LETTER

doi:10.1038/nature16454

Single-chip microprocessor that communicates directly using light

Chen Sun[1,2]*, Mark T. Wade[3]*, Yunsup Lee[1]*, Jason S. Orcutt[2]†*, Luca Alloatti[2], Michael S. Georgas[2], Andrew S. Waterman[1], Jeffrey M. Shainline[3]†, Rimas R. Avizienis[1], Sen Lin[1], Benjamin R. Moss[2], Rajesh Kumar[3], Fabio Pavanello[3], Amir H. Atabaki[2], Henry M. Cook[1], Albert J. Ou[1], Jonathan C. Leu[2], Yu-Hsin Chen[2], Krste Asanović[1], Rajeev J. Ram[2], Miloš A. Popović[3] & Vladimir M. Stojanović[1]

Data transport across short electrical wires is limited by both bandwidth and power density, which creates a performance bottleneck for semiconductor microchips in modern computer systems—from mobile phones to large-scale data centres. These limitations can be overcome[1–3] by using optical communications based on chip-scale electronic–photonic systems[4–7] enabled by silicon-based nanophotonic devices[8]. However, combining electronics and photonics on the same chip has proved challenging, owing to microchip manufacturing conflicts between electronics and photonics. Consequently, current electronic–photonic chips[9–11] are limited to niche manufacturing processes and include only a few optical devices alongside simple circuits. Here we report an electronic–photonic system on a single chip integrating over 70 million transistors and 850 photonic components that work together to provide logic, memory, and interconnect functions. This system is a realization of a microprocessor that uses on-chip photonic devices to directly communicate with other chips using light. To integrate electronics and photonics at the scale of a microprocessor chip, we adopt a 'zero-change' approach to the integration of photonics. Instead of developing a custom process to enable the fabrication of photonics[12], which would complicate or eliminate the possibility of integration with state-of-the-art transistors at large scale and at high yield, we design optical devices using a standard microelectronics foundry process that is used for modern microprocessors[13–16]. This demonstration could represent the beginning of an era of chip-scale electronic–photonic systems with the potential to transform computing system architectures, enabling more powerful computers, from network infrastructure to data centres and supercomputers.

The electro-optic system on a chip (Fig. 1) contains a dual-core RISC-V instruction set architecture[17] (ISA) microprocessor and an independent 1 MB bank of static random access memory that is used for memory. The on-chip electro-optic transmitters and receivers enable both the microprocessor and the memory to communicate directly to off-chip components using light, without the need for separate chips or components to host the optical devices. The chip was fabricated using a commercial high-performance 45-nm complementary metal–oxide semiconductor (CMOS) silicon-on-insulator (SOI) process[18]. No changes to the foundry process were necessary to accommodate photonics and all optical devices were designed to comply with the native process-manufacturing rules. This 'zero-change' integration enables high-performance transistors on the same chip as optics, reuse of all existing designs in the process, compatibility with electronics design tools, and manufacturing in an existing high-volume foundry.

The process includes a crystalline-silicon layer that is patterned to form both the body of the electronic transistors and the core of the optical waveguides. A thin buried-oxide layer separates the crystalline-silicon layer from the silicon-handle wafer (Extended Data Fig. 1). Because the buried-oxide layer is <200 nm thick, light propagating in crystalline-silicon waveguides will evanescently leak into the silicon-handle wafer, resulting in high waveguide loss. To resolve this, we perform selective substrate removal on the chips after electrical packaging to etch away the silicon handle under regions with optical devices (Extended Data Fig. 2). We leave the silicon handle intact under the microprocessor and memory (which dissipate the most power) to allow a heat sink to be contacted, if necessary. Substrate removal has a negligible effect on the electronics[13] and the processor is completely functional even with a fully removed substrate.

Silicon-germanium (SiGe) is present, although in low germanium mole fractions, in advanced CMOS processes to enhance hole mobility and transistor performance via compressive strain engineering of p-channel transistors[18]. Selecting a 1,180-nm wavelength band for the optical channel enables the use of photodetectors built using this SiGe (ref. 19). Silicon is transparent at 1,180 nm and no adverse effects are observed. At these wavelengths, the optical propagation loss in silicon-strip waveguides is 4.3 dB cm^{-1} (losses at industry-standard wavelengths of 1,300 nm and 1,550 nm are 3.7 dB cm^{-1} and 4.6 dB cm^{-1}, respectively[13]). The receiver circuit[20] resolves photocurrent produced by the illuminated photodetector into digital ones and zeros. The receiver sensitivity in optical modulation amplitude (OMA) is −5 dBm for a bit error ratio better than 10^{-12}.

The electro-optic transmitter consists of an electro-optic modulator and its electronic driver. The modulator is a silicon micro-ring resonator with a diameter of 10 μm, coupled to a waveguide. We dope the structure with the n-well and p-well implants used for transistors to form radially extending p–n junctions, interleaved along the azimuthal dimension[21,22], taking the form of a 'spoked ring'. The ring exhibits a sharp, notched-filter optical transmission response, with a stop-band at the resonant wavelength of the ring (λ_0). Applying a negative voltage across the junctions depletes the ring of free carriers (electron and hole concentrations), while a small positive voltage refills the carriers. A change in carrier concentration influences the refractive index of the ring waveguide as a result of the carrier plasma dispersion effect[23], which, in turn, shifts λ_0. Electro-optic modulation (on-off keying) is achieved by changing the voltage applied across the junction to move the λ_0 stop-band in and out of the laser wavelength (λ_L). The modulator has a loaded quality factor of approximately 10,000, and a voltage swing of only 1 V_{pp} (where V_{pp} is the peak-to-peak voltage) across the modulator achieves on:off ratios of 6 dB at an insertion loss of 3 dB for non-return-to-zero binary data. The low voltage, near-zero quiescent current, and low capacitance (15 fF, including wiring capacitance) result in an energy-efficient modulator driven by a standard CMOS logic

[1]University of California, Berkeley, Berkeley, California 94720, USA. [2]Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA. [3]University of Colorado, Boulder, Boulder, Colorado 80309, USA. †Present addresses: IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, USA (J.S.O.); National Institute for Standards and Technology, Boulder, Colorado 80305, USA (J.M.S.).
*These authors contributed equally to this work.

