JOM Journal of Optical Microsystems

RESEARCH PAPER

Thermal modeling of hybrid three-dimensional integrated, ring-based silicon photonic–electronic transceivers

David Coenen,[a,b,*] Minkyu Kim,[b] Herman Oprins,[b] Yoojin Ban,[b] Dimitrios Velenis,[b] Joris Van Campenhout,[b] and Ingrid De Wolf[a,b]

[a]KU Leuven, Department of Materials Engineering, Leuven, Belgium

[b]imec, Leuven, Belgium

ABSTRACT. Co-packaged optics for high performance computing or other data center applications requires dense integration of silicon photonic integrated circuits (PICs) with electronic integrated circuits (EICs). This work discusses the impact of three-dimensional (3D) hybrid integration on the thermal performance of Si ring-based photonic devices in wavelength-division multiplexing PICs. A thermal finite element model of the EIC-PIC assembly is developed and calibrated with thermo-optic device measurements, before and after integration of an electrical driver on top of the PIC by means of microbump flip-chip bonding. Both measurements and simulations of the thermal tuning efficiency and crosstalk between silicon photonic devices show that the EIC can have a significant impact on the thermal performance of the integrated heaters in the PIC by acting as an undesired heat spreader. This heat spreading lowers the heater efficiency with 43.3% and increases the thermal crosstalk between the devices by up to 44.4% compared with a PIC-only case. Finally, it is shown that these negative thermal effects of 3D integration can largely be mitigated by a thermally aware design of the microbump array and the back-end-of-line interconnect, guided by the calibrated thermal simulation model.

Keywords: silicon photonics; co-packaged optics; thermal modeling

Paper 23016SS received Jun. 28, 2023; revised Nov. 14, 2023; accepted Nov. 20, 2023; published Dec. 6, 2023; corrected Mar. 11, 2024.

1 Introduction

The increase in internet data traffic has been a major topic of discussion in recent years. With the proliferation of connected devices and the increasing use of the internet for a wide range of activities, the demand for data transmission has skyrocketed.[1] This trend is expected to continue in the coming years, with some estimates projecting that global internet traffic will reach unprecedented levels. Furthermore, the recent development of large generative artificial intelligence (AI) models (e.g., ChatGPT) showcases the need for sufficient networking bandwidth for high-performance AI/machine learning (ML) clusters that are used for training these models. As the demand for networking bandwidth continues to rise, AI/ML compute clusters have turned to optical transceivers as a means of transmitting large amounts of data at high data rates, while maintaining low power (<5 pJ/bit), low bit-error ratio ($<10^{-12}$), and low link latency.[2] So, traditional Cu connects are being replaced by pluggable optical interconnects for interconnecting multi-XPU servers. As these systems scale up to thousands of interconnected XPUs, the optical interconnects will move to the package level to support the chip-to-chip

*Address all correspondence to David Coenen, david.coenen@imec.be

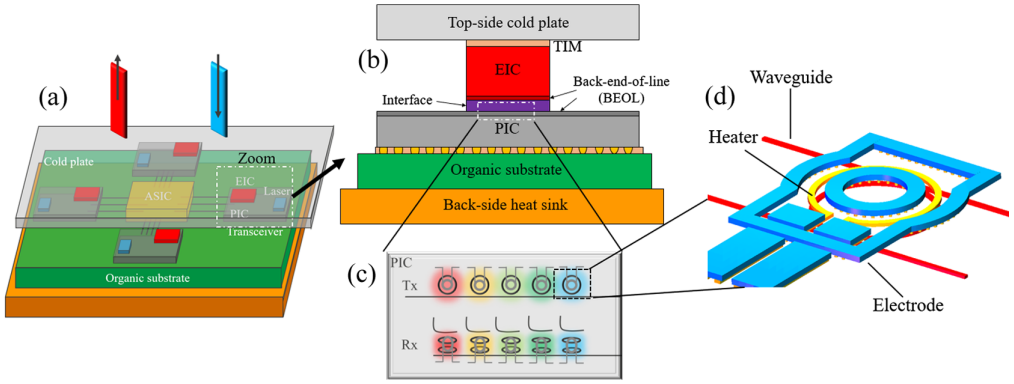

Fig. 1 (a) CPO conceptual drawing. (b) The optical transceiver module consists of a PIC with 3D integrated driver EIC and laser. The EIC-PIC is bonded face-to-face (F2F). The whole assembly has a back-side heat sink and a top-side cold plate. (c) The transceiver utilizes DWDM architecture with ring modulators (d) on the transmit side (Tx) and ring filters on the receive side (Rx), which are tuned with a heater.

bandwidth and input/output (I/O) density requirements. This industry trend toward the use of co-packaged optics (CPO) in which the transceiver and network switch are integrated into a single package has been reported in numerous studies recently.[3–5] This trend has been driven by the need for more compact and cost-effective solutions, as well as the desire to improve performance and reliability.[6]

In Fig. 1(a), a conceptual drawing of a CPO module is shown. At the center is the host integrated circuit (IC) (e.g., an XPU or network switch), which is connected through the packaging substrate to the optical modules around it. In this work, we focus on the optical module, which consists of an assembly of a photonic integrated circuit (PIC), a laser,[7,8] and an electronic integrated circuit (EIC) containing electrical drivers, trans-impedance amplifiers, and control circuits.[2] More specifically, the focus of this work is the thermal interaction between the PIC-EIC [Fig. 1(b)] and the impact on the performance of the photonic devices. For the PIC, a dense wavelength-division multiplexing (DWDM)-based architecture is assumed; it consists of arrays of ring-based photonic devices (ring modulators and ring filters). The operation principle of ring-based silicon photonic devices is based on resonance, making them very compact but also highly susceptible to temperature changes.[9] A change of 0.5°C in ambient temperature can already degrade the signal quality of the optical link. For this reason, integrated heaters are used in the devices to allow for locally controlling the temperature.[10–12] Because the power consumption of these heaters should be minimized, the heater efficiency is an important metric in the design of the heater. Thus, any factor that influences the heater efficiency is important to consider. In this paper, we investigate the thermal impact of the three-dimensional (3D) integration of the EIC on top of the PIC by means of a flip-chip μbump assembly. The novelty of this work is the characterization of the thermal performance of the ring-based devices before and after EIC integration, along with detailed models that allow us to interpret the results and develop mitigation strategies.

