Building 3D integrated circuits with electronics and photonics

Standfirst: The three-dimensional integration of electronic and photonic integrated circuits could solve critical input/output limitations in existing computing chips, and create larger, more complex chips for application in future data centres and high-performance systems.

Chao Xiang & John E. Bowers

In modern data centres and disaggregated computing systems, power consumption due to data movement has become a critical issue, limiting system scale-out and the development of higher data processing capabilities[1]. Integrated photonic technologies — which have the intrinsic advantages of optical connectivity, including low latency, low loss, high integrity and high-bandwidth data transmission — could provide an answer. Applications such as data centre network switches already rely on optical communications, but a move from pluggable optics to co-packaged optics could reduce metal interconnect lengths and thus avoid high power consumption and signal degradation at high frequencies[2]. Co-packaged optics could also address the input/output (I/O) data bottleneck associated with chip-to-chip communication, allowing electronic integrated circuits to perform beyond the limits of single chips with high-throughput inter-chip connectivity.

The requirements of these systems are, however, changing, driven by increasing demands for artificial intelligence (AI) training and inference, high-performance computing, and hyperscale data centre interconnects. And to meet inter-chip interconnect demands at various distances, the role and architecture of both the electronics and photonics in such packages need to be defined and generalized towards powerful electronic integrated circuits with reliable and efficient optical I/Os (Fig. 1).

How 3D integration can help

With photonic integrated circuits, interfacing driver electronics are required for purposes such as signal generation, modulation, detection, and resonance tuning and locking. However, integration strategies can vary across different architectures depending on the availability of chip fabrication process and packaging technologies. One such strategy is three-dimensional (3D) integration, which leverages vertical space to extend on-chip capabilities beyond what is possible with two-dimensional (2D) integration[3].

A benefit of this 3D integration is ultra-short metal interconnects, which replace the long metallic traces that are routed on the chip surface in 2D integration. The short distances in these vertical interconnects are determined by the size of bump connections or 3D interposer layer thicknesses (which are used in through-silicon via technology). Short-distance metal interconnects can dramatically reduce signal loss at high frequencies and remove the need for retimers (which transmit a new copy of a signal) or digital signal processing, providing additional power budget savings.

Future 3D integrated circuits that have logic and/or memory layers are likely to require a photonics I/O layer or module for terabit-per-second-scale or higher planar connectivity at distances greater than millimetres. Moving the photonics closer to the electronics will be advantageous for this due to superior low-loss optical connectivity at such distances. Close electronic–photonic interfaces can also mitigate electronic integration constraints at high driving frequencies and help unlock the full

potential of photonic integrated circuits in terms of performance metrics including bandwidth and power consumption. For instance, optical modulator bandwidth and receiver signal-to-noise ratio is inversely proportional to the parasitic capacitance of driving circuits, and to achieve a low driver capacitance the electrical wiring needs to be short.

However, photonic integrated circuits generally require 28-nm node or older complementary metal–oxide–semiconductor (CMOS) processes, while the power consumption of electronic integrated circuits generally scales with lower processing nodes. Such node mismatch means that the route to monolithic integration of electronics and photonics on the same chip will be challenging, and potentially economically impracticable.

3D integrated optical and electronic modules can provide close electronic interfaces for photonic integrated circuits, and — unlike monolithically integrated photonics and electronics — can support different process nodes. Such modules can thus serve as standalone optical I/O elements that work with different host integrated circuits with advanced nodes.

Furthermore, the photonics and electronics elements in future intra-package interconnects could be merged into more complex active silicon interposers that provide both electrical and optical signal redistribution functions. These interposers must meet stringent interconnect requirements for both electrical bandwidth (signal rate per wavelength) and optical bandwidth (wavelength multiplexing), as outlined by standards such as the universal chiplet interconnect express (UCIe)[4]. The use of advanced packaging (such as micro-bumps and hybrid copper bonding) permits a high density of electronic integrated circuit and photonic integrated circuit combinations in their respective optimal forms: that is, process nodes.

On the electronics side, advanced packaging using microbumps can lead to a bond pad pitch of 10 to 50 µm. Through-silicon vias and hybrid copper bonding can lower these values to less than 10 µm. Such dense electronics can provide sub-femtofarad per square micrometre capacitance — as a result, direct driving of modulators or photodetectors eliminates the 50 Ohm load and thus improves the efficiencies of these photonic–electronic devices.

On the photonics side, higher-order multiplexing and higher data rate-per-lane approaches are being developed to boost optical connectivity[5]. Higher-order multiplexing relies on a high number of multiplexers, demultiplexers, modulators and photodetectors that all require electronics for control, as well as optical-to-electronic conversion. The higher data rate-per-lane approach, on the other hand, sets a limit on the tolerable capacitance, resistance, and inductance of the driving electronics in order to meet bandwidth requirements. Both approaches thus require dense integration with electronics with fine pitches to increase channel numbers and allow higher driving efficiencies. 3D integration naturally fits these demands.

