LETTER

https://doi.org/10.1038/s41586-018-0028-z

Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip

Amir H. Atabaki[1,11]*, Sajjad Moazeni[2,11], Fabio Pavanello[3,6,7,11], Hayk Gevorgyan[4], Jelena Notaros[3,8], Luca Alloatti[1,9], Mark T. Wade[3,10], Chen Sun[2,10], Seth A. Kruger[5], Huaiyu Meng[1], Kenaish Al Qubaisi[4], Imbert Wang[4], Bohan Zhang[4], Anatol Khilo[4], Christopher V. Baiocco[5], Miloš A. Popović[4], Vladimir M. Stojanović[2] & Rajeev J. Ram[1]

Electronic and photonic technologies have transformed our lives—from computing and mobile devices, to information technology and the internet. Our future demands in these fields require innovation in each technology separately, but also depend on our ability to harness their complementary physics through integrated solutions[1,2]. This goal is hindered by the fact that most silicon nanotechnologies—which enable our processors, computer memory, communications chips and image sensors—rely on bulk silicon substrates, a cost-effective solution with an abundant supply chain, but with substantial limitations for the integration of photonic functions. Here we introduce photonics into bulk silicon complementary metal–oxide–semiconductor (CMOS) chips using a layer of polycrystalline silicon deposited on silicon oxide (glass) islands fabricated alongside transistors. We use this single deposited layer to realize optical waveguides and resonators, high-speed optical modulators and sensitive avalanche photodetectors. We integrated this photonic platform with a 65-nanometre-transistor bulk CMOS process technology inside a 300-millimetre-diameter-wafer microelectronics foundry. We then implemented integrated high-speed optical transceivers in this platform that operate at ten gigabits per second, composed of millions of transistors, and arrayed on a single optical bus for wavelength division multiplexing, to address the demand for high-bandwidth optical interconnects in data centres and high-performance computing[3,4]. By decoupling the formation of photonic devices from that of transistors, this integration approach can achieve many of the goals of multi-chip solutions[5], but with the performance, complexity and scalability of 'systems on a chip'[1,6–8]. As transistors smaller than ten nanometres across become commercially available[9], and as new nanotechnologies emerge[10,11], this approach could provide a way to integrate photonics with state-of-the-art nanoelectronics.

Sustained innovations in electronics, predominantly in CMOS, have transformed computing, communications, sensing and imaging. More recently, silicon photonics has been leveraging the CMOS infrastructure to address the growing demands for optical communications for internet and data centre networks[3,6,7]. This convergence of photonics with CMOS promises to transform electronic–photonic technologies, enabling processor and memory chips with high-bandwidth optical input/output[1,4], communications chips with high-fidelity optical signal processing[2,12], and highly parallel optical biochemical sensors for blood analysis[13] and gene sequencing[14]. To make these a reality, photonic devices need to be integrated with a variety of nanoelectronic functions (digital, analogue, memory, storage and so on) on a single silicon die (chip).

Monolithic (that is on a single chip) integration of photonic devices in close proximity to electronic circuits is crucial for two main reasons: it allows us to achieve the required levels of performance, scalability and complexity simultaneously for electronic–photonic systems; and substantially accelerates system-level innovation by enabling a cohesive design environment and device ecosystem to realize entire 'systems on a chip'. In fact, the accelerated progress in recent years in electronics is a direct result of such a system-on-a-chip approach and the addition of new functions and components to CMOS to create new monolithic device platforms, such as wireless communications and radar imaging chips (through the addition of inductors and transmission lines[15]) and image sensors (through silicon photodiodes[16]).

The greatest challenge towards the integration of photonic circuits into CMOS has been the lack of a semiconductor material with suitable optical properties for realizing active and passive photonic functions in bulk CMOS, which is the dominant manufacturing platform for microelectronic chips (every Intel, Apple and Nvidia CPU/GPU, all computer memory and flash storage, and so on). As a result, all efforts so far to integrate photonics into CMOS have been limited to silicon-on-insulator (SOI) substrates[1,6–8]. These processes are cost-prohibitive for many applications (for example, computer memory) and have a limited supply chain for high volume markets. The same photonic integration challenge also exists for the leading CMOS technologies below 28-nm transistor nodes—fin field effect transistor (FinFET) and thin-body fully depleted SOI[17] (TBFD-SOI)—where the crystalline silicon layers are too thin (less than 20 nm) to support photonic structures with sufficient optical confinement. To address these integration challenges, we have developed a photonic platform using an optimized polycrystalline silicon (polysilicon) film that could be deposited on silicon oxide islands that are ubiquitous in CMOS (used to isolate transistors) even in the most recent technologies using FinFET and TBFD-SOI[17] (Fig. 1a).

Deposited electronic and photonic devices on glass have already affected many fields: thin-film transistors have enabled today's display technologies, and photonic platforms with thin-film components on glass have been commercially deployed in optical communications systems[18]. However, deposited photonic components have been restricted to passive functions (for example, filters and delay lines) lacking light detection and modulation. A variety of materials, including amorphous and polycrystalline silicon[19–21], polymer-based devices[22] and chalcogenides[23], have been deposited on glass in the attempt to realize active photonic components. Nevertheless, the integration of a fully functional photonic platform (that is, passive functions, optical modulators and detectors) and its integration with CMOS nanoelectronics is yet to be demonstrated. In this work, we have integrated a fully functional polysilicon photonic platform with a 65-nm bulk CMOS process through the addition of a few extra processing steps without affecting the transistors' native performance, and demonstrated large-scale monolithic electronic–photonic systems.

Figure 1a shows transistor structures in today's three dominant deeply scaled CMOS processes. The silicon oxide shallow trench

[1]Massachusetts Institute of Technology, Cambridge, MA, USA. [2]University of California, Berkeley, Berkeley, CA, USA. [3]University of Colorado, Boulder, Boulder, CO, USA. [4]Boston University, Boston, MA, USA. [5]Colleges of Nanoscale Science and Engineering, State University of New York (SUNY) Polytechnic Institute, Albany, NY, USA. [6]Present address: Photonics Research Group, Ghent University-IMEC, Ghent, Belgium. [7]Present address: Center for Nano- and Biophotonics, Ghent University, Ghent, Belgium. [8]Present address: Massachusetts Institute of Technology, Cambridge, MA, USA. [9]Present address: Institute of Electromagnetic Fields (IEF), ETH Zurich, Zurich, Switzerland. [10]Present address: Ayar Labs, Inc., Emeryville, CA, USA. [11]These authors contributed equally: Amir H. Atabaki, Sajjad Moazeni, Fabio Pavanello. *e-mail: atabaki@mit.edu