2 Methodology

2.1 Samples and Experimental Characterization

The focus of this work is to study the effect of the EIC-PIC 3D assembly on the thermal performance of the PIC, so the experimental test vehicle is limited to these components (Fig. 2). In this figure, the PIC is visible, with the EIC flip-chipped[13] on top. In Fig. 2(c), a cross section is shown of the 3D interconnect between the two dies, which consists of a 50 μm pitch Cu-Sn μbump array and an epoxy-based underfill material. The PIC die size is 5×5 mm^2, and the EIC die size is 1.5×1 mm^2. The EIC-PIC assembly is mounted on a printed circuit board (PCB) and connected with wire bonding for electrical I/O for testing and a fiber array for optical I/O for testing. The transmitter (Tx) part of the PIC consists of an array of ring modulators with a 100 μm pitch. First, experimental characterization of the heater efficiency and thermal crosstalk

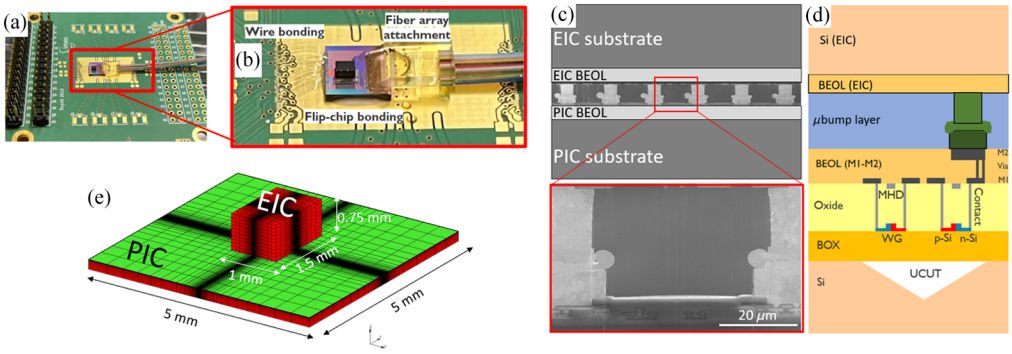

Fig. 2 Picture of test vehicle: (a), (b) PIC with electrical driver (EIC) hybrid integrated with flip-chip bonding. Cross-section scanning electron microscope (SEM) picture of (c) 3D die stack and (d) sketch: MH and WG. (e) Finite element mesh.

for the ring modulators is carried out. Those parameters are measured for devices with- and without a substrate undercut (UCUT), which is an additional design feature in which a part of the Si substrate below the device is removed with the intent of improving the heater efficiency.[10] Furthermore, the heater efficiency is measured before and after EIC integration to quantify its thermal impact. The heater efficiency is measured by tracking the ring resonance wavelength shift by applying power to the heater. This results in a figure of merit in units of wavelength per unit power, or nm/mW typically. Next, thermal crosstalk measurements are carried out, quantifying the parasitic heating of a ring modulator from the heaters of the neighboring rings in the array. This is measured by activating the heater in the first channel and measuring the resonance wavelength shift in all neighboring rings.

2.2 Finite Element Thermal Model

A finite element thermal model is made of a PIC with EIC flip-chipped on top. The model is develop using the commercial finite element solver MSC Marc.[14] The finite element model is shown in Fig. 2(e). Using the cross section in Fig. 2(d), the thermal simulation is explained. The simulation input is Joule heat generated in the metal heater (MH; made of tungsten), situated above the waveguide (WG). The total amount of dissipated heat is set to 1 mW, such that the simulation results are normalized per unit of power and can easily be converted to different units. The main point of interest is the resulting WG temperature of the ring modulator, which is extracted from the simulation result as the volumetric average temperature in the Si ring WG. This metric is used to compare the simulation with the measurement result of the heater efficiency. The thermal boundary conditions are natural convection ($h = 20$ W/m^2-K) on the top-facing surfaces, adiabatic sidewalls, and a thermal resistance ($R_{th} = 8$ K/W[15]) on the bottom face of the PIC, representing the PCB below the die. The material properties used in the simulations are summarized in Table 1. More details on the finite element modeling can be found in Ref. 10. Note that only the EIC-PIC assembly is included in the computational domain; all other parts of the package (Fig. 1) are represented as equivalent boundary conditions.[16] Also, the finite element model is made to mimic the test setup in Fig. 2; consequently, there is no top side cold plate. In Fig. 2(c), a cross section of the interface layer (μbumps) is shown. This is a relatively complex layer, with in-plane anisotropy and cross-plane thermal conductivity that depends on the used underfill and metallic composition of the bumps. Two different ways of modeling this layer are compared. First, the actual geometry of the bumps is used in the model (i.e., geometric model). In Fig. 3, the ring modulator is shown along with an array of 12 μbumps around the device. We opt for including only 12 μbumps because the neighboring μbumps have the largest

Table 1 Thermal conductivity values for thermal simulations,[10] UF = underfill in μbump layer.

	Si	SiO_2	Cu	W	UF	Cu-Sn	BEOL (EIC)
Thermal conductivity [W/(m-K)]	150	1	400	100	0.18	65	1.5

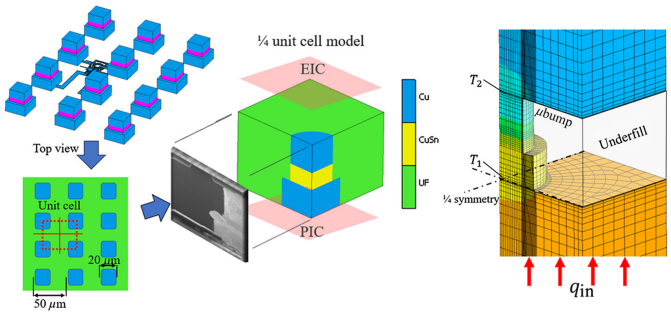

Fig. 3 μbumps modeling: individual geometric or with equivalent material properties through unit cell approach. The equivalent out-of-plane thermal conductivity is calculated using the temperature drop across the μbump $\Delta T = T_1 - T_2$.

impact on the ring modulator. In the actual test chip, the PIC-EIC interface is fully filled with μbumps (1.5×1 mm^2 area). Including the actual geometry of the μbumps will yield the most accurate results, but it is quite inefficient because of the large mesh size and hence increased central processing unit (CPU) time for simulation. The second method of modeling the μbump layer is by replacing it with an equivalent material with properties that match the combination of μbumps and underfill. To achieve this, the symmetry of the layer is exploited, and a separate, quarter-slice model of the bump is made (Fig. 3). With this highly simplified model, the effective in-plane and cross-plane thermal conductivity can be estimated.[17] Finally, these equivalent material properties are used in the full μbump layer (i.e., equivalent model). The equivalent out-of-plane thermal conductivity of the μbump layer is calculated using the simulated temperature drop across the μbump $\Delta T = T_1 - T_2$ (Fig. 3), given as

$$k_{z,\text{eff}} = \frac{q_{\text{in}}}{\Delta T \cdot t} = 9.38 \text{ W/(m} - \text{K)}, \tag{1}$$

where q_{in} is the input heat (see Fig. 3) and t is the μbump thickness (≈ 20 μm).

