

Three-dimensional photonic integration for ultra-low-energy, high-bandwidth interchip data links

Received: 7 June 2024

Accepted: 27 January 2025

Published online: 21 March 2025

 Check for updates

Stuart Daudlin¹, Anthony Rizzo^{1,2,3}, Sunwoo Lee⁴, Devesh Khilwani⁴, Christine Ou⁴, Songli Wang¹, Asher Novick¹, Vignesh Gopal¹, Michael Cullen¹, Robert Parsons¹, Kaylx Jang¹, Alyosha Molnar⁴ & Keren Bergman¹✉

Artificial intelligence (AI) hardware is positioned to unlock revolutionary computational abilities by leveraging vast distributed networks of advanced semiconductor chips. However, a barrier for AI scaling is the disproportionately high energy and chip area required to transmit data between the chips. Here we present a solution to this long-standing overhead through dense three-dimensional (3D) integration of photonics and electronics. With 80 photonic transmitters and receivers occupying a combined chip footprint of only 0.3 mm², our platform achieves an order-of-magnitude-greater number of 3D-integrated channels than prior demonstrations. This enables both high-bandwidth (800 Gb s⁻¹) and highly efficient, dense (5.3 Tb s⁻¹ mm⁻²) 3D channels. The transceiver energy efficiency is showcased by a state-of-the-art 50 fJ and 70 fJ per communicated bit from the transmitter and receiver front ends, respectively, operating at 10 Gb s⁻¹ per channel. Furthermore, the design is compatible with commercial complementary metal–oxide–semiconductor foundries fabrication on 300-mm-sized wafers, providing a route to mass production. Such ultra-energy-efficient, high-bandwidth data communication links promise to eliminate the bandwidth bottleneck between spatially distinct compute nodes and support the scaling of future AI computing hardware.

Light, as a medium for communication, has the unique ability to transmit volumes of data with minimal energy loss. This capability not only sparked the revolution of internet-based communication over fibre-optic networks but also holds the potential to greatly expand computing power beyond current capabilities. Specifically, artificial intelligence (AI) is poised to dramatically transform the computational landscape if provided with more efficient data communication between the nodes in computer networks^{1,2}. A critical bottleneck to the

full implementation of light-based communication is the conversion of electrical data from inside a computer chip to optical data. At present, the data are stored densely in these semiconductor chips in compute nodes but are sent out of the chip through centimetre-long electrical wires before finally interfacing with optical transmitters in the form of pluggable optical transceivers. The limited number of these electrical channels and their wire lengths result in data rates that are not scalable without accounting for a substantial amount of energy consumption^{3,4}.

¹Department of Electrical Engineering, Columbia University, New York, NY, USA. ²Air Force Research Laboratory Information Directorate, Rome, NY, USA.

³Thayer School of Engineering, Dartmouth College, Hanover, NH, USA. ⁴Department of Electrical and Computer Engineering, Cornell University, Ithaca, NY, USA. ✉e-mail: bergman@ee.columbia.edu

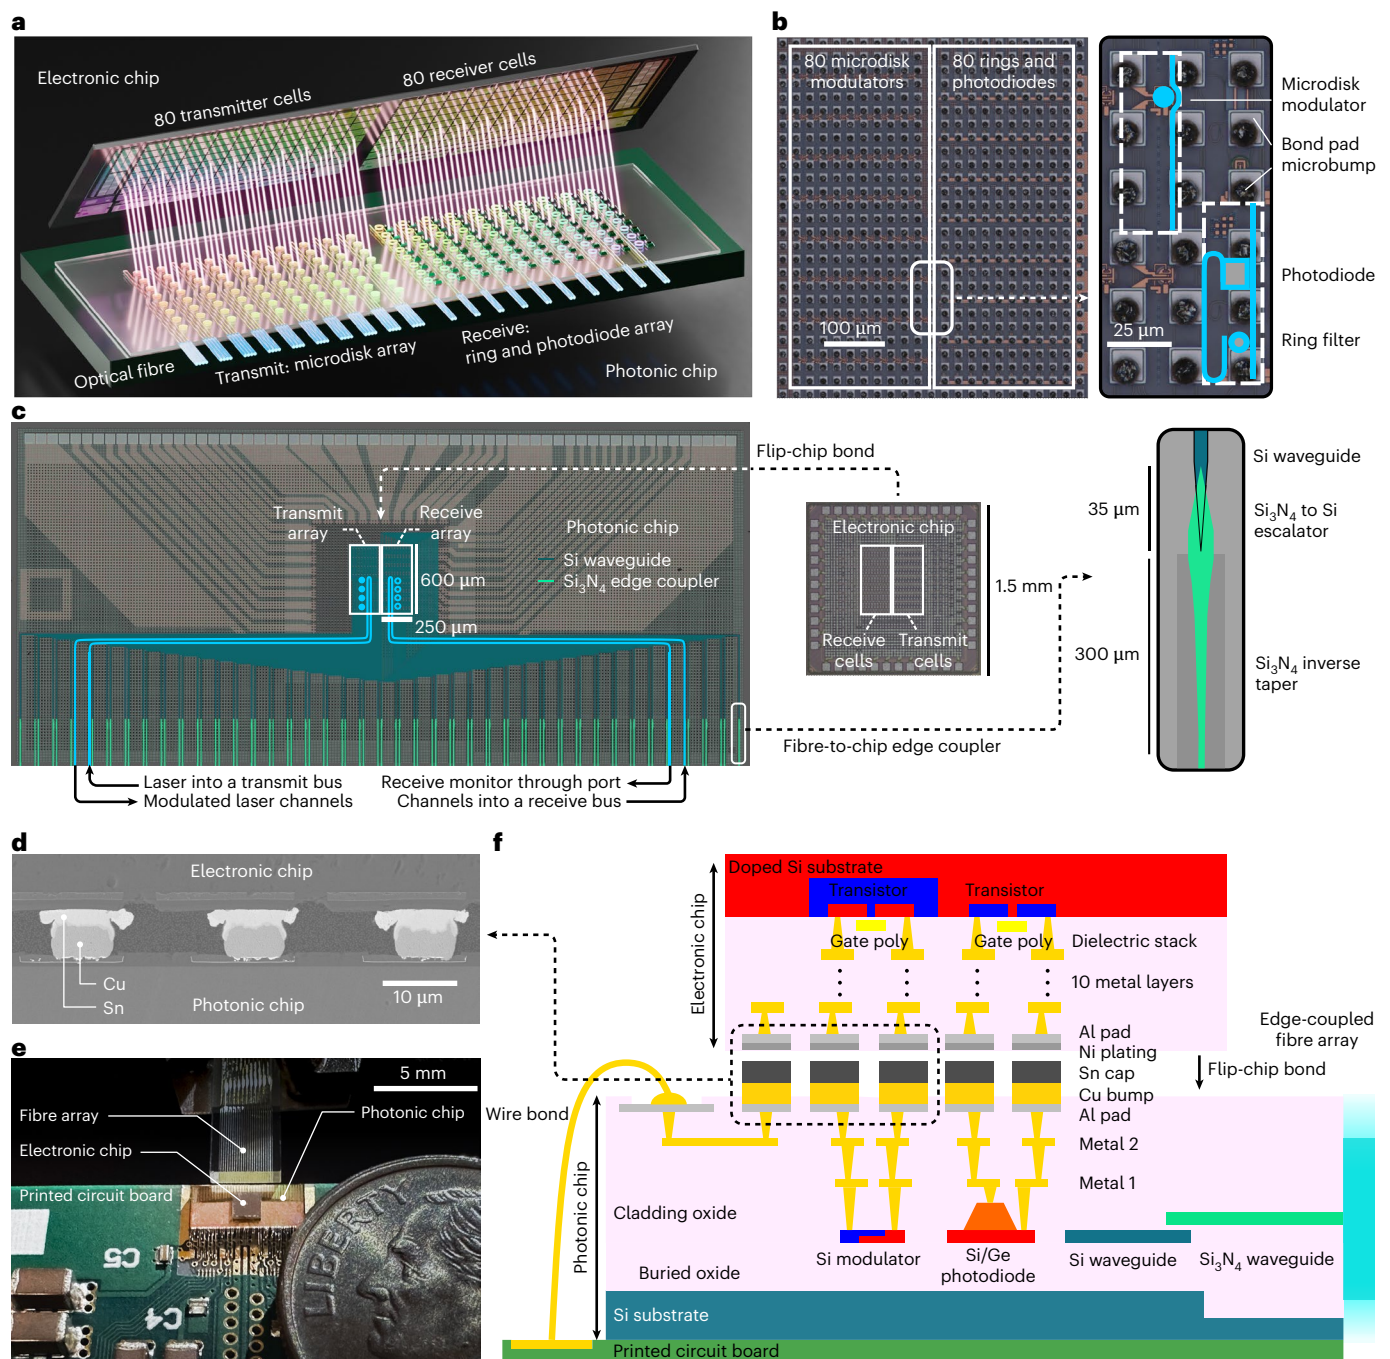


Fig. 1 | 3D-integrated photonic-electronic transceiver. **a**, An illustration of the 3D-integrated photonic-electronic system combining arrays of electronic cells with arrays of photonic devices. **b**, A microscope image of the 80-channel photonic device arrays with an inset of two transmitter and two receiver cells. **c**, Microscope images of the photonic and electronic chips. The active photonic circuits occupy an area outlined in white, while the outer photonic chip area is used to fan out the optical/electrical lanes for fibre coupling and wire bonding. The blue overlay shows a four-channel transmitter and receiver waveguide path; the disk and ring overlays are not to scale. An inset shows a diagram of

