

Myths and rumours of silicon photonics


Tom Baehr-Jones, Thierry Pinguet, Patrick Lo Guo-Qiang, Steven Danziger, Dennis Prather and Michael Hochberg

Low-cost manufacturing, high yields and seamless on-chip integration with electronics are often touted as the guaranteed benefits of silicon photonics, but is this really the case? Michael Hochberg and colleagues explain that the situation is much more complex in reality.

The field of silicon photonics has expanded substantially in recent years and become increasingly diverse in its applications. Today, silicon-based platforms support the realization of a wide variety of devices, including high-speed modulators¹ and detectors², low-loss waveguides³ and other passive components. One theme that has emerged as the field has grown is that silicon, as a platform, does not provide best-in-class devices for every task.

This isn't at all surprising, and in fact it mirrors a development in the electronics industry: silicon does not dominate consumer electronics because CMOS transistors provide the highest performance in every metric. In microelectronics, silicon dominates largely because of the huge investments that have been made in manufacturing infrastructure to create chip-scale systems containing billions of components. Silicon was chosen initially because it was the platform that allowed yield to be improved quickest, due to the convenient and uniform growth of thermal oxide. Today, CMOS circuits commonly consist of billions of elements; no other platform provides similar capabilities in terms of yield and complexity. But for individual devices, alternative mature material systems such as group III–V semiconductors, silicon-on-insulator and silicon-on-sapphire offer far better performance in some metrics than generic bulk silicon. One effort that has attempted to address this trade-off is being funded by the Defense Advanced Research Projects Agency in the USA, where a number of teams⁴ are attempting to integrate group III–V semiconductor-based transistors (as well as optical components) directly on top of a CMOS platform.

A worrying dynamic now emerging in silicon photonics is the desire to obtain best-in-class performance at the expense



M. HOCHBERG

Silicon photonic components must operate at much lower voltages and energy-per-bit metrics before it will be worth integrating them on-chip with a CPU.

of process- or material-compatibility with other devices. Furthermore, engaging in iterative process modifications to improve the performance of a single device is often orthogonal to the efforts needed to turn devices into working integrated systems. Although process iteration is a core activity for device-level development, this kind of activity is very different from the kind of work required to build complex on-chip systems, where yield and variation are critical parameters.

It should not be forgotten that the key benefits of silicon photonics include the ability to integrate photonics with electronics, as well as the simplicity and potentially low cost of manufacturing. Being able to reuse the manufacturing infrastructure developed for electronics grants huge advantages in terms of cost, time-to-market and complexity scaling. Furthermore, chip-scale integration will

make it possible to create new form factors and functions that simply could never be addressed with bulk devices. The bottom line is that individual silicon devices will probably not outperform single devices based on other material platforms, except in a few particular areas. The applications that will benefit most from the silicon platform are those that require many devices to be strung together into a complex system, just like in electronics.

There are many broad generalizations regarding silicon photonics. Our goal in this Commentary is to explore and perhaps dispel some of these generalizations, based on our experience in building and commercializing silicon photonic systems.

“Silicon means low cost”

This is emphatically false in the case of small production volumes. The cost of low-volume manufacturing in silicon turns out

to be pretty similar to that of other material systems. Silicon only provides a defining 'win' on cost when it provides higher device yields or pushes costs down by sharing processes and facilities or exploiting high production volumes.

The per-gate costs for high-volume silicon electronics are staggeringly low, and continue to decrease. But this comes at a price: up-front investments in fabrication infrastructure run into billions of dollars, and tooling costs for a specific design can easily reach tens of millions of dollars. For the chips that emerge from this infrastructure to be cheap, it is necessary to divide these costs over many millions of units. In photonics, although up-front costs can be significantly reduced through the use of shuttle runs and processes that promote shared foundry infrastructure, they cannot be eliminated completely.


There are approaches to silicon photonic fabrication that do not rely on a standard foundry but instead utilize a captive fabrication facility. However, the ultimate costs and production capacity in such an approach are likely to be similar to building a captive facility for other material systems such as lithium niobate, as the equipment sets are quite similar. The main differences would emerge only from the cost of the starting material, as well as from differences in process tolerance and, consequently, yield.