3 Experimental Results

3.1 Heater Efficiency

The heater efficiency of the ring modulators is measured for four different cases: with and without UCUT and with and without EIC flip-chipped on top of the PIC. The measurement results are shown in Fig. 4. The addition of UCUT in the design increases the heater efficiency, as reported in Ref. 10. Furthermore, the impact of the EIC on the heater efficiency is also significant: for the case with UCUT, a decrease by 44% is observed. The experimental results are summarized in Table 2 and compared with simulation results, which are further discussed in Sec. 4.1. Note that the result of the thermal simulations is the average WG temperature (in units of K/mW), which can be converted to a resonance wavelength shift (in units of pm/mW)[10] to compare with the

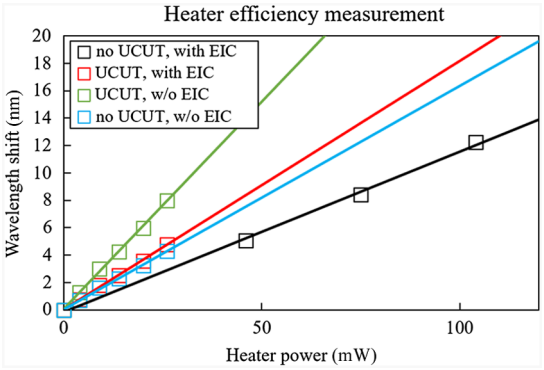

Fig. 4 Measurement of the ring modulator heater efficiency.

Table 2 Ring modulator heater efficiency from experiments (pm/mW) and simulations (equivalent model) (K/mW). The value between brackets is obtained by the geometric model.

	No UCUT		UCUT	
	No EIC	With EIC	No EIC	With EIC
Experiment (pm/mW)	165	131	312	180
Simulation (K/mW)	3.07	2.14	5.44	2.86 (3.08)
Simulation (pm/mW)	192	134	340	179
Difference (%)	16.3	2.10	8.97	0.55

experimental heater efficiency. The conversion can be done using the temperature-dependence of the ring resonance wavelength, given as

$$\frac{\Delta\lambda}{\Delta P} = \frac{\Delta T}{\Delta P} \times \frac{\Delta\lambda}{\Delta T}, \tag{2}$$

where $\Delta\lambda/\Delta T = 62.54$ pm/K for the rings measured in the O-band wavelength range.

3.2 Thermal Crosstalk

In the second experiment, the thermal crosstalk between the ring modulators is analyzed. This is done as follows: in the array of four ring modulators, only the heater of the first device is actuated. The optical transmission spectrum of each device is measured for difference values of heater power in the first device of the array. Due to thermal crosstalk, it is expected that there will be a small shift in the resonance wavelength in the neighbor's transmission spectra, even though their heater is not actuated. The measurements are conducted on the test setup (Fig. 2) with passive natural convection cooling on the die top side and with the EIC bonded on top of the PIC. In Sec. 4.2, simulations are carried out with and without EIC to quantify its impact on thermal crosstalk. Furthermore, the boundary condition on the top side of the assembly is replaced with an equivalent thermal resistance that models an actively cooled cold plate, which is more representative of a real CPO package (Fig. 1). In Fig. 5, the measured optical transmission spectra of the ring modulators show the increasing applied heater power. The spectra of all four devices are plotted together. The four ring modulators are positioned on the same bus WG; hence, four resonance peaks are visible in each wavelength sweep (e.g., the green profile in Fig. 5). The first ring modulator, which is actuated, exhibits a clear resonance wavelength red-shift. The first, second, and third neighbor resonance wavelengths are indicated. Their wavelength shift is not clearly

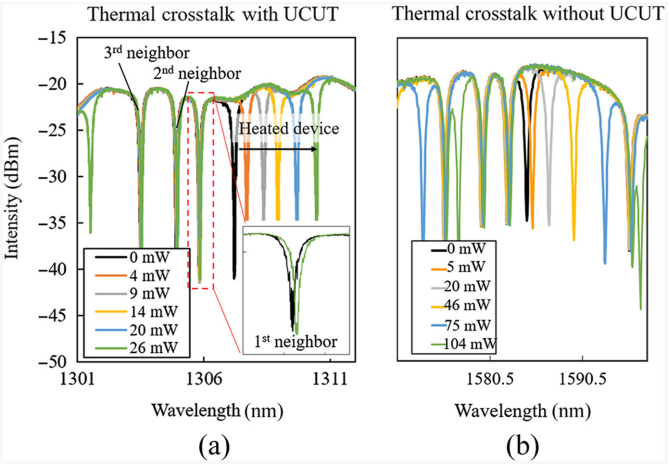

Fig. 5 Thermal crosstalk measurement with UCUT (a) and without UCUT (b). Both results are with EIC integration. Note that the ring modulators with UCUT (a) are designed for O-band operation and without UCUT (b) for C-band operation.

Table 3 Thermal crosstalk experiment and simulation (equivalent model) of EIC-PIC assembly, in percentages.

	Crosstalk (%)			
	Experiment		Simulation	
Distance (μm)	No UCUT	UCUT	No UCUT	UCUT
100	2.62	1.65	2.32	1.68
200	2.33	1.43	2.06	1.49
300	1.94	1.32	1.92	1.39

visible, so an extra inset is added for the first neighbor. For the devices with UCUT [Fig. 5(a)], up to 26 mW of heater power is applied, which corresponds to a maximum resonance wavelength shift of 8.112 nm, given a heater efficiency of 312 pm/mW (see Table 2). For the devices without UCUT [Fig. 5(b)], these numbers are104 mW of heater power and 17.16 nm resonance wavelength shift, respectively. Because the devices without UCUT are expected to have a heater efficiency, which is approximately half compared with the devices with UCUT, it is important to increase the heater power for devices without UCUT for crosstalk characterization. By doing so, it is ensured that a sufficient resonance wavelength shift is obtained. Because the thermal crosstalk is in the 1% order of magnitude, the measured shift in resonance wavelength from the different neighbors can be estimated as $\Delta\lambda_{\text{neighbor}}(26\ \text{mW}) = 8.112\ \text{nm} \times 0.01 = 0.08112\ \text{nm}$. This is a small value compared with the width of the resonance peak of the ring modulator; consequently, it is important to apply sufficient power to the heater. Otherwise, the crosstalk could fall below the measurement precision of the setup. Based on this wavelength shift, thermal coupling is defined as