the fibre-to-chip edge coupler, consisting of a silicon nitride (Si_3N_4) inverse taper and escalator to silicon. **d**, A scanning electron microscope image of the bonded electronic and photonic chip cross-section. **e**, An image of the wire-bonded transceiver die bonded to a printed circuit board and optically coupled to a fibre array with a US dime for scale. **f**, A cross-sectional diagram of the electronic and photonic chips and their associated material stacks. Both chips consist of a crystalline silicon substrate, doped-silicon devices and metal interconnection layers.

To overcome this bottleneck, electrical channels must be condensed and converted into optical signals within a compact area.

Previously, intensive efforts have produced chip-scale transmitters and receivers (transceivers) towards this goal but have been marked by a lack of efficiency or scale. These works build on the field of integrated photonics, a technology that aggregates a multitude of

optical components onto a single integrated chip. In particular, silicon is highly appealing as a material platform for integrated photonics since it can leverage the tremendous investment in the complementary metal–oxide–semiconductor (CMOS) infrastructure used to fabricate microelectronics chips⁵. The silicon photonics technology platform includes devices such as microresonator-based modulators^{6,7}, filters

and germanium photodiodes⁸ that are compact, efficient in their electrical-to-optical and optical-to-electrical conversions and scalable to many wavelength channels^{9,10}. So far, the largest of these systems is composed of 64 channels of photonics and electronics on a single chip and achieves 240 fJ per communicated bit by the transmitter^{11–13}. However, this system has receiver energy consumption above 1,000 fJ per bit and has a limited density from the lateral arrangement of photonics and electronics on the same two-dimensional chip. While this monolithic integration of CMOS transistors alongside photonic devices on the same chip may appear highly appealing^{14,15}, this configuration ‘freezes’ transistors at a given node size and, thus, cannot benefit from the further energy efficiency, size and speed gains of moving to more advanced CMOS nodes. Alternatively, three-dimensional (3D) integration combines a more efficient, leading edge CMOS node electronic chip and a separate photonic chip to improve on these limitations. Ongoing 3D efforts have demonstrated sub-200 fJ per bit powers from transmitter^{16–19} and receiver front ends^{17,18,20}, but the chip-to-chip bond spacings are either substantially larger than the devices themselves^{16,17,19,20} or rely on emerging hybrid bonding technology¹⁸. Furthermore, 3D-integrated transceiver demonstrations have yet to show more than eight channels¹⁶.

Here, we present a transceiver that is a large-scale 3D system, featuring 80 channels on a single chip, and is at the forefront of energy and area efficiency, consuming only 120 fJ per bit from its electro-optic front ends in a high-density platform. The transceiver achieves low energy consumption through low-capacitance 3D connections between photonics and codesigned CMOS electronic circuits. Although each channel has a relatively low data signalling rate of 10 Gb s⁻¹, this allows the receiver electronics to operate with high sensitivity and minimal energy consumption. The large array of channels compensates for the low per-channel data rates, delivering a high aggregate data rate of 800 Gb s⁻¹ in a compact transmitter or receiver area of only 0.15 mm² (5.3 Tb s⁻¹ mm²). In addition, having many low-data-rate channels relaxes signal processing and time multiplexing of the low-data-rate streams native to the processor^{3,21}. Furthermore, wavelength-division multiplexing sources for these numerous data streams are becoming available with the advent of chip-scale microcombs^{9,22}. Our demonstration unlocks the tremendous potential of light as a high-bandwidth and energy-efficient medium for interchip communication, offering an immediate solution to the pressing challenge of AI scaling.

Results

We implement this high-density transceiver using compact photonic devices and dense, 28-nm-node codesigned electronic circuits; however, the total density ultimately depends on the 3D bond spacings. To address this, we use a high-density bonding process using copper pillar bumps. An electroplating process is used to form bumps on the photonic chip with copper pedestals capped with a layer of tin. The copper–tin bumps are then bonded to a nickel-plated electronic chip under a thermal and compressive force bringing the chips together. Figure 1a illustrates the bonding of the electronic and photonic chips. Figure 1b shows a microscope image of the photonic device array after microbumping. This dense photonic device array resides within the photonic test chip, which Fig. 1c shows with the electronic test chip and

