



Heterogeneous integration in silicon photonics: opportunities and challenges: opinion

ROEL BAETS^{1,2,*}  AND ABDUL RAHIM^{1,2}

¹Photonics Research Group, Ghent University – IMEC, Technologiepark-Zwijnaarde 126, B-9052, Gent, Belgium

²Center of Nano- and Biophotonics, Ghent University, Technologiepark-Zwijnaarde 126, B-9052, Gent, Belgium

*roel.baets@ugent.be

Abstract: The application span of silicon photonics is rapidly evolving from high speed transceivers for data- and telecommunication to a broad range of functionalities for many different markets, especially in the sensing and computing space. As a result, the demand for new building blocks and enhanced performance is accelerating and diversifying. Heterogeneous integration of new materials, chips and thin-film chipllets is becoming of key importance in this context. But the implementation of industrial supply chains for this diverse need will be challenging and may require a new supply chain model with dedicated standardization and test methods at the interface between the actors involved. This opinion article discusses opportunities and challenges associated with heterogeneous integration in silicon photonics, in particular with respect to future market growth and the design of process flows for heterogeneous integration.

© 2023 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

Silicon photonics is the field of Photonic Integrated Circuits (PICs) that can be fabricated in the fabs of the microelectronics industry. It takes advantage of many decades of massive investment in ever more performant fabrication tools and processes on large silicon wafers. The field was founded in the 1980's with reports of silicon-based waveguides [1] operating at O and C band (1.3/1.6 μm wavelengths), that are popular in fiber optics, soon followed by reports of modulation of light in such waveguides by building p-n junctions in the waveguide and exploiting the plasma dispersion effect [2]. In the early nineties, a high-index-contrast silicon waveguide was created by patterning a Silicon-On-Insulator (SOI) wafer [3]. This type of SOI-waveguide would become – and still is today – one of the dominant waveguide modalities in the field of silicon photonics [4].

In the early years 2000, it was demonstrated that the mainstream patterning technique of microelectronics – deep ultraviolet (DUV) lithography – was capable of producing low-loss SOI-waveguides [5]. In the same period the first practical structures to couple light from a single-mode optical fiber to the chip (and back) were demonstrated [6]. Also the first waveguide-integrated photodiodes were demonstrated through the epitaxial growth of germanium on silicon [7]. With that, the main ingredients of what would become silicon photonics PIC-platforms were on board: low-loss waveguides, coupling structures to optical fiber, optical modulators and detectors, both with high speed operation capability. We analyzed the R&D efforts of the last 25 years in silicon photonics to plot the dashed lines in Fig. 1. These dashed lines represent the main trends of the evolution of the performance of SOI-based silicon photonics between 2000 and 2025.

Our research has shown that today there are approximately eight Complementary Metal-Oxide-Semiconductor (CMOS) foundries, four Integrated Device Manufacturers (IDMs) and approximately 20 research institutes around the world that have developed a mature SOI-PIC process flow with these ingredients on 200 or 300 mm wafers. The technology platforms of these fabs act as a mature vehicle to realize ultracompact and cost-effective high-speed transceivers (with aggregate data rate of 100-400 Gb/s through advanced modulation and multiplexing methods) for datacenters and telecommunication networks that convert high speed digital electrical data




Fig. 1. The performance evolution for key SOI-based silicon photonics building blocks (dashed lines). Integrating new materials boosts the performance of SOI-based silicon photonics (solid lines). The points in the figure are based on Ref. [5,7,10,11], and [27,–34]. The performance trends are plotted for telecom wavelengths. The dots on the dashed and solid lines represent the early and state-of-the-art results for the respective trends.

into optical signals for transport over optical fiber and then back into electrical data [8–11]. It is estimated that several million of such transceivers based on silicon photonics are being manufactured each year [8]. In recent years research efforts have boosted the performance of such transceivers further to symbol rates beyond 100 GBaud [8,9].

