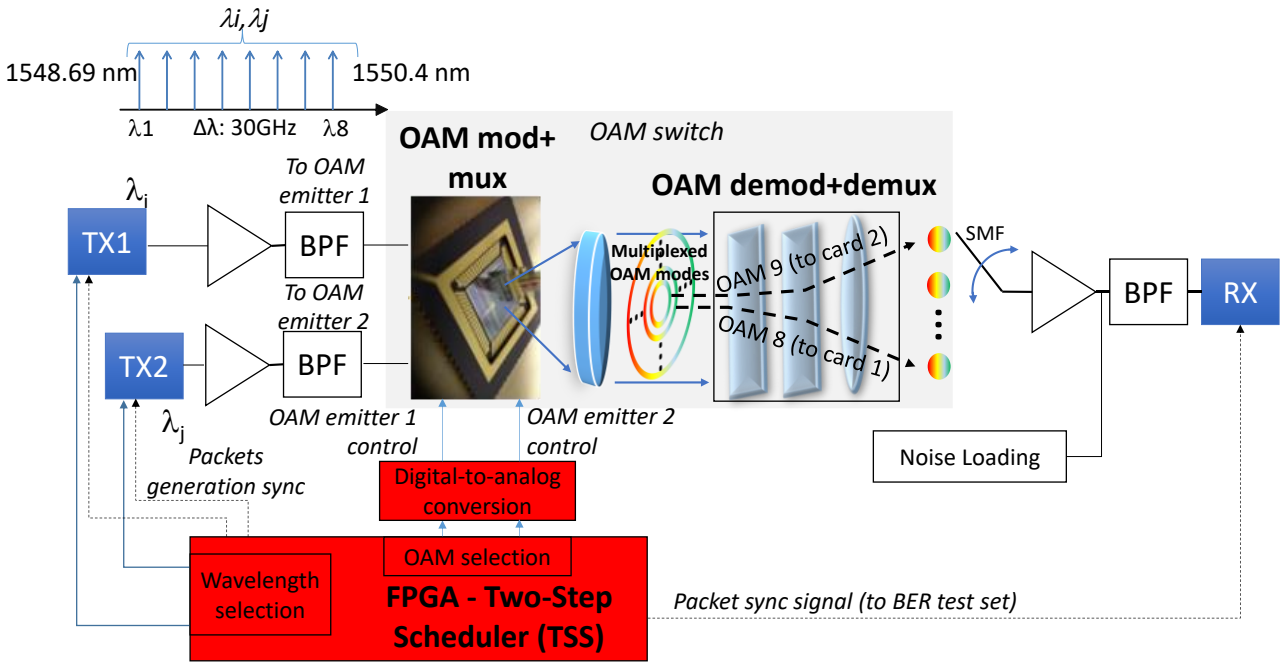


Fig. 8: Experimental setup for the implementation of FPGA controlled OAM-wavelength switch. TX: transmitter; BPF: band pass filter; RX: receiver.

terms of maximum operating frequency and scheduling duration.

Fig. 6 shows the maximum operating frequency of the investigated implementations. The graph shows that iSLIP-iSLIP (LQF-LQF) achieves the highest (lowest) maximum operating frequency since iSLIP exploits only binary logic, which requires less combinational logic elements in every clock cycle. Furthermore, the use of iSLIP in the first step boosts the attainable frequency when LQF is used in the other step.

The scheduling duration can be computed by dividing the worst-case number of clock cycles to complete both steps by the maximum operating frequency. On an interconnection network with $I$ ports, iSLIP needs at most $3log2(I) + 2$ clock cycles, while LQF requires up to $I + 2$ clock cycles [31]. Fig. 7 shows the scheduling duration of the investigated 8Port-4Ccard TSS implementation. The fastest schedulers are iSLIP-iSLIP and iSLIP-LQF due to the higher maximum operating frequency sizes, however LQF-LQF is still very fast as the scheduling duration is around 320 ns.