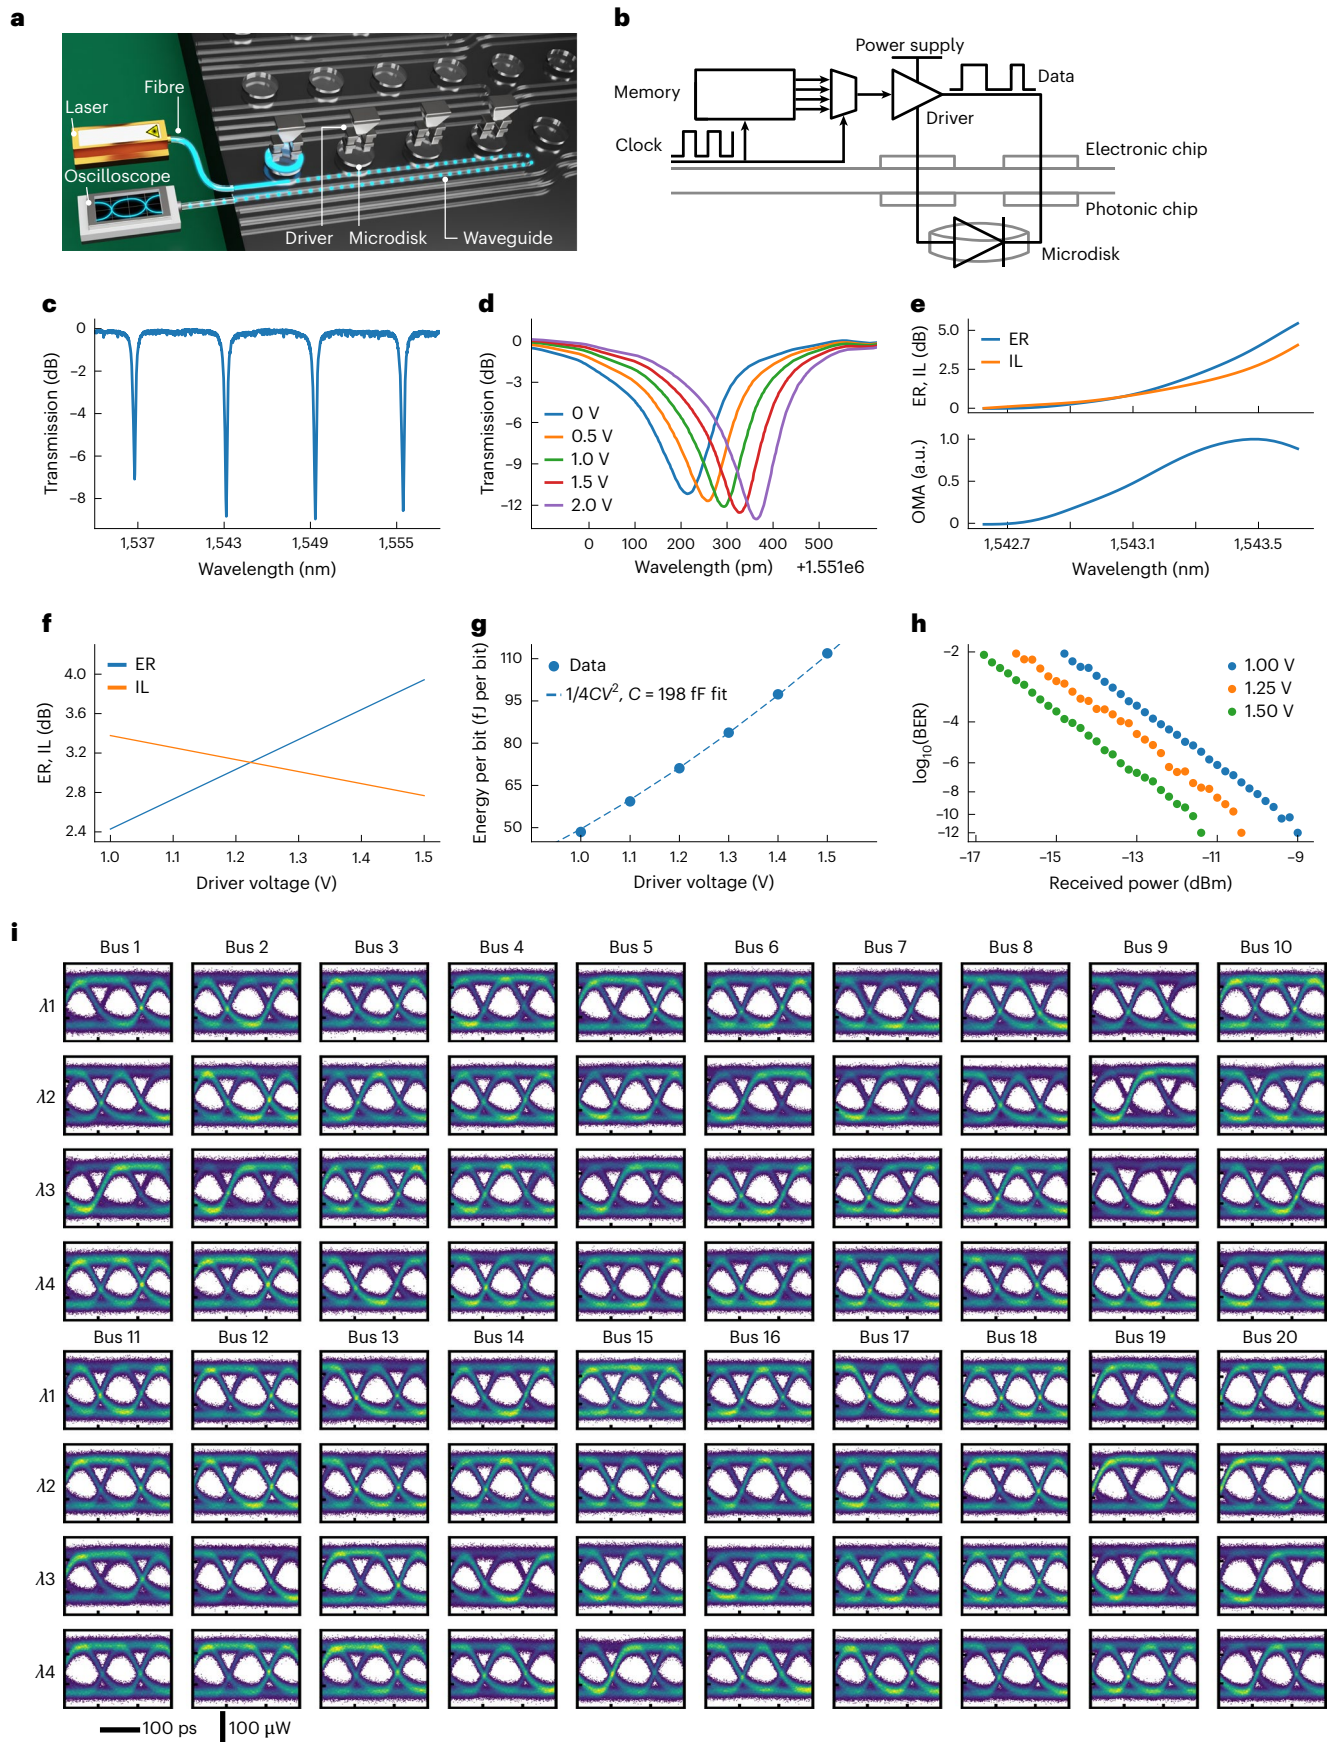
a diagram of the fibre-to-chip coupler (see Supplementary Note 1 for coupler details). We push the limits of this bonding technology by using a 15 µm spacing and 10 µm bump diameters (25 µm pitch) in an array of 2,304 bonds. This process balances two potential failure modes for such close spacing: excessive tin causing flow and electrical shorting to adjacent bonds during bonding and insufficient tin leading to brittle bonds²³. We test our bonding process using cross-sectional scanning electron micrographs of the bonds, displayed in Fig. 1d, and by measuring the force needed to separate the bonded chips. The cross-section analysis reveals that tin does not flow to adjacent bonds, while the shear test demonstrates a robust 2.1 kg (114.9 MPa) force required to separate the bonded chips. The modelling and measurements show each bond has a 10 fF capacitance (Methods). Figure 1e shows the assembled transceiver, which is wire bonded to a printed circuit board and optically coupled to a fibre array (Methods). Figure 1f illustrates the material layers in the electronic chip, photonic chip and bonds and shows how fibre and wire bonds interface with the transceiver. This bonding technique provides an ideal platform to achieve the required density for chip-to-chip data communication links.

The 3D-integrated chip contains an array of 80 transmitter cells and 80 receiver cells; these cells are organized into 20 waveguide buses with four wavelength channels per bus. A blue overlay in Fig. 1c outlines a single transmitter and receiver bus. Each transmitter cell has a local memory in the electronic chip that stores a pseudo-random bit sequence. A periodic clock signal triggers the electronics, and the transmitter cell electronics send out the programmed bit sequences as voltage pulses incident on the photonic modulator electrodes. These voltage pulses blue shift the microdisk resonance from a blocking to a non-blocking state and, thus, modulate an on-resonance laser line. Figure 2a illustrates the transmitter experiment, and Fig. 2b shows a schematic of the transmitter cell, while Fig. 2c shows the spectrum of the modulator bus with four microdisk resonances. After the transmitter characterization, we test the receivers, which function similarly: in each receiver cell, the wavelength channels carry signals on the photonic chip, and microrings selectively drop wavelengths onto each respective photodiode. The electronic chip then amplifies the photocurrent generated by the photodiode and writes the data into the local memory of each receiver cell, as illustrated in Fig. 3a. Figure 3b shows a schematic of the receiver cell. For performance characterization, an on-chip circuit compares the receiver memory to expected data and keeps an error count that is periodically read out of the chip. This architecture of transmitters and receivers fills the array of channels in the dense area provided by the bonding process.

The transmitter cell within the 80-cell array consumes 50 fJ per bit when driving the microdisks with a 1 V swing. This power is dynamic, equal to $\frac{1}{4} CV^2$, where C is the capacitance being charged or discharged during a bit transition, and V is the charged-to-discharged voltage²⁴. The vertical p–n junction microdisk enables a low voltage drive by featuring a higher overlap of the p–n depletion region and the optical whispering gallery mode of the disk compared with lateral junctions⁷ and results in an electrical-to-optical response of 75 pm resonance shift per applied volt (Fig. 2d). We further characterize this response with the dynamic insertion loss (IL, the power of a ‘1’ bit divided by the power before the modulator) and extinction ratio (ER, the power of a ‘1’ bit divided by the power of a ‘0’ bit).

Fig. 2 | Transmitter characterization and performance. **a**, An illustration of the transmitter experimental test setup showing a single wavelength laser channel modulated by the transmitter and measured on an oscilloscope. **b**, The transmitter cell circuit schematic. **c**, An optical spectrum of the four-channel transmitter bus. **d**, A resonance shift of a representative microdisk as a function of reverse-bias voltage applied to the vertical p–n junction. **e**, The transmitted signal ER (power of ‘1’ bit divided by power of ‘0’ bit), IL (power of ‘1’ bit divided by power before the modulator) and normalized OMA (normalized power of ‘1’ bit

minus power of ‘0’ bit) with a 1.5 V driver voltage. **f**, The ER and IL at maximum OMA for a range of driver voltages. **g**, Measured energy consumption of the transmitter array for a range of driver voltages and a $\frac{1}{4} CV^2$ fit, where C is the capacitance charged by the driver voltage V . **h**, A BER measurement of the modulated signal input to a commercial receiver at 1, 1.25 and 1.5 V driver voltages; the received power is the average signal power at the commercial receiver. **i**, Eye diagrams for all 80 modulators on the photonic chip at 10 Gb s⁻¹ per modulator and 1 dBm input laser power.