$$C[\%] = 100\% \cdot \frac{\Delta\lambda_{\text{neighbor}}}{\Delta\lambda_{\text{heated}}}, \tag{3}$$

where $\Delta\lambda_{\text{neighbor}}$ is the wavelength shift of the (unheated) neighbor ring modulator and $\Delta\lambda_{\text{heater}}$ is the wavelength shift of the heated ring modulator. The thermal coupling, or crosstalk, is measured for three neighboring ring modulators, each with 100 μm pitch. The coupling coefficients [Eq. (3)] are thus known at distances of 100 to 200 to 300 μm from the heated ring modulator. This data are used for the calibration of the finite element model in Sec. 4.2 and are summarized in Table 3.

4 Simulation Results

4.1 Heater Efficiency

4.1.1 *Equivalent versus real geometry approach*

The first part of the simulation results is the comparison between the equivalent model for the μbumps and the geometric model. In Fig. 6, a temperature contour plot of the ring modulator is shown. The reference case is the cross section of the bare PIC, without EIC [Fig. 6(b)]. There is a clear horizontal temperature gradient as the heat is forced to flow around the UCUT before it can leak into the Si substrate below. When the EIC is added, the temperature contour changes. In Fig. 6(c), the simulation result is shown for the equivalent model. The μbump layer, consisting of the μbumps and underfill, is replaced by material with an equivalent thermal conductivity, calculated with the unit cell model (Fig. 3). In Fig. 6(d), the result is shown for the geometric model. In this case, the actual μbumps geometry is modeled. The local temperature field above the UCUT is different than the equivalent model. However, the main point of interest is the WG: the average WG temperature determines the heater efficiency. In Fig. 7(a), a horizontal, linear temperature profile through the WG is shown for the two model types (geometric and equivalent). In Fig. 7(b), a circular temperature profile along the ring circumference is shown.

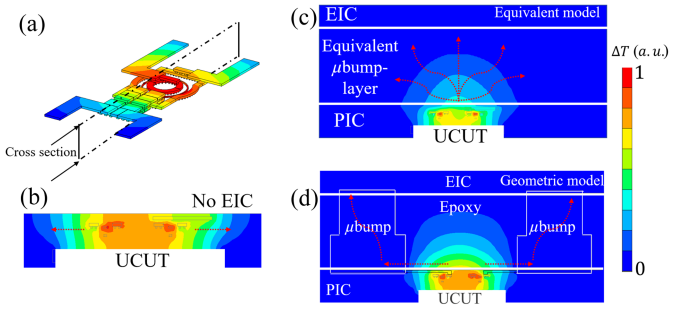

Fig. 6 Simulation result: (a) temperature contour plot of ring modulator cross section, (b) without EIC and (c), (d) with EIC integration.

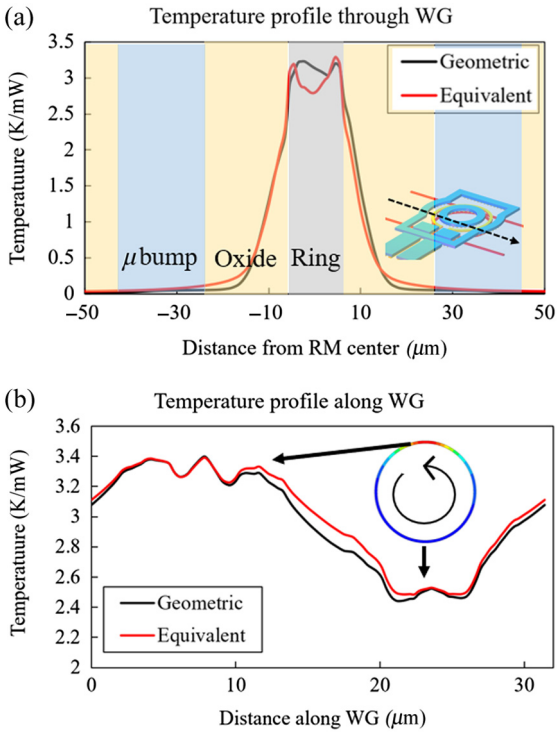

Fig. 7 (a) Linear temperature profile through WG and (b) circular temperature profile along the WG circumference.

From this, it can be concluded that, for the purpose of simulating the WG temperature, both the equivalent and geometric models produce similar results.

4.1.2 *Model validation*

The heater efficiency simulated with the equivalent model, that is, the average WG temperature, is now compared with the experimental results in Fig. 8. The experimentally observed drop in heater efficiency due to the flip-chip integration of the EIC (42.4% with UCUT and 20.6% without UCUT) is well captured in the model. The cause of this drop can be retrieved from the temperature contour plots in Fig. 6: the flip-chip bonded EIC acts as a heat spreader for the local heat generation in the heater of the ring modulator (RM). In the reference case without EIC [Fig. 6(b)], the heat transfer is forced horizontally by the presence of the UCUT below the RM. If EIC is flip-chipped on top, heat spreading occurs in the EIC, lowering the overall heater efficiency. By the addition of UCUT, the heat flow is forced horizontally through the SiO_2 surrounding the device. However, if the EIC is added, a secondary parallel heat transfer path is created, lowering the overall thermal resistance and partially offsetting the efficiency gain by UCUT.

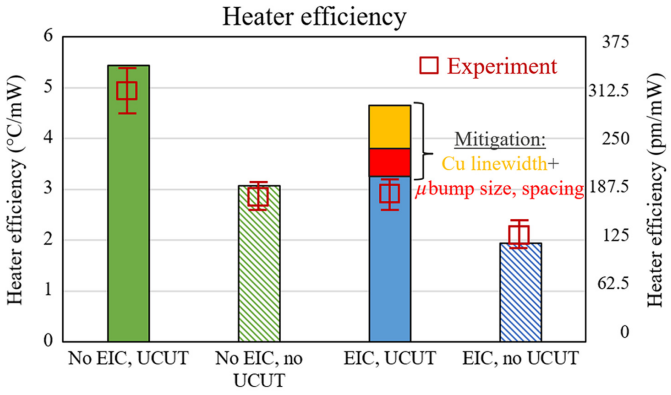

Fig. 8 Measured versus simulated heater efficiency for four cases: with and without UCUT/EIC. Red and orange impact on the heater efficiency with EIC and UCUT is obtained by a thermal modeling study (Sec. 5).