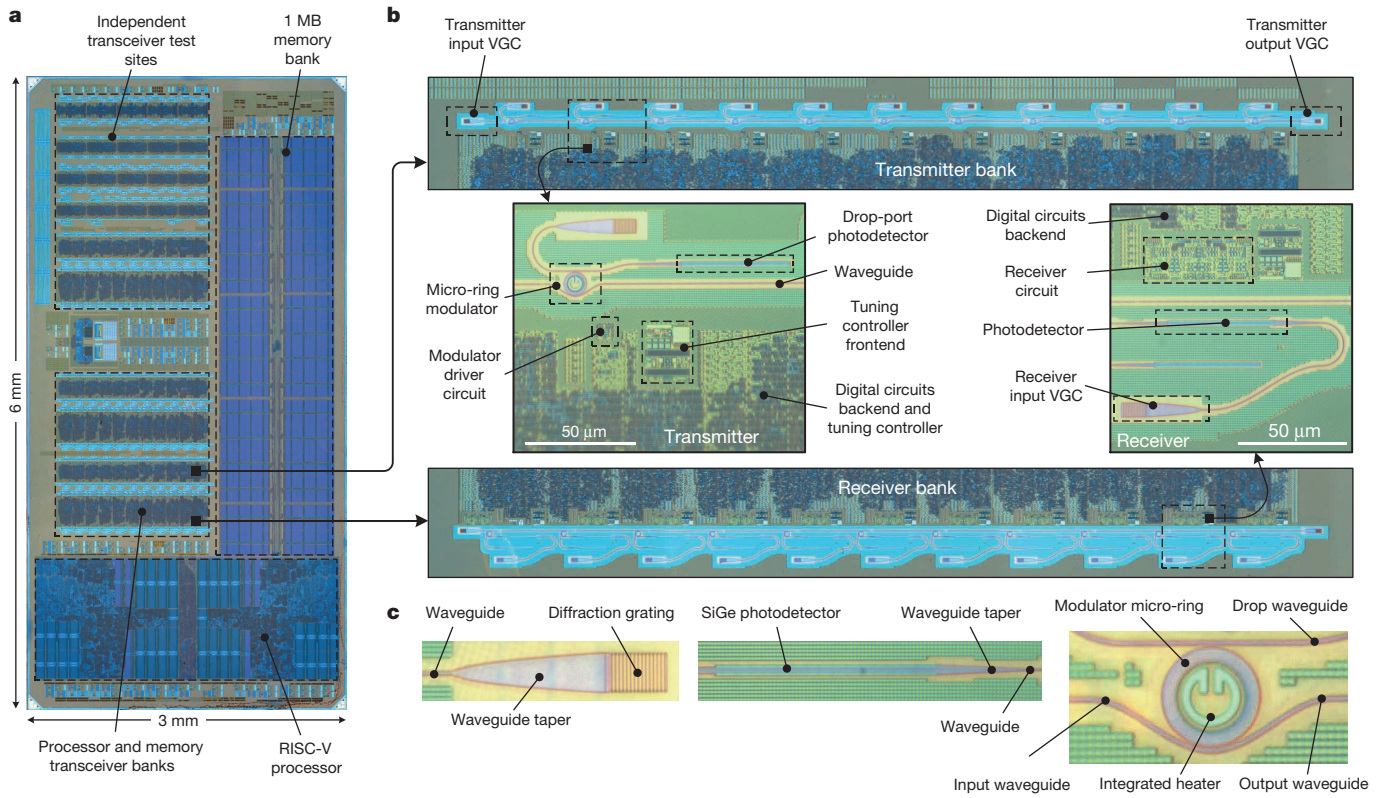

Figure 1 | The electro-optic system on a chip. a, Die photo of the 3 mm × 6 mm chip showing the locations and relative sizes of the processor, memory, and transceiver banks, imaged from the backside of the chip. **b**, The processor transmitter and receiver banks (the memory transmitter and receiver banks are identical) with close-ups of individual transmitters and receivers sites. **c**. Micrographs of the grating coupler, photodetector, and resonant micro-ring modulators (left to right).

inverter at gigabit data rates using the same 1-V nominal supply that powers digital electronics.

As a resonant device, the modulator is highly sensitive to variations in the thickness of the crystalline-silicon layer within and across SOI wafers[24] as well as to spatially and rapidly temporally varying thermal environments created by the electrical components on the chip[25,26]. Both effects cause λ_0 to deviate from the design value, necessitating tuning circuitry. We embedded a 400-Ω resistive microheater inside the ring to efficiently tune λ_0 and added a monitoring photodetector weakly coupled to the modulator drop port. When light resonates in the modulator ring, a small fraction of it couples to and illuminates the photodetector. This generates photocurrent proportional to the amount of resonating light, which is maximized when $\lambda_0 = \lambda_L$ (modulator is directly on resonance). Taking advantage of the densely integrated electronics, we designed a digital controller that monitors the photocurrent and controls the power to the microheater to keep λ_0 locked to λ_L under thermal variations[20]. When λ_0 has a large offset from λ_L, such as during chip power-up, and when no photocurrent feedback is available, the controller 'sweeps' λ_0 by stepping the power output of the heater up or down. This sweep works to reduce the λ_0-to-λ_L offset until sufficient photocurrent to begin the main feedback loop is obtained. The controller achieves initial lock ($\lambda_0 = \lambda_L$) within 7 ms and has a tracking time constant of 13 µs after lock-on. This system provides up to 3 nm of change in λ_0 and can compensate temperature swings of 60 K (ref. 20), aided by the superior thermal isolation afforded by selective substrate removal.