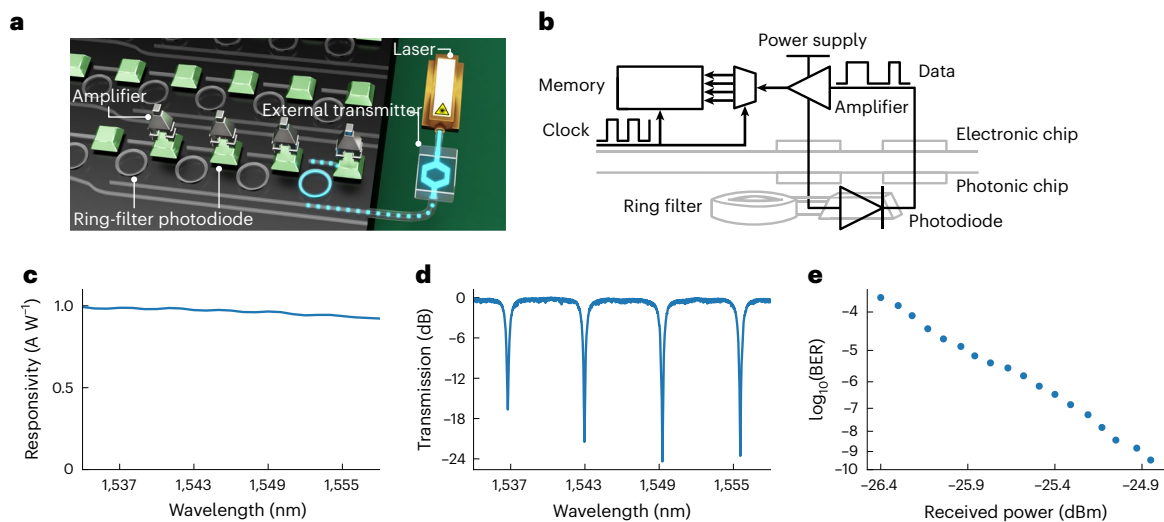


Fig. 3 | Receiver characterization and performance. **a**, An illustration of the receiver test setup showing a laser line modulated by a commercial transmitter and received by the 3D-integrated photonic-electronic receiver. **b**, A receiver cell circuit schematic. **c**, A responsivity (light to electrical current conversion

efficiency) measurement of the photodiode. **d**, An optical spectrum of a receiver bus. **e**, A BER test of a receiver cell using a commercial transmitter signal; the received power is the average signal power at the photodiode.

Figure 2e shows these metrics, captured from the modulated signal output of a transmitter cell driven at 1.5 V. In this measurement, the laser wavelength moves into the shifting resonance and the optical modulation amplitude (OMA, the power of '1' bit minus the power of '0' bit) of the output signal increases, reaching a maximum at 2.5 dB IL and 4 dB ER. Figure 2f shows the ER and IL at maximum OMA for driver voltages between 1 and 1.5 V. These high ERs and low ILs per volt enable a reduced V in $\frac{1}{4} CV^2$. The capacitance sources include the microdisk p–n junction (128 fF), bond pads (10 fF) and capacitances within the driver circuit (61 fF), combining for a total expected capacitance of 199 fF (Methods). These capacitances exhibit low values through microdisk compactness, miniaturized bonds and careful design in the 28 nm electronic chip technology. Figure 2g shows the transmitter energy consumption as all 80 modulators are transmitting data with drive voltages ranging between 1 and 1.5 V. The $\frac{1}{4} CV^2$ model fits 198 fF total capacitance per cell using this data, aligning closely with the expected 199 fF from the independently measured and modelled devices. Next, we record eye diagrams for each of the 80 transmitters on the chip with a drive at 10 Gb s^{−1} per transmitter and 1 dBm laser power before the modulator (Fig. 2i). As every transmitter modulates, we measure current and voltage on the transmitter power supply for the previous energy consumption (Fig. 2g). With no optical amplification, the oscilloscope receiver is the limiting factor of the eye qualities with an 8 μW input-referred root mean square noise (denoted as input-referred noise here on). All 80 eye diagrams in the array are open and uniform, which confirms the high yield of the bonding process and validates our many-channel approach. As a further confirmation of transmitter signal quality, a bit-error ratio (BER) test with a reference receiver demonstrates error-free performance (BER < 10^{−12}) down to a receiver noise-determined power for 1-, 1.25- and 1.5-V modulator drives (Fig. 2h). In the following section, our on-chip receiver shows a dramatically reduced input-referred noise of 480 nW. The 80-channel photonic transmitter array outputs an aggregate data rate of 800 Gb s^{−1} and occupies an area of 0.15 mm², demonstrating an unprecedented bandwidth density of 5.3 Tb s^{−1} mm².

The receiver cell consumes 70 fJ per bit when receiving a 10 Gb s^{−1} signal at −24.85 dBm average power with a 4×10^{-10} BER. The receiver spends energy as a static biasing of the electronic amplifier. The photodiode is a vertical p–silicon, i–germanium and n–germanium diode that efficiently converts optical signals to electrical current with

an efficiency of 1 A W^{−1} (Fig. 3c). The capacitance of this photodiode is crucial since receiver noise is proportional to the amplifier input capacitance and is compensated by static biasing power (Methods and Supplementary Note 3). Minimizing this noise is critical for reducing laser power sourced into the link and improving energy efficiency. With a measured photodiode capacitance of 17 fF and pad capacitances of 10 fF, the simulated input-referred noise is 300 nW. A BER test is used to evaluate this performance; a signal from an ideal modulator is used to send a 10 Gb s^{−1} data stream into the chip from which on-chip circuits measure errors in the received bits. On the photonic chip, a ring resonator on a four-channel bus filters the modulated signal to a photodiode. Figure 3d shows the four-channel bus spectrum. We then gradually reduce the signal power while counting errors on the electronic chip, obtaining the BER curve in Fig. 3e. The receiver exhibits a 4×10^{-10} BER at a −24.85 dBm average input optical power. This reveals a sensitivity of −22.3 dBm OMA using the 13 dB ER signal and an input-referred noise of 480 nW. We record the energy consumption of the receiver from its power supply at −24.85 dBm input optical power to be 70 fJ per bit. This result, along with the transmitter cell performance, demonstrates that both the receiver and transmitter consume less than 100 fJ per bit while supporting a massive bandwidth transfer of 800 Gb s^{−1} from the dense array.

The transmitter and receiver cells independently demonstrate a combined 120 fJ per bit; we next connect them and validate their combined performance. Optical fibre (4 m of SMF28) connects two separate transceivers as a complete data communication link, with one transceiver functioning as a transmitter and the other as a receiver (Fig. 4a). A polarization controller is the only component between the transmitter and receiver. A shared clock synchronizes the two electronic chips, and programmable clock delays in each receiver cell align the transmitted data with the receiver sampling point. A laser diode array provides four wavelength channels at −5 dBm power per channel, which we feed into a bus on the transmitter chip. Figure 4b shows the spectrum of these laser channels. Individual transmitter cells simultaneously modulate each wavelength at 8 Gb s^{−1} at a 1.5 V drive, resulting in open eye diagrams for each channel (Fig. 4c). The signal powers are too low for detection by the diagnostic oscilloscope receiver, so we amplify the signal before the oscilloscope and normalize the eye diagrams. However, the signals do not require amplification for the electronic-photonic receiver due to its high sensitivity.