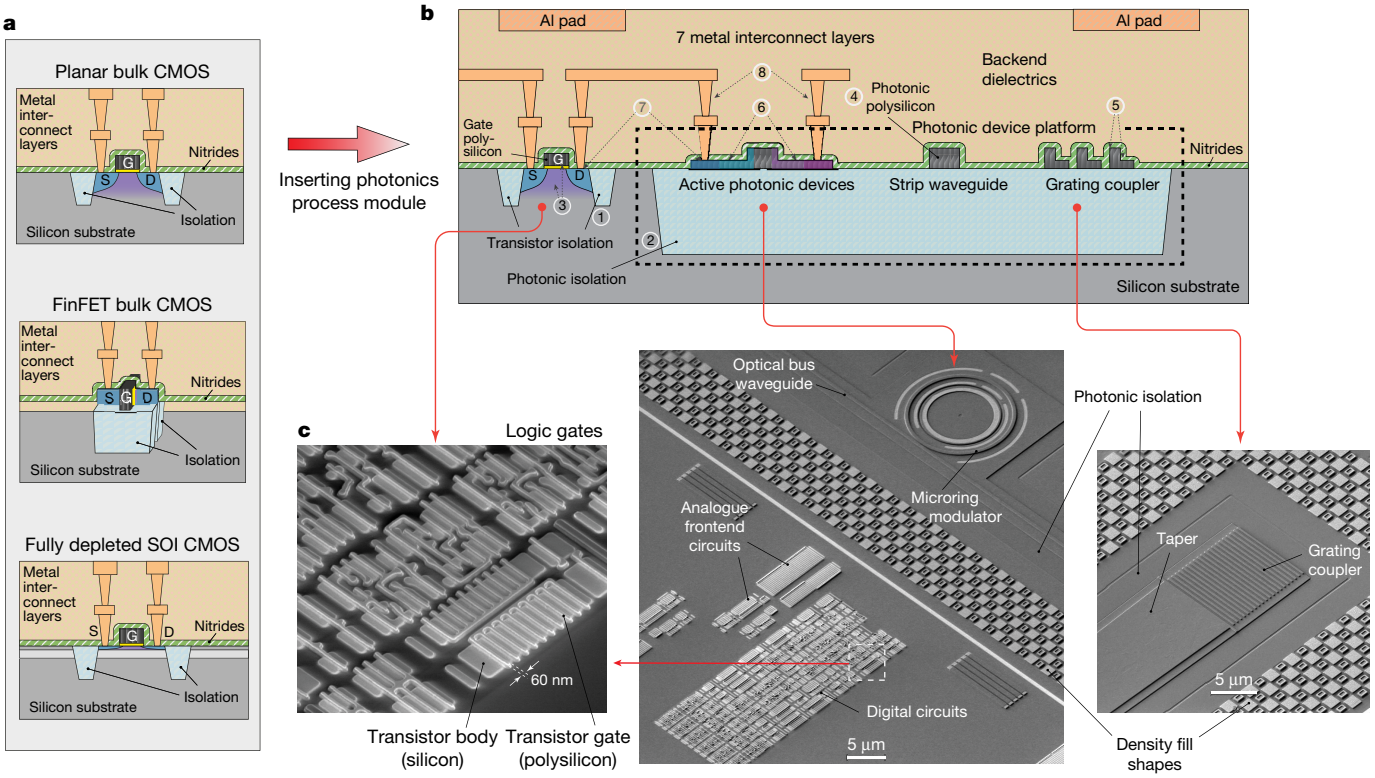

Fig. 1 | Photonic integration with nanoscale transistors. a, Illustration of three major deeply scaled CMOS processes: planar bulk CMOS, FinFET bulk CMOS, and fully depleted SOI CMOS. **b**, Integration of a photonics process module into planar bulk CMOS with photonic devices implemented in an optimized polysilicon film (220 nm) deposited on a photonic trench filled with silicon oxide (about 1.5 µm). The numbers indicate major fabrication steps in the order appearing in the process: (1) and (2), transistor and photonic isolation fabrication; (3) transistor frontend fabrication up to source/drain implant, including gate definition; (4) deposition, annealing and polishing of photonic polysilicon film; (5) polysilicon full and partial etching for forming strip and ridge photonic structures; (6) doping implants (P and N) for active photonics; (7) high doping implants (P++ and N++) and salicidation for both electronic and photonic devices; and (8) metallization. **c**, Scanning electron micrographs of different photonic and electronic blocks in our monolithic platform.

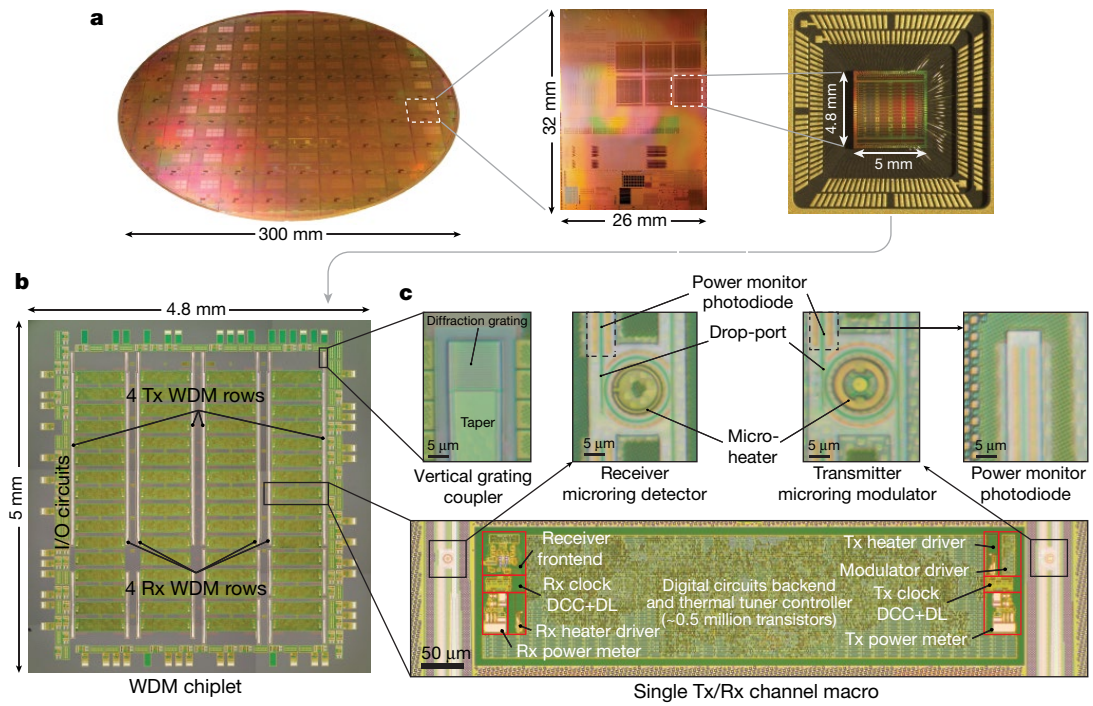

Fig. 2 | Monolithic electronic–photonic platform in 65-nm bulk CMOS. a, Photograph of a fully fabricated 300-mm wafer with monolithic electronics and photonics, and close-ups of a reticle on this wafer, and a packaged WDM chiplet. **b**, Micrograph of a WDM chiplet with four transmitter (Tx) and receiver (Rx) rows. **c**, Close-up of a single transceiver macro and its photonic and electronic circuit components, such as the grating coupler, optical modulator and detector, and power monitor photodiodes. I/O, input/output; DCC, duty-cycle corrector; DL, delay line.