We use the direct chip-to-chip optical connectivity of the microprocessor chip to build a photonically connected main memory system for the microprocessor (Fig. 2). The microprocessor chip optically communicates to the 1 MB memory array located remotely on a second identical chip an arbitrary distance away. The microprocessor sends requests (a 'read' or 'write'), the memory address (location in memory to read or write), and write data (for write requests) via the microprocessor-to-memory (P → M) link. The memory-to-microprocessor (M → P) link returns read data for read requests. A field programmable gate array (FPGA) provides the peripheral functionality of a motherboard, completing a user controllable computer.

For both P → M and M → P links, the laser light first couples into an electro-optic transmitter; laser light arriving in a single-mode fibre couples into an on-chip waveguide through a vertical grating coupler (VGC). The optical modulator, driven by circuits, modulates light in the waveguide and imprints it with on-off keyed binary data from the source. The light then exits the chip through a second vertical grating into a single-mode fibre bound for the other chip. Once there, the light couples into the receive site through a VGC, illuminates a receive photodetector, and is resolved back by the receiver circuit into binary data for the destination. The communication between the microprocessor and memory is full-duplex. Both P → M and M → P links run at 2.5 Gb s^{-1}, providing an aggregate 5 Gb s^{-1} of memory bandwidth. Our demonstration uses only one wavelength of light; each additional wavelength increases the memory bandwidth by 5 Gb s^{-1} for a total potential aggregate bandwidth of 55 Gb s^{-1} without the need to use additional fibres.

A single 1,183-nm continuous-wave off-chip solid-state laser acts as the light source, with output power split 50/50 to share it across both the P → M and M → P links. To overcome the 4–6 dB coupling losses through each VGC due to unoptimized grating couplers, we insert an optical amplifier, which provides about 9 dB of gain, to obtain sufficient optical power at the receiver to resolve the signal. Using the optimized VGCs with losses of 1.2 dB (ref. 27) that exist as standalone test devices elsewhere on the same chip would eliminate the need for optical amplifiers in future design iterations.

To verify functionality of the photonically connected memory in the computer, we ran a combination of terminal-based and graphical programs (see Fig. 3b for an excerpt). To run a program, the control FPGA

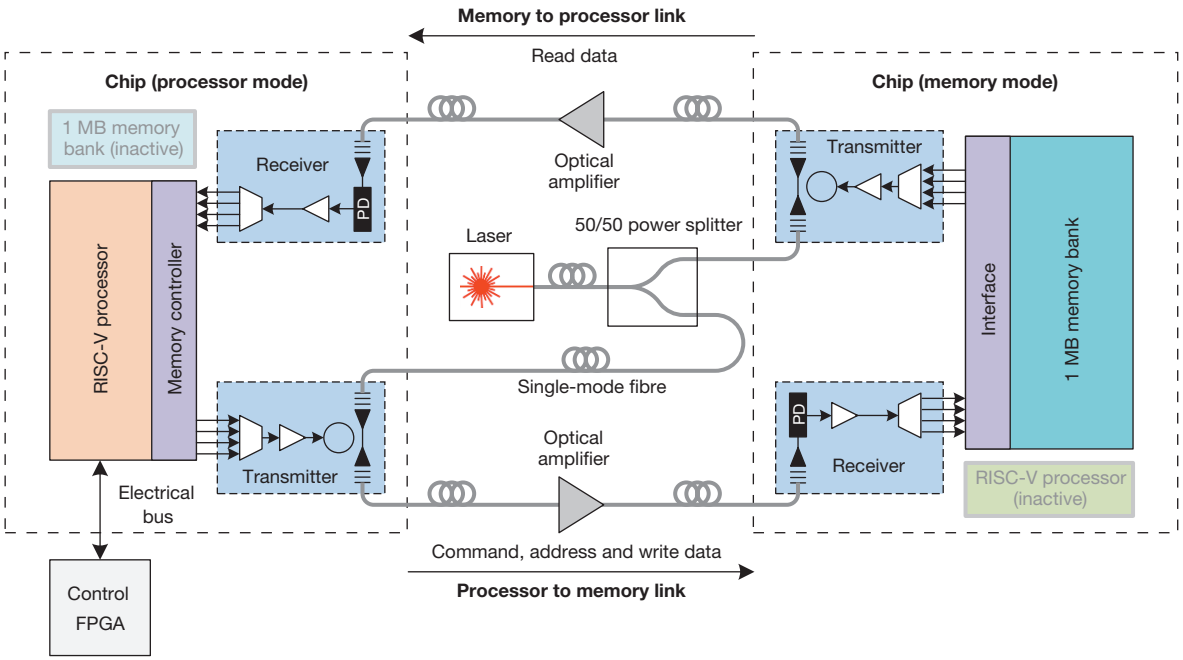

Figure 2 | Block diagram of the optical memory system. The system uses one chip acting as the processor and the other acting as memory, connected by a full-duplex optical link with a round-trip distance of 20 m by fibre. PD, photodetector.

first performs direct memory access through the memory controller to write all of the program's instructions into memory. Once the program is fully loaded, the FPGA issues a 'reset' signal to the processor and the processor begins execution of the program by fetching the first program instruction from memory (from address 0x00002000). During program execution, the processor writes and reads program data to and from memory, in addition to reading the instructions from the memory. The control FPGA handles the printing of terminal outputs and acts as a display driver that reads from the frame buffer residing in memory to display a screen to the user. In all cases, the P → M and M → P optical links handle all communications to and from memory (which holds all the program instructions and data).