Getting to a situation where fixed costs are not dominant in silicon photonics will mean building a product that requires hundreds of wafers to be run every month. Mask sets commonly cost many millions of dollars. To reduce costs to the bare minimum, large semiconductor fabs will often run many thousands of product wafers per week. Nothing in silicon photonics is anywhere near this scale. To date, the highest-volume products in silicon photonics, to the best of our knowledge, are the chips produced by Luxtera. Hopefully over the next few years, a wide range of high-volume products will begin to emerge from across the community.

"Silicon makes it easy to get high yields"

Being able to access the yield-control methodologies and tools that are common in advanced microfabrication facilities provides huge benefits in terms of improving device-level yields. Complex systems achieve high yields by implementing high manufacturing volumes, thereby allowing for very precise statistical-process control. The high complexity and yield of silicon electronics is enabled by the large volumes of statistical-process control data that are generated by running huge numbers of wafers.

As a community, our understanding of what drives photonic device yield is still



Increasing component count for on-chip silicon photonic systems, based on recent literature and using the methodology of Hochberg *et al.*²³. Recent results from Luxtera, Kotura, Intel, Bell Labs and the Yoo lab (University of California, Davis) are also included.

quite primitive. Inline photonic test systems do not yet exist, meaning that photonic parameters must often be inferred from metrology and electronic test data. Today, transistor models⁵ often use hundreds of parameters to create faithful and predictive results. Photonic components are every bit as complex in many cases, but our models are comparatively primitive. Finding ways of rapidly and accurately measuring the second-order properties of photonic components (for example, the static coherence length of waveguides, electro-optic noise figures and nonlinearities) will be a key area for innovation over the coming years.

That being said, the tools available for reaching high yields in silicon are extraordinarily powerful. By accessing the same front-end tools used to make transistors, it is common to fabricate entire 200 mm wafers populated with hundreds of thousands of working devices⁶. Although achieving this goal requires careful process design and metrology strategies, the tools used to yield extremely consistent device performance already exist. We suggest that future papers in silicon photonics should include yield and variation information whenever practical, in order to start developing a cultural understanding of the relevant issues in the field.

"The goal is monolithic integration with advanced CMOS transistors"

The history of electronics suggests that the monolithic (single-chip) integration of multiple technologies is generally driven by strong economic and technical factors. One great modern example of process convergence — the integration of bipolar transistors with CMOS (BiCMOS) — emerged in 1992⁷. This technology combines the ultrahigh-radiofrequency performance of

heterojunction bipolar transistors with the low-power operation of CMOS, but at the expense of significantly increased process complexity (and therefore cost). It is worth noting that these BiCMOS fabrication processes operate at linewidths around 130 nm and are therefore many generations behind today's most advanced techniques, which operate at 28 nm and below.

True monolithic integration of photonic devices (modulators, detectors, waveguides and sources, for example) with cutting-edge 28 nm (or smaller) CMOS processes is a very challenging task. Making process modifications to support such integration will fundamentally change the models for the transistors, thus requiring them to be re-characterized. In addition, transistor performance will almost certainly be degraded. Deciding not to change the process at all is one option⁸, and researchers have shown that some photonic functionality can be integrated with minimal post-processing. So far, however, the performance of photonic devices is far from what can be achieved with process modification. Lastly, real estate in a highly scaled CMOS process is extremely expensive; as photonic components are much larger than transistors, the economic barrier to bringing them into an advanced front-end process is significant.

Many of the benefits of monolithic integration can be achieved through multichip integration with through-silicon vias, interposers or front-to-front bonding⁹, probably with little in the way of technical or yield penalties. This is a technology that is already in large-scale production and is now rapidly being driven to maturity in the electronics industry. The history of system-on-a-chip and system-in-package development in the electronics world suggests that a diversity of processes will continue to be valuable for silicon photonics (such as with radiofrequency, power, logic and memory in the electronics world) and that heterogeneous process integration with photonic silicon is going to remain the dominant paradigm for the immediate future.