While an annual volume of several million products is substantial from an optics and photonics perspective and even so from an optical communication perspective, it is small from a semiconductor industry perspective. A 200 mm wafer can harbor approximately thousand 5×5 mm chips – a size large enough for a basic optical transceiver function. So it takes only 50 wafer lots with 20 wafers each to manufacture a million such chips. And half so many in case of 300 mm wafers. This is to be contrasted with typical wafer capacities of semiconductor companies, that often surpass 100,000 wafer starts per month (WSPM). So silicon photonics is small business for most of the industrial players in the field, even if the sales price of a silicon photonics wafer is often a multiple of that of an electronics wafer (for the same technology node). So one can wonder whether the current market size of silicon photonics is sufficiently large to guarantee the sustainability of the supply chain.

There is hope for substantial growth. Market analysts predict double digit growth figures for the transceiver market [12], and technology leaders are keen in deploying silicon photonics solutions for AI applications [13]. But, more importantly, there is the expectation that silicon photonics will move in the direction of becoming a generic technology for many applications and markets, some of which may even be high-volume consumer markets. These applications include microwave photonics (serving for example 5 G wireless networks) [14], Light Detection and Ranging (LIDAR) [15], neuromorphic computing [16], quantum computing [17], Augmented

Reality (AR)/Virtual Reality (VR), environmental and industrial sensing, medical sensing, and more.

Our market research has shown that there are at least 100 companies around the world that are in the process of developing silicon-photonics-based products for these markets. Worth noting in this context is that the size of silicon photonics chips for applications such as beam-forming and computing is typically much larger than the 5×5 mm chip mentioned earlier and can in certain cases even fill an entire reticle. This obviously boosts the wafer manufacturing volumes. Time will tell to what degree all the developments towards new markets will be successful, but it seems safe to predict that the wafer manufacturing will grow at least tenfold and possibly hundredfold in the next decade. At that point the making of silicon photonics chips could easily become a thriving business.

With the proliferation of silicon photonics into many different markets there is also an increasing demand for diversity in the needs and requirements. The most obvious driver for diversity is wavelength. Not all applications can operate in the optical communication windows of 1.3 or 1.55 μm . This is particularly true in the sensing space. Some applications require operation in the visible or the very near infrared (below 1.2 μm), others require operation at a wavelength beyond 2 μm or even in the mid infrared.

Next, there is performance. The SOI-based silicon photonics platform has limitations with respect to the performance of the basic components, either because of fundamental physical limitations of the materials involved or because of imperfections in fabrication. Nonlinear absorption and free carrier absorption [18] are in the first category, sidewall roughness and waveguide thickness and width inaccuracy [18,19] are in the second.

Next there is functionality. Passive optical functions along with modulation and detection are possible, but other functions, such as light sources, optical isolators, non-volatile building blocks and ultra-compact modulators with low-power dissipation have not been possible in most of the mainstream platforms.

As a response to this diversity the research community has developed an extremely rich set of activities, all with the aim of extending silicon photonics with new functions or better components, while respecting as much as possible the original “credo” of silicon photonics, which is to manufacture PICs on silicon wafers with the toolset and infrastructure of the microelectronics world. Invariably this has led to the addition of new materials to silicon photonics platforms or the replacement of one material by another. A prominent example is the use of silicon nitride as the waveguide core material, instead of silicon. This has allowed to extend the spectral range of operation towards visible wavelengths as well as to improve the performance of passive building blocks (see Fig. 1), while at the same time active modulation and detection functions were no longer possible. Not surprisingly, in some platforms one has added a silicon nitride waveguide layer on top of an SOI waveguide layer, thereby combining the best of two worlds and extending the design freedom in optical wiring. Silicon nitride is an “easy” material in this context because it is well known in a microelectronic fab and it can be integrated monolithically on silicon wafers through chemical vapour deposition methods.

