Research Article

Antonio La Porta, Jonas Weiss*, Roger Dangel, Daniel Jubin, Norbert Meier, Folkert Horst and Bert Jan Offrein

Broadband and scalable optical coupling for silicon photonics using polymer waveguides

https://doi.org/10.1515/aot-2017-0064
Received October 4, 2017; accepted January 24, 2018; previously published online February 21, 2018

Abstract: We present optical coupling schemes for silicon integrated photonics circuits that account for the challenges in large-scale data processing systems such as those used for emerging big data workloads. Our waveguide based approach allows to optimally exploit the on-chip optical feature size, and chip- and package real-estate. It further scales well to high numbers of channels and is compatible with state-of-the-art flip-chip die packaging. We demonstrate silicon waveguide to polymer waveguide coupling losses below 1.5 dB for both the O- and C-bands with a polarisation dependent loss of <1 dB. Over 100 optical silicon waveguide to polymer waveguide interfaces were assembled within a single alignment step, resulting in a physical I/O channel density of up to 13 waveguides per millimetre along the chip-edge, with an average coupling loss of below 3.4 dB measured at 1310 nm.

Keywords: coupling; integration; optics; packaging; silicon photonics.

1 Introduction

Emerging workloads for cognitive applications and machine and deep learning on large and unstructured data, among other, require compute systems with very low latency access to very large memory pools and high-bandwidth node-to-node communication [1]. Whether implemented in on premise data centres or in the cloud, the node-local bandwidth and density requirements start to stress traditional, copper-based interconnect schemes. At very high frequencies and data rates, they suffer from attenuation, signal crosstalk and bandwidth-limitation induced intersymbol interference. Though mitigation schemes are available, they introduce inherent power- and channel-density penalties [2]. Optical transmission technologies have long solved most of these limitations for long distance data communication and are currently also heavily deployed in large computing systems to connect between data centres, individual racks and down to drawer-to-drawer interconnections [3]. To facilitate the next step and directly optically connect to the package of the compute engines, optical technologies are required that fit well into the existing assembly and packaging ecosystem and provide a clear path towards tighter integration and scaling in the future, similar to what copper-based interconnect technology did up to today. For this, various optical technologies were proposed [4–6]. Among them, integrated silicon (Si) photonics provides the tightest integration between the optical functions and electronics on a single Si die, at a competitive cost by being fully compatible with complementary metal-oxide semiconductor (CMOS) technology. While the necessary Si photonics building blocks are available, a remaining challenge is system-level assembly and optical connectorisation in a way to support very high numbers of optical I/Os. Direct fibre-to-chip coupling has been demonstrated; however, it constrains existing packaging schemes and is density-limited by the mechanical dimensions of the fibres employed. We overcome these limitations by optically coupling the Si photonics chip into polymer waveguides (PWGs), which can be routed at high densities to the Si die chip-edges.

In this paper, we give a brief overview of system-level interconnect requirements in Section 2, will discuss different possible integration schemes in Section 3, elaborate on and show test results of our Si waveguide-to-PWG coupling technology in Section 4 and summarise in Section 5.

***Corresponding author: Jonas Weiss,** IBM Zurich Research Laboratory, Ruschlikon, Switzerland, e-mail: jwe@zurich.ibm.com
Antonio La Porta, Roger Dangel, Daniel Jubin, Norbert Meier, Folkert Horst and Bert Jan Offrein: IBM Zurich Research Laboratory, Ruschlikon, Switzerland