Unless silicon photonic devices shrink very quickly, the cost of full front-end integration with every successive node on the semiconductor roadmap, in terms of lost area for transistors, will increase. This does not negate the value of silicon photonics, but it certainly suggests that front-end integration into the most advanced processes is not necessarily going to be the universal winning strategy.

Although there are certainly technical advantages to monolithic integration, the development of high-density interconnect

technologies by the electronics industry has dramatically mitigated the penalty for multichip integration. Furthermore, given the costs of modifying high-end CMOS processes, multichip solutions are likely to become dramatically more economical in the near future for a wide variety of applications.

“Silicon means cheap and simple packaging”

In the electronics world, packaging is remarkably cheap because large up-front investments in infrastructure can be amortized across a large volume of product. The precision that is routinely required (and achieved) in the high-end packaging of electronic devices is similar to what is needed to connect two single-mode optical fibres. However, in order for this infrastructure to be successfully reused in silicon photonics, the optics community must figure out how to attach fibres without the need for active alignment. So far, the only silicon photonic products to integrate a light source have used package-level integration, which involves bonding the light source to the back-end dielectric of a CMOS process¹⁰. Many new approaches are being explored, including full front-end integration¹¹, front-end bonding¹² and multichip system-in-package approaches¹³, although none of these are yet in commercial production.

There are two distinct problems here: attaching the fibre, and integrating the light source. Because the packaging costs of photonics are relatively high compared with electronics, we can expect that many of the early products in silicon photonics will be expensive and address low-volume markets. Some of these applications may not even require off-chip optical connections if the light sources can be integrated — one can easily imagine chips for analog signal processing or biosensing whose inputs and outputs are electrical but whose on-chip connections are photonic. Thus, if the laser co-packaging problem is solved, the problem of attaching fibres would be essentially avoided.

The packaging and die-attach costs will continue to represent a substantial fraction of the cost of any finished photonic device, at least until the chip value is dramatically increased. This will continue to be true until silicon photonic products are being sold on the scale of millions of units per year, which will provide the economic drive to invest in dramatically reducing packaging costs.

“On-chip data interconnects are a near-term and dominant application”

Although integration with the most advanced CMOS processes is an extremely challenging goal that drives fundamental

research in the field, it is highly unlikely that this will become practical in the near future for high-volume, low-cost applications such as central processing units (CPUs). Silicon photonic components that are developed without the constraints of cutting-edge (<28 nm) monolithic process integration are now competing with III–V vertical-cavity surface-emitting lasers for many-metre-long optical links. It is widely recognized that silicon photonic components must operate at much lower voltages and energy-per-bit metrics before it will be worth integrating them on-chip with a CPU¹⁴. Our recent work has shown that non-resonant silicon junction-based electro-optic modulators can operate at 25 Gbit s⁻¹ with drive voltages of much less than 1 V, making them compatible with high-performance CMOS electronics for the first time¹⁵. However, silicon is still a long way from being a winning technology for millimetre-scale data links.

As a trend, photonic technologies are generally displacing electronics for progressively shorter links. To address the millimetre-scale links required within a CPU, we would have to leapfrog over the challenges of centimetre-to-metre scale links, which are already very substantial. Moving information with photons tends to be more attractive at longer distances and higher bandwidths because of the low loss of optical fibres and the high power consumption and cable size of electrical links. Competing at ultrashort on-chip length scales is a goal that remains well in the future for the field. Other applications, such as biosensing¹⁶, light detection and ranging¹⁷, radiofrequency oscillators¹⁸, sensors¹⁹, analog components and systems-on-chip²⁰, test equipment, hyperspectral imaging²¹ and inertial sensing²², are more likely to reach commercial scale in the near term.