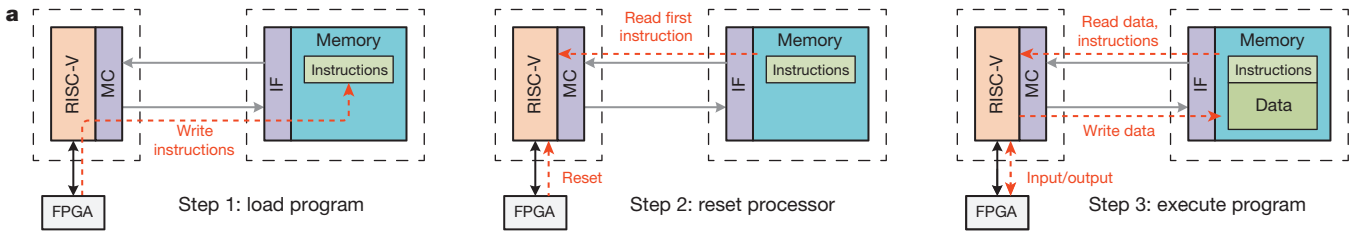

b

```
------------------------------------------------------------
>> ./fesvr-zedboard-head.1MB +divisor=1 +hold=1 ./pk -p ./hello.photonics
------------------------------------------------------------
CPU reset complete
uncore slowio divisor=1, hold=1
host_clk frequency = 15.64 MHz
cpu_clk frequency  = 31.27 MHz
hello world with photonics!
------------------------------------------------------------
>> ./fesvr-zedboard-head.1MB +divisor=1 +hold=1 ./pk -p ./stream.256KB
------------------------------------------------------------
CPU reset complete
uncore slowio divisor=1, hold=1
host_clk frequency = 15.64 MHz
cpu_clk frequency  = 31.27 MHz
-------------------------------------------------------------
STREAM version $Revision: 5.10 $
-------------------------------------------------------------
This system uses 8 bytes per array element.
-------------------------------------------------------------
Array size = 32768 (elements), Offset = 0 (elements)
Memory per array = 0.2 MiB (= 0.0 GiB).
Total memory required = 0.8 MiB (= 0.0 GiB).
Each kernel will be executed 10 times.
 The *best* time for each kernel (excluding the first iteration)
 will be used to compute the reported bandwidth.
-------------------------------------------------------------
Your clock granularity/precision appears to be 1 microseconds.
Each test below will take on the order of 799 microseconds.
   (= 799 clock ticks)
Increase the size of the arrays if this shows that
you are not getting at least 20 clock ticks per test.
-------------------------------------------------------------
WARNING -- The above is only a rough guideline.
For best results, please be sure you know the
precision of your system timer.
-------------------------------------------------------------
Function    Best Rate MB/s  Avg time     Min time     Max time
Copy:           640.3     0.000823     0.000819     0.000831
Scale:          551.3     0.000954     0.000951     0.000964
Add:            584.8     0.001350     0.001345     0.001364
Triad:          585.9     0.001351     0.001342     0.001364
-------------------------------------------------------------
Solution Validates: avg error less than 1.000000e-13 on all three arrays
-------------------------------------------------------------
```

c

Figure 3 | Processor optical demonstration. a, Program loading and execution. MC is the memory controller in the processor; IF is the memory interface of the memory bank. **b**, Successful execution of the 'Hello world!' basic functionality test and the STREAM[28] memory benchmark, two examples of terminal-based programs. **c**, Screen capture of the output of a three-dimensional teapot-rendering application running on the processor.

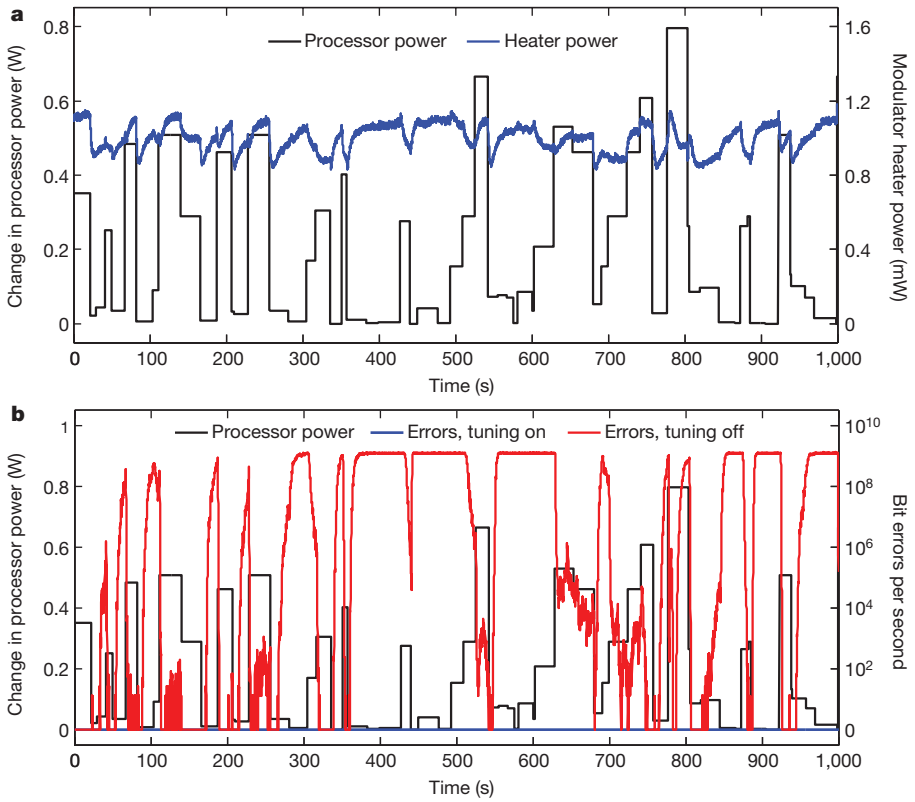

Figure 4 | Thermal-tuning stress test of the P → M link. a, Modulator heater output power with tuning switched 'on', overlaid on the power trace for the processor. The thermal-tuning controller changes the heater power output to adapt to the changes in temperature created by the changes in processor power. **b**, Measured bit errors per second versus time with the thermal-tuning controller switched 'on' and 'off', overlaid on the power trace for the processor. The link with the tuning controller 'on' has no bit errors over the entire interval (a total of 2.5 Tb transmitted and received).