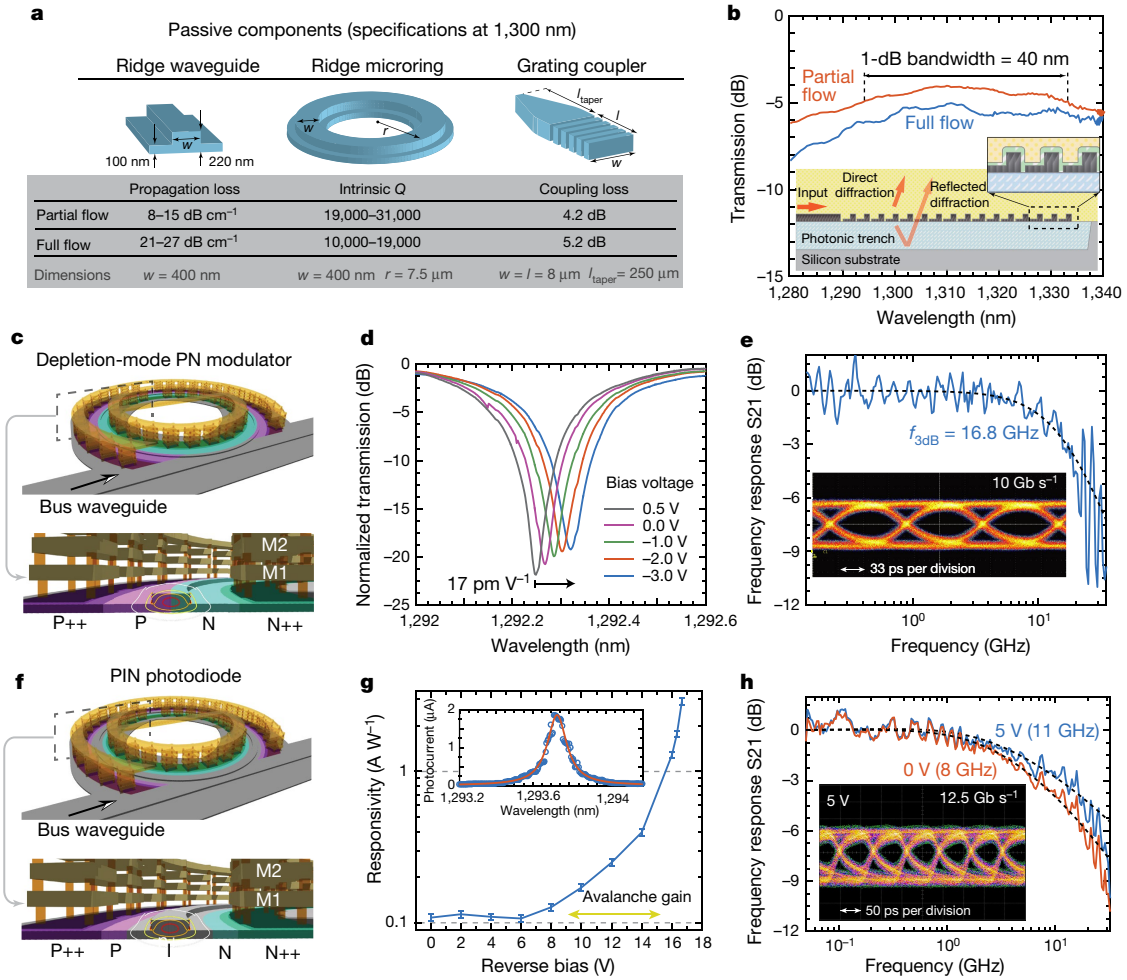

Fig. 3 | Photonics platform performance. a, Passive component specifications at 1,300 nm for partial- and full-flow wafers. **b**, Transmission spectrum and the longitudinal cross-section of the grating coupler (inset). **c**, Microring modulator three-dimensional layout. **d**, Transmission spectrum of a modulator resonance with loaded *Q*-factor of 5,000. **e**, Modulator electro-optic frequency response (S21) and the eye-diagram obtained with 2 V peak-to-peak drive voltage (V_{pp}), extinction ratio of 5.7 dB and insertion loss of 3.8 dB. **f**, Microring photodiode three-dimensional layout. M1 and M2 are the first and second metal layers in the process. P++ and N++ are the high doping regions under metal contacts and I is the intrinsic region in the photodiode. **g**, Responsivity versus reverse bias voltage. Avalanche gain is observed at biases above 8 V. **h**, Photodiode frequency response (S21) under 0 V and 5 V reverse bias with 3-dB bandwidths of 8 GHz and 11 GHz, respectively. The inset shows the eye diagram obtained under 5 V bias.

isolation for transistors in advanced CMOS nodes is too thin to support low-loss optical waveguides on top of this layer owing to light leakage into the substrate. We address this issue by locally adding a thicker silicon oxide photonic isolation layer (about 1.5 μm) with a fabrication process very similar to shallow trench isolation. An optimized polysilicon film (220 nm thick) with low optical propagation loss and high carrier mobility is then deposited on this layer, and is used for passive photonic components, free-carrier plasma dispersion modulators[24,25], and photodetectors that make use of the absorption by defect states at polysilicon grain boundaries[26,27]. Photonic isolation layer fabrication and polysilicon film deposition are followed by two etching steps (full and partial, for strip and ridge structures) and two doping implants (N-type and P-type for modulators and detectors) to form our photonics process module that is inserted into the CMOS fabrication process flow. Figure 1b shows the cross-sectional drawing of three representative photonic components in our polysilicon photonic platform next to a transistor in a planar bulk CMOS process.

The photonics process module is inserted in the middle of transistor processing, after gate definition, but before source and drain implants (see numbers in Fig. 1b for the fabrication order). With this approach, all of the high-temperature photonics processing takes place before the definition of the source, drain and channel of transistors. This eliminates the need for re-optimizing the source and drain implants and anneal processes that would otherwise be needed because of the sensitivity of deeply scaled transistors to the source and drain doping profiles[28]. Also, this approach allows us to reuse some of the frontend processing steps (high-doping implants, and silicide formation) for active photonic components to minimize the number of photolithography masks. In doing so, the entire fabrication development is shifted to the photonics side, because low-loss photonic structures have to be implemented while transistor gate features already exist on the same level. This necessitates careful optimization of the polysilicon film deposition, polishing, and etching steps to achieve low (<1 nm) surface and sidewall roughness for low-loss and high-performance devices (see Methods for process details).

This optimized photonics process was integrated with an entire commercial 65-nm bulk CMOS process with seven metal interconnect layers, featuring transistors with three different threshold voltages and two gate oxide thickness variants. Figure 1c shows bird's eye scanning electron micrographs of our monolithic platform with photonic components next to transistors with 60 nm channel length. This platform is fabricated on 300-mm-diameter wafers (the largest size in production at present) in a CMOS foundry located at the Colleges for Nanoscale Sciences and Engineering, SUNY Polytechnic Institute, Albany, New York. Figure 2a shows a photo of a fully fabricated wafer, and close-up photos of the entire reticle and one packaged chiplet composed of

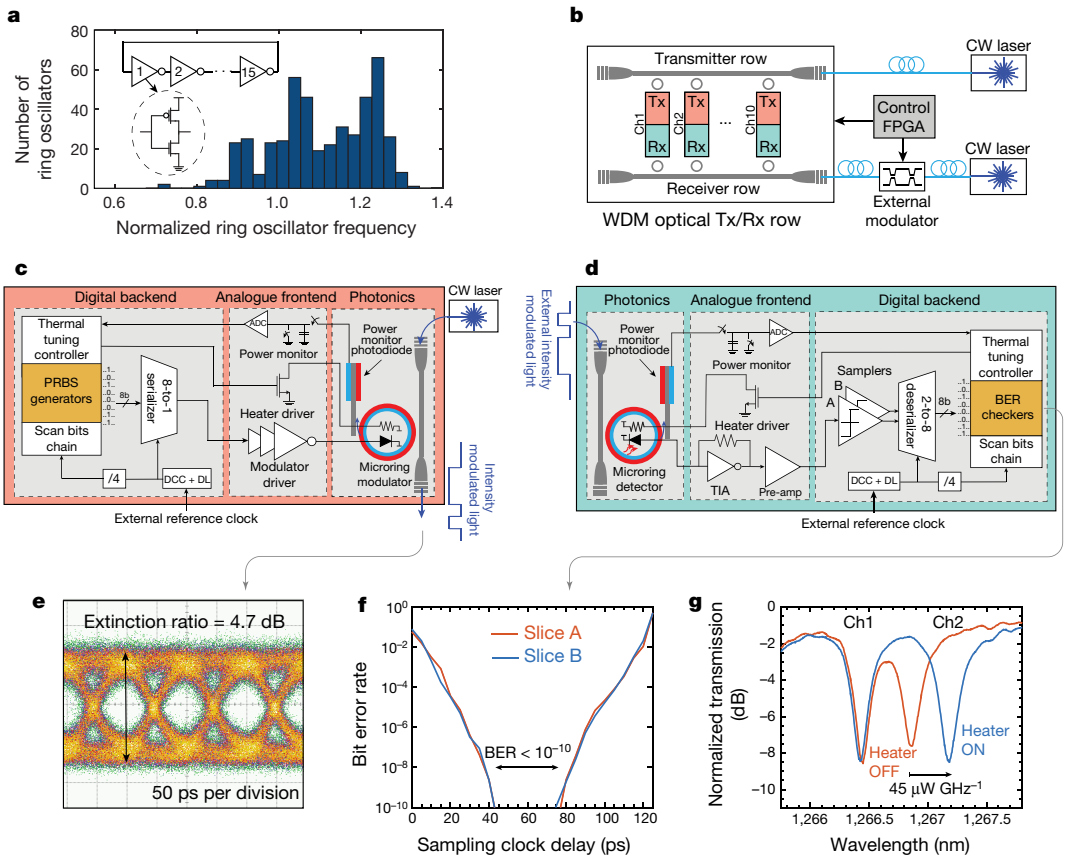

Fig. 4 | Electro-optical testing of WDM transceiver chips. a, Histogram of measured frequencies of 485 test ring oscillators normalized to the frequency obtained from simulation with the native CMOS process design kit models. FPGA, field-programmable gate array. **b**, Block diagram of one WDM transmit–receive row in our test setup. **c**, Block-diagram of one transmitter channel. **d**, Block diagram of one receiver channel. TIA, transimpedance amplifier. **e**, 10 Gb s^{-1} transmit eye diagram using the on-chip pseudorandom binary sequence (PRBS) generator. **f**, 7 Gb s^{-1} receiver bathtub curve obtained from the on-chip bit-error-rate (BER) tester by sweeping the delay between the clocks for the receiver and external transmitter (sampling clock delay). **g**, Thermal tuning of one WDM channel using the integrated microheater.

several wavelength division multiplexing (WDM) photonic transceivers. The micrographs of the transceiver chiplet, transmitter and receiver circuit blocks, and individual photonic components are shown in Fig. 2b and c. We were able to build a library of passive and active photonic components (waveguides, microring resonators, vertical grating couplers, high-speed modulators, and avalanche photodetectors) with a performance similar or better than previous demonstrations on polysilicon[20,21,26], next to circuit blocks composed of millions of transistors operating at native CMOS process specifications.