## IV. Implementation of the OAM-wavelength packet switch with FPGA control

In this section we report the implementation of the FPGA-controlled OAM-wavelength packet switch suitable for 4-OAM (i.e. cards), 8-wavelength (i.e., ports per card) operation.

### A. Experimental Setup

The experimental setup is shown in Fig. 8. Two 20Gb/s on-off keying (OOK) transmitters (TX) are exploited to generate, upon FPGA control, streams of optical packets of duration 100 μs. A guard time of 20 μs is inserted to accommodate for switch reconfiguration time which is about 10 μs in the 10%-90% output signal amplitude variation interval [32]. To enable BER measurements, packet data is filled with a pseudorandom bit sequence (PRBS) $2^{31}$-1. The input wavelengths are chosen from a comb of 8 wavelengths (λ1, …, λ8, see inset of Fig. 8), allowing for the addressing of 8 ports per card, with a separation of 30 GHz (0.24 nm), which is adequate to support 20 Gb/s WDM channels. The packet streams are amplified, filtered, and then fed to two ports of the OAM-wavelength switch, which is composed by the cascade of an OAM mod+mux and an OAM demod+demux. The OAM mod+mux is implemented with an integrated silicon-on-insulator (SOI) chip with concentric OAM emitters [24]. Each OAM emitter generates OAM beams orthogonal to the chip plane. The chip is packaged into a 16mm×16mm square ceramic chip carrier with bonded electrical wires for OAM modes thermal tuning. The tuning power is 18mW per OAM mode. The integrated OAM multiplexer total loss is 13 dB, 7 dB due to the emission efficiency and 6 dB due to the fibre-to-waveguide coupling. The vertically emitted OAM beams are directed to a compact OAM demod+demux, implemented using two cascaded refractive element followed by a lens [28][29]. The demultiplexer loss is 3 dB.

The packet-based operation of the multiplane switch is enforced by an Altera DE2-115 FPGA. This FPGA hosts the two-step scheduling algorithm, commands the transmitters (TX1, TX2), and drives the OAM mod+mux assigning the proper voltage to the OAM emitters in order to tune the OAM order of the emitted mode (i.e., to select the card, as order 8 is employed to reach card 1 and order 9 to reach card 2).

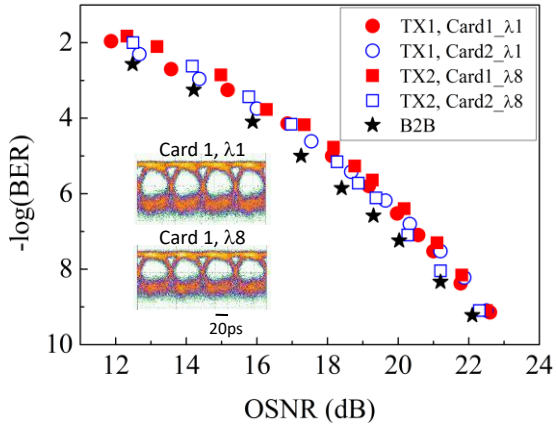


Fig. 9: BER vs. OSNR for packet stream at λ1 and λ8 for destination Card1 and Card2.

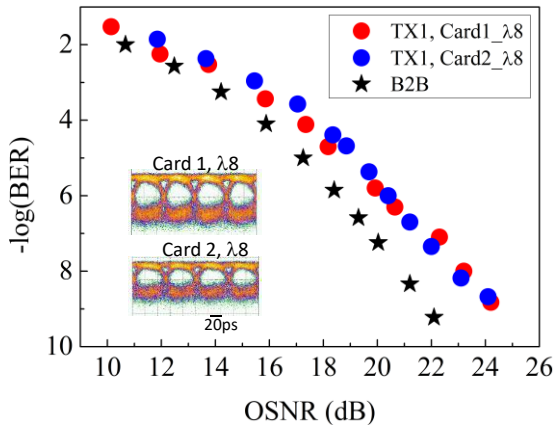


Fig. 10: BER vs. OSNR for packet stream at same wavelengths (λ8) for Card1 and Card2.