The processor clock frequency is locked to 1/80th of the aggregate bit rate of the P → M link (corresponding to a clock frequency of 31.25 MHz at 2.5 Gb s^{-1}) when demonstrating the processor using the optical link, the result of a decision that simplified engineering efforts during chip design. When operating in non-optical mode—by electrically communicating to the 1 MB bank of memory local to the same chip, or memory connected to the control FPGA by time-multiplexing memory data over the control interface—the processor can run at a maximum speed of 1.65 GHz. A demonstration of the system running these programs is provided in Supplementary Video 1.

To evaluate the robustness of the optical links and ring tuning control against thermal perturbations, we create a synthetic processor power trace by changing the voltage and frequency operating points of the processor (Fig. 4) over a 1,000-s period. The changes in processor power are representative of the behaviour of a processor as it runs different loads, affecting the chip temperature. The difference in temperature between the highest and lowest temperatures (processor at maximum and minimum power, respectively) is approximately 8 K. The thermal-tuning circuitry controls the output of the microheater integrated with the ring modulator to keep the resonant device locked to the laser wavelength, which keeps the link free of bit errors despite changes in temperature produced by the processor. With the tuning circuitry disabled, the same link experiences a number of bit errors depending on the processor power draw. The effect of thermal perturbations on the system during the execution of a program is shown in Supplementary Video 1.

Our demonstration of an electronic–photonic microprocessor chip could enable advances in very-large-scale integrated circuit (VLSI) technology, by adding nanophotonics as a new design dimension. Tailoring photonic devices to be integrated directly with electronics in an advanced-node CMOS process enabled a fully functioning electronic–photonic system on a single chip to be produced in a high-volume electronics foundry. The level of integration allowed on-chip thermal-tuning control systems to guarantee robust operation of compact and energy-efficient, but also thermally sensitive, optical resonator devices, addressing one of the key remaining challenges for nanophotonic circuits adoption in VLSI technology.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 17 August; accepted 10 November 2015.

1. Goodman, J. W., Leonberger, F. J., Kung, S.-Y. & Athale, R. A. Optical interconnections for VLSI systems. *Proc. IEEE* **72,** 850–866 (1984).
2. Miller, D. A. Rationale and challenges for optical interconnects to electronic chips. *Proc. IEEE* **88,** 728–749 (2000).
3. Young, I. *et al.* Optical I/O technology for tera-scale computing. *IEEE J. Solid-State Circ.* **45,** 235–248 (2010).
4. Vantrease, D. *et al.* Corona: system implications of emerging nanophotonic technology. In *Proc. 35th Annual International Symposium on Computer Architecture (ISCA '08)* 153–164, http://dx.doi.org/10.1109/ISCA.2008.35 (IEEE Computer Society, 2008).
5. Shacham, A., Bergman, K. & Carloni, L. P. Photonic networks-on-chip for future generations of chip multiprocessors. *IEEE Trans. Comput.* **57,** 1246–1260 (2008).
6. Batten, C. *et al.* Building manycore processor-to-DRAM networks with monolithic CMOS silicon photonics. *IEEE Micro* **29(4),** 8–21 (2009).
7. Beamer, S. *et al.* Re-architecting DRAM memory systems with monolithically integrated silicon photonics. In *Proc. 37th Annual International Symposium on Computer Architecture (ISCA '10)* 129–140, http://doi.acm.org/10.1145/1815961.1815978 (ACM, 2010).
8. Xu, Q., Schmidt, B., Pradhan, S. & Lipson, M. Micrometre-scale silicon electro-optic modulator. *Nature* **435,** 325–327 (2005).
9. 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, http://www.osapublishing.org/abstract.cfm?URI=OFC-2008-OMK7 (Optical Society of America, 2008).
10. Assefa, S. *et al.* CMOS integrated nanophotonics: enabling technology for exascale computing systems. In *Optical Fiber Communication Conference* OMM6, https://www.osapublishing.org/abstract.cfm?uri=OFC-2011-OMM6 (Optical Society of America, 2011).