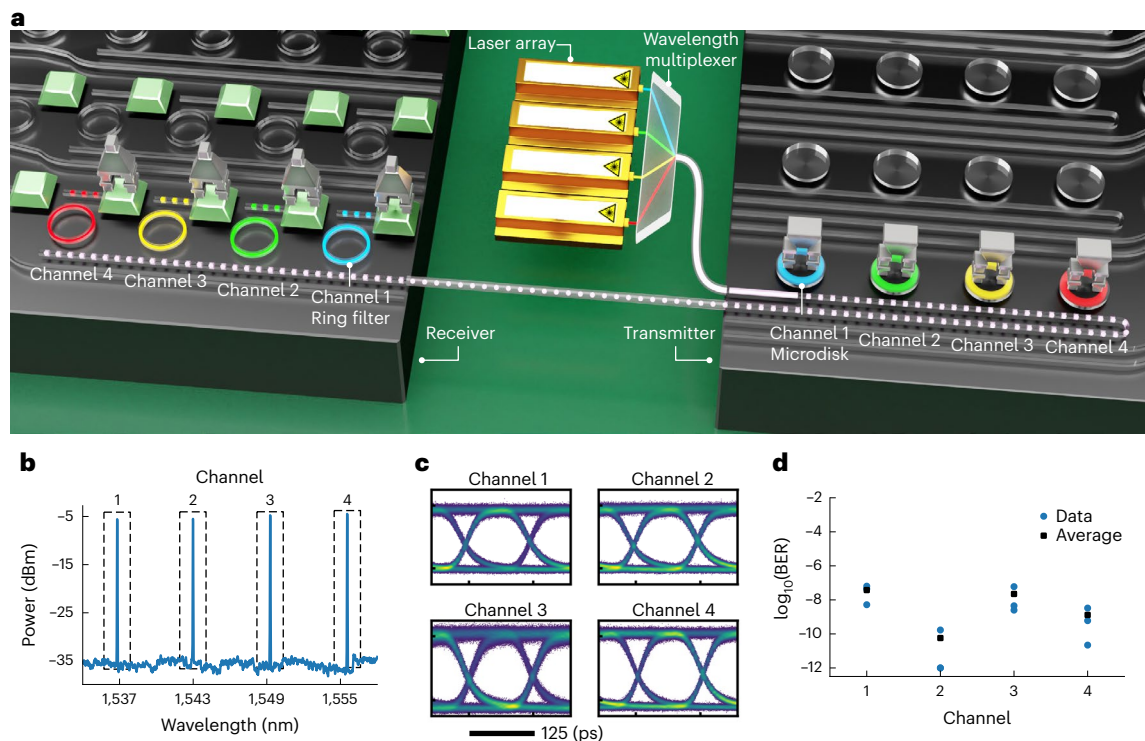


Fig. 4 | Transmitter to receiver data communication link demonstration.

a, An illustration of the transmitter to receiver link showing a 3D-integrated photonic-electronic transmitter modulating four laser lines and a separate photonic-electronic receiver converting the four data channels back to the electrical domain. The illustration excludes a polarization controller between the

transmitter and receiver. **b**, A spectrum of the link laser source. **c**, Eye diagrams of the four channels after the transmitter. The diagrams do not include a power scale bar, as an amplifier is used before the oscilloscope, but it is not used in the optical link. **d**, BERs of the data channels after the receiver.

The average power per channel at each receiver photodiode is -19.5 dBm. On-chip error counters record errors in each channel in 1-min intervals, revealing a maximum recorded BER of 6×10^{-8} and a minimum count of no errors in the interval, denoted as 10^{-12} , in Fig. 4d. This result shows that the transmitters and receivers within the 3D integration can form a complete low-power, high-bandwidth link needed for next-generation computing systems.

Discussion

Integrated photonic chips present a promising low-power platform to address the data transfer demands of AI computing. Here, we realized this promise by demonstrating a scaled-up array of 80 channels on a single electronic-photonic, densely 3D-integrated transceiver. This multichip module consumes minimal energy by virtue of the large number of channels, cutting-edge low-capacitance bonding technology, codesigned electronic-photonic circuits and advanced devices used. While our demonstrated system achieves record performance in terms of energy efficiency and bandwidth density, the performance can be further improved in future implementations. Although the microdisks used in this demonstration exhibit high performance, resonant modulators can be developed with lower capacitance^{7,25} and a higher electro-optical response^{7,26}, both of which would decrease the dynamic power of the transmitter. Similarly, on the receiver side, lower capacitance photodiodes^{27,28} could reduce the power and noise of the receiver architecture. However, miniaturizing photodiodes requires considering a loss of responsivity²⁹, presenting complex link-level trade-offs. In addition, the energy consumption of the electronic circuits can be further reduced by moving to a more advanced CMOS node. While our demonstrated bonding technology is approaching the limit of how closely spaced tin bonds can be made, further density scaling could be realized through the development of hybrid bonding¹⁸. However, after achieving the low-capacitance value

of the bonds demonstrated in this work, pursuing a further reduction in bond capacitance would yield diminishing returns in terms of energy efficiency.

Reduced chip-to-fibre optical losses can improve the loss budget of our demonstrated link, and laser powers may be as low as 47 fJ per bit with a distributed feedback laser and 30 fJ per bit for a scalable, high channel count comb laser (Supplementary Note 1). Furthermore, silicon resonators are sensitive to temperature and fabrication variation, requiring thermal control circuits^{30,31} and, for minimized power, reduced heat leakage into the environment using methods such as a silicon substrate removal around resonators^{19,32,33}. Detailed wafer-scale resonance variation data and approximate thermal energy contributions across a range of scenarios provide a best case thermal energy consumption of 176 fJ per bit (Supplementary Note 2). A dense wavelength-division multiplexing architecture may be used to reduce this power penalty²¹. Moreover, the photonic circuits are highly polarization-sensitive and require polarization-maintaining fibre or the addition of polarization management circuits^{29,34,35}. Finally, while we demonstrate high-bandwidth density, a higher per-fibre bandwidth and photonic chip edge bandwidth density can be achieved with wavelength scaling by cascading more arrayed channels onto fewer waveguide buses^{9,10}. This architecture can combine with chip-scale frequency combs to generate hundreds of wavelength channels⁹.

While the potential impact of this technology is evident for the advancement of energy-efficient AI computing, its use may extend to far reaching applications. These low-power, massively parallel optical links could enable pervasive device connectivity and transform computing by streamlining resource allocation through optically linked, disaggregated and reconfigurable computing and memory resources^{36–39}, revolutionizing the computing landscape over the next decade and beyond.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41566-025-01633-0>.