For many other materials, in particular crystalline materials, the integration is much more demanding and one needs to resort to “hybrid” or “heterogeneous” integration methods. Semantically speaking, the precise distinction between “hybrid” and “heterogeneous” is not fully resolved, but in what follows we will use the term “heterogeneous”. In our view, heterogeneous integration refers to a non-monolithic wafer-scale integration technique where the same silicon wafer hosts the “native” silicon photonics devices and the newly integrated materials, wafers, chips or chiplets. A building block that is missing in most platforms – and that is key to some of the applications – is an integrated light source or an optical amplifier. Therefore a large-scale research effort has been spent over the past 10-15 years to develop methods to integrate III-V semiconductor lasers or optical amplifiers on silicon photonics wafers [8,20]. A wide range of

methods is being used for this purpose, including: (1) attachment of micro-optic benches (MOB) on the silicon photonics wafer [21], (2) flip-chip mounting of III-V dies [22], (3) die-to-wafer or wafer-to-wafer bonding of unprocessed epitaxial layer stacks [23], (4) micro-transfer printing of processed thin-film III-V chipllets [24] and (5) monolithic direct epitaxy of III-V onto silicon [25]. The first and the third of these techniques have matured up to the point of industrial manufacturing in decent volume.

Apart from the III-V semiconductors (InP, GaAs-, GaN-, or GaSb-based, depending on target wavelength), there is a long list of other materials that one may want to integrate heterogeneously on silicon photonics wafers so as to create new functions or enhance performance. In recent years there has been considerable interest in electro-optic materials such as lithium niobate (LiNbO_3), barium titanate (BTO) and poled organic polymers to bring pure phase modulation at ultra-high speed into the world of silicon photonics [26]. But there are many other “exotic” materials that are being explored, such as graphene and other 2D-materials, rare-earth-doped oxides, magneto-optic materials such as Ce:YIG, piezoelectric materials such as Lead Zirconate Titanate (PZT), phase change materials such as $\text{Ge}_2\text{Sb}_2\text{Te}_5$ (GST), liquid crystals and more. Figure 1 (full lines) shows the trend of the performance boost enabled by the integration of new materials with SOI-based silicon photonics platforms.

But here comes the caveat. Diversity brings dilution of effort. This is not a major problem at the research level but, in our view, it is a substantial barrier for the development and maintenance of a manufacturing process flow, especially when talking about process flows in wafer-level fab environments. Such a development only makes economic sense if the manufacturing volume reaches a break-even point. While it is difficult enough to reach break-even for the more generic silicon photonics platforms, it will be more challenging for “niche” material combinations, especially in those cases where the heterogeneous integration of a given material has a major impact on the rest of the process flow, for example due to the thermal budgets involved. Another factor of worry is that some of the materials mentioned earlier are not very well accepted in a CMOS-environment, because of uncertainty with respect to the impact of contamination on other process steps and other wafers. But the fact remains that for each material combination that is considered for heterogeneous integration there is a good rationale from an application point of view.

Do we face a blocking chicken-and-egg problem here? Some of the promising markets cannot grow because there is no manufacturing environment for the heterogeneous integration that is required for that market, while at the same time dedicated manufacturing capacity is not being invested into because the market still looks too small and risky. And even if there is a potential large-volume market, the path to that market passes via low-volume and medium-volume. This blocking situation may be the fate for certain cases but there may be smart ways to mitigate the problem.

As discussed earlier there are many technologies for heterogeneous integration [21–24] and the non-recurrent engineering cost for developing a manufacturing supply chain and process flow can be very different. The worst-case scenario is that a heterogeneous integration step of material A would lead to substantial modifications of the generic process flow for the silicon photonics wafer, while for another material B those modifications are once more entirely different. Ideally one would hope to use, as much possible, generic process flows that allow for heterogeneous integration of materials A, B, C.. with minimal customization of the flow depending on material. This would suggest that the heterogeneous integration happens best late in the process flow, at the back-end-of-line level, possibly even in an environment that is less risk averse with respect to contamination issues. In other words we may see an evolution in the field of silicon photonics in which the wafer-level processing is no longer done in one fab, but where it involves multiple fabs. There are those that execute a generic front-end-of-line process flow that fits many purposes; there are those that prepare wafers, chips or chipllets in special materials; and finally there are

those that execute the actual heterogeneous integration and associated wafer-level postprocessing. We believe that such a model demands a great need for standardization of the “interfaces” and their specification and testability. Today this is still poorly developed. It is also far from trivial because the testability of the objects involved (wafers, naked chips, naked chipllets) is somewhat limited and requires dedicated tools. Our recommendation therefore is to gear up research on methods that allow to specify and test the relevant interface parameters so as to create a trustable and reliable supply chain for a wide variety of heterogeneous silicon photonics chips, ready to enable many markets.