11. Buckwalter, J., Zheng, X., Li, G., Raj, K. & Krishnamoorthy, A. A monolithic 25-Gb/s transceiver with photonic ring modulators and Ge detectors in a 130-nm CMOS SOI process. *IEEE J. Solid-State Circ.* **47,** 1309–1322 (2012).
12. Dupuis, N. *et al.* 30Gbps optical link utilizing heterogeneously integrated III-V/Si photonics and CMOS circuits. In *Optical Fiber Communications Conference,* Th5A.6, http://dx.doi.org/10.1364/OFC.2014.Th5A.6 (Optical Society of America, 2014).
13. Orcutt, J. S. *et al.* Open foundry platform for high-performance electronic-photonic integration. *Opt. Express* **20,** 12222–12232 (2012).
14. Takahashi, O. *et al.* Migration of Cell broadband engine from 65nm SOI to 45nm SOI. In *International Solid-State Circuits Conference (ISSCC 2008) Dig.Tech. Pap*. 86–597, http://dx.doi.org/10.1109/ISSCC.2008.4523069 (IEEE, 2008).
15. IBM Blue Gene team. Design of the IBM Blue Gene/Q Compute chip. *IBM J. Res. Dev.* **57,** 1–13 (2013).
16. Wendel, D. *et al.* The implementation of POWER7™: a highly parallel and scalable multi-core high-end server processor. In *International Solid-State Circuits Conference (ISSCC 2010) Dig. Tech. Pap*. 102–103, http://dx.doi.org/10.1109/ISSCC.2010.5434074 (IEEE, 2010).
17. Waterman, A., Lee, Y., Patterson, D. A. & Asanović, K. *The RISC-V Instruction Set Manual, Volume I: User-Level ISA, Version 2.0*. Technical Report No. UCB/EECS-2014-54, http://www.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-54.html (EECS Dept., Univ. California, Berkeley, 2014).
18. Narasimha, S. *et al.* High performance 45-nm SOI technology with enhanced strain, porous low-k BEOL, and immersion lithography. In *International Electron Devices Meeting (IEDM '06)* 1–4, http://dx.doi.org/10.1109/IEDM.2006.346879 (IEEE, 2006).
19. Alloatti, L., Srinivasan, S., Orcutt, J. & Ram, R. Waveguide-coupled detector in zero-change complementary metal–oxide–semiconductor. *Appl. Phys. Lett.* **107,** 041104 (2015).
20. Sun, C. *et al.* A 45nm SOI monolithic photonics chip-to-chip link with bit-statistics-based resonant microring thermal tuning. In *2015 Symposium on VLSI Circuits* C122–C123, http://dx.doi.org/10.1109/VLSIC.2015.7231348 (IEEE, 2015).
21. Shainline, J. M. *et al.* Depletion-mode carrier-plasma optical modulator in zero-change advanced CMOS. *Opt. Lett.* **38,** 2657–2659 (2013).
22. Wade, M. T. *et al.* Energy-efficient active photonics in a zero-change, state-of-the-art CMOS process. In *Optical Fiber Communication Conference* Tu2E.7, http://dx.doi.org/10.1364/OFC.2014.Tu2E.7 (Optical Society of America, 2014).
23. Soref, R. A. & Bennett, B. Electrooptical effects in silicon. *IEEE J. Quantum Elect.* **23,** 123–129 (1987).
24. Selvaraja, S. K., Bogaerts, W., Dumon, P., Van Thourhout, D. & Baets, R. Subnanometer linewidth uniformity in silicon nanophotonic waveguide devices using CMOS fabrication technology. *IEEE J. Sel. Top. Quantum Elect.* **16,** 316–324, (2010).
25. Padmaraju, K., Chan, J., Chen, L., Lipson, M. & Bergman, K. Thermal stabilization of a microring modulator using feedback control. *Opt. Express* **20,** 27999–28008 (2012).
26. Sun, C. *et al.* A monolithically-integrated chip-to-chip optical link in bulk CMOS. *IEEE J. Solid-State Circ.* **50,** 828–844 (2015).
27. Wade, M. *et al.* 75% efficient wide bandwidth grating couplers in a 45 nm microelectronics CMOS process. In *2015 Optical Interconnects Conference* 46–47, http://dx.doi.org/10.1109/OIC.2015.7115679 (IEEE, 2015).
28. McCalpin, J. D. STREAM: sustainable memory bandwidth in high performance computers. http://www.cs.virginia.edu/stream/ (1995).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank S. Twigg, Q. Nguyen, and M. Moreto Planas for help with processor infrastructure, A. Srinivasan for help with photodetector characterization, and S. Han for help with chip photos. This work was supported by DARPA POEM award HR0011-11-C-0100, led by J. Shah and DARPA PERFECT award HR0011-12-2-0016, led by J. Cross. We thank M. Casper, J. Torneden, and the team at the Kansas City Plant for their support of our design submissions over the years leading up to this work. Support is also acknowledged from the Berkeley Wireless Research Center, UC Berkeley ASPIRE Lab, MIT CICS, National Science Foundation, FCRP IFC, Trusted Foundry, Intel, Santec, and NSERC. 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 C.S. developed the thermal tuning circuitry, designed the memory bank, implemented the 'glue-logic' between various electronic components, and performed top-level assembly of electronics and photonics. M.T.W. optimized modulator designs for thermal tuning, designed the grating couplers, and performed top-level assembly of photonics regions used in our demonstration. C.S. and Y.L. designed the system-level architecture and demonstrated the processor with photonic input/output. Y.L. wrote and/or adapted the test programs for the processor demonstration. Y.L. and A.S.W. developed the RISC-V ISA and processor implementation. J.S.O. created the CAD infrastructure for photonic layouts, designed the photodetector used in our demonstration, and assembled initial photonic layouts and passive devices. L.A. improved the CAD infrastructure, developed new rules for design rule checking, and contributed new photodetector designs. C.S., M.T.W., Y.L., and L.A. contributed to chip verification and testing. M.S.G. designed and implemented the receiver circuit. J.M.S. designed, implemented, and tested the original version of the modulator. R.R.A. performed the physical implementation of the processor and designed the chip and adapter printed circuit boards. S.L. developed the selective substrate removal process and contributed to the thermal tuning method. B.R.M. assisted with chip implementation and performed initial substrate removal experiments. R.K. assisted in the rework of new grating coupler designs. F.P. contributed to layout and analysis for couplers and modulators. A.H.A. created new photodetector designs. H.M.C. and A.J.O. assisted with processor design. J.C.L. and Y.-H.C. contributed components in the transceiver regions. V.M.S., M.A.P., R.J.R., and K.A. supervised the project.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare competing financial interests: details are available in the online version of the paper. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to V.M.S. (vlada@berkeley.edu), M.A.P. (milos.popovic@colorado.edu), R.J.R. (rajeev@mit.edu) or K.A. (krste@berkeley.edu).

METHODS

Chip implementation. The key chip characteristics are summarized in Extended Data Table 1. Photonic devices were prepared in Cadence Virtuoso (an industry-standard design tool for frontend electronics) in conjunction with mixed-signal electronics[29]. Digital electronics were implemented using a combination of digital-synthesis and place-and-route tools from Synopsys and Cadence. All photonic and electronic designs conform to the CMOS manufacturing rules (more than 5,000 rules) of IBM's commercial 45-nm thin buried-oxide SOI process (12SOI), with physical verification performed using Mentor Graphics Calibre.

Chip fabrication. The chips were fabricated through the standard 12SOI process flow. We submitted our design for mask aggregation through the Trusted Access Program Office (TAPO) shuttle run, with the chip mask set treated as if it were an ordinary electronics design. The physical design dimensions, including the cross-sectional layer type and thickness information not reported here, are provided as part of the standard electronic design kit that is made available to IBM foundry customers under a non-disclosure agreement. A subset of process and performance information regarding this process can be found in various official IBM publications on electronic CMOS process development[18,30,31].