References

- Wang, G. et al. ZeRO++: extremely efficient collective communication for giant model training. Preprint at <https://arxiv.org/abs/2306.10209> (2023).
- Wu, Z. et al. Peta-scale embedded photonics architecture for distributed deep learning applications. *J. Lightwave Technol.* **41**, 3737–3749 (2023).
- Miller, D. A. Attojoule optoelectronics for low-energy information processing and communications. *J. Lightwave Technol.* **35**, 346–396 (2017).
- Lee, B. G., Nedovic, N., Greer, T. H. & Gray, C. T. Beyond CPO: a motivation and approach for bringing optics onto the silicon interposer. *J. Lightwave Technol.* **41**, 1152–1162 (2022).
- Hochberg, M. & Baehr-Jones, T. Towards fabless silicon photonics. *Nat. Photonics* **4**, 492–494 (2010).
- Xu, Q., Schmidt, B., Pradhan, S. & Lipson, M. Micrometre-scale silicon electro-optic modulator. *Nature* **435**, 325–327 (2005).
- Timurdogan, E. et al. An ultralow power athermal silicon modulator. *Nat. Commun.* **5**, 4008 (2014).
- Michel, J., Liu, J. & Kimerling, L. C. High-performance Ge-on-Si photodetectors. *Nat. Photonics* **4**, 527–534 (2010).
- Rizzo, A. et al. Massively scalable Kerr comb-driven silicon photonic link. *Nat. Photonics* **17**, 781–790 (2023).
- Rizzo, A. et al. Petabit-scale silicon photonic interconnects with integrated kerr frequency combs. *IEEE J. Quantum Electron.* **29**, 1–20 (2022).
- Wade, M. et al. A bandwidth-dense, low power electronic–photonic platform and architecture for multi-Tbps optical I/O. In *2018 European Conference on Optical Communication* 1–3 (IEEE, 2018).
- Sun, C. et al. Teraphy: an O-band WDM electro-optic platform for low power, terabit/s optical i/o. In *2020 IEEE Symposium on VLSI Technology* 1–2 (IEEE, 2020).
- Wade, M. et al. An error-free 1 Tbps WDM optical i/o chiplet and multi-wavelength multi-port laser. In *Optical Fiber Communication Conference F3C–6* (Optica Publishing Group, 2021).
- Sun, C. et al. Single-chip microprocessor that communicates directly using light. *Nature* **528**, 534–538 (2015).
- Atabaki, A. H. et al. Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip. *Nature* **556**, 349–354 (2018).
- Zheng, X. et al. Ultralow power 80 Gb/s arrayed cmos silicon photonic transceivers for WDM optical links. *J. Lightwave Technol.* **30**, 641–650 (2011).
- Rakowski, M. et al. Hybrid 14nm finfet-silicon photonics technology for low-power Tb/s/mm² optical I/O. In *2018 IEEE Symposium on VLSI Technology* 221–222 (IEEE, 2018).
- Chang, P.-H. et al. A 3D integrated energy-efficient transceiver realized by direct bond interconnect of co-designed 12 nm finfet and silicon photonic integrated circuits. *J. Lightwave Technol.* **41**, 6741–6755 (2023).
- Ban, Y. et al. Highly optimized O-band Si ring modulators for low-power hybrid CMOS–SiPho transceivers. In *2023 Optical Fiber Communications Conference and Exhibition* 1–3 (Optica Publishing Group, 2023).
- Saeedi, S., Menezes, S., Pares, G. & Emami, A. A 25 Gb/s 3D-integrated cmos/silicon-photonics receiver for low-power high-sensitivity optical communication. *J. Lightwave Technol.* **34**, 2924–2933 (2016).
- Georgas, M., Leu, J., Moss, B., Sun, C. & Stojanović, V. Addressing link-level design tradeoffs for integrated photonic interconnects. In *2011 IEEE Custom Integrated Circuits Conference* 1–8 (IEEE, 2011).
- Gaeta, A. L., Lipson, M. & Kippenberg, T. J. Photonic-chip-based frequency combs. *Nat. Photonics* **13**, 158–169 (2019).
- Li, Z. et al. Scaling solder micro-bump interconnect down to 10 um pitch for advanced 3D IC packages. In *2021 IEEE 71st Electronic Components and Technology Conference* 451–456 (IEEE, 2021).
- Miller, D. A. Energy consumption in optical modulators for interconnects. *Opt. Express* **20**, A293–A308 (2012).
- Sun, J. et al. A 128 Gb/s PAM4 silicon microring modulator with integrated thermo-optic resonance tuning. *J. Lightwave Technol.* **37**, 110–115 (2018).
- Gevorgyan, H., Khilo, A., Wade, M. T., Stojanović, V. M. & Popović, M. A. Miniature, highly sensitive MOSCAP ring modulators in co-optimized electronic-photonic CMOS. *Photonics Res.* **10**, A1–A7 (2022).
- Lischke, S. et al. High bandwidth, high responsivity waveguide-coupled germanium pin photodiode. *Opt. Express* **23**, 27213–27220 (2015).
- Chen, H. et al. –1 V bias 67 GHz bandwidth Si-contacted germanium waveguide pin photodetector for optical links at 56 Gbps and beyond. *Opt. Express* **24**, 4622–4631 (2016).
- Chrostowski, L. & Hochberg, M. *Silicon Photonics Design: From Devices to Systems* (Cambridge Univ. Press, 2015).
- Sun, C. et al. A 45 nm CMOS-SOI monolithic photonics platform with bit-statistics-based resonant microring thermal tuning. *IEEE J. Solid-State Circuits* **51**, 893–907 (2016).
- Hattink, M., Dai, L. Y., Zhu, Z. & Bergman, K. Streamlined architecture for thermal control and stabilization of cascaded DWDM micro-ring filters bus. In *Optical Fiber Communication Conference W2A–2* (Optica Publishing Group, 2022).
- Masood, A. et al. Comparison of heater architectures for thermal control of silicon photonic circuits. In *10th International Conference on Group IV Photonics* 83–84 (IEEE, 2013).
- Rizzo, A. et al. Ultra-efficient foundry-fabricated resonant modulators with thermal undercut. In *CLEO: Science and Innovations SF2K–6* (Optica Publishing Group, 2023).
- Ma, Y. et al. Symmetrical polarization splitter/rotator design and application in a polarization insensitive WDM receiver. *Opt. Express* **23**, 16052–16062 (2015).
- Park, A. H., Shoman, H., Ma, M., Shekhar, S. & Chrostowski, L. Ring resonator based polarization diversity WDM receiver. *Opt. Express* **27**, 6147–6157 (2019).
- Zhu, Z. et al. Photonic switched optically connected memory: an approach to address memory challenges in deep learning. *J. Lightwave Technol.* **38**, 2815–2825 (2020).
- Jouppi, N. et al. TPU v4: an optically reconfigurable supercomputer for machine learning with hardware support forembellings. In *Proc. 50th Annual International Symposium on Computer Architecture* 1–14 (2023).
- Gonzalez, J. et al. Optically connected memory for disaggregated data centers. *J. Parallel Distrib. Comput.* **163**, 300–312 (2022).
- Michelogiannakis, G. et al. A case for intra-rack resource disaggregation in HPC. *ACM Trans. Archit. Code Optim.* **19**, 1–26 (2022).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share

adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025

Methods

Transceiver assembly

Separate CMOS foundries are used to fabricate the electronic and photonic chips. The photonic chips were fabricated through the American Institute for Manufacturing Integrated Photonics on a custom 300-mm silicon-on-insulator wafer. The American Institute for Manufacturing Integrated Photonics process design kit includes the microdisks, ring filters and photodiodes⁴⁰. The microdisks and ring filters have quality factors of 5,800 and 2,800, respectively, and both have a free spectral range of 25.6 nm for a channel spacing of 6.4 nm. An example of the photodiode design can be found in ref. 41. The electronic chips were fabricated through Taiwan Semiconductor Manufacturing Company Limited on a shared multiproject wafer in a 28-nm CMOS process node. Both chips feature a square array of 2,304 aluminium pads at a 25- μm pitch on their top surfaces. These pads connect to lower metal layers within each chip and the devices on the silicon layer.

The photonic and electronic chips are then processed postfabrication before bonding their pad arrays together. In this step, we core the 300-mm photonic wafer to a 200-mm wafer, and a wafer-level process is used to bump its pads with electroplated layers of copper and tin. The electronic chips, received as individual 1.6 mm² units from a shared wafer, are unsuitable for wafer-level photolithography-based processes. Alternatively, we adopt a chip-level process of electroless nickel plating, followed by an additional layer of immersion gold plating to prevent nickel oxidation. After dicing the bumped photonic wafer into 6.5 mm \times 3 mm chips, a thermo-compression bond is used to connect the bumped photonic chips to the plated electronic chips.

To power and operate the transceiver, we create electrical connections to the electronic chip through the bonds to the photonic chip. The metal layers on the photonic chip wire these connections to large electrical pads on an exposed edge of the photonic chip. Wire bonds connect these pads to a printed circuit board, which connects to: (1) a microcontroller that programs the electronic chip, (2) power sources that supply the electronic chip voltage rails and (3) a radio-frequency (RF) clock generator for the 5 GHz clock of the electronic chip. This clock line is a coplanar RF waveguide on the printed circuit board with a matched impedance to 50 Ω RF cables. The optical fibres couple light to waveguide buses through silicon nitride edge couplers; these couplers are on the side of the photonic chip that is opposite the wire-bond pads. A micropositioner was used to align a standard single mode fibre v-groove array with the edge couplers, which are spaced at a 127 μm pitch. The assembly procedures of photonic wafer bumping and bonding, electronic chip plating and wire bonding are conducted at Micross AIT, CVI and Cornell University, respectively.

Capacitance models and measurements

We identify several sources of capacitance that affect energy efficiency: (1) chip pads, (2) bump parasitics, (3) the microdisk junction, (4) the photodiode junction and (5) electronic driver capacitance. Extended Data Fig. 1a depicts these capacitance sources. These capacitances are determined through electrostatic simulations, circuit model simulations and empirical measurements. Focusing initially on the bumps and photonic chip pads, an electrostatic solver (Ansys Maxwell) is used to simulate the photonic chip pad-to-substrate capacitance (4 fF) and bump-to-bump parasitic capacitance (<1 fF). Similarly, we use a circuit model simulator (Cadence Virtuoso) to determine the electronic chip pad capacitance, which we find to be 6 fF. The simulator is also used to calculate the effective capacitance of the driver. We simulate the transmitter driving a capacitor (C_{load}), subtract $\frac{1}{4}C_{\text{load}}V^2$ from the total simulated energy consumption of the transmitter, and treat this adjusted driver-only power as $\frac{1}{4}C_{\text{driver}}V^2$. From this simulation, the effective capacitance of the driver is calculated to be 61 fF.

Experimental methods are used to determine the capacitances of the microdisk and photodiode junctions, along with a validation of the simulated pad capacitance. A vector-network analyser

(Keysight P5007A) is used to record an electrical RF reflection from probed devices. We fit the magnitude and phase of the reflected wave across RF frequencies to a reflection from a lumped complex impedance. Extended Data Fig. 1b shows the imaginary impedances of the devices and their associated fitted capacitor impedances. During these measurements, an RF bias tee provides a d.c. reverse-bias voltage to the device. An electronic calibration module (Keysight N4693D) is used to calibrate out the vector-network analyser response of the RF cable up to the output of the bias tee. The bias tee connects to a 25- μm pitch RF probe (FormFactor InfinityXT), and this probe response is calibrated out of the results by de-embedding its unlanded response as an electrical open. After this calibration, measurements of the photonic pad and bumped photonic pad show capacitances of 3 and 4 fF, respectively. Subsequent measurements are used to first de-embed these pad responses and then measure the photodiode and microdisk capacitances. The measured photodiode capacitance is 17 fF.