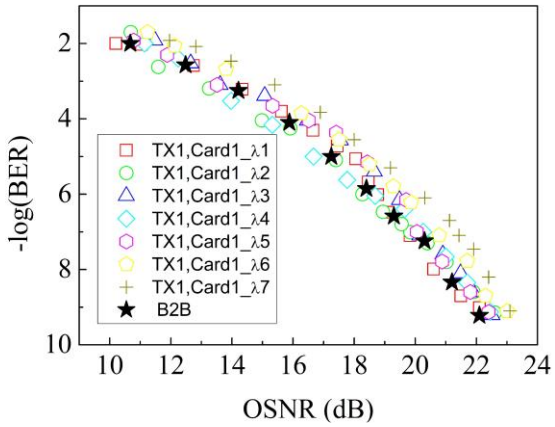


Fig. 11: BER vs. OSNR for packet stream of TX1 at different wavelengths (λ1-λ7), while TX2 is set at λ8.

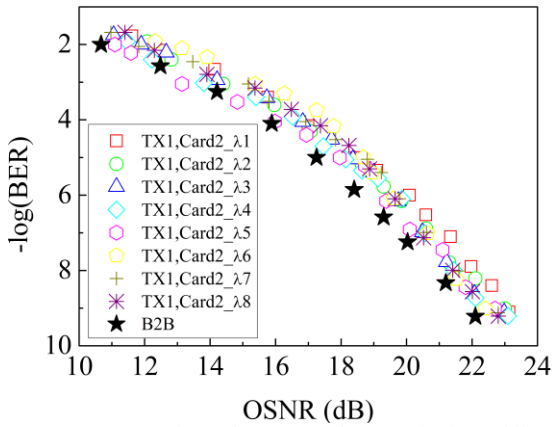


Fig. 12: BER vs. OSNR for packet stream of TX1 at (λ1-λ7), while TX2 is set to same wavelengths as TX1, but at different OAM order.

In order to preserve the integrated OAM multiplexer from thermal breakdown risks, the DAC output voltage was set to limit the OAM tuning over two modes. However, tuning up to 4 OAM modes has already been demonstrated in [24], thus enabling a 4-card switch. The signal at the output of the OAM demod/demux is filtered with a tunable band pass filter, in order to properly select the wavelength. The receiver is composed by a 40 GHz bandwidth photoreceiver followed by a bit error rate tester gated by a packet synchronization signal coming from the FPGA. The power of the received packets is equalized for all the streams coming from the different transmitters. Since the spatial position of the signal at the OAM demod+demux output depends on the order of the selected OAM mode (card), the fiber position at the RX input is set accordingly to the OAM mode under test.

### *B. Performance*

In a first experiment only two wavelengths were considered (λ1 and λ8), implementing two main switch configurations. In the first configuration, the packets from TX1 and TX2 are set at two different wavelengths, thus addressed to different destination ports, while the OAM order (i.e., the destination card) of the packet streams is changed according to the decisions of the FPGA-based scheduler. The performance is shown in Fig. 9. The BER for the packets at destination card 1 and 2 show that the performance is similar for both the considered output cards and input ports (TX1, TX2), with a penalty with respect to B2B < 1 dB. In the second configuration, the worst case scenario was analyzed, in which the packets from both transmitters (TX1 and TX2) are set at the same wavelengths (λ8), while again the OAM order of the packets streams from both ports is controlled by the FPGA according to the computed packet schedule. The measurements are shown in Fig. 10 for card 1 and card 2. The BER curves of the signal coming from TX1 shows a penalty with respect to B2B < 2 dB, with an increase with respect to previous case due to the crosstalk induced by the packet stream coming from TX2 at the same wavelength, but at different OAM order since the scheduler prohibits output port collisions. The eye diagrams of the received signals at the switch output, shown in the insets in Fig. 9 and Fig. 10, look open.