Electrical packaging. The chips from the foundry are bumped with controlled-collapse chip connection (C4) solder balls. The chips are then flip-chip mounted (the chip's substrate is exposed on top) to an 8-layer FR4 printed circuit board through C4 solder reflow. This forms all 249 electrical connections (including power and ground) from the chip to the printed circuit board. Epoxy encapsulation is added to the mounted chips for additional mechanical support and to protect the mounted chips. These steps are typical for an electrical chip package and were performed by CVInc.

Patterned substrate removal of a packaged chip. The electrically packaged samples are first backside-ground to thin the chip substrate down to 100–150 μm (performed by Aptek Industries). We then clean the backside surface with isopropyl alcohol and an N_2 air gun. We next apply Kapton tape over the substrate regions that we do not wish to remove (over the processor and the 1 MB memory bank). Afterwards, the chips are placed in a chamber that supplies XeF_2 gas to isotropically etch the silicon substrate, removing it as the volatile product SiF_4. We use a pulsed-etch technique, in which etch steps of 120 s were interleaved with 60-s periods during which we pump out the reaction products. The pressure used in the chamber is 3.4 Torr. Because electronics are unaffected by the substrate removal, the very coarse feature definition provided by tape and hand alignment is sufficient. On average, the substrate removal process takes 10–30 cycles (depending on the thickness after the backside grind) with a success yield of 80% (defined as having a working processor after substrate removal). We stop the etch when the substrate over the desired etch region has disappeared when inspected by eye. The steps above are easily implementable at wafer scale in high-volume manufacturing using standard photolithographic techniques[32], which can also improve uniformity and yield of the post-processing as well as the resolution and alignment of the etch regions.

Optical testing. The 1,183-nm laser is a quantum dot DFB (distributed feedback) laser available from QDLaser. We used lensed fibres available from Oz Optics with a spot size of 5 μm and a working distance of 26 μm to couple light into the VGCs through the chip backside (after substrate removal). The spot size is matched to the 5-μm mode-field diameter of the VGCs. We used 3-axis positioner stages (Thorlabs NanoMax) to position and align fibres over the grating couplers of the test sites. The shown demonstrations require a total of 3 fibres coupled to each chip. Minimum fibre-to-coupler insertion loss was achieved by angling the fibres at 19° off-normal from the surface of the chip. To adjust the polarization of the input light, we use 3-paddle manual polarization controllers from Thorlabs (although these can be avoided if using polarization-maintaining fibres). For this first demonstration, we chose the manual fibre alignment approach to freely couple into any of the hundreds of optical test sites located throughout the chip. To make a permanent fibre-attach, we could leverage commercial optical packaging techniques for VGCs, such as through horizontal fibre array blocks with angle-cleaved fibres[33] or through vertical fibre array pigtails[9,34].

Processor testing. The control FPGA is a Zedboard FPGA, providing an intermediate hardware interface between the electrical links of the processor and an ethernet connection to the laboratory control computer. The individual cores incorporate a 64-bit scalar core, floating-point unit, vector accelerator, and private caches[35]. Programs are compiled from C source code using a GCC-based C compiler targeted for the RISC-V ISA. The implementations of the RISC-V processor and the software compilation stack are available at http://bar.eecs.berkeley.edu/projects.html. Details of the RISC-V ISA standard are found at http://www.riscv.org. The full system is stable and can execute an arbitrary number of programs. A representative set of programs tested on this processor is as follows.

- 'Memory test'. The control FPGA writes to and reads from every location in memory through direct memory access to verify that the memory interface is fully functional and that all bits are correct. The processor is idled for this test.
- 'Hello world!' A program that asks the processor to print out a single line of text to the terminal, which is sent to the control FPGA to be displayed to the user.
- 'STREAM'. A popular memory benchmarking application[28]; the outputs of the program are printed to the terminal and displayed to the user.
- 'Teapot renderer'. A program that pixel shades a three-dimensional teapot using the Blinn–Phong shading model and outputs the rendered image. The location and colour of the light source illuminating the teapot in the rendered image is controlled by the user using the keyboard. The processor performs all calculations and writes the image to the frame buffer in memory using the optical links. It then reads the content of the frame buffer over the optical link and sends it to the control FPGA to display it as an image to the user.
- 'Linux'. A full Linux operating system. Once Linux boots, the user is free to run any program, including 'python', 'top', or file system operations (the file system behaviour is coordinated by the control FPGA). This test uses memory connected to the control FPGA and not the optically connected memory, because the memory footprint of the Linux kernel is too big to fit in the 1 MB memory bank.

The 1:80 ratio between processor clock frequency and the P → M link throughput was chosen to keep processor frequency reasonable if the links operated at higher data rates than anticipated at design time and when all wavelengths in the P → M link are active. For example, if the P → M link supported an 80 Gb s^{-1} aggregate data rate, then the processor needs to operate at 1.0 GHz, which is well within its abilities. Alternatively, if the ratio was 1:10, then the processor would need to operate at 8 GHz, which is impractical.

Transmitter and receiver circuit specifications. At the 2.5 Gb s^{-1} operating point used in our demonstration, the transmitter uses the 1-V digital supply, which corresponds to a transmitter energy of 20 fJ per bit and achieves an insertion loss of 3 dB at an on:off ratio of 6 dB for non-return-to-zero binary data. The modulator is effectively 'driverless' insofar as no analogue driver electronics are needed to bridge between digital logic and the optical modulator, owing to the efficiency of the latter. The thermal tuning for the modulator ring consumes a fixed 192 fJ per bit for the control circuit and 0–2.5 mW for the heater power, dependent on the tuned range (for the heater output power of 1.5 mW in Supplementary Video 1, this corresponds to 600 fJ per bit). More detailed transmitter and thermal tuner descriptions have been reported previously[20]. The receiver has a 10^{-12} bit-error-rate sensitivity (OMA) of −5 dBm up to 5 Gb s^{-1}, degrading to −3.8 dBm at 8 Gb s^{-1}, and −0.8 dBm at 10 Gb s^{-1}. At 2.5 Gb s^{-1}, the receiver energy efficiency is 496 fJ per bit, improving to 297 fJ per bit at 10 Gb s^{-1}. Summing up, we report a total circuit energy efficiency of 1.3 pJ per bit at 2.5 Gb s^{-1}—a power consumption of 3.25 mW. The bandwidth density of the transceivers is approximately 300 Gb s^{-1} mm^{-2} of chip area. The key specifications are summarized in Extended Data Table 2.