2 Interconnect requirements for compute systems

High-density and high-speed interconnect technologies have been established and matured over the last few decades to serve the increasing demand for bandwidth and channel-count of emerging compute systems [7]. While speeds beyond 20 Gbps per lane/differential pair can be reached with state-of-the-art on-board electrical links, it is much more challenging to provide such speeds for densely spaced links and/or over distances beyond a few centimetres. In Figure 1, we aim at identifying specific data communication bottlenecks within a data centre and assess the potential for optical technologies. As mentioned, very short data links, e.g. between processors or to the memory on the same card, can still be built over printed circuit board (PCB)/package copper traces by exploiting powerful equalisation schemes [7]. Relatively low bandwidth links are found between storage and backup media and respective controllers. At these low data rates, links can also be implemented cost efficiently with electrical 'legacy' technologies. Large modern workloads tend to stretch beyond the drawer, or the rack boundary, either by requiring more memory, more compute or both. Consequently, intra data centre networks are much more heavily loaded and so are their endpoints, the network interface controllers (NIC). At the NIC, the highest concentration of high-speed signals for longer reach is found. Thus, the highest growth potential and the position where optics can outperform copper-based interconnects are at these edges of the nodes and within the data centre network, which on a global scale constitutes the majority of all IP traffic [8].

3 Optical chip-level communication interfaces

Since more than a decade, optics is used in data centres to connect from the card/node edge to the intra data centre network, which is very much in line with the observation made in the previous section. The technology of choice here is pluggable active optical cables (AOC) [9]. A scheme of a typical compute node is depicted in Figure 2. The advantage of AOC technology is that it completely decouples system design from any optics constraint and provides a lot of design flexibility. However, drawbacks of such integration schemes are as follows: First, all high-speed electrical signals need to be transmitted over relatively large distances to/from the card edge, consuming board real estate and power. Second, there is constraint by the drawer cross section, so very little volume is available for the electro-optical (EO) conversion. The aggregated data rate into and out of the node are thus fundamentally limited by the height and width of the node drawer.

To increase the bandwidth density and lower the power consumption, closer integration of the optical

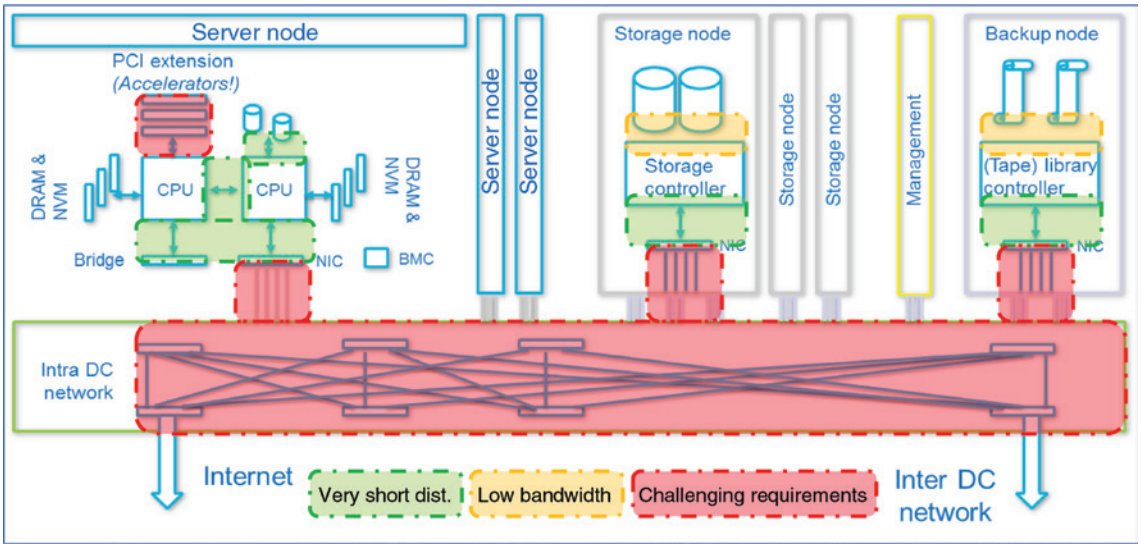

Figure 1: High-level server architecture. Very short distance links (green) and low-bandwidth links (yellow) will likely not be limited by electrical interconnect technologies in the near future. Longer and high-bandwidth links and the intra data centre network (red) can often no longer be efficiently implemented electrically and will benefit the most from optical data communication.