The measured microdisk capacitance decreases with an increasing reverse bias, as anticipated from a p–n junction. Extended Data Fig. 1c shows the measured capacitance of the four microdisks across reverse biases. The energy spent per bit transition is the integral of this capacitance function weighted by the difference of supply voltage and output voltage²⁴. A midpoint Riemann sum of the microdisk capacitances, weighted by reverse-bias voltage subtracted from 1 V, between 0 V and 1 V bias voltage, yields an effective disk junction capacitance for dynamic energy consumption of 128 fF (averaged across the four microdisks).

A summation of the pad, driver and junction capacitances results in a 199 fF transmitter capacitance. This result is in excellent agreement with the capacitance directly measured from the energy consumption of the transmitter (198 fF), validating the capacitance models and measurements.

Transmitter characterization

Each transmitter result is experimentally measured from the 3D-integrated photonic-electronic chip. An exception is the d.c. electro-optic response measurement, for which a separate photonic chip with the same modulator design is used. We apply a voltage to a probe to reverse bias the modulator at varying d.c. voltages and use an optical spectrum analyser (Keysight 8164B) to record each response. The remainder of the transmitter, receiver and link characterizations use an optical switch (Polatis 1000n 24 \times 10). This switch optically connects equipment and devices under test and minimizes fibre mating cycles. This approach streamlines the measurement process and eliminates potential power discrepancies that might stem from fibre mating inconsistencies. An optical spectrum analyser (Aragon BOSA 400) is used to measure the transmitter bus spectrum. For dynamic data transmission, a microcontroller is used to write a different PRBS6 pattern into the 64-bit registers in each of the 80 transmitter cells. Next, all modulators transmit this data simultaneously as the electronic chip is clocked, and all data registers are driven out of the chip to their respective microdisks. The 64-bit pattern transmitted by each modulator repeats indefinitely as the chip clock is running. In this state, we record the eye diagrams of each modulator, dynamic characteristics of the modulators and transmitter energy consumption from the electronic driver array voltage rail. The asymmetry between upper and lower regions of the eye diagrams may be due to uneven sizing of the P-channel metal-oxide-semiconductor (PMOS) and N-channel metal-oxide-semiconductor (NMOS) transistors in the transmitter driver. A narrow linewidth tunable laser (Santec TSL-210) is used as the light source in these measurements. Laser light travels through a fibre polarization controller and then into the chip. An oscilloscope (Tektronix DSA8300 with an 80C01 Optical Sampling Module) is used to receive modulated light for dynamic characterization and eye diagrams. In the BER test, modulated light initially passes through a variable optical attenuator (VOA) before reaching a commercial receiver

(Thorlabs RXM40AF). The commercial receiver converts the optical signals into electrical signals that are read by a BER tester (Anritsu MUI95040A). We sweep the received optical power with the VOA and record errors from the BER tester to construct the transmitter BER curves.

Receiver characterization

An ideal modulation source and an on-chip bit-error checker circuit are used to characterize the receiver cell performance. Separately, a tunable laser (Keysight 8164B) and a d.c. electrical probe landed on photodiode pads are used to measure the photodiode responsivity; the probe applies a reverse-bias voltage and senses photocurrent from a known input laser power. Next, for the dynamic characterization, we use an ideal modulation source (Thorlabs MX35E) consisting of an internal laser and a lithium niobate Mach–Zehnder modulator. A pulsed pattern generator (PPG, Anritsu MUI95020A) is used to drive the modulator with a repeating 64-bit PRBS6 pattern. The signal travels through fibre and a polarization controller before coupling into the photonic chip. Voltage is applied to a doped-silicon resistor adjacent to the ring filter to generate heat and tune the ring-filter resonance to the desired wavelength channel. The ring resonator drops the signal to a photodiode, which converts it from light into photocurrent for the electronic chip to then amplify. Timing circuits continuously write the received bits into a 64-bit long memory in a cycle. For timing, a programmable timing offset circuit in the electronic chip and a timing offset of the PPG align the incoming data to the receiver sampling point. A split clock source synchronizes the receiver chip and PPG clock frequencies. As the final step, readouts from the serial programming port display the saved received bits and confirm data reception. However, the serial port cannot update fast enough to give a BER in a short time frame. Instead, an on-chip error-counter circuit in each cell compares the received memory with preprogrammed expected bits, and if there is a discrepancy, it adds an error to an ongoing count. The readouts from the serial port display this count, and we obtain a BER curve as we sweep signal power using a VOA inside the ideal modulation source.

Link demonstration

The link demonstration combines the previously described experimental setups of the transmitter and receiver. A microcontroller sets the sent bits in a transmitter chip, and a second microcontroller reads the received bits in a separate receiver chip. Four channels of data transmit through the link at 8 Gb s^{-1} per channel and serial port readouts from the receiver record these data, along with an ongoing error count for each channel. A shared clock signal synchronizes the two transceivers and a programmable delay block in each receiver cell is used to delay the receiver sampling point with respect to the transmitter clock. A distributed feedback laser array (Thorlabs PRO8) is used as the four-channel optical source for the link. An arrayed waveguide grating multiplexes these different wavelengths of light from the laser array onto a single fibre. We place polarization controllers before the transmitter chip and before the receiver chip to ensure optimal coupling into the fundamental transverse electric mode of each waveguide. The fibres in the link are standard SMF28 fibre. An optical switch is used to direct light from the transmitter to an erbium-doped fibre amplifier and an oscilloscope for each eye diagram; the switch then directs light back to the 3D-integrated receiver for BER measurements without amplification. The optical losses in this link amount to 14.5 dB. These losses are from several sources: three chip-to-fibre interfaces at 3 dB each account for 9 dB, the modulation IL is 2.5 dB, a modulation penalty accounts for 1.5 dB (this is the difference between the optical power of a '1' bit and the average power) and an extra 1.5 dB of power is lost through the optical switch and fibre connectors throughout the link. Here, we use 8 Gb s^{-1} per channel to accommodate the timing circuits after our receiver. However, in the transmitter section, we

show error-free operation while transmitting 10 Gb s^{-1} data using a commercial clock and data recovery circuit. A 1.5 V drive is used due to the limitations of the receiver input current digital-to-analogue converter (DAC). Transmitter data at a 1 V drive has a lower extinction ratio and a higher average input current to the TIA, which exceeds the DAC circuit upper limit.

Electronic circuit architectures

Transmitted data starts as bits in the memory of each transmitter cell. Timing circuits running on a 5-GHz input clock (half transmitted data rate) generate memory read addresses and two-to-one multiplexer select signals. The circuits in the data path operate at a voltage supply of 1 V, except for the driver, which operates between 1 and 1.5 V. Extended Data Fig. 2a shows the driver circuit. Inside of the design, high-speed 1 V transistors in a cascode configuration prevents transistor junction breakdown from supply voltages exceeding 1 V. The main driver branch (M5–M8) has wide transistors to reduce the switching delay on modulator capacitance ($C_{\text{microdisk}} = 128 \text{ fF}$). A capacitor ($C_{\text{coupling}} = 183 \text{ fF}$) ensures a high switching speed while the auxiliary branch (M1–M4) holds the d.c. voltage level.

The receiver circuit senses a modulated photodiode current, amplifies it to digital levels at the supply rail voltage and deserializes the signal before writing it into internal memory. Extended Data Fig. 2b shows the amplifier circuit, which uses an inverter-based transimpedance amplifier (TIA) as an initial gain stage followed by an equalizer and inverters. A programmable current DAC at the amplifier input supplies a current (I_{DAC}) that cancels the d.c. offset of the photodiode current. The TIA stage has a high feedback resistance for a high gain ($R_f = 18.6 \text{ k}\Omega$). This resistor equates to a lower input resistance ($R_{\text{in}} = 2.1 \text{ k}\Omega$) from the Miller effect; however, it combines with the input capacitance ($C_{\text{photodiode}} = 17 \text{ fF}$, $C_{\text{pad}} = 10 \text{ fF}$) for a low-frequency pole. As a remedy, an active inductor circuit in the subsequent equalizer stage cancels out the TIA pole ($R_{\text{eq}} = 3.1 \text{ k}\Omega$, $C_{\text{eq}} = 33.6 \text{ fF}$). After the equalizer, ensuing inverters ensure the output swings between 0 and 1 V. An isolated, 1-V power supply of the receiver amplifiers mitigates supply noise.