4.2 Thermal Crosstalk

Section 4.2 discusses the effect of 3D hybrid integration on thermal crosstalk between the different ring modulators. In Sec. 3.2, the experimental results were shown (Fig. 5 and Table 3). In our setup, the thermal crosstalk was measured at three different distances from the heated ring modulator: 100, 200, and 300 μm. At these locations, the relative thermal crosstalk is extracted and compared with the simulation results in Fig. 9, with and without UCUT. The simulation result is extracted in a way which is similar to the method that was used for the experimental coupling coefficients in Eq. (3): the coupling is defined as the temperature $T(x)$ at distance x relative to the temperature in the heated WG T_{ref}, in the PIC WG plane: $C(x) = T(x)/T_{\text{ref}} \cdot 100\%$. The first conclusion from the simulations is that the addition of UCUT lowers the thermal crosstalk, which is also observed experimentally. This is mainly caused by the higher reference temperature T_{ref}, that is, the higher WG temperature of the heated ring modulator. The UCUT itself has little impact on the absolute temperature at a large distance from the heated device.[18] Also, for large distances (>200 μm), the thermal crosstalk is quasi-constant, similar to the results reported in Ref. 19. To assess the impact of the flip-chip EIC integration on the PIC, the calibrated FE model is simulated with and without EIC [Fig. 9(b)] for the case with UCUT. The addition of the EIC increases the thermal crosstalk by 44.4%. Based on the previous conclusions about the heat spreading in the EIC, this result is expected. The flip-chip integration of the EIC causes additional lateral heat transfer, which in turn enhances thermal crosstalk in the PIC.

All reported crosstalk values are for the case without an active cooling solution on the EIC-PIC assembly top side. This allows a direct comparison with the experiments (Fig. 2). However, a more realistic thermal packaging approach for CPO-based transceivers is based on liquid cooling as shown schematically in Fig. 1(a). This scenario can be simulated by adapting the thermal boundary condition on the EIC top side by replacing the natural convection boundary condition

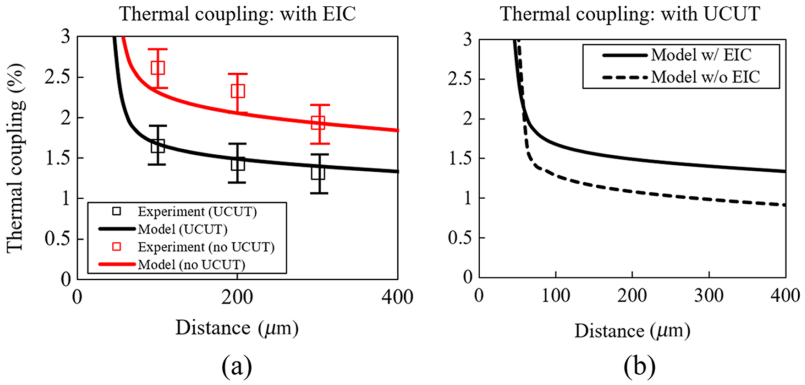

Fig. 9 (a) Thermal crosstalk measurement and simulation result with EIC. (b) Simulated crosstalk with and without EIC.

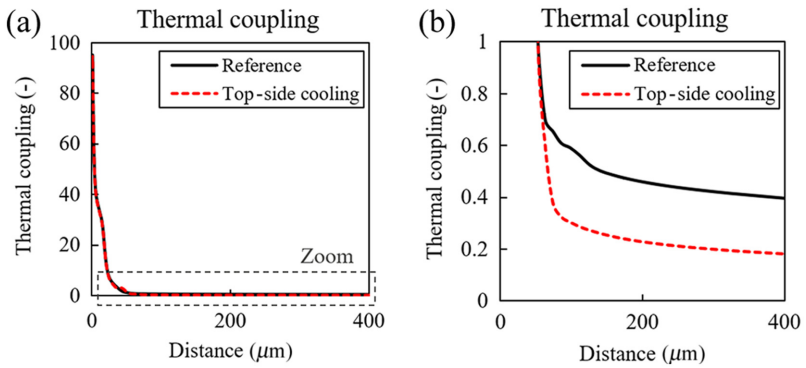

Fig. 10 (a), (b) Simulated thermal crosstalk for the reference case (test setup) and for a real package case (top-side cold plate).

with a thermal resistance of 0.1 K/W. The simulation results are shown in Fig. 10. The reference case corresponds to the device with UCUT and EIC flip-chipped on top, with natural convection cooling. The thermal crosstalk for the reference case and the top side cooling case is plotted in Fig. 10(a), and there appears to be no difference between both. However, in Fig. 10(b), the same data are shown, but zoomed-in around 1%. Here, the impact of the top side cooling is clear. The long range thermal crosstalk is greatly reduced (49.9%). This result shows a pathway for minimizing the effects of thermal crosstalk in heater-based PIC architectures. Furthermore, the impact of the top side cooling on the heater efficiency is negligible (−0.56%).

5 Mitigation Strategies

In this section, a set of design changes is proposed to mitigate the unwanted heat spreading effect of the EIC. The main strategy for avoiding undesired lateral heat transfer is by simply increasing the EIC-PIC interface thermal resistance and reducing the heat spreading in the PIC back-end-of-line (BEOL) layers. To illustrate this, the heater efficiency is simulated in the function of the bond layer thermal resistance (Fig. 11). This value is defined as the thickness-normalized thermal resistance, given as

$$R_{\text{bond}} = \frac{t}{k_{z,\text{eff}}}, \tag{4}$$

where t is the interface thickness and $k_{z,\text{eff}}$ is the effective out-of-plane thermal conductivity. For the reference case, this results in a bond layer thermal resistance of ∼2 mm^2-K/W. From Fig. 11, it is clear that a high thermal resistance value is preferred, which is the opposite of typical 3D (EIC) die stacks in which a good thermal contact is preferred.[20]

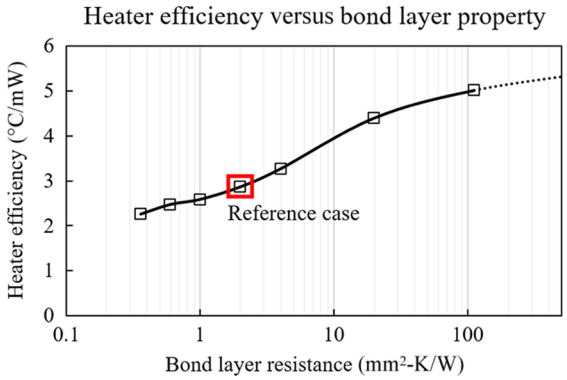

Fig. 11 Heater efficiency versus bond layer thermal resistance.