In a second experiment, the number of considered wavelengths was increased to 8. The performance measured for the packets transmitted from TX1 to card 1 are shown in Fig. 11 for the different wavelengths (i.e., destination port). The packets from TX2 are set at λ8. The BER is almost similar for all the wavelengths (ports), with a penalty with respect to B2B < 1 dB. the cascade of integrated OAM multiplexer and OAM demultiplexer operates a filtering function. The cascade of OAM multiplexer and demultiplexer works as an adapted receiver, improving the BER OSNR penalty, as reported also in [33][34]. The measurements in the

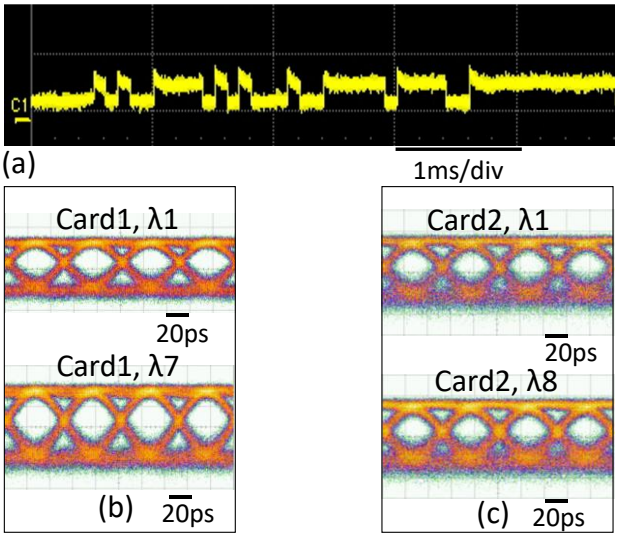


Fig. 13: (a) TX1 output packet stream at λ1 for card1 (b) Eye diagrams for TX1 at λ1 and λ7 while TX2 set to λ8 (c) Eye diagrams for TX1 at λ1 and λ8 while TX2 set to same wavelengths as TX1.

switch configurations where both the transmitters (TX1 and TX2) are at the same wavelength are shown in Fig. 12. Also in this case the FPGA-based scheduler sets the OAM order for the packets from TX1 and TX2 at different values (i.e., destination cards) to avoid output port collisions. The BER curves, plot for the destination card 2, show a penalty < 2 dB. This penalty is slightly higher than the one of the first configuration due to the crosstalk induced by the packet streams from TX2, now set at the same wavelength of TX1. The crosstalk comes from the not ideal OAM mode purity after generation [32] and the suboptimal OAM demultiplexing. Fig. 13(a) shows the output packet stream for TX1 (card 1, λ1) taken with a real-time oscilloscope. The eye diagrams of the signal at the output of TX1 for the first configuration (different wavelengths for TX1 and TX2) is shown in Fig. 13(b) for λ1 and λ7. Fig. 13(c) shows the eye diagram for the second configuration (same wavelengths for TX1 and TX2) at λ1 and λ8. The eyes in Fig. 13(c) show a higher noise due to the crosstalk.

The switch bandwidth limitation is mainly given by the integrated OAM multiplexer, since the OAM demultiplexer is composed of just two refractive elements, i.e., 3D engineered poly-methyl-methacrylate (PMMA) surfaces [29], thus covering the whole C-band. Nevertheless OAM multiplexers with a wider bandwidth, e.g., > 1 nm, can be utilized [35].

The OSNR penalty is 2 dB when the signal crosses twice the switch in a loopback configuration. This suggests that the OAM-wavelength switch can be cascaded a few times, and thus employed in a fat-tree data center network architecture.

## V. Conclusions

In this work we demonstrated an FPGA-controlled multiplane switch architecture exploiting OAM and wavelength as switching domains.

The switch is implemented by cascading an integrated silicon tunable OAM modulator+multiplexer with a refractive element-based OAM demodulator+demultiplexer. For this demonstration the two-step scheduler framework was implemented in FPGA, assessed through emulation, validated with real-time packet operations, and then utilized in the packet-based switching experiments.