Link specifications. In the 2.5 Gb s^{-1} P → M and M → P links used in our demonstration, the transmitter input VGC, transmitter output VGC, and receiver input VGC contribute 4 dB, 4 dB, and 6 dB of link insertion loss, respectively. The 1,183-nm laser outputs 9.2 dBm such that 5.2 dBm (50/50 split, with an approximately 1-dB excess loss of the splitter) is incident upon each of the input transmit VGC of each link (P → M and M → P). At this laser power level, the OMA of each transmitter is −7 dBm, with an average optical power of −9 dBm. Each amplifier adds 9 dB of optical gain, completing the P → M and M → P links each with an extra 1-dB link margin. A chip iteration incorporating 1.2-dB-loss VGCs[27] into the P → M and M → P link would remove 10.4 dB of excess insertion loss. These devices were high-risk test structures on the current chip and so were not placed in the P → M and M → P transceivers. Using these couplers at the same input laser power level as before, both links could complete, without an amplifier, with an extra 2.4 dB of link margin. The 1,183-nm laser was made by QDLaser and uses 55 mA of pump current at a laser diode bias of 1.3 V to output 9.2 dBm (8.3 mW) of power. This corresponds to a power use of 71.5 mW and a wall-plug efficiency of 11.6%. The laser is shared across both P → M and M → P links and the total wall-plug energy efficiency (laser and circuit) is 15.6 pJ per bit. The laser has a threshold of 29 mA and a slope efficiency of 0.32 mW mA^{-1}.

Potential for improved performance. The chip demonstrated here is a first working research prototype, and the current achieved performance is by no means representative of the absolute performance limits of this technology. We describe a few known ways to improve performance in the following.

(1) The current modulator design uses a mid-level p-implant (10^{17}–10^{18} cm^{-3}) for p-contacts as opposed to a p+ implant, creating high series contact resistance

that limits its bandwidth. Future design iterations will use the p+ implant to improve device bandwidth. Moreover, the modulators used only two out of several different doping implants available in the process for different transistors and transistor thresholds. Substantial improvements may be possible with other available implants.

(2) The current detector is absorption-length limited[19] and resonating the detector can improve sensitivity without an increase in the device size. Resonant detectors, implemented as a spoked-ring cavity in a manner similar to that of the modulator, exist on the same chip as standalone devices in the independent-device and transceiver regions. If incorporated with processor and memory transceivers in a future chip, they would improve sensitivity by approximately 6 dB (to an OMA sensitivity of −11 dBm), which would be competitive with state-of-the-art integrated receivers. In addition, the current receiver circuit design is very conservative and could be optimized to further improve the sensitivity by 6 dB. The circuit could also be placed closer to the photo-detector to minimize wiring capacitance.

(3) Our demonstration uses the laser at a power level far below that for peak efficiency, which is 16% at 30 mW. Operation of the current laser at the peak-efficiency power level and sharing of the output power across multiple links on the chip or, alternatively, usage of a laser optimized for the given output power are techniques for improving the energy efficiency of the link, even without any device improvements.

Applicability to CMOS processes with bulk silicon substrates. CMOS processes using a bulk silicon substrate lack a patternable crystalline silicon layer, which motivates the use of alternative devices in polycrystalline silicon and a small number of process changes[36]. However, some guiding principles of zero-change integration, such as reuse of existing transistor mask levels, repurposing of transistor materials for optics, and compact integration using silicon micro-rings, can be applied to minimize changes to the process frontend, which are the most harmful to process-native electronics. These concepts have been applied successfully in practice to enable functional photonics in bulk[26,36], although at a far smaller scale than is demonstrated here.

Code availability. The source code for the processor is available at http://bar.eecs.berkeley.edu/projects.html. Other test applications (such as STREAM) can be found at the respective cited references.

29. Orcutt, J. S. & Ram, R. J. Photonic device layout within the foundry CMOS design environment. *IEEE Photonic. Tech. Lett.* **22,** 544–546 (2010).
30. Kalla, R., Sinharoy, B., Starke, W. J. & Floyd, M. Power7: IBM's next-generation server processor. *IEEE Micro* **30,** 7–15 (2010).
31. Lee, S. *et al.* Record RF performance of 45-nm SOI CMOS technology. In *International Electron Devices Meeting (IEDM 2007)* 255–258, http://dx.doi.org/10.1109/IEDM.2007.4418916 (IEEE, 2007).
32. Roger, A. Breaking a new sound barrier: it's a mic-on-a-chip. *Electron. Des.* **54,** http://electronicdesign.com/analog/breaking-new-sound-barrier-its-mic-chip (2006).
33. Pavarelli, N., Lee, J. S. & O'Brien, P. A. Packaging challenges for integrated silicon photonic circuits. In *Proc. SPIE* Vol. 9133 (eds Vivien, L. *et al.*) 91330F, http://dx.doi.org/10.1117/12.2058559 (SPIE, 2014).
34. Kopp, C. *et al.* Silicon photonic circuits: on-CMOS integration, fiber optical coupling, and packaging. *IEEE J. Sel. Top. Quant. Elect.* **17,** 498–509 (2011).
35. Lee, Y. *et al.* A 45nm 1.3GHz 16.7 double-precision GFLOPS/W RISC-V processor with vector accelerators. In *40th European Solid State Circuits Conference (ESSCIRC 2014)* 199–202, http://dx.doi.org/10.1109/ESSCIRC.2014.6942056 (IEEE, 2014).
36. Meade, R. *et al.* Integration of silicon photonics in bulk CMOS. In *2014 Symposium on VLSI Technology Dig. Tech. Pap.* 228–229, http://dx.doi.org/10.1109/VLSIT.2014.6894427 (IEEE, 2014).