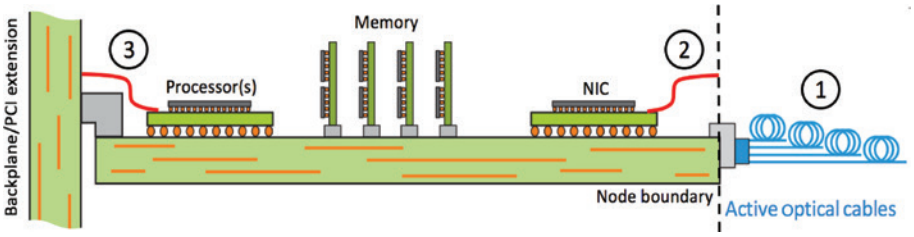

Figure 2: Compute node scheme. Active optical cables are attached at the edge (1). High bandwidth-length-density requirements are also found between the NIC and the node edge (2) and from processors to PCI(e) extension cards and accelerators (3).

transceivers with the processing modules (i.e. processor chips and switch chips) is required. Performing the EO conversion within the system, closer to the processor/NIC packages, will allow for (much) shorter electrical signals, thus better signal quality and lower power consumption. Such tight integration also allows to fully exploit the cross section of the node chassis edge to feed through optical signals at much higher densities/numbers than is possible with the electrical connectors of the AOCs. Discrete mid-board and NIC carrier co-packaged multi-mode transceivers, based on vertical cavity surface emitting laser (VCSEL), were successfully used in high-performance data centres and computing systems [9, 10]. Despite their advantages, they always induce extra cost through separate packaging, reduced board real-estate efficiency through the use of bulky electrical connector or socket interfaces and even constrain cooling airflow or cable wiring. Additionally, for bit rates higher than 25 Gbps, the maximum reach of links based on VCSEL-based multi-mode transceivers drops below the requirements of larger data centres.

The combination of two technologies, the CMOS Si photonics technology integrating single-mode (SM) waveguides and the optical PCB technology with optical PWGs, offers a path to tight integration of SM optical devices with electronics.

Integrated CMOS-based Si photonics combines electrical and optical functions on a single Si chip, typically built on silicon-on-insulator (SOI) layer stacks. It provides bandwidth scalability by the fundamental scale-out properties of integrated technology and by introducing additional functions such as wavelength division multiplexing or higher-order modulation formats. All necessary Si photonics building blocks have been already demonstrated [11–13]. In comparison to VCSEL-based technology, Si photonics is a SM platform operating in the telecom wavelength windows around 1.31 or 1.55 µm, (O- and C-bands, respectively). SM operation provides a larger bandwidth-distance product and enables kilometre-range SM fibre links with channel data rates of ≥25 Gbps.

Optical PCB technology is an extension of the established electrical PCB technology with the addition of optical PWGs processed on or embedded in the PCB stack. The combination of PWGs with copper lines in a high-density laminate offers a versatile platform with EO interconnects for signal routing and a scalable EO assembly [14].

Our technology supports compute node architectures that include Si photonics transceivers co-packaged with the processors and application-specific integrated circuits directly on the same carrier substrate. PWGs are used to distribute the optical signals between the various subsystems across the optical PCB. This was partially reported in [15], and it is schematically summarised in Figure 3.

The Si photonics chip packaging and the interface between the Si photonics waveguides on the Si photonics chip and the optical fibres are a major challenge. The large difference of the numerical apertures between optical fibres and Si photonics waveguides leads to high optical coupling loss for direct butt coupling. Optical mode adaption is required at the optical chip I/Os. There exist two well-known approaches for direct fibre coupling: out-of-plane coupling (e.g. vertical grating couplers) and edge coupling. With vertical grating couplers, fibre-to-waveguide coupling losses as low as 1 dB have been reported [16]. However, vertical grating couplers are polarisation sensitive and limited in their usable wavelength range. Edge coupling fibre interfaces using