Figure 3a summarizes the performance of passive components measured on partial-flow (passive photonics only) and full-flow (active and passive photonics with electronics) wafers, at a wavelength of 1,300 nm. We achieved a propagation loss of approximately 10 dB cm^{-1} for ridge and strip waveguides, and a loaded quality factor (Q-factor) of >20,000 for microring resonators on partial-flow wafers. Full-flow wafers exhibit higher loss, but this issue did not have much effect on the performance of our optical transceivers: the 20 dB cm^{-1} waveguide loss results in a loss of 3 dB across the 10-lambda (that is, wavelength) WDM rows, and the loaded Q-factor of 10,000 of microring resonators is close to optimal for resonant modulators and detectors of bandwidth 10–20 GHz (see Methods for further discussion). Waveguide loss and resonator Q-factor are two times better at 1,550 nm (Extended Data Fig. 1), but all optical transceivers were initially designed at 1,300 nm. Grating couplers for coupling light into and out of the chip are designed using both the partial- and full-etch steps to construct a periodic L-shaped geometry (Fig. 3b). The measured grating transmission, shown in Fig. 3b, indicates a peak efficiency of −4.2 dB for the partial flow and −5.2 dB for the full flow, with 1-dB bandwidth of around 40 nm (see Methods for discussions of further device improvements).

Depletion-mode resonant modulators and defect-based photodetectors were implemented using ridge microring structures with lateral PN and PIN diode junctions, respectively (where P, N and I refer to P-type, N-type and Intrinsic). Two mid-level doping implants, P-type and N-type with concentrations of 6×10^{18} cm^{-3}, are optimized for these devices. Device micrographs and designs are shown in Figs. 2c and 3c and f (see Methods for discussion on implants).

The modulator uses a lateral PN junction to modulate light through the modulation of the resonance wavelength using the free-carrier plasma dispersion effect[29] (Fig. 3d). By operating the PN junction under full depletion, a 3-dB bandwidth of 16.8 GHz (Fig. 3e) and digital modulation at 10 Gb s^{-1} is achieved with only 2 V V_{pp} modulation signal (inset to Fig. 3e).

The defect-based photodetector has a responsivity of 0.11 A W^{-1} (quantum efficiency of 10%) near 1,300 nm under very low bias voltages (Fig. 3g) using a resonant design that enhances the weak absorption in polysilicon (inset to Fig. 3g). We also observed avalanche gain for the first time in polysilicon photodetectors[30] at bias voltages above 8 V (Fig. 3g), leading to a responsivity of 1.3 A W^{-1} at bias 16 V with a noise equivalent power of 0.27 pW Hz$^{-1/2}$. This device has a 3-dB bandwidth of more than 8 GHz under reverse bias voltages above 0 V, reaching 11 GHz for a bias of 5 V (Fig. 3h). More results on photodetectors are given in Extended Data Fig. 2.

To examine the performance of transistors after introducing the photonics module into the CMOS process, electrical ring oscillators composed of 15 equally sized inverting stages were used inside all electronic–photonic blocks to probe the speed of transistors, as well as the intra- and inter-die variations. The fastest transistors (low threshold voltage) in the process with gate lengths of 55 nm were used for the ring

oscillator design. Figure 4a shows the histogram of the normalized frequency of ring oscillators relative to the nominal frequency (2.33 GHz, single stage delay of 14.3 ps), simulated with the original process design kit provided by the foundry. This distribution is within the standard range of native CMOS processes, confirming that our photonic process module does not degrade the performance of transistors.

As a first demonstration of monolithic electronic–photonic systems in this platform, we have implemented high-bandwidth photonic WDM transceivers. We designed a total of six chiplets, each containing four stand-alone WDM transmitter and receiver rows, each supporting up to 16 channels. Different designs for resonant modulators and detectors were used in WDM rows. The chiplets are diced and wire-bond packaged with 100 pads to provide direct-current (d.c.) supplies, bias signals and high-speed clocks for electro-optical testing (Fig. 2a). By integrating all of the analogue and digital blocks, signal generation and error estimation of transceivers can be performed on the chip.

The transmitter is composed of a full digital backend that generates a pseudorandom binary sequence signal, a serializer with an 8 to 1 ratio, and finally an inverter chain that drives the microring modulator (Fig. 4c). On the receiver side, a transimpedance amplifier analogue frontend converts and amplifies the received photo-current into a voltage signal, and a pair of double-data-rate samplers converts the signal into the digital domain[31] (Fig. 4d). These bits are deserialized and fed into a bit-error-rate checker on the chip. The generated pseudorandom binary sequence signal and bit-error-rate data are monitored via on-chip scan chains to measure the functionality and performance of the transceivers. Overall, approximately 30,000 logic gates (about 0.5 million transistors) from digital standard cells have been used in each transceiver channel.

The operation of one channel of a 10-lambda WDM transceiver is shown in Fig. 4e and f. Modulators were operated in the depletion mode (voltage swing from 0 to −1.5 V) at 10 Gb s^{-1} data-rate with an extinction ratio of 4.7 dB (Fig. 4e). The receiver achieved a bit-error rate better than 10^{-10} with −3 dBm input optical power at 7 Gb s^{-1}, as shown in Fig. 4f. The speed of receivers is limited at present by the long on-chip clock distribution network and could be further improved by integrating local clock generators[31]. Thermal tuning controllers and heater drivers are also included in the transceivers, to adjust for microring resonance fluctuations due to temperature and process variations[32]. Using the digital-to-analogue converters in the thermal tuning controllers and heater drivers, we measured a tuning efficiency of 45 μW GHz^{-1} for integrated microheaters on microring modulators and detectors (Fig. 4g). The total electrical energy consumption of the transmitter and receiver including the serializer and deserializer was 100 fJ b^{-1} and 500 fJ b^{-1}, respectively. The transceiver achieved a bandwidth density of 180 Gb s^{-1} mm^{-2} with 10% of the effective area occupied by photonics, which can be reduced to 5% by optimizing the floorplan. Incorporating this photonics platform in advanced sub-10-nm technology nodes with higher transistor densities[33] would lead to >2 Tb s^{-1} mm^{-2} bandwidth densities meeting the needs of next-generation systems on a chip.

The optical transceivers in bulk CMOS demonstrated here are an important milestone towards multi-terabytes-per-second optical interconnects for direct integration with logic and memory to improve the performance of computing systems, at present limited by the chip input/output bandwidth. This photonic platform and integration approach illustrates how adding photonic functions onto a variety of substrates could enable the next generation of systems on a chip for computing, communications, imaging and sensing.

Online content

Any Methods, including any statements of data availability and Nature Research reporting summaries, along with any additional references and Source Data files, are available in the online version of the paper at https://doi.org/10.1038/s41586-018-0028-z

Received: 5 October 2017; Accepted: 19 February 2018;
Published online 18 April 2018.