The proposed FPGA-controlled OAM-wavelength switch was tested for different configurations taking into consideration different destination cards (OAM modes) and different ports (wavelengths). Successful transmission of packet streams at 20 Gb/s with bit error rate < $10^{-9}$ without exploitation of forward-error-correction codes has been observed for all the configurations, demonstrating the feasibility of a 32-port switch with 8 ports and 4 cards. Since Gaussian signals are present at the input and the output ports of the switch, they can be coupled to single mode fibres, making the switch suitable to be interfaced to standard networking devices.

The presented results demonstrate that an optical switch exploiting the OAM of light and wavelength as switching domains can support packet traffic and can be driven by commercial FPGAs. The switch could be employed in data-center architectures in parallel to electronic switches to route relatively large (e.g., elephant) traffic flows. In order to move the switch towards practical applications further improvements in the packaging are needed, i.e., reducing the distance from the integrated OAM multiplexer and demultiplexer by exploiting microlenses for beam collimation, and using a fibre array at the OAM demultiplexer output to collect all the demultiplexed beams.

In the proposed OAM-wavelength switching architecture the number of integrated devices scales linearly with the number of ports, thus being competitive with respect to single-plane integrated implementations that scale with the square of the number of ports. The reduction of the number of integrated elementary switches limits the technological issues in terms of fabrication yield, making the solution practical also for a large number of switch ports.

## References

[1] J. Shuja, K. Bilal, S. Madani, M. Othman, R. Ranjan, P. Balaji, and S. Khan, "Survey of Techniques and Architectures for Designing Energy-Efficient Data Centers," IEEE Syst. J. 10, 507-519, June 2016.

[2] H. Cho, P. Kapur, and K.C. Saraswat, "Power comparison between high-speed electrical and optical interconnects for interchip communication," J. Lightwave Technol., v. 22, n. 9, p. 2021, 2004.

[3] D. A. B. Miller, "Device requirements for optical interconnects to silicon chips," Proc. IEEE, vol. 97, no. 7, pp. 1166–1185, July 2009.

[4] H. Cho, P. Kapur, and K. Saraswat, "Power comparison between high-speed electrical and optical interconnects for interchip communication," J. Lightwave Technol., vol. 22, no. 9, pp. 2021–2033, Sept. 2004.

[5] A. Benner, "Cost-effective optics: enabling the exascale roadmap," in Proc. 17th IEEE Symp. on High Perf. Interconn. (HOTI), pp. 133-137 (2009).

[6] http://www.polatis.com/series-7000-384x384-port-software-controlled-optical-circuit-switch-sdn-enabled.asp.

[7] A. Wonfor, H. Wang, R. V. Penty, and I. H. White, "Large Port Count High-Speed Optical Switch Fabric for Use Within Datacenters", IEEE J. Opt. Comm. Netw., vol. 3, no. 8, p. A32, 2011.

[8] K. Tanizawa, K. Suzuki, M. Toyama, M. Ohtsuka, N. Yokoyama, K. Matsumaro, M. Seki, K. Koshino, T. Sugaya, S. Suda, G. Cong, T. Kimura, K. Ikeda, S. Namiki, and H. Kawashima, "32×32 Strictly Non-

Blocking Si-Wire Optical Switch on Ultra-Small Die of 11×25 mm$^2$", in OFC Tech. Dig., paper M2B.5 (2015).

[9] C. Guillemot et al., "Transparent optical packet switching: The European ACTS KEOPS project approach," J. Lightw. Technol., vol. 16, no. 12, pp. 2117-2134, Dec. 1998.

[10] L. Xu, H. G. Perros, and G. Rouskas, "Techniques for Optical Packet Switching and Optical Burst Switching," IEEE Commun. Mag., pp. 136-142, Jan. 2001.

[11] Z. Pan et al., "Advanced optical-label routing system supporting multicast, optical TTL, and multimedia applications," J. Lightw. Technol., vol. 23, no. 10, pp. 3270-3281, Oct. 2005.