3D multilayer photonic chips

3D integration can also lead to gains in photonic functionalities. Vertical stacking of photonic layers can, in particular, allow different material layers to work seamlessly on one chip[6]. These layers can be stacked by monolithic deposition, heteroepitaxial growth, or heterogeneous bonding, depending on the material properties. The approach can, for instance, be used to add signal generation and detection capabilities to conventional integrated photonic platforms such as silicon nitride and thin-film lithium niobate. Such multi-layered 3D structures can also be effective in minimizing process-

induced degradation[7]. The approach can thus be used to combine active and passive devices with optimized performance on a single 3D chip — a crucial step in the development of complex photonic integrated circuit functionalities for applications such as light detection and ranging (LIDAR)[8], neuromorphic networks[9] and quantum photonic computing[10].

With 3D integration, the number of photonic devices can break the limits of mask reticle sizes and advance density scaling; a similar trend has already occurred in multi-chip heterogeneously-packaged electronics. Furthermore, due to the diffraction limit and variation in optical mode confinement, different waveguide platforms offer different device footprints: silicon waveguides offer a bending radius down to a few micrometres, whereas with silicon nitride waveguides, a bending radii below around 100 μm can result in strong radiation loss due to the much weaker mode confinement. Circuit density mismatch thus exists across the integration platforms. 3D integration can, however, allow mismatched footprints in different layers, and avoids wasting the 2D planar space in layers comprised of dense photonic circuits.

Despite the potential advantages of 3D integrated electronic and photonic circuits, several challenges remain to be addressed. The challenges are focused in four key areas: density scaling, thermal management, energy efficiency, and laser integration.

Density scaling

The density of 3D electronic–photonic integrated circuits rely heavily on packaging technologies, including the pitch density of electronic interfaces and the shoreline density of optical interfaces (which defines the number of parallel spatial optical channels around a chip with limited size). These optical interfaces connect with optical fibres or waveguides, depending on I/O bandwidth density and reach distance. Progress in electronic 3D integrated circuits and optical packaging solutions are thus both critical in maintaining a compact package and reducing the cost of interconnects using high channel-count wavelength division and parallel spatial multiplexing. For instance, the development of vertical broadband optical couplers for fibres or fibre array units could relax shoreline density limitations. Another key development is to replace optical fibres with planar optical waveguides and corresponding couplers for high-density interconnects at centimetre or shorter distances.

Due to the size mismatch that exists between the photonic elements and transistors, effective fan-in and fan-out of electrical traces is needed in a 3D electronic–photonic integrated circuits to save space. When stacking multiple wafers together using through-silicon vias, a keep-out zone is normally required to balance wafer bowing but this requires additional footprints. Moreover, evanescent couplers with low crosstalk support efficient inter-layer transitions in 3D photonic integrated circuits, but the coupler lengths normally remain above 10 μm. Efficient couplers with short lengths are required to further reduce the chip footprint.

Thermal management

Similar to 3D electronic integrated circuits, thermal management is a major challenge for 3D electronic–photonic integrated circuits[11]. With high-density bonding, the generated heat could alter the photonic device performance, especially for phase-sensitive elements (such as ring resonances and laser/amplifier efficiencies). The corresponding thermal management imposes limits on the power budget and needs to be accounted for at the initial design stages. Efficient locking electronic circuits with low control power consumption are also needed for the compensation of temperature-

induced wavelength drift and power loss. Furthermore, excessive heat generation can affect the mechanical reliability of such 3D electronic–photonic integrated circuits, which will require special thermal dissipation techniques and mechanical reliability tests.

Energy efficiency

The potential energy efficiency of electronic–photonic integrated circuits is the main driver for 3D integration. However, power consumption is currently tens of picojoule per bit, and to achieve the desired sub-picojoule per bit energy efficiency both the electronic and photonic power consumption have to be optimized. The respective packaging will also play a role here.

New optical link architectures could help. For example, a dramatic reduction in laser power consumption is expected if distributed feedback laser arrays can be replaced by multi-wavelength comb sources in a wavelength-division-multiplexing system. But comb generation efficiencies and spectral manipulation techniques must first improve[12]. Furthermore, the laser wall-plug-efficiency and the coupling efficiency with downstream components are two important parameters regarding the power consumption from laser sources in an optical link. In addition, low-power-consumption modulators, high-responsivity avalanche photodetectors and low-insertion-loss passives could play key roles in reducing the link power consumption, especially when new architectures require fewer multiplexer/demultiplexer elements. The control and locking circuits for these components, and related temperature management, also consume additional power, which require improved architectures and electronic designs.