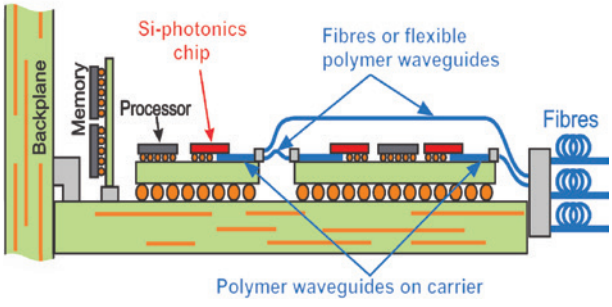

Figure 3: Integration of Si photonics chip on processor package.

subwavelength grating structures at the Si waveguide end and Si substrated removal have coupling losses as low as 1.5 dB [17]. Both approaches are limited to low fibre count I/O interfaces, e.g. Si photonics stand-alone transceiver modules [18].

Direct-fibre coupling constrains packaging schemes and is density limited by the diameter of the fibres attached to the chips. In our approach, we exploit optical fibre mode matched PWGs integrated on the board and/or carrier as an intermediate media between the Si photonics waveguides and the fibres. The latter can now be attached at the edge of the card (see Figure 3). This scheme offers high optical I/O channel density at the chip edge and a robust mechanical interface, decoupled from the chip. It is compatible with standard electrical assembly techniques (solder reflow and flip-chip bonding) of the Si photonics chip and thus removes the cost overhead of dedicated EO packaging technology. Previous related work spans from Si waveguide-to-PWG coupling of multiple optical I/Os, over scalable optical interfaces based on flip-chip packaging [19, 20], to PWG directly processed on Si photonics chip [21]. In the next section, we provide implementation details of our optical interface between Si photonics waveguides and PWGs.

4 Waveguide-based Si photonics chip coupling

We established a SM PWG technology operating in the O- and C-bands. A siloxane-based polymer from Dow Corning Corporation (Midland, MI, USA) was selected for its excellent optical properties and easy processing. The material is thermally very stable [14], thus compatible with flip-chip bonding for system-level integration as we envision in Figure 3. A core-cladding index contrast $\Delta n \approx 0.005–0.008$ provides SM operation of square-shaped waveguides of 6–8 μm edge length in the O- and C-bands and low-loss optical coupling to standard SM fibres, which we measure to be 0.5 dB. A detailed report on the processing, integration, properties and application of our PWG technology is provided in [15].

The coupling between Si waveguides and PWGs is realised by the adiabatic transformation of the optical super-mode in the Si/PWG system. For this, we taper the Si waveguide end at the edge of the Si photonics chip. The optical mode evolves adiabatically along the coupler from a full confinement in the Si waveguide to almost complete localisation in the PWG core (and vice versa). Figure 4 shows the Si/PWG coupling system at the edge of the Si photonics chip that can be attached to the carrier by flip-chip bonding and the evolution of the optical super-mode intensity along the coupler [15].

For our experiments, we realised a Si waveguide/PWG system as shown in Figure 4D. The PWG refractive index contrast between core and cladding was $\Delta n = n_{core} - n_{clad} = 0.0055$. The Si waveguide taper consists of three linear segments and has a Si width that varies from 350 nm waveguide width to 90 nm Si tip width, as depicted in Figure 4C. Our taper design has the transverse electrical (TE) mode transfer from PWG-to-Si waveguide at 170 nm Si taper width. Thus, the smallest taper slope was designed to be between 140 and 300 nm Si width (see [22]), and the total taper lengths L_c range from 0.5 to 2.0 mm. The Si waveguides were structured on an SOI wafer with a 2 μm thick buried oxide and a top-Si thickness of 170 nm. The latter is to allow single mode operation in the O- and C-bands and to match the localisation of the transverse magnetic (TM) and TE mode transfer; see Figure 4C centre intensity plot. On a separate glass wafer, PWGs were fabricated with a lower polymer cladding of 25 μm. Subsequently, the processed SOI and PWG wafers were diced into test-unit chips, aligned and assembled to