1. Sun, C. et al. Single-chip microprocessor that communicates directly using light. *Nature* **528**, 534–538 (2015).
2. Khilo, A. et al. Photonic ADC: overcoming the bottleneck of electronic jitter. *Opt. Express* **20**, 4454–4469 (2012).
3. Rumley, S. et al. Silicon photonics for exascale systems. *J. Lightwave Technol.* **33**, 547–562 (2015).
4. Miller, D. A. Rationale and challenges for optical interconnects to electronic chips. *Proc. IEEE* **88**, 728–749 (2000).
5. Boeuf, F. et al. A multi-wavelength 3D-compatible silicon photonics platform on 300 mm SOI wafers for 25 Gb/s applications. In *Electron Devices Meeting (IEDM)* 13–3 (IEEE, 2013).
6. Assefa, S. et al. Monolithic integration of silicon nanophotonics with CMOS. In *Photonics Conference (IPC)* 626–627 (IEEE, 2012).
7. Narasimha, A. et al. A 40-Gb/s QSFP optoelectronic transceiver in a 0.13 μm CMOS silicon-on-insulator technology. In *Optical Fiber Communication Conference (OMK7)* (Optical Society of America, 2008).
8. Awny, A. et al. A 40 Gb/s monolithically integrated linear photonic receiver in a 0.25 μm BiCMOS SiGe:C technology. *IEEE Microw. Wirel. Compon. Lett.* **25**, 469–471 (2015).
9. Ieong, M., Doris, B., Kedzierski, J., Rim, K. & Yang, M. Silicon device scaling to the sub-10-nm regime. *Science* **306**, 2057–2060 (2004).
10. del Alamo, J. A. Nanometre-scale electronics with III–V compound semiconductors. *Nature* **479**, 317–323 (2011).
11. Desai, S. B. et al. MoS_2 transistors with 1-nanometer gate lengths. *Science* **354**, 99–102 (2016).
12. Ghelfi, P. et al. A fully photonics-based coherent radar system. *Nature* **507**, 341–345 (2014).
13. Mudumba, S. et al. Photonic ring resonance is a versatile platform for performing multiplex immunoassays in real time. *J. Immunol. Methods* **448**, 34–43 (2017).
14. Quail, M. A. et al. A tale of three next generation sequencing platforms: comparison of Ion Torrent, Pacific Biosciences and Illumina MiSeq sequencers. *BMC Genomics* **13**, 341 (2012).
15. Zolfaghari, A., Chan, A. & Razavi, B. Stacked inductors and transformers in CMOS technology. *IEEE J. Solid-State Circuits* **36**, 620–628 (2001).
16. Fossum, E. R. CMOS image sensors: electronic camera-on-a-chip. *IEEE Trans. Electron Dev.* **44**, 1689–1698 (1997).
17. Magarshack, P., Flatresse, P. & Cesana, G. UTBB FD-SOI: A process/design symbiosis for breakthrough energy-efficiency. In *Proc. Conf. Design, Automation and Test in Europe* 952–957 (EDA Consortium, 2013).
18. Wörhoff, K., Heideman, R. G., Leinse, A. & Hoekman, M. TriPleX: a versatile dielectric photonic platform. *Adv. Opt. Technol.* **4**, 189–207 (2015).
19. Selvaraja, S. K. et al. Low-loss amorphous silicon-on-insulator technology for photonic integrated circuitry. *Opt. Commun.* **282**, 1767–1770 (2009).
20. Sun, C. et al. A monolithically-integrated chip-to-chip optical link in bulk CMOS. *IEEE J. Solid-State Circ.* **50**, 828–844 (2015).
21. Preston, K., Manipatruni, S., Gondarenko, A., Poitras, C. B. & Lipson, M. Deposited silicon high-speed integrated electro-optic modulator. *Opt. Express* **17**, 5118–5124 (2009).
22. Dangel, R. et al. Polymer-waveguide-based board-level optical interconnect technology for datacom applications. *IEEE Trans. Adv. Packag.* **31**, 759–767 (2008).
23. Eggleton, B. J., Luther-Davies, B. & Richardson, K. Chalcogenide photonics. *Nat. Photon.* **5**, 141–148 (2011).
24. Xu, Q., Schmidt, B., Pradhan, S. & Lipson, M. Micrometre-scale silicon electro-optic modulator. *Nature* **435**, 325–327 (2005).
25. Shainline, J. M. et al. Depletion-mode carrier-plasma optical modulator in zero-change advanced CMOS. *Opt. Lett.* **38**, 2657–2659 (2013).
26. Mehta, K. K. et al. Polycrystalline silicon ring resonator photodiodes in a bulk complementary metal-oxide-semiconductor process. *Opt. Lett.* **39**, 1061–1064 (2014).
27. Atabaki, A. H., Meng, H., Alloatti, L., Mehta, K. K. & Ram, R. J. High-speed polysilicon CMOS photodetector for telecom and datacom. *Appl. Phys. Lett.* **109**, 111106 (2016).
28. Nishi, Y. & Doering, R. *Handbook of Semiconductor Manufacturing Technology* (CRC Press, 2000).
29. Soref, R. & Bennett, B. Electrooptical effects in silicon. *IEEE J. Quantum Electron.* **23**, 123–129 (1987).
30. Ackert, J. J. et al. 10 Gbps silicon waveguide-integrated infrared avalanche photodiode. *Opt. Express* **21**, 19530–19537 (2013).
31. Settaluri, K. T. et al. Demonstration of an optical chip-to-chip link in a 3D integrated electronic-photonic platform. In *41st European Solid-State Circuits Conf. (ESSCIRC)* 156–159 (IEEE, 2015).
32. Sun, C. et al. A 45 nm CMOS-SOI monolithic photonics platform with bit-statistics-based resonant microring thermal tuning. *IEEE J. Solid-State Circ.* **51**, 893–907 (2016).
33. Yang, S. et al. 10 nm high performance mobile SoC design and technology co-developed for performance, power, and area scaling. In *Symp. on VLSI Technology* T70–T71 (IEEE, 2017).

Acknowledgements We thank E. Timurdogan and M. Byrd for help with high-speed photonic measurements, and A. Nikiforov of the Boston University Photonics Center for assistance with sample imaging and analysis. This work was supported by DARPA POEM award HR0011-11-C-0100 and contract HR00111190009, led by J. Shah. We also acknowledge support from the Microsystems Technology Laboratories (MTL) at MIT, the Berkeley Wireless Research Center, the Center for Nanoscale Sciences and Engineering and the

Packard Foundation. The views expressed are those of the authors and do not reflect the official policy or position of the DoD or the US Government.

Author contributions A.H.A., S.M. and F.P. all contributed equally to this work. A.H.A. designed and tested avalanche photodetectors, tested photonic devices during process development and designed WDM receiver rows. S.M. designed the mixed-signal electronics for transceiver chiplets, tested optical transceivers and electronics, and performed top-level assembly of electronics and photonics on the chip. F.P. designed and tested optical modulators, designed WDM transmitter rows, and performed top-level assembly of photonics in the transceiver chips. A.H.A., S.M. and F.P. all contributed to the chip verification. H.G. tested optical modulators and loss test structures on full flow wafers. J.N. designed and tested the original version of grating couplers. L.A. developed the CAD infrastructure for the photonics layout and chip verification, and contributed to the layout of waveguide test structures. M.T.W. designed and contributed to the layout of optical passive devices. C.S. contributed to the design of WDM optical transceivers and thermal tuning circuits. S.A.K. ran process experiments for propagation loss improvement, performed CMOS and photonics fabrication compatibility optimization, fabricated full flow wafers, and performed inline electrical and optical testing. H.M. contributed to the design of avalanche photodetectors. K.A.K. tested grating couplers on full-flow wafers. I.W. performed device metrology and post processing of optical waveguide and resonator loss data. B.Z. performed grating coupler simulations and contributed to the new grating coupler designs. A.K. analysed device metrology and measurement data for studying and optimizing photonic performance. C.V.B. developed the 65-nm CMOS-compatible photonics design kit, designed methodology for the in-line electrical tests, metrology and optical photonics testing, worked with the mask house with building the photolithography masks and managed the overall process development and wafer fabrication. M.A.P., V.M.S. and R.J.R. supervised the project.

Competing interests C.S., M.T.W., R.J.R., M.A.P. and V.M.S. are involved in developing silicon photonic technologies at Ayar Labs.

Additional information
Extended data is available for this paper at https://doi.org/10.1038/s41586-018-0028-z.
Reprints and permissions information is available at http://www.nature.com/reprints.
Correspondence and requests for materials should be addressed to A.H.A.
Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

METHODS

Chip implementation. Photonic device layouts were developed and drawn in Cadence Virtuoso (an industry-standard design tool for frontend electronics in conjunction with mixed-signal electronics[34]). Digital electronics were implemented using a combination of digital synthesis and place and route tools from Cadence. All photonic and electronic designs conform to the 65-nm CMOS technology manufacturing rules (more than 5,000 rules). New design rules were added to the original CMOS rules for the new photonics masks that were added to the process. The most critical mask rules with the introduction of photonics into the process are density rules. This is important because photonics and electronics occupy separate regions on the chip, but their respective masks have to maintain a certain maximum and minimum density of shapes across the whole design area. These density rules were met by custom fill shapes designed by our team. Density fill shapes can be seen in the scanning electron micrographs in Fig. 1c. The physical design verification was performed using Mentor Graphics Calibre (https://www.mentor.com/products/ic_nanometer_design/verification-signoff/).

Fabrication. Designs were fabricated on 300-mm wafers in the fabrication facility at Colleges for Nanoscale Sciences and Engineering, State University of New York, Albany, New York. The photonics passive-only wafers (partial flow) were fabricated on silicon wafers with 1.5-μm-thick SiO_2 under-cladding blankets with the whole CMOS backend dielectric stack as the over-cladding. Partial-flow wafers were fabricated for photonics process optimization before integration with electronics. For full-flow wafers (passive and active photonics with electronics), the deep photonic trench is first fabricated by etching the trench in the silicon substrate and filling it with SiO_2 by chemical vapour deposition followed by a planarization step. At this point, the wafer goes through the CMOS frontend process up to the source and drain formation. Photonic device fabrication is then followed by the deposition, annealing and planarization of a 220-nm photonic polysilicon layer. Using two reactive ion etching steps, one full (etching the entire 220 nm depth of polysilicon film) and one partial (120 nm deep), strip and ridge photonic structures are formed. This is followed by the photonic mid-level doping implants. From here, electronics and photonics share the rest of the fabrication process, including the high-doping implants (for the transistor source and drain, and the photonic modulator and detector ohmic contacts), nitride liners (silicide block, and etch-stop for the first via), silicide formation and metallization. There are a total of seven metal interconnect layers in this process, with the first four having a lithography resolution of less than 100 nm.