Laser integration

There is still debate over the preferred placement method for lasers, and whether to use remote laser sources or on-chip lasers. There have been considerable foundry developments on integrating lasers with silicon photonics in recent years[13]. And with the increase in the number of laser sources needed, integrated on-chip lasers also offer advantages in terms of cost and power consumption. This is due to the ease of scalable fabrication and minimized power coupling loss to downstream components. However, the heat generated by the lasers could be a concern for 3D integration and heat dissipation methods need to be specifically designed. Due to the elevated temperature in 3D chips, integrated lasers also need to be able to operate efficiently at higher temperatures. Concerns about laser reliability in 3D integration can though be largely mitigated if additional lasers are used as backups on the same chip, an approach that is not available with disaggregated lasers[14]. By addressing these issues, 3D electronic–photonic integrated circuits with integrated lasers could provide the highest density and energy efficiency.

Chao Xiang[1] and John E. Bowers[2]
[1]Department of Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong, China
[2]Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA, USA

e-mail: cxiang@eee.hku.hk

Competing interests

The authors declare no competing interests.

References

1. Margalit, N. et al. *Appl. Phys. Lett.* **118**, 220501 (2021).
2. Cheng, Q., Bahadori, M., Glick, M., Rumley, S. & Bergman, K. *Optica* **5**, 1354–1370 (2018).
3. Sabry Aly, M. M. et al. *Proceedings of the IEEE* **107**, 19–48 (2019).
4. Sharma, D. Das, Pasdast, G., Qian, Z. & Aygun, K. *IEEE Trans Compon Packaging Manuf Technol* **12**, 1423–1431 (2022).
5. Dai, D., Bauters, J. & Bowers, J. Light *Sci. Appl.* **1**, e1 (2012).
6. Zhang, Y., Samanta, A., Shang, K. & Ben Yoo, S. J. *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1–10 (2020).
7. Xiang, C. et al. *Nature* **620**, 78–85 (2023).
8. Poulton, C. V et al. *Opt. Lett.* **42**, 4091–4094 (2017).
9. Huang, C. et al. *Nat. Electron.* **4**, 837–844 (2021).
10. Wang, J., Sciarrino, F., Laing, A. & Thompson, M. G. *Nat. Photon.* **14**, 273–284 (2020).
11. Tu, K.-N. *Microelectronics Reliability* **51**, 517–523 (2011).
12. Gaeta, A. L., Lipson, M. & Kippenberg, T. J. *Nat. Photon.* **13**, 158–169 (2019).
13. Zhou, Z. et al. *Elight* **3**, 1 (2023).
14. Fathololoumi, S. et al. *2022 IEEE Symposium on High-Performance Interconnects (HOTI)* https://doi.org/10.1109/HOTI55740.2022.00016 (2022).

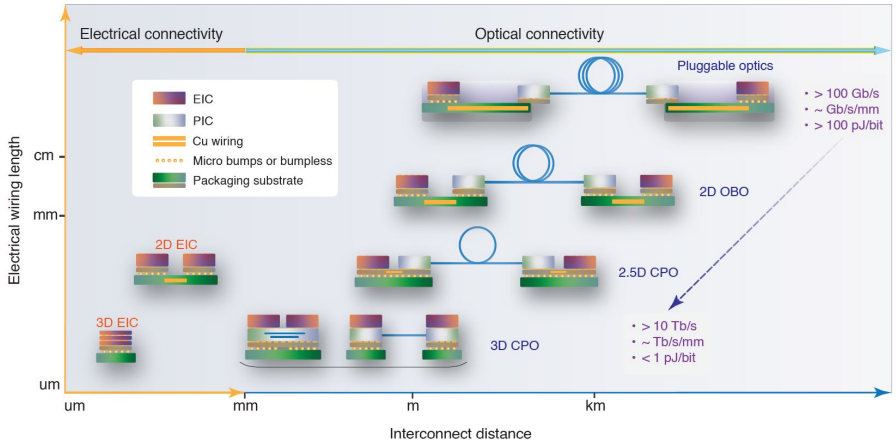

Figure 1. Interconnect distances and electrical wiring length dependence for connectivity solutions. Shorter-reach interconnects move photonic integrated circuits (PICs) from the ensemble perimeter to be close to electronic integrated circuits (EIC) through different packaging techniques. These packaging forms generally include on-board optics (OBO) that place separate EIC and PIC on a common PCB, and 2.5D/3D co-packaged optics that packages EIC and PIC together. The difference in 2.5D and 3D lies in that whether a 2D interposer substrate is required in such packages. It needs to be noted that packaging methods EIC and PIC can vary with substrate, and microbumps/hybrid copper bonding are used here as general indications. For 3D CPOs, both optical waveguide and optical fibre based connectivity are included with a difference in the PIC and interconnect distance. The figure highlights the technological trend seen in connectivity solutions in terms of interconnect bandwidth, bandwidth density and energy efficiency.