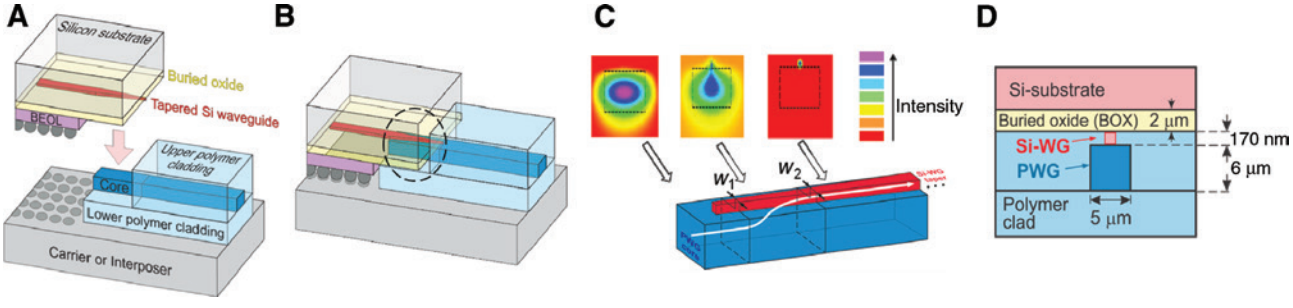

Figure 4: (A) Flip-chip bonding of Si photonics chip onto PWG on carrier. (B) Physical contact between a tapered Si waveguide and PWG enables low-loss adiabatic optical coupling. (C) Optical super-mode intensity evolution in the adiabatic Si waveguide/PWG system. For Si taper widths $w \approx w_2$, the mode is confined in the Si waveguide. For Si taper widths $w \approx w_1$, the mode is localised in the PWG core. (D) Cross-sectional scheme of the Si waveguide/PWG system used for adiabatic coupling.

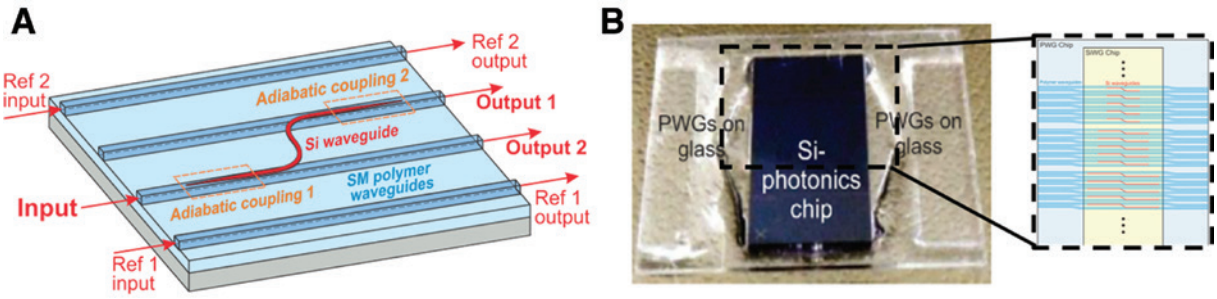

Figure 5: (A) Layout of an S-shape coupler test unit with two identical adiabatic couplers and four PWGs to determine coupling efficiency. (B) Pictures of the Si photonics chip assembled onto a PWG-on-glass chip (left) and schematic of the multiple test-unit arrangement for large optical I/O channel count test (right).

each other with a flip-chip bonder. Thus, the silicon waveguide and the PWG core were brought into direct contact, realising the system in Figure 4D. The fixation of the bonded assembly was finalised by UV-curing photosensitive polymer cladding material, which acted as filler and optical glue. The UV-curing process provides the required Δn as reported above.

Figure 5A shows an assembled test unit used for coupling loss measurements. When no Si waveguide is in contact with the PWG core, the light is almost fully localised in the PWG core because of the relatively high refractive index contrast between the PWG core (1.516) and the BOX (1.450). In this case, the propagation loss of the optical mode is calculated to be <0.05 dB/cm. However, because of the intrinsic material loss of the current polymer material formulation and the nonideal patterning of the PWG core sidewalls, a total propagation loss of 1.0 dB/cm was measured.