Each wafer quadrant on full-flow wafers received a separate mid-level doping implant concentration for photonic active components (modulators and detectors) of $[1, 2, 3, 6] \times 10^{18}\,cm^{-3}$. Owing to the presence of a large density of defects in polysilicon, the carrier activation occurs only after the majority of defect states are occupied, whose onset occurs for a doping concentration[35] of roughly $10^{18}\,cm^{-3}$. This necessitates careful optimization of photonic mid-level P and N doping concentrations to balance loss, modulator efficiency and device series resistance, which affects the speed of both modulators and detectors. By using a separate doping concentration in each quadrant of the wafer, we tested the performance of modulators and detectors as a function of doping concentration. From the results of doping splits in an earlier fabrication run for optimizing the photonics process, we expected the optimal doping concentration to be close to $3 \times 10^{18}\,cm^{-3}$. However, in the full-flow run, owing to an increase in optical loss caused by polishing residues, microring Q-factors dropped by a factor of two. This required larger wavelength shifts in modulators to compensate for the broadened resonance lineshape to achieve the same level of modulation depth. Therefore, we observed the best overall performance for a P and N implant concentration of $6 \times 10^{18}\,cm^{-3}$. The results presented in this paper are for devices receiving this implant concentration.

Fabrication results. The mask density rules ensure that material density during polishing, etching and lithography is within an acceptable range over the entire reticle and wafer to eliminate pattern-dependent results and achieve a high fabrication yield. Nevertheless, the maximum density range for each layer (polysilicon, metals, and so on) is desirable for more design flexibility. In this fabrication run, we faced unforeseen issues with photonic trench planarization, owing to a large density gradient of photonic trenches across the reticle field. This caused dielectric residues on the wafer after photonic trench planarization. We also experienced metal residues after the fabrication of the first via contact. Both of these issues led to a factor-of-two degradation in the passive photonic performance in full-flow runs ($20\,dB\,cm^{-1}$ versus $10\,dB\,cm^{-1}$ for waveguide loss, and 10,000 versus 20,000 for microring Q-factor at 1,300 nm). Both of these issues were resolved through modified design rules, and optimized fabrication processes for the next fabrication run.

Optical testing. Tunable lasers from Agilent Technologies and Santec were used for the optical characterization. Standard single-mode fibres (SMF28) were used to couple light into and out of the chip using grating couplers. The width of the grating couplers is matched to the mode size of the SMF28 fibre. We used 3-axis positioner stages (Thorlabs NanoMax) to position and align fibres over the grating couplers of the test sites. Minimum fibre-to-coupler insertion loss was achieved by angling the fibres at 15° off-normal from the surface of the chip. Waveguide losses were estimated by measuring transmitted optical power for four different waveguide lengths (50 μm, 1 mm, 5 mm and 15 mm), and fitting the propagation loss to the transmission measurements. Microring Q-factors were estimated by fitting a Lorentzian function to microrings at close to the critical coupling condition (the condition for zero transmission at resonance). Fibre-to-chip grating coupler efficiencies were extracted from the transmission measurement data, by fine-tuning the fibre angle at the input and output to achieve minimum transmission loss and subtracting the waveguide loss connecting the two identical grating couplers at the input and output of the test structure. Electro-optical frequency response (S21) of optical modulators and photodetectors were measured using an Agilent Vector Network Analyzer (VNA, 8722D). For modulator bandwidth testing, the modulator was driven by VNA Channel 1 with the bias voltage applied using a Bias Tee (SHF BT 65), and the modulator output was detected by a high-bandwidth photodiode (Discovery DSC30-3-2010) which was connected to Channel 2 of the VNA for S21 measurement. For photodetector bandwidth testing, VNA Channel 1 was used to drive an external lithium niobate modulator (JDSU 10022054), and the photodetector was biased using SHF BT 65, whose radio-frequency output was connected to VNA Channel 2 for S21 measurement. We used high-bandwidth radio-frequency probes (Cascade Infinity, 50 μm pitch) for high-speed testing. Eye diagrams for the modulators and detectors were obtained using a similar setup, with a pseudorandom binary sequence pattern generator (Picosecond Programmable Pattern Generator, SDG Model 12072) and a high-bandwidth oscilloscope (Agilent Technologies 86108B Precision Waveform Analyzer). The microring photodetector responsivity was measured by dividing the device photocurrent by the input optical power, which was estimated by measuring the optical power in the input fibre before entering the chip and accounting for the fibre-to-chip grating coupler, and excess waveguide losses.

Electrical testing. Chips were assembled in ceramic packages (CPG20809) with 100 wirebond connections. These packages were plugged into a socket on a host printed circuit board, which delivers supplies, bias signals, the high-speed clock (from an Agilent 81142 A pulse generator), and scan control signals from an Opal-Kelly FPGA. Scan commands for each measurement are set in Python scripts on a computer and then sent to the FPGA to configure the chip for each particular experiment. To read out the ring-oscillator frequencies, the output of the oscillator is fed into an asynchronous digital divider (divide by 8) and the divided clock runs a digital counter block. The oscillator's frequency was then estimated by scanning out the counter's value. The transmitter's eye diagram was captured via an external Ortel photoreceiver with 10-GHz bandwidth on a digital communication analyser (DCA) oscilloscope by running the on-chip pseudorandom binary sequence modules at a 5-GHz external clock frequency. On-chip clock adjustment circuits composed of a duty-cycle corrector and a delay line were used to synchronize the timing of different transceivers. The receiver's bit-error-rate test was performed by first programming a KC705 Xilinx FPGA with the same pseudorandom binary sequence coefficients used for our on-chip bit-error-rate checkers. The output of the FPGA was then amplified using a high-voltage modulator driver (JDSU H301), which then drives an external JDSU MZI optical modulator. Modulated light is amplified using a semiconductor optical amplifier (Thorlab BOA1130) and is coupled into the chip. The bit-error-rate bathtub curves in Fig. 4f show the measured bit-error rate at each time delay point between the clock fed into the chip and the FPGA reference clock.

Device design and discussion. The reticle area was divided into two main sections: the test device section with photonic and electrical test structures and the electronic-photonic microsystem section (WDM chiplets). The test area included waveguide loss, microring Q-factor, grating coupler efficiency, and sheet resistance test structures as well as individual modulators and detectors. A variety of parameters in every device was swept to find optimal designs and to extract information about the quality of the fabrication (for example, estimating surface and sidewall roughness, doping activation, and so on). Photonic test devices were designed for both 1,300 nm and 1,550 nm wavelengths. Overall, approximately 1,000 test devices were laid out on the reticle. The microsystems included 6 transceiver chiplets, each 4.8 mm × 5 mm. Since the same set of masks was used for process optimization and system implementation, there were some uncertainties about the performance of photonic passives and actives. This required sweeping modulator and detector parameters in the WDM transceiver designs to cover a large enough range of device performance.

All photonic components were designed using two etch steps (partial and full). The thickness of the polysilicon layer (220 nm) and the depth of the partial etch (120 nm) were chosen to optimize the overall performance of the whole platform, including the efficiency of grating couplers, radiation loss of the microring resonators, and series resistance of the modulators and detectors. The optimal width of the single-mode waveguides and diameter of the high-Q microrings were around 450 nm and 15 μm, respectively, for 1,300 nm operation. The high doping regions for ohmic contacts were about 1 μm away from the centre of the ridge waveguides

to avoid free-carrier loss. The width of the intrinsic region for the photodiode is 800 nm in the reported device. Microheaters were incorporated in the microring modulators and detectors for tuning their resonances to the desired laser wavelength. We used silicided P+ polysilicon resistors for the microheaters. The fabricated microheaters had a resistance of about 100 Ω.

The thickness of the photonic trench was optimized to maximize the grating coupler efficiency. A thickness of approximately 1.5 μm results in constructive interference of the main beam diffracted upwards with the beam diffracted downward and reflected from the surface of the silicon substrate. This improved the directionality and coupling efficiency of the grating couplers.

Device micrographs. The micrographs of the chip and discrete photonic components in Fig. 2 are taken from the back side of the die, subsequent to complete removal of the silicon substrate using XeF_2 gas after mounting the die on a carrier substrate.

Improvement of photonic performance. We have already taken the necessary steps to improve the waveguide loss in our full-flow wafers by fixing the photonic isolation planarization issue on the next fabrication run. Hence, we expect to achieve the same passive photonics performance in our fully integrated wafers as that reported in partial-flow runs in this work. This means improvements of a factor of two in waveguide loss (10 dB cm^{-1}) and microring *Q*-factor (20,000), simply by resolving the photonic isolation planarization issue.