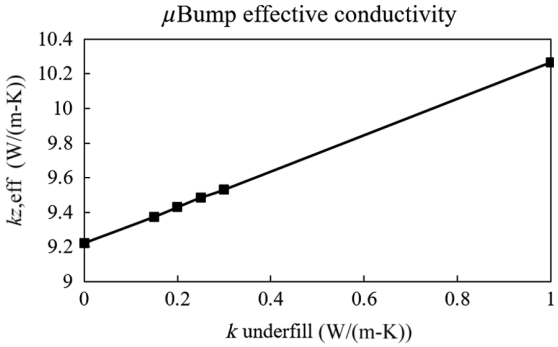

Fig. 12 Effective thermal conductivity of the μbump layer as a function of the underfill thermal conductivity.

5.1 Underfill

To maximize the interface thermal resistance, there are two degrees of freedom that can be changed: the metallic connections and the underfill. The impact of a low thermal conductivity underfill on the overall effective thermal conductivity can be simulated with the equivalent μbump model. The result is shown in Fig. 12. The reference underfill thermal conductivity is $k_{\mathrm{UF}} = 0.18$ W/m-K,[21] a standard value for an unfilled epoxy underfill. A simulation with $k_{\mathrm{UF}} = 0$ W/m-K results in a decrease of the effective thermal conductivity from 9.38 to 9.22 W/m-K. This indicates that the effective thermal conductivity of the underfill-μbump array is insensitive to the type of underfill used in dense μbump arrays. In other words, the thermal conduction path between EIC-PIC is dominated by the metallic connections that have a high thermal conductivity.

5.2 Metallic Connections: Cu Lines and μBumps

Two aspects of the metallic connections are investigated by means of a thermal modeling study: (1) the metallic linewidth of the connections to the μbumps in the PIC and (2) the local μbump size and placement around the ring modulator. First, the linewidth is varied in the Cu connections in the M2 BEOL layer from the ring modulator to the μbumps. The reference linewidth is 2.4 μm, which is decreased down to 0.4 μm in the model. The impact on the heater efficiency is plotted in Fig. 14. Up to 30.1% heater efficiency gain is possible for the smallest linewidth. Changing the Cu linewidth obviously impacts the electrical performance of the ring modulator. Mainly, the parasitic resistance between the device and its driver will increase. The impact on the 3 dB bandwidth can be estimated using literature data on equivalent circuit parameters. The interconnect resistance (~1 Ω[22]) is typically much smaller compared with the series resistance of the WG pn-junction (200[23,24] to 350[22] Ω). The electro-optic time constant is $\tau = \mathrm{RC}$, and the 3 dB bandwidth is $f_{3\,\mathrm{dB}} = 1/(2\pi\tau)$, which results in a 1.6% to 2.9% drop in electro-optic bandwidth for a six times higher interconnect resistance in the case of the smallest Cu linewidth. Also, decreasing the Cu linewidth poses a concern for reliability, more specifically electromigration failure.[25] A smaller cross-section results in larger current density. The peak current during the charging of the WG capacitance can be estimated with $I = V/R \cdot e^{-t/\mathrm{RC}}$. Assuming a −2 V ring modulator bias, the peak current amplitude is ~5.5 mA, which is much lower to the expected current for the heater connections of the device (~30 mA) for full free spectral range tuning; furthermore, this peak current is only maintained very briefly and is not constant in time. This all points to the fact that the radio frequency (RF)-connections of the ring modulator are more robust against electromigration failure compared with the heater connections.

Second, the impact of adapting the μbump design is simulated. Two design parameters are varied: the μbump size and placement around the ring modulator. The μbump size is reduced by 50%, and the simulation result is shown in Figs. 13(a) and 13(b). Reducing the size results in two thermal effects: first, the thermal resistance increases because of the lower metallic volume fraction in the μbump layer (desirable effect). Second, if the μbump pitch is kept constant, a smaller μbump results in a larger spacing between the ring modulator and μbump, which limits unwanted horizontal heat spreading. These effects combined result in a +12.1% heater efficiency (Fig. 14). Now, we explore further the effect of increasing the spacing between the μbump and ring

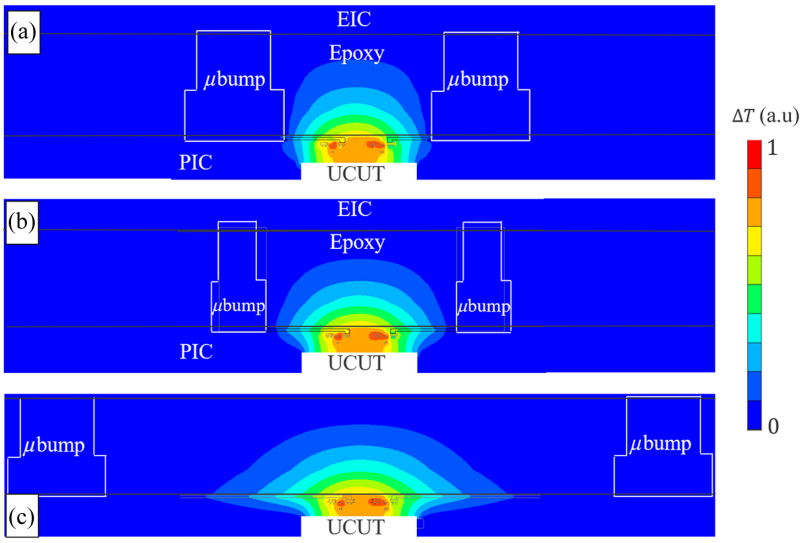

Fig. 13 Cross section of ring modulator and μbumps, with simulated and normalized temperature contours. (a) the reference design, (b) the adapted design with −50% μbump size, and (c) the increased gap between ring modulator and μbumps.