The Si photonics chip was designed with a large number of test sites to prove the scalability of our Si

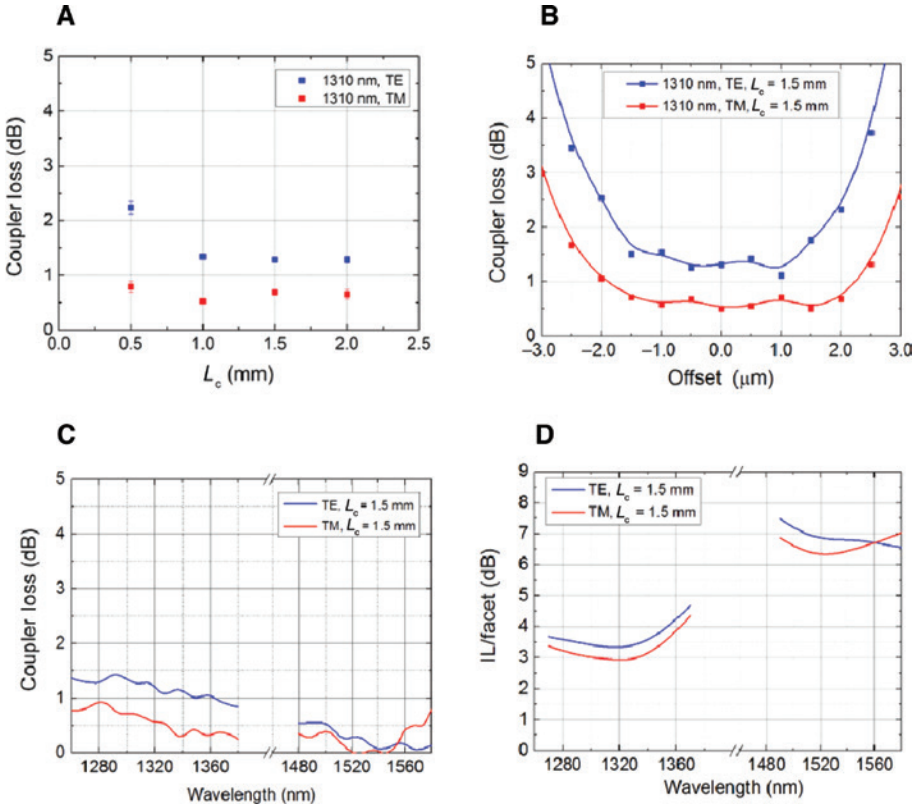

Figure 6: (A) Coupler loss vs. coupler length L_c, at $\lambda = 1310$ nm. (B) Misalignment tolerance for $\lambda = 1310$ nm and $L_c = 1.5$ mm. (C) Coupling loss vs. wavelength for $L_c = 1.5$ mm. (D) IL/facet in the O- and C-bands.

photonics interface concept to a high number of optical I/O channels. Figure 5B shows an assembled chip with 152 optical I/Os and a schematic of the test unit with different taper lengths L_c.

The assembled couplers were characterised by injecting and collecting the light into the input and output PWGs, respectively, using SM fibres. The coupler loss for both TE and TM modes was obtained by comparison of the transmitted optical power through the S-shape Si waveguide optical path with the optical transmitted power through the straight reference PWG. Figure 6A shows measured coupler loss of <0.6 dB for TM and <1.4 dB for TE at $\lambda = 1310$ nm wavelength for taper lengths $L_c \geq 1.0$ mm. The larger loss of the TE mode is attributed to the Si waveguide sidewall roughness and to its more abrupt transition from the PWG to the Si waveguide along the taper [23]. The lateral alignment tolerance of the Si photonics chip-to-PWG process assembly was evaluated using a Si-to-PWG test unit designed with varying Si waveguide to PWG horizontal misalignment. An additional 1 dB loss was measured for ±2 µm lateral misalignment for $L_c = 1.5$ mm, as shown in Figure 6B. Figure 6C shows coupler loss of <1.5 dB over the entire O-band and C-band for both polarisations for designs with $L_c = 1.5$ mm. The lower coupling loss measured for the C-band is due to the smaller influence of Si waveguide imperfections at longer wavelengths. Finally, we evaluated the overall fibre-to-PWG-to-Si waveguide loss contribution (IL/facet). The IL/facet is shown in Figure 6D for the O- and C-bands and for both polarisations. The IL/facet was found to be <3 dB and <3.5 dB at 1310 nm for TM and TE, respectively. The IL/facet wavelength dependency and the higher IL/facet in the O- and C-bands are due to the long PWG propagation length of ~1.5 cm in the test unit, which we selected for ease of handling and greater flexibility to accommodate different test chips.