The TIA dominates the receiver amplifier energy consumption, and its energy per bit is the static biasing power divided by the data rate. However, the TIA design introduces trade-offs between noise, bandwidth and power. Equation (1) shows how the receiver signal-to-noise ratio, SNR, relates to receiver energy per bit (E/bit_{RX}), input signal (I) input capacitance (C) and channel bandwidth (BW). Supplementary Note 3 provides a derivation of equation (1). This relationship sets a boundary on channel data rate scaling. With constant SNR and C , the design can expand BW with an increase in the input signal. In this context, the energy per bit remains constant, with the growing BW balancing out added laser power. This could imply an indefinite data rate scaling. However, a rise in BW necessitates wider TIA transistors, which subsequently contribute to the input capacitance, C . This sequence results in a degradation of SNR at the receiver for high BWs, establishing a cap on the energy-efficient per-channel data rate. To achieve higher data rates without compromising energy, the focus should be on parallel data communication across multiple channels. Similar conclusions have been made in other studies, which advocate for parallel channels operating at moderate data rates^{3,21}.

$$\text{SNR} \approx \left(\frac{I}{\text{BW}} \right)^2 \frac{E/\text{bit}_{\text{RX}}}{C^2}. \quad (1)$$

Data availability

The data used to generate the plots in Figs. 2–4 are available via figshare at <https://doi.org/10.6084/m9.figshare.28217315.v1> (ref. 42). All data that support the other findings of this study are available from the corresponding author upon reasonable request.

References

40. Fahrenkopf, N. M. et al. The AIM Photonics MPW: a highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits. *IEEE J. Sel. Top. Quantum Electron.* **25**, 1–6 (2019).
41. Byrd, M. J. et al. Mode-evolution-based coupler for high saturation power Ge-on-Si photodetectors. *Opt. Lett.* **42**, 851–854 (2017).
42. Daudlin, S. et al. 3D photonic integration for ultra-low energy, high bandwidth inter-chip data links. *figshare* <https://doi.org/10.6084/m9.figshare.28217315.v1> (2025).

Acknowledgements

This work was supported by the US Defense Advanced Research Projects Agency (DARPA) under PIPES (grant no. HRO011920014) and by the US Advanced Research Projects Agency-Energy (ENLITENED grant no. DE-AR000843). This work was also supported in part by the Center for Ubiquitous Connectivity (CUBiC), sponsored by Semiconductor Research Corporation (SRC) and Defense Advanced Research Projects Agency (DARPA) under the JUMP 2.0 program. S.D. acknowledges support by the National Science Foundation GRFP (grant no. DGE-1644869). We thank G. Keeler for leading the PIPES programme, N. Abrams and M. Hattink for helpful discussions and the engineering teams at AIM/SUNY Poly Photonics, Micross AIT and CVI for their roles in the transceiver fabrication and assembly.

Author contributions

S.D. designed the photonic chip and developed the 3D bonding process. S.D. led transceiver testing and photonic device

analysis with assistance from A.R., S.W., A.N. and V.G. A.R. compiled chip designs for the custom photonic wafer run. S.L., D.K., C.O. and A.M. designed the electronic chip and conducted bring-up tests of the electronic chip. S.D. and M.C. designed the printed circuit boards. R.P. and K.J. gathered wafer-level microdisk fabrication variation data. A.M. and K.B. supervised the project.

Competing interests

The authors declare no competing interests.

Additional information

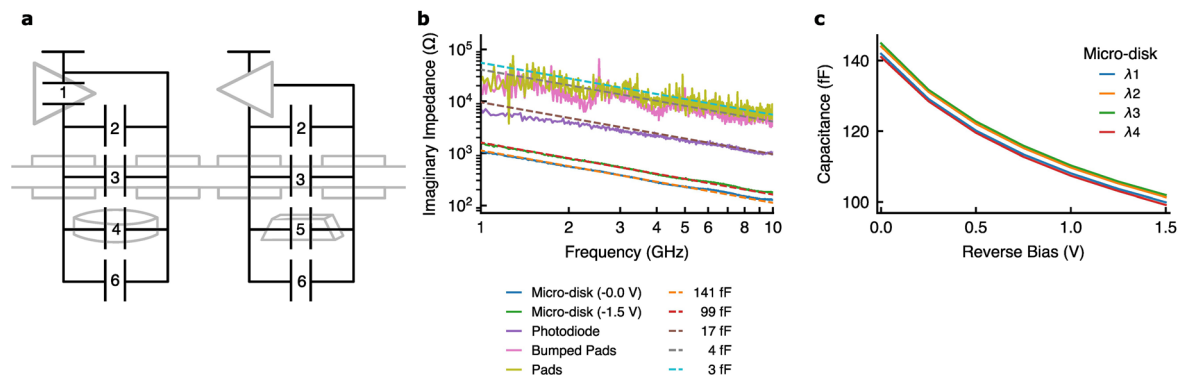
Extended data is available for this paper at <https://doi.org/10.1038/s41566-025-01633-0>.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41566-025-01633-0>.

Correspondence and requests for materials should be addressed to Keren Bergman.

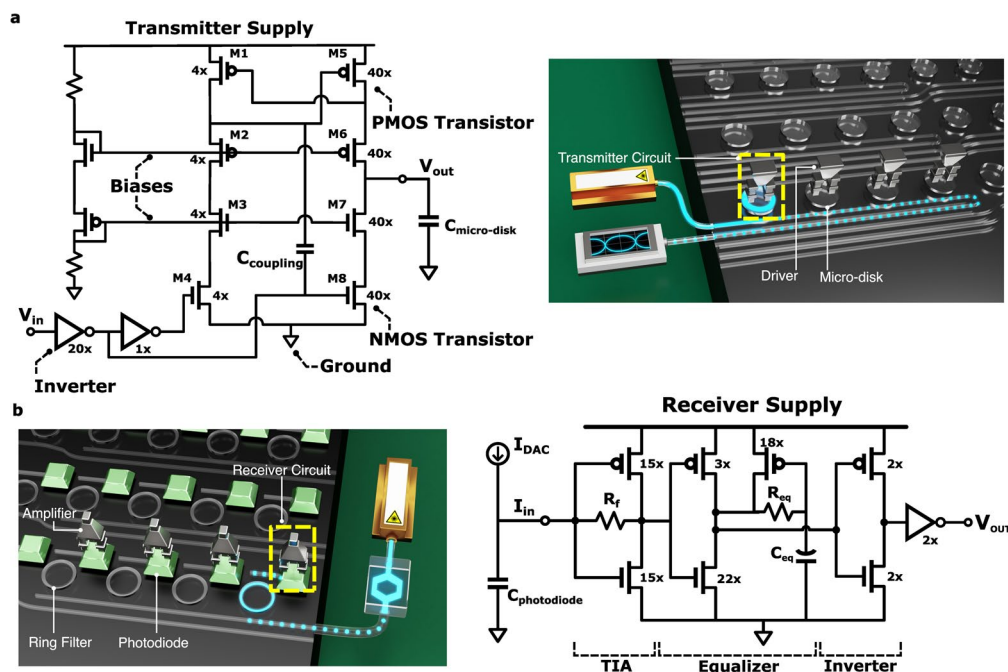
Peer review information *Nature Photonics* thanks David Moss for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.



Extended Data Fig. 1 | Device capacitances. **a**, Schematic of the transmitter and receiver capacitance sources; 1: electronic driver circuit, 2: electronic chip pads, 3: bump parasitics, 4: micro-disk PN junction, 5: photodiode PIN junction,

6: photonic chip pads. **b**, Measured imaginary impedances of devices and their fitted capacitor impedances. **c**, Measured micro-disk junction capacitance as a function of reverse bias voltage.



Extended Data Fig. 2 | Transmitter and receiver electronic circuit schematics. **a**, Transmitter driver circuit schematic and driver highlighted in the transmitter illustration; data as V_{in} drives the micro-disk as V_{out} . **b**, Receiver circuit schematic and the receiver highlighted in the receiver illustration; amplification

converts I_{in} to V_{out} . In the schematics, a multiplier labels each transistor width as a multiple of a '1x' transistor. The '1x' transistor has a width of 500 nm and length of 30 nm in the transmitter and 300 nm by 30 nm in the receiver.