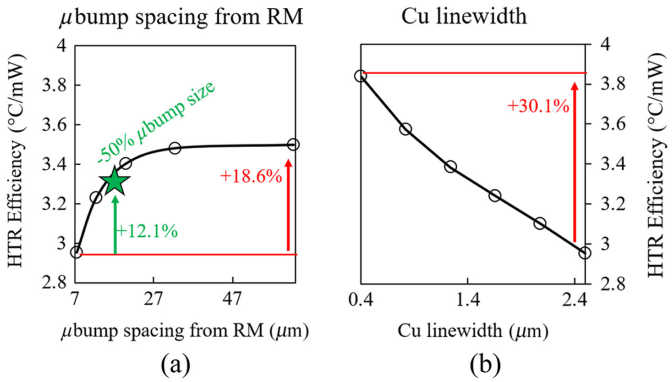

Fig. 14 Simulation result: (a) the heater efficiency as a function of the gap between the ring modulator and μbump. The green data point is obtained for −50% size μbumps at the original pitch (50 μm); see Fig. 13(b). Decreasing μbump size effectively increases the μbump-RM gap. (b) The heater efficiency as a function of the Cu linewidth.

modulator. The reference μbump size is kept and the spacing of the two closest μbumps is increased. In Fig. 15, this concept is shown. This results in an irregular μbump pattern, which introduces additional integration challenges such as the loss of periodicity in the array, which are not discussed further. In Fig. 15, the heat flux through the μbumps is shown for the reference case and for the case with 20 μm additional spacing between the ring modulator and μbumps. For the second case, a significant decrease in μbump heat flux is observed because the local heat sinking effect is decreased. This effect is further illustrated in Fig. 13, where a cross section of the model is shown with temperature contours. In Fig. 13(a), the reference design is shown. In Fig. 13(c), the design with added μbump spacing is shown. The effect on the heater efficiency in function of μbump spacing is plotted in Fig. 14. Moving the μbumps away from the device increases the heater efficiency up to 18.6%. After a spacing of ∼20 μm between the edge of the RM and the μbump, there are diminishing returns for further spacing increases.

Furthermore, changing the material composition of the μbumps or the EIC itself could also be done for further thermal improvements, but this is left for follow-up research. In conclusion regarding the thermal improvements, in combining smaller μbumps with larger spacing between the RM and μbumps together with decreased Cu linewidth, an overall heater efficiency gain of

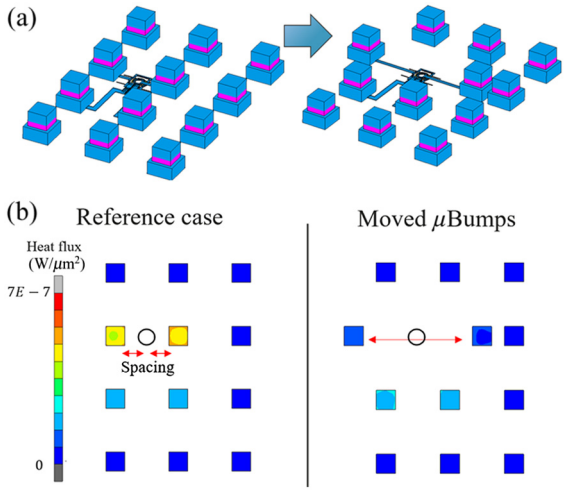

Fig. 15 (a) Model showing the ring modulator and 4 × 3 array of μbumps in its close vicinity and increased gap between the ring modulator and its μbumps. (b) Top view of the heat flux through μbumps for the reference case and moved μbumps case.

49.2% is obtained. This improvement is relative to the reference case with EIC bonded on top of the PIC. In absolute terms, the heater efficiency is increased from 3.08 to 4.60 K/mW (Fig. 8), which indicates that the thermal penalty on the heater efficiency due to the EIC stacking is reduced from 43% to 20% though the proposed mitigation solutions. The efficiency loss due to heat spreading in the EIC can be offset partially, but not fully. For further improvements, more drastic design changes, such as increased UCUT area, are required.

6 Conclusion

In this work, the thermal impact of 3D hybrid photonic–electronic integration is investigated for a ring-based DWDM transceiver architecture. Because of the thermal sensitivity of the ring modulators, they require active thermal tuning for stable operation. The efficiency of the integrated heaters and the thermal crosstalk between them are two main points of interest. Experiments are carried out with and without EIC flip-chipped on top of the PIC with μbumps, for ring modulators with and without substrate UCUT. The experimental results are used for the validation of finite element models. From the analysis of the results, we conclude the following:

1. On the heater efficiency: this metric is decreased with 43.3% after EIC integration with μbumps. This has a significant impact on overall thermal tuning energy consumption. From simulation, it is concluded that this is caused by the heat conduction and spreading inside the EIC. A thermally aware design with an additional gap between the ring modulator and μbumps and a smaller Cu linewidth can largely offset the loss in heater efficiency.
2. On the thermal crosstalk: the calibrated finite element model is used to simulate thermal crosstalk with and without EIC. The addition of the EIC increases thermal crosstalk with 44.4%. Also, if the transceiver is cooled by a liquid cooled cold plate on the top side, thermal crosstalk can be lowered with a factor of 2.

The integration of the EIC is instrumental for low electrical parasitics and good RF performance of the transceiver. However, in this work, it is concluded that the thermal behavior of this 3D EIC-PIC integration needs to be carefully monitored as the close integration of the EIC, without thermal mitigation strategies, causes a heat spreading effect resulting in a reduced heater efficiency and increased thermal coupling.

Disclosures

The authors declare no conflicts of interest related to this work.

Code and Data Availability

The data that support the findings of this study are available from the corresponding author, David Coenen, upon reasonable request.

Acknowledgments

This work was carried out as part of imec's industry affiliation R&D program "optical I/O."