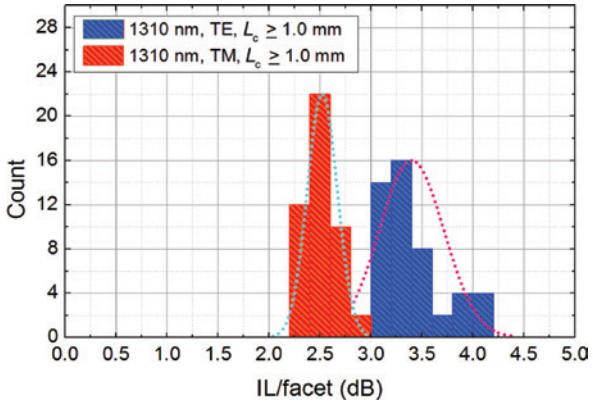

Figure 7: TE- and TM-dependant IL/facet count for a total of 50 optical I/Os with $L_c \geq 1.0$ mm.

We demonstrated that more than 100 functional optical Si waveguide to PWG interfaces can be assembled in a single flip-chip bonding step. Out of these 100, we considered the 50 PWG/Si waveguide interfaces with $L_c \geq 1.0$ mm and which were designed for perfect horizontal alignment, in our statistical analysis. An average $\langle \text{IL/facet} \rangle_{TE} = 3.4$ dB (with standard deviation $\sigma_{TE} = 0.32$ dB) and an average $\langle \text{IL/facet} \rangle_{TM} = 2.5$ dB (with $\sigma_{TM} = 0.15$ dB) were found at 1310 nm, as shown in Figure 7.

5 Summary

Low-loss adiabatic optical coupling between Si waveguides and fibre-matched PWGs is reported. Our Si photonics chip-level coupling scheme facilitates broadband interfaces to SM fibres and is scalable to very high optical I/O channel counts. Si waveguide-to-PWG coupling losses below 1.5 dB could be achieved for both the O- and C-bands, with a relaxed 1 dB lateral misalignment tolerance of about ±2 µm. More than 100 optical Si waveguide to PWG interfaces were assembled in a single alignment step, and average $\langle \text{IL/facet} \rangle_{TE}$ of 3.4 dB and $\langle \text{IL/facet} \rangle_{TM}$ of 2.5 dB were measured at 1310 nm, clearly highlighting scalability potential of our technology.

Acknowledgements: The authors acknowledge Dow Corning Corporation (Midland, MI, USA) for developing and providing the optical polymers. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 688172 (STREAMS) and no. 688572 (WIPE). It was also supported by the Swiss National Secretariat for Education, Research and Innovation (SERI) under contract nos. 15.0339 and 15.0309. The opinions expressed and arguments employed herein do not necessarily reflect the official views of the Swiss Government.

References

[1] J. Weiss, R. Dangel, J. Hofrichter, F. Horst, D. Jubin, et al., Optical Interconnects for Disaggregated Resources in Future Datacenters, 2014 The European Conference on Optical Communication (ECOC) (Cannes, France, 2014).
[2] D. A. B. Miller and H. M. Ozaktas, J. Parallel Distrib. Comput. 41, 42–52 (1997).
[3] Y. A. Vlasov, IEEE Commun. Mag. 50, 67–72 (2012).
[4] C. R. Doerr, IEICE Trans. Electron. 96, 950–957 (2013).
[5] Y. Hibino, MRS Bull. 28, 365–371 (2003).
[6] X. Zhang, A. Hosseini, X. Lin, H. Subbaraman, R. T. Chen, IEEE J. Sel. Top. Quant. 19, 196–210 (2013).