[12] M. Funabashi, Z. Pan, Z. Zhu, and S. J. B. Yoo, "Nanosecond Guard Time Packet-by-Packet Burst-Mode Optical 3R Regeneration in an Optical-Label Switching Router," IEEE Photon. Technol. Lett., vol. 18, no. 9, pp. 1091-1093, May 2006.

[13] H. Huang, G. Xie, Y. Yan, N. Ahmed, Y. Ren, Y. Yue, D. Rogawski, M. J. Willner, B. I. Erkmen, K. M. Birnbaum, S. J. Dolinar, M. P. J. Lavery, M. J. Padgett, M. Tur, and A. E. Willner, "100 Tbit/s free-space data link enabled by three-dimensional multiplexing of orbital angular momentum, polarization, and wavelength", Opt. Lett. 39, 197-200, 2014.

[14] R. Gaudino, G. Castillo, F. Neri, and J. Finochietto, "Can simple optical switching fabrics scale to terabit per second switch capacities?" J. Opt. Commun. Netw., vol. 1, no. 3, pp. B56–B69, Aug. 2009.

[15] O. Liboiron-Ladouceur, I. Cerutti, P. G. Raponi, N. Andriolli, and P. Castoldi, "Energy-efficient design of a scalable optical multiplane interconnection architecture," IEEE J. Sel. Top. Quantum Electron., vol. 17, no. 2, pp. 377–383, Mar./Apr. 2011.

[16] O. Liboiron-Ladouceur, P. G. Raponi, N. Andriolli, I. Cerutti, M. Shafiqul Hai, and P. Castoldi, "Scalable Space–Time Multi-plane Optical Interconnection Network Using Energy- Efficient Enabling Technologies [Invited]", J. Opt. Comm. Netw., vol. 3, n. 8, pp. A1-A11, Aug. 2011.

[17] M. Scaffardi, M. N. Malik, E. Lazzeri, G. Meloni, F. Fresi, L. Potì, N. Andriolli, I. Cerutti, C. Klitis, L. Meriggi, M. Sorel, and A. Bogoni, "A Silicon Microring Optical 2x2 Switch Exploiting Orbital Angular Momentum for Interconnection Networks up to 20Gbaud", J. Lightwave Technol., 35, 3142-3148, 2017.

[18] J. Wang, J.-Y. Yang, I. M. Fazal, N. Ahmed, Y. Yan, H. Huang, Y. Ren, Y. Yue, S. Dolinar, M. Tur and A. E. Willner, "Terabit free-space data transmission employing orbital angular momentum multiplexing", Nature Photonics, vol. 6, pp. 488-496, 2012.

[19] W. Xuli, "Orbital Angular Momentum Multiplexing in 0.1-THz Free-Space Communication via 3D Printed Spiral Phase Plates", Proc. CLEO, paper Stu2F.2 (2014).

[20] J. Wang, J.-Y. Yang, I. M. Fazal, N. Ahmed, Y. Yan, H. Huang, Y. Ren, Y. Yue, S. Dolinar, M. Tur, and A. E. Willner, "Terabit free-space data transmission employing orbital angular momentum multiplexing", Nat. Photonics, v. 6, p. 488, 2012.

[21] Q. Xiao, C. Klitis, S. Li, Y. Chen, X. Cai, M. Sorel, and S. Yu., "Generation of photonic orbital angular momentum superposition states using vortex beam emitters with superimposed gratings", Optics Express, 24, 3168 (2016).

[22] N. K. Fontaine, C. R. Doerr, and L. L. Buhl, "Efficient multiplexing and demultiplexing of free-space orbital angular momentum using photonic integrated circuits", in OFC Tech. Dig., paper OTu1I.2 (2012).

[23] C. Quin, B. Guan, R. Scott, R. Proietti, N. K. Fontaine, T. Su, C. Ferrari, M. Cappuzzo, F. Klemens, B. Keller, M. Earnshaw, and S. J. B. Yoo, "Demonstration of Orbital Angular Momentum State Conversion using Two Hybrid 3D Photonic Integrated Circuits", in OFC Tech. Dig., paper Th4A (2014).