Conclusion

As more and more silicon photonic devices reach the market, we will have to work together as a community to understand and address the fundamental technical and economic challenges associated with CMOS integration. A variety of different fabrication and packaging solutions, dictated by varying requirements of cost, performance and production volume, are likely to be needed, rather than a single ‘winning’ integration approach. Silicon photonics is an incredibly promising technology, but until the community addresses a number of critical issues it will be no different than a host of other custom, low-volume optical technologies. Silicon photonics will not be a world-changing

technology until we can win against existing approaches at a system level, in terms of both cost and performance. Building new and improved devices is only one part of this problem; system-level integration has not yet received enough attention within the silicon photonic community. Recent work is certainly beginning to show a very healthy shift to systems-level analysis and towards more complex chip-scale systems, both in academia and in industry. As a whole, the community must focus on understanding the technical and economic challenges associated with silicon fabrication and the packaging of integrated photonic circuits. None of these are challenges are insurmountable, as long as we acknowledge and address them. □

Tom Baehr-Jones is at the Department of Electrical Engineering, University of Washington, Seattle, Washington 98195-2500, USA. Michael Hochberg and Dennis Prather are in the Department of Electrical and Computer Engineering, University of Delaware, 140 Evans Hall, Newark, Delaware 19716, USA. Steven Danziger is at BAE Systems North America, 9300 Wellington Rd, Manassas, Virginia 22209, USA. Thierry Pinguet is at Luxtera, Carlsbad, California 92011, USA. Patrick Lo Guo-Qiang is at the Institute of Microelectronics, Agency for Science, Technology and Research (A*STAR), Singapore. e-mail: hochberg@udel.edu

References

1. Reed, G. T., Mashanovich, G., Gardes, F. Y. & Thomson, D. J. *Nature Photon.* **4**, 518–526 (2010).
2. DeRose, C. T. *et al. Opt. Express* **19**, 24897–24904 (2011).
3. Tien, M.-C., Bauters, J. F., Heck, M. J. R., Blumenthal, D. J. & Bowers, J. E. *Opt. Express* **18**, 23562–23568 (2010).
4. Kazior, T. E. *et al. Compound Semiconductor Integrated Circuit Symp.* <http://dx.doi.org/10.1109/csics.2010.5619670> (2010).
5. Chan, M., Hui, K. Y., Hu, C. & Ko, P. K. *IEEE Trans. Electron. Devices* **45**, 834–841 (1998).
6. Ding, R. *et al. IEEE Photon. Conf.* <http://dx.doi.org/10.1109/pho.2011.6110876> (2011).
7. Cressler, J. D. *Silicon Heterostructure Devices* Ch. 1 (CRC, 2008).
8. Orcutt, J. S. *et al. Opt. Express* **19**, 2335–2346 (2011).
9. Motoyoshi, M. *Proc. IEEE* **97**, 43–48 (2009).
10. Mekis, A. *et al. IEEE J. Sel. Top. Quant. Electron.* **99**, 1–12 (2010).
11. Sun, X. C., Liu, J. F., Kimerling, L. C. & Michel, J. *IEEE J. Sel. Top. Quant. Electron.* **16**, 124–131 (2010).
12. Fish, G., Nicholes, S., Kaman, V. & Fang, A. *IEEE Photonics Conf.* <http://dx.doi.org/10.1109/pho.2011.6110438> (2011).
13. Narasimha, A. *et al. Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference* paper OMK7 (OSA, 2008).
14. Miller, D. A. B. *Proc. IEEE* **97**, 1166–1185 (2009).
15. Baehr-Jones, T. *et al.* Preprint at <http://arxiv.org/abs/1203.0767> (2012).
16. Shang, J. *et al. Langmuir* **28**, 3338–3344 (2012).
17. Doyle, J. K. *et al. Opt. Express* **19**, 21595–21604 (2011).
18. Capmany, J. & Novak, D. *Nature Photon.* **1**, 319–330 (2007).
19. Schmidt, B., Almeida, V., Manolatu, C., Preble, S. & Lipson, M. *Appl. Phys. Lett.* **85**, 4854–4856 (2004).
20. Yoo, S. J. B. *et al. in IEEE Trans. Terahertz Sci. Tech.* **2**, 167–176 (2012).
21. Kasprzak, J., Patton, B., Savona, V. & Langbein, W. *Nature Photon.* **5**, 57–63 (2011).
22. Sanders, S. J., Strandjord, L. K. & Mead, D. *Optical Fiber Sensors Conference Technical Digest* 5–8 (2002).
23. Hochberg, M. & Baehr-Jones, T. *Nature Photon.* **4**, 492–494 (2010).