[7] P. A. Francese, T. Toifl, M. Braendli, C. Menolfi, M. Kossel, et al., 10.6 Continuous-Time Linear Equalization with Programmable Active-Peaking Transistor Arrays in a 14nm FinFET 2mW/Gb/s 16Gb/s 2-Tap Speculative DFE Receiver, IEEE ISSCC (San Francisco, CA, USA, 2015).
[8] CISCO, Cisco Global Cloud Index: Forecast and Methodology, 2015–2020 (2016). Available at: https://www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.pdf.
[9] InfiniBand Trade Association, InfiniBand® Roadmap. Available at: http://www.infinibandta.org/content/pages.php?pg=technology_overview, Accessed: 7/Feb/2018.
[10] A. F. Benner, D. M. Kuchta, P. K. Pepeljugoski, R. A. Budd, G. Hougham, et al., Optics for High-Performance Servers and Supercomputers, OFC Conference (San Diego, CA, USA, 2010).
[11] S. Assefa, S. Shank, W. Green, M. Khater, E. Kiewra, et al., A 90nm CMOS Integrated Nano-Photonics Technology for 25Gbps WDM Optical Communications Applications, IEDM (San Francisco, CA, USA, 2012).
[12] C. R. Doerr, Front. Phys. 3, 37 (2015).
[13] H. Subbaraman, X. Xu, A. Hosseini, X. Zhang, Y. Zhang, et al., Opt. Express 23, 2487–2511 (2015).
[14] R. Dangel, F. Horst, D. Jubin, N. Meier, J. Weiss, et al., J. Lightwave Technol. 31, 3915–3926 (2013).
[15] R. Dangel, J. Hofrichter, F. Horst, D. Jubin, A. La Porta, et al., Opt. Express 23, 4736–4750 (2015).
[16] A. Mekis, S. Gloeckner, G. Masini, A. Narasimha, T. Pinguet, et al., IEEE J. Sel. Top. Quant. 17, 597–608 (2011).
[17] T. Barwicz, N. Boyer, A. Janta-Polczynski, J.-F. Morisette, Y. Thibodeau, et al., A Metamaterial Converter Centered at 1490 nm for Interfacing Standard Fibers to Nanophotonic Waveguides, OFC, 2016 (Anaheim, CA, USA, 2016).
[18] P. De Dobbelaere, S. Abdalla, S. Gloeckner, M. Mack, G. Masini, et al., Si Photonics Based High-Speed Optical Transceivers, ECOC (Amsterdam, The Netherlands, 2012).
[19] I. M. Soganci, A. La Porta, and B. J. Offrein, Opt. Express 21, 16075–16085 (2013).
[20] T. Barwicz, Y. Taira, S. Takenobu, N. Boyer, A. Janta-Polczynski, et al., Optical Demonstration of a Compliant Polymer Interface between Standard Fibers and Nanophotonic Waveguides, OFC (CA, USA, 2015).
[21] A. La Porta, J. Weiss, R. Dangel, D. Jubin, N. Meier, et al., Silicon Photonics Packaging for Highly Scalable Optical Interconnects, ECTC (San Diego, CA, USA, 2015).
[22] A. La Porta, R. Dangel, D. Jubin, N. Meier, D. Chelladurai, et al., Scalable Optical Coupling between Silicon Photonics Waveguides and Polymer Waveguides, ECTC (Las Vegas, NV, USA, 2016).
[23] A. La Porta, R. Dangel, D. Jubin, F. Horst, N. Meier, et al., Optical Coupling between Polymer Waveguides and a Silicon Photonics Chip in the O-band, OFC (Anaheim, CA, USA, 2016).