Acknowledgments. The authors thank Tom Vanackere for assisting in drawing figures. The authors thank ePIXfab – the European silicon photonics alliance (<https://epixfab.eu>) for providing the data presented in this article.

Disclosures. The authors declare no conflicts of interest.

Data availability. No data were generated or analyzed in the presented research.

References

1. R. Soref and J. Lorenzo, “Single-crystal silicon: a new material for 1.3 and 1.6 μm integrated-optical components,” *Electron. Lett.* **21**(21), 953–954 (1985).
2. R. Soref and J. Lorenzo, “All-silicon active and passive guided-wave components for $\lambda = 1.3$ and 1.6 μm ,” *IEEE J. Quantum Electron.* **22**(6), 873–879 (1986).
3. J. Schmidtchen, A. Splett, B. Schuppert, *et al.*, “Low loss single mode optical waveguides with large cross-section in silicon-on-insulator,” *Electron. Lett.* **27**(16), 1486–1488 (1991).
4. X. Chen, M. Milosevic, S. Stankovic, *et al.*, “The emergence of silicon photonics as a flexible technology platform,” *Proc. IEEE* **106**(12), 2101–2116 (2018).
5. P. Dumon, W. Bogaerts, V. Wiaux, *et al.*, “Low-loss SOI photonic wires and ring resonators fabricated with deep UV lithography,” *IEEE Photonics Technol. Lett.* **16**(5), 1328–1330 (2004).
6. D. Taillaert, P. Bienstman, and R. Baets, “Compact efficient broadband grating coupler for silicon-on-insulator waveguides,” *Opt. Lett.* **29**(23), 2749–2751 (2004).
7. D. Ahn, C. Hong, J. Liu, *et al.*, “High performance, waveguide integrated Ge photodetectors,” *Opt. Express* **15**(7), 3916–3921 (2007).
8. S. Shekhar, W. Bogaerts, L. Chrostowski, *et al.*, “Silicon photonics – roadmapping the next generation,” *arXiv*, arXiv:2305.15820, (2023).
9. J. Witzens, “High-speed silicon photonics modulators,” *Proc. IEEE* **106**(12), 2158–2182 (2018).
10. X. Wu, R. Kumar, D. Huang, *et al.*, “A single-chip high-speed silicon photonic transmitter with integrated laser and micro-ring modulator,” *2023 IEEE Silicon Photonics Conference (SiPhotonics)*, Washington, DC, USA, (2023), 1–2.
11. K. Li, S. Liu, D. J. Thomson, *et al.*, “Electronic–photonic convergence for silicon photonics transmitters beyond 100 Gbps on–off keying,” *Optica* **7**(11), 1514–1516 (2020).
12. <https://www.yolegroup.com/press-release/silicon-photonics-to-soi-and-beyond/>.
13. <https://www.eenewseurope.com/en/tsmc-teams-with-nvidia-broadcom-on-silicon-photonics/>.
14. R. Sabella, “Silicon photonics for 5 G and future networks,” *IEEE J. Sel. Top. Quantum Electron.* **26**(2), 1–11 (2020).
15. C. V. Poulton, M. J. Byrd, P. Russo, *et al.*, “Coherent LiDAR With an 8,192-element optical phased array and driving laser,” *IEEE J. Sel. Top. Quantum Electron.* **28**(5: Lidars and Photonic Radars), 1–8 (2022).
16. B. J. Shastri, A. N. Tait, T. F de Lima, *et al.*, “Artificial intelligence and neuromorphic computing,” *Nat. Photonics* **15**(2), 102–114 (2021).
17. T. Rudolph, “Why I am optimistic about the silicon-photonic route to quantum computing,” *APL Photonics* **2**(3), 030901 (2017).
18. A. Rahim, E. Ryckeboer, A. Z. Subramanian, *et al.*, “Expanding the silicon photonics portfolio with silicon nitride photonic integrated circuits,” *J. Lightwave Technol.* **35**(4), 639–649 (2017).
19. K. Okamoto, “Wavelength-division-multiplexing devices in thin SOI: Advances and prospects,” *IEEE J. Sel. Top. Quantum Electron.* **20**(4), 248–257 (2014).
20. D. Thomson, A. Zilkie, J. E. Bowers, *et al.*, “Roadmap of silicon photonics,” *J. Opt.* **18**(7), 073003 (2016).
21. I. L. Bundalo, P. E. Morrissey, A. Annoni, *et al.*, “PIXApp photonics packaging pilot line – development of a silicon photonic optical transceiver with pluggable fiber connectivity,” *IEEE J. Sel. Top. Quantum Electron.* **28**(3), 1–11 (2022).
22. A. Marininis, S. Hänsch, H. Sar, *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).
23. Z. Wang, A. Abbasi, D. Utsav, *et al.*, “Novel light source integration approaches for silicon photonics,” *Laser Photonics Rev.* **11**, 4 (2017).
24. J. Zhang, B. Haq, J. O’Callaghan, *et al.*, “Transfer-printing-based integration of a III-V-on-silicon distributed feedback laser,” *Opt. Express* **26**(7), 8821–8830 (2018).
25. Y. De Koninck, C. Caer, D. Yudistira, *et al.*, “GaAs nano-ridge laser diodes fully fabricated in a 300 mm CMOS pilot line,” *arXiv*, arXiv:2309.04473 (2023).