The loss improvement will also affect the performance of active photonics, resulting in a higher quantum efficiency for photodetectors by reducing the fraction of photons lost by scattering. We expect a 50%–100% improvement in detector responsivity (0.15–0.2 A W^{-1} in the linear mode, and 2–2.6 A W^{-1} with avalanche gain). The improvement in the microring modulator *Q*-factor leads to a sharper resonance feature, which consequently reduces the drive voltage and lowers the transmitter power consumption.

The waveguide loss can be improved even further to about 6 dB cm^{-1} across the entire telecom and datacom bands (1,250–1,700 nm) by optimizing the polysilicon film. At present, waveguide loss at shorter wavelengths (1,300 nm) is twice as high as at 1,550 nm (Fig. 1). The similarity of the detector quantum efficiency (11%) at 1,300 nm and 1,550 nm suggests that the scattering and absorption losses are increasing proportionately at shorter wavelengths. As optical modes at shorter wavelengths are more confined in the waveguide core and exhibit less scattering by the sidewall roughness, the dominant mechanism for the increase in scattering loss at shorter wavelengths is most probably the increase in scattering by polysilicon grain boundaries. This source of scattering can be reduced by optimizing polysilicon deposition and anneal conditions such that the scattering correlation length is reduced.

The bandwidth of the modulators and detectors can also be improved by optimizing the doping profiles. In the present run, we employed the same source and drain implants for the low-resistance regions (P++ and N++) of our modulators and detectors. These implants have concentrations that are designed for the shallow source and drain of deeply scaled transistors and are not high enough to provide low series resistance for the 100-nm-thick polysilicon regions in our photonic devices. By using dedicated implant masks and high implant dosages, and by further optimizing dimensions in all doping profiles, we expect to improve the bandwidth of modulators by 25%–50% to 20–24 GHz based on the estimates for the contributions of mid-level (P and N) and high-level doping (P++ and N++) regions to the series resistance. For the detectors, we also expect an improvement in bandwidth by reducing the width of the intrinsic region from 800 nm to reduce the transit time, which is at present limiting the speed of the device. We expect that a 50% reduction of the intrinsic region width to 400 nm would not substantially affect the *Q*-factor and responsivity of detectors, while reducing the depletion width and transit time. A careful optimization of the intrinsic region width combined with the reduction of the *RC* time constant (resistance × capacitance) by adjusting the implant conditions can improve the bandwidth of detectors by 25%–50% to 14–16.5 GHz in the linear mode and to 10–12 GHz in the avalanche mode.

We also expect an improvement in the grating coupler efficiency in the next fabrication run. In the present run, a 30-nm photolithography bias caused the dimensions of the grating couplers to be smaller than the nominal design values. This caused the grating coupler efficiency to drop from −1.8 dB in design to −5 dB in the fabricated devices. We are addressing this issue by optimizing the lithography step and pre-biasing the photolithography mask. Also, by using a nonuniform grating design[36], we expect to improve the mode matching of the grating coupler to the Gaussian mode profile of the optical fibre. This will enable us to improve the coupler efficiency to below −1 dB.

Electronic–photonic systems on glass. The present work was aimed at integrating photonics into bulk CMOS technologies. However, a fully functional deposited photonic platform on glass transcends any one particular substrate or application. All of our partial-flow photonics-only silicon wafers were covered by a blanket of 1.5 μm-thick plasma-enhanced chemical vapour deposition silicon oxide, on which we have fabricated optical waveguides, resonators, modulators and photodetectors. These thin-film integrated photonic devices, along with the thin-film transistors that are currently used in display panels, could enable electronic–photonic systems on glass. These systems can be fabricated on low-cost large-area substrates such as metal foils, transparent glass or even flexible substrates as long as they are covered with roughly 1 μm of glass. Such a platform can enable a variety of new systems and applications that current electronic–photonic technologies cannot address owing to substrate size or cost limitations. For example, several space and astrophysics applications, such as laser communications and astronomical spectroscopy, require large-area optics and detectors. Also, many optical phased array applications (lidars, augmented reality headsets, and so on) could benefit greatly from large-area integrated photonic circuits. An electronic–photonic platform on glass, enabled by the deposited polysilicon photonics demonstrated in this work, could address these application areas. The performance of photonics on this platform would be similar to the devices we have reported on partial-flow wafers in this paper.

Data availability. The main data supporting the findings of this study are available within the article. Extra data are available from the corresponding author upon request.

34. Alloatti, L., Wade, M., Stojanovic, V., Popovic, M. & Ram, R. Photonics design tool for advanced CMOS nodes. *IET Optoelectron.* **9**, 163–167 (2015).
35. Seto, J. Y. W. The electrical properties of polycrystalline silicon films. *J. Appl. Phys.* **46**, 5247 (1975).
36. Notaros, J. et al. Ultra-efficient CMOS fiber-to-chip grating couplers. In *Optical Fiber Communications Conf. Exhib. (OFC)* 1–3 (IEEE, 2016).

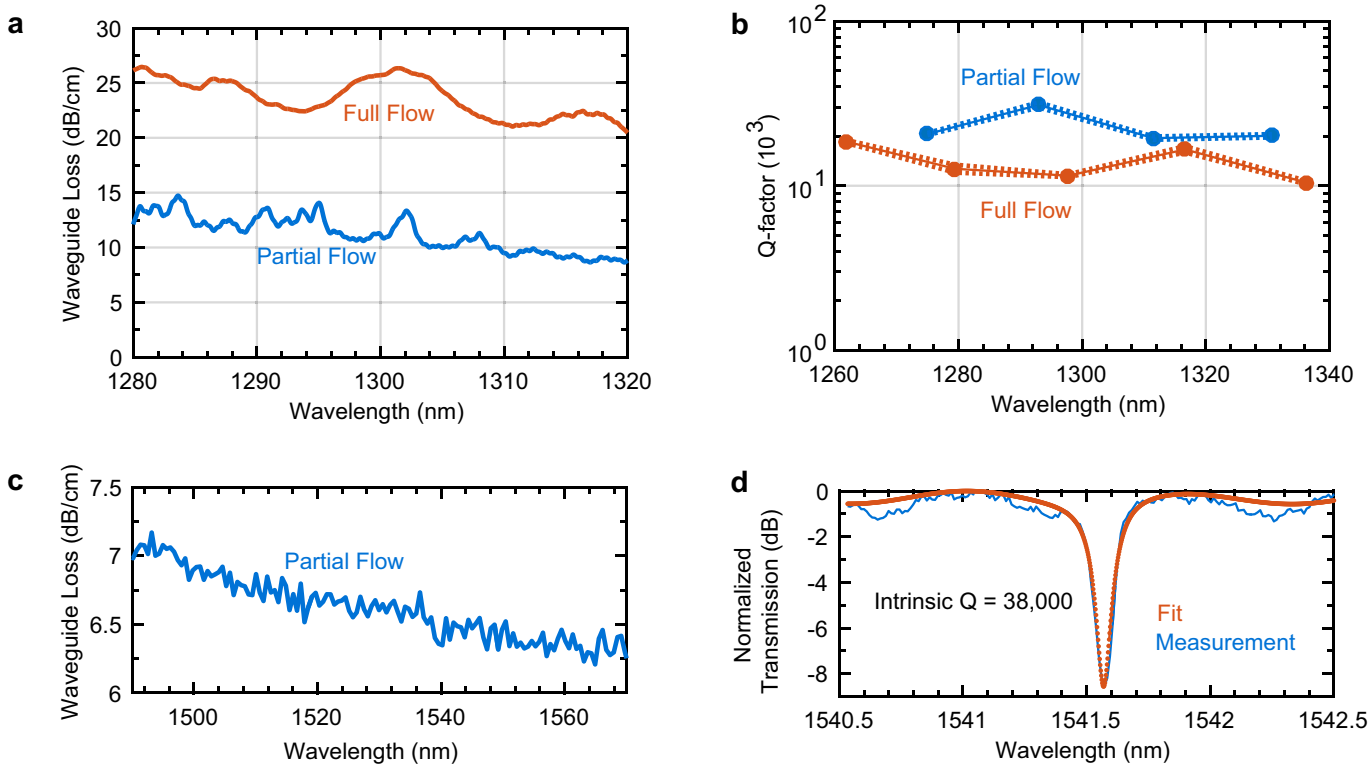

Extended Data Fig. 1 | Passive photonic performance at 1,300 nm and 1,550 nm. a, Waveguide propagation loss at 1,300 nm. Waveguide loss drops with wavelength because of a combination of lower absorption and scattering by polysilicon. **b**, *Q*-factor of a 15-μm-diameter microring resonator. **c**, Waveguide propagation loss at 1,550 nm. **d**, One resonance of a 17-μm-diameter microring near 1,540 nm with a *Q*-factor of 38,000.