References

1. Cisco, "Cisco Annual Internet Report—Cisco Annual Internet Report (2018–2023) White Paper," https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html (accessed January 2023).
2. Y. Ban et al., "Highly optimized O-band Si ring modulators for low-power hybrid CMOS-SiPho transceivers," in *Opt. Fiber Commun. Conf. (OFC)*, San Diego, CA, (March 2023).
3. R. Nagarajan et al., "2.5D heterogeneous integration for silicon photonics engines in optical transceivers," *IEEE J. Sel. Top. Quantum Electron.* **29**(3), 1–10 (2023).
4. B. G. Lee et al., "Beyond CPO: a motivation and approach for bringing optics onto the silicon interposer," *IEEE J. Lightwave Technol.* **41**(4), 1152–1162 (2022).
5. CPO JDF, "Co-packaged optic assembly guidance document," http://www.copackagedoptics.com/wp-content/uploads/2020/05/CPO-Assembly-Guidance-Doc-V1.0-FINAL.pdf (accessed January 2023).
6. J. S. Lee et al., "Meeting the electrical, optical, and thermal design challenges of photonic-packaging," *IEEE J. Sel. Top. Quantum Electron.* **22**(6), 409–417 (2016).
7. D. Coenen et al., "Thermal characterization of hybrid, flip-chip InP-Si DFB lasers," *MDPI Micromach.* **14**(2), 381 (2023).
8. A. Marinins et al., "Wafer-scale hybrid integration of InP DFB lasers on Si photonics by flip-chip bonding with sub-300 nm alignment precision," *IEEE J. Sel. Top. Quantum Electron.* **29**(3), 1–11 (2023).
9. K. Padmaraju and K. Bergman, "Resolving the thermal challenges for silicon microring resonator devices," *Nanophotonics* **4**(4–5), 269–281 (2014).
10. D. Coenen et al., "Thermal modelling of silicon photonic ring modulator with substrate undercut," *IEEE J. Lightwave Technol.* **40**(13), 4357–4363 (2022).
11. A. Masood et al., "CMOS-compatible Tungsten heaters for silicon photonic waveguides," in *9th Int. Conf. Group IV Photonics (GFP)*, pp. 234–236 (2012).
12. I. Shubin et al., "Integration, processing and performance of low power thermally tunable CMOS-SOI WDM resonators," *Opt. Quantum Electron.* **44**, 589–604 (2012).
13. E. Beyne et al., "The 3-D interconnect technology landscape," *IEEE Des. Test* **33**(3), 8–20 (2016).
14. MSC Software Corporation, "Marc: advanced non-linear simulation solution," https://www.mscsoftware.com/product/marc (accessed 5 April 2022).
15. H. Oprins et al., "Thermal aspects of silicon photonic interposer packages," in *21st Int. Conf. Thermal, Mech. and Multi-Phys. Simul. and Exp. in Microelectron. and Microsyst. (EuroSim Conf.)* (2020).
16. H. Oprins et al., "Thermal modeling and model validation for 3D stacked ICs," in *Handbook of 3D Integration: Design, Test, and Thermal Management*, P. D. Franzon, E. J. Marinissen, and M. S. Bakir, Eds., Chap. 18, Vol. **4**, Wiley (2019).
17. H. Oprins et al., "Experimental characterization of the vertical and lateral heat transfer in three-dimensional stacked die packages," *ASME J. Electron. Packag.* **138**(1), 010902 (2016).
18. D. Coenen et al., "Analysis of thermal crosstalk in photonic integrated circuit using dynamic compact models," *IEEE Trans. Components Packag. Manuf. Technol.* **12**(8), 1350–1357 (2022).
19. S. Biasi et al., "On the effect of the thermal crosstalk in a photonic feed-forward neural network based on silicon microresonators," *Front. Phys.* **10**, 1093191 (2022).
20. H. Oprins et al., "3D wafer-to-wafer bonding thermal resistance comparison: hybrid Cu/dielectric bonding versus dielectric via-last bonding," in *19th IEEE Intersoc. Conf. Thermal and Thermomech. Phenom. in Electron. Syst. (ITherm)*, pp. 219–228 (2020).
21. H. Oprins et al., "Transient analysis based thermal characterization of die-die interfaces in 3D-ICs," in *13th InterSociety Conf. on Thermal and Thermomech. Phenom. in Electron. Syst. (ITherm)*, pp. 1395–1404 (2012).
22. M. Shin et al., "A linear equivalent circuit model for depletion-type silicon microring modulators," *IEEE Trans. Electron. Devices* **64**(3), 1140–1145 (2017).
23. J. Rhim et al., "Verilog-a behavioral model for resonance modulated silicon micro-ring modulator," *Opt. Express* **23**(7), 8762–8772 (2015).
24. M. Kim et al., "Large-signal SPICE model for depletion-type silicon ring modulators," *Photonics Res.* **7**(9), 948–954 (2019).
25. D. Coenen et al., "Electromigration performance improvement of metal heaters for Si photonic ring modulators," *IEEE Trans. Device Mater. Reliab.* **22**(3), 417–423 (2022).

David Coenen received his MSc degrees in aerospace engineering and mechanical energy engineering from KU Leuven and Ghent University, in 2018 and 2020, respectively. He received his PhD in materials engineering from KU Leuven in 2023, focusing on thermal management of Si photonic optical transceivers. He is currently a researcher at imec, with research interests in Si photonics and advanced cooling solutions.

Minkyu Kim received his BS degree and PhD in electrical and electronics engineering from Yonsei University, Republic of Korea, in 2015 and 2021, respectively. His doctoral thesis concerned the monolithic silicon photonics transmitter design with a ring modulator and its temperature controller. Since 2021, he has been an R&D engineer at imec, where he is currently working on high-speed silicon photonics modulator design as well as I/O system design.

Herman Oprins is a principal member of technical staff and R&D team leader at imec, where he is leading the thermal modeling and characterization team. He received his MSc degree and PhD in mechanical engineering from the KU Leuven, Belgium. In 2003, he joined imec, where he has been involved in the thermal experimental characterization, thermal modeling, and thermal management solutions ranging from the device level to the chip level to the system level.

Yoojin Ban received her ME degree in electrical and electronics engineering from Yonsei University, Seoul, Republic of Korea, in 2015. Then she was responsible for the Si modulator devices and Si Photonics optical transceiver at imec for 6 years. Since 2021, she has been a program manager of the industry-affiliation R&D program on optical I/O at imec, which targets the development of an industrially viable short-reach optical interconnect technology based on silicon photonics.

Dimitrios Velenis is the leader of the 3D and silicon photonics device and components group at imec. He has been with imec for more than 15 years, working on 3D and silicon photonics interconnects. He received his MSc degree and PhD from the University of Rochester. Previously, he worked as an assistant professor at Illinois Institute of Technology and as a research associate at the University of Rochester. He has contributed to more than 70 papers in journals and conference proceedings.

Joris Van Campenhout is a fellow silicon photonics and director of the optical I/O R&D program at imec. Prior to joining imec in 2010, he was a post-doctoral researcher at IBM's TJ Watson Research Center (USA). He received his PhD in electrical engineering from Ghent University, Belgium, in 2007.

Ingrid De Wolf received her PhD in physics from KU Leuven, Belgium, in 1989. In the same year, she joined imec, Belgium, where she worked in the field of microelectronics reliability. From 1999 to 2014, she headed the reliability research group. She has (co)-authored 14 book chapters and more than 350 peer-reviewed journal articles. She is a fellow at imec, IEEE senior member, and a full professor in materials engineering of the KU Leuven.