[24] N. Zhang, M. Scaffardi, M. N. Malik, V. Toccafando, C. Kliti, M. P. J. Lavery, G. Meloni, F. Fresi, E. Lazzeri, D. Marini, J. Zhu, X. Cai, S. Yu, L. Poti, G. Preve, A. Bogoni, M. Sorel, "4 OAM × 4 WDM Optical Switching Based on an Innovative Integrated Tunable OAM Multiplexer", in OFC Tech. Dig., paper Th3H.1 (2018).

[25] M. Scaffardi, N. Andriolli, M. N. Malik, N. Jhang, E. Lazzeri, C. Klitis, M. P. J. Lavery, M. Sorel, A. Bogoni, "Multiplane Orbital Angular Momentum and Wavelength Switch based on Integrated Tunable Vortex Emitters", in OFC Tech. Dig., paper Th2A.44 (2018).

[26] P. G. Raponi, N. Andriolli, I. Cerutti, and P. Castoldi, "Two-step scheduling framework for space-wavelength modular optical interconnection networks," IET Commun., vol. 4, no. 18, pp. 2155–2165, Dec. 2010.

[27] J. Borromeo, M. Malik, N. Andriolli, N. Zhang, C. Klitis, M. Lavery, G. Preve, V. Toccafondo, R. Reyes, P. Castoldi, M. Sorel, A. Bogoni, M. Scaffardi, "First Demonstration of an FPGA-Controlled Multiplane OAM-Wavelength Packet Switch", in OFC Tech. Dig., paper W2A.35 (2019).

[28] G. C. G. Berkhout, M. Lavery, J. Courtial, M. W. Beijersbergen, and M. J. Padgett, "Efficient sorting of orbital angular momentum states of light," Physical Review Lett., 105, 153601 (2010).

[29] M. P. J. Lavery, D. J. Robertson, G. C. G. Berkhout, G. D. Love, M. J. Padgett, and J. Courtial, "Refractive elements for the measurement of the orbital angular momentum of a single photon," Opt. Express 20, 2110-2115 (2012).

[30] N. McKeown, "The iSLIP scheduling algorithm for input-queued switches," IEEE/ACM Trans. Netw., vol. 7, no. 2, 1999.

[31] I. Cerutti, J. A. Corvera, S. M. Dumlao, R. Reyes, P. Castoldi, and N. Andriolli, "Simulation and FPGA-based implementation of iterative parallel schedulers for optical interconnection networks," IEEE/OSA J. Opt. Commun. Netw., vol. 9, no. 4, pp. C76–C87, Apr. 2017.

[32] M. J. Strain, X. Cai, J. Wang, J. Zhu, D. B. Phillips, L. Chen, M. L. Garcia, J. L. O'Brien, M. G. Thompson, M. Sorel, S. Yu, "Fast electrical switching of orbital angular momentum modes using ultra-compact integrated vortex emitters," Nat. Commun., 5, 4856 (2014).

[33] A. Parini, G. Bellanca, A. Annoni, F. Morichetti, A. Melloni, M. J. Strain, M. Sorel, M. Gay, C. Pareige, L. Bramerie, and M. Thual, "BER evaluation of a passive SOI WDM router," IEEE Photon. Technol. Lett. 25(23), 2285–2288 (2013).

[34] F. Gambini, P. Pintus, S. Faralli, M. Chiesa, G. B. Preve, I. Cerutti, and N. Andriolli, "Experimental demonstration of a 24-port packaged multi-microring network-on-chip in silicon photonic platform," Opt. Express 25, 22004-22016 (2017).

[35] Q. Xiao, C. Klitis, H. Chen, S. Li, Y. Chen, J. Zhu, X. Cai, M. Sorel, and S. Yu, "A coaxially integrated photonic orbital angular momentum beam multiplexer", in Proc. 13th Int. Conf. on Group IV Photonics (GFP), pp. 114-115 (2016).