26. A. Rahim, A. Hermans, B. Wohlfeil, *et al.*, “Taking silicon photonics modulators to a higher performance level: state-of-the-art and a review of new technologies,” *Adv. Photonics* **3**(02), 2 (2021).
27. T. Horikawa, D. Shimura, H. Okayama, *et al.*, “A 300-mm silicon photonics platform for large-scale device integration,” *IEEE J. Sel. Top. Quantum Electron.* **24**(4), 1–15 (2018).
28. W. D. Sacher, Y. Huang, G.-Q. Lo, *et al.*, “Multilayer silicon nitride-on-silicon integrated photonic platforms and devices,” *J. Lightwave Technol.* **33**(4), 901–910 (2015).
29. M. He, M. Xu, Y. Ren, *et al.*, “High-performance hybrid silicon and lithium niobate Mach–Zehnder modulators for 100 Gbit s⁻¹ and beyond,” *Nat. Photonics* **13**(5), 359–364 (2019).
30. Ansheng Liu, Richard Jones, Ling Liao, *et al.*, “A high-speed silicon optical modulator based on a metal–oxide–semiconductor capacitor,” *Nature* **427**(6975), 615–618 (2004).
31. M. Eppenberger, A. Messner, B. I. Bitachon, *et al.*, “Resonant plasmonic micro-racetrack modulators with high bandwidth and high temperature tolerance,” *Nat. Photonics* **17**(4), 360–367 (2023).
32. S. Lischke, A. Peczek, J. S. Morgan, *et al.*, “Ultra-fast germanium photodiode with 3-dB bandwidth of 265 GHz,” *Nat. Photonics* **15**(12), 925–931 (2021).
33. D. Schall, E. Pallecchi, G. Ducournau, *et al.*, “Record high bandwidth integrated graphene photodetectors for communication beyond 180 gb/s,” *2018 Optical Fiber Communications Conference and Exposition (OFC)*, pp. 1–3 (2018).
34. M. Koepfli, M. Baumann, S. Giger, *et al.*, “High-speed graphene photodetection: 300 ghz is not the limit,” *2021 Conference on Lasers and Electro-Optics Europe & European Quantum Electronics Conference (CLEO/Europe-EQEC)*, Munich, Germany, p. 1, (2021).