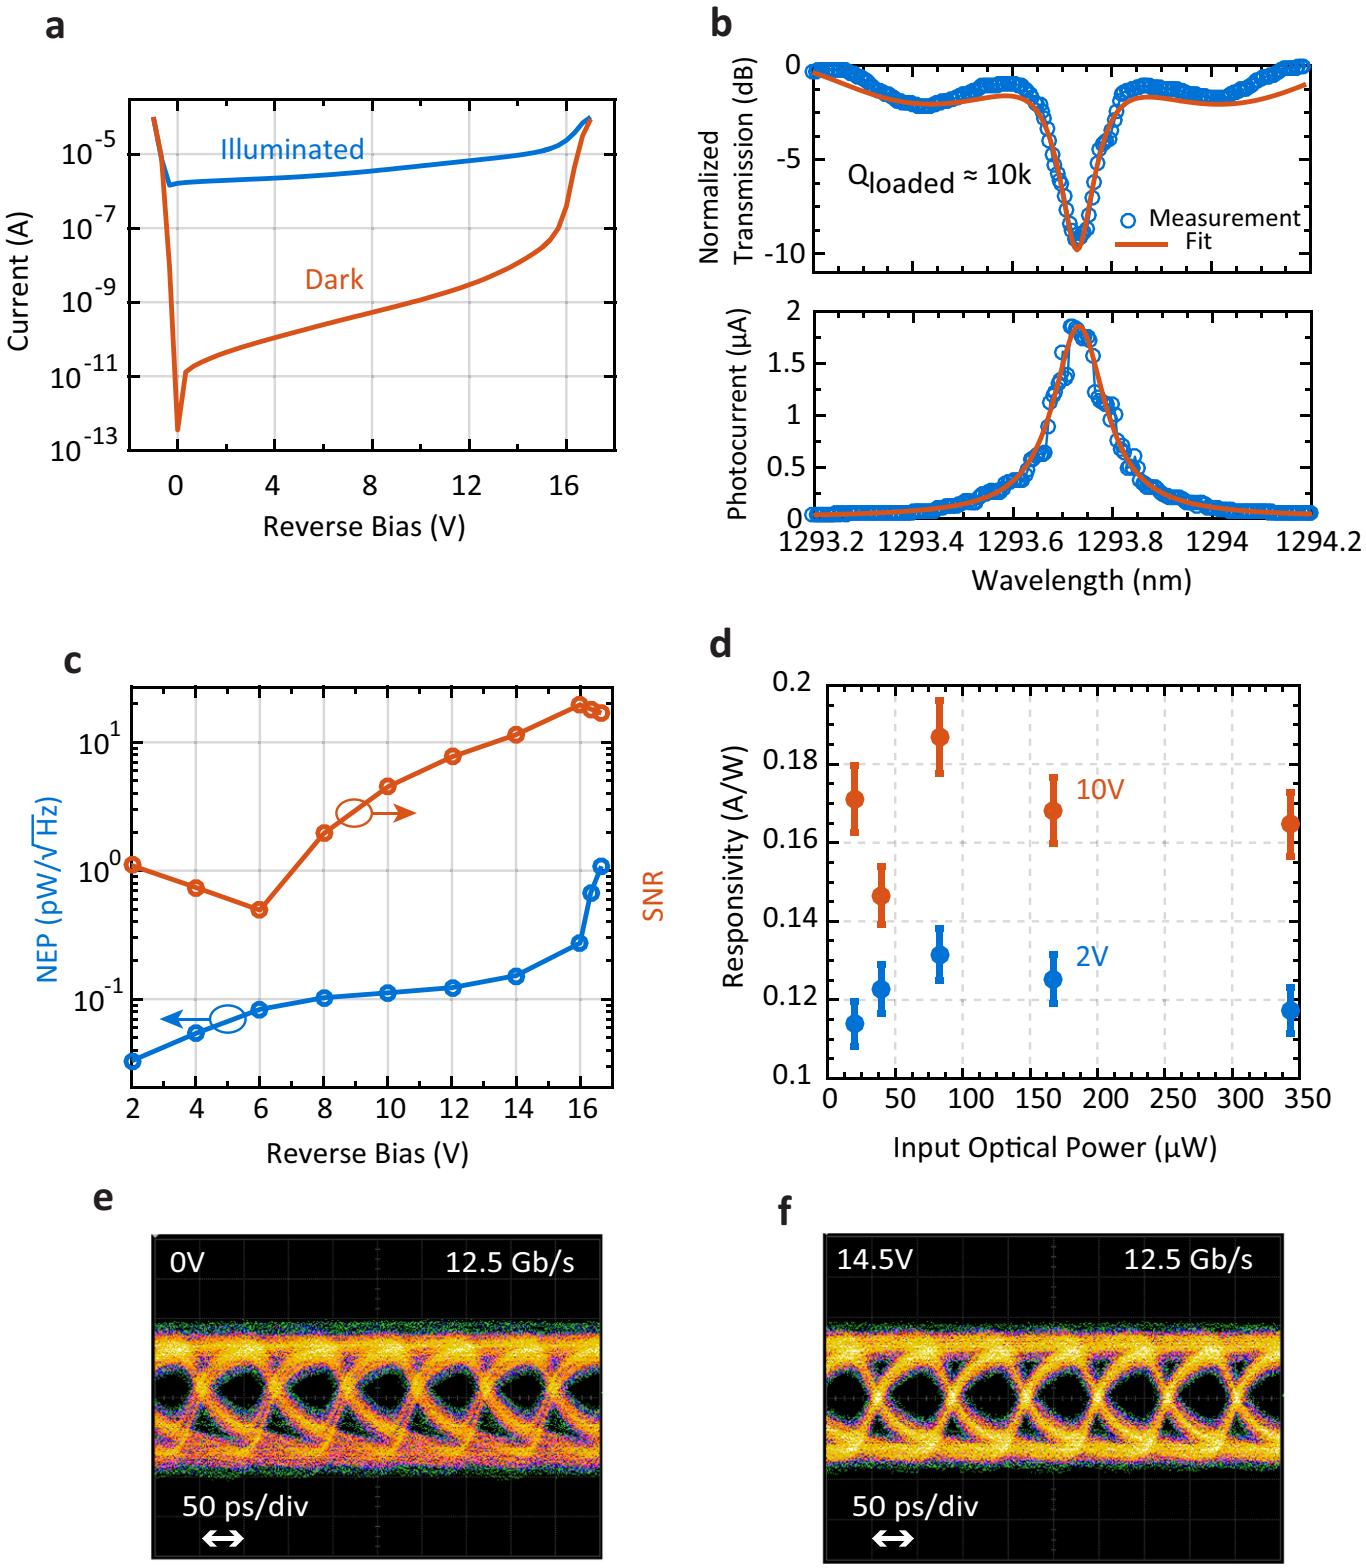

Extended Data Fig. 2 | Polysilicon avalanche photodetector. a, Current–voltage curve of the microring photodiode under dark and illumination for an input optical power of 20 μW. Dynamic range is about 60 dB and about 10 dB at 0 V and 16 V, respectively. **b**, One microring photodetector resonance (top) and the corresponding photo-current (bottom) as the wavelength is swept across the resonance. The loaded *Q*-factor (Q_{loaded}) of the microring is about 10,000. The fit is obtained through least-squares optimization of a model that includes a Lorentzian resonance for the microring and accounts for the reflections from the end facets of the chip to model the Fabry–Perot resonances observed in the transmission curve. **c**, Noise equivalent power (NEP, blue curve) of the photodiode estimated from the dark-current shot noise, which dominates the detector noise. Avalanche gain is 13 at 16 V bias, with an noise equivalent power of 0.27 pW $Hz^{-1/2}$. The simulated signal-to-noise ratio (SNR) (red curve) at the output of the optical receiver, assuming an optical signal of 1 μW, and a receiver circuit input-referred noise spectral density of 1 pA $Hz^{-1/2}$. **d**, The responsivity of the photodetector versus input optical power, showing minimal power dependency. The error bar is estimated based on a ±5% error in estimating the optical power in the waveguide before coupling into the detector. This error comes from variations in fibre to chip coupling efficiency owing to fibre-grating coupler misalignment. **e**, **f**, Eye diagrams at 12.5 Gb s^{-1} at 0 V and 14.5 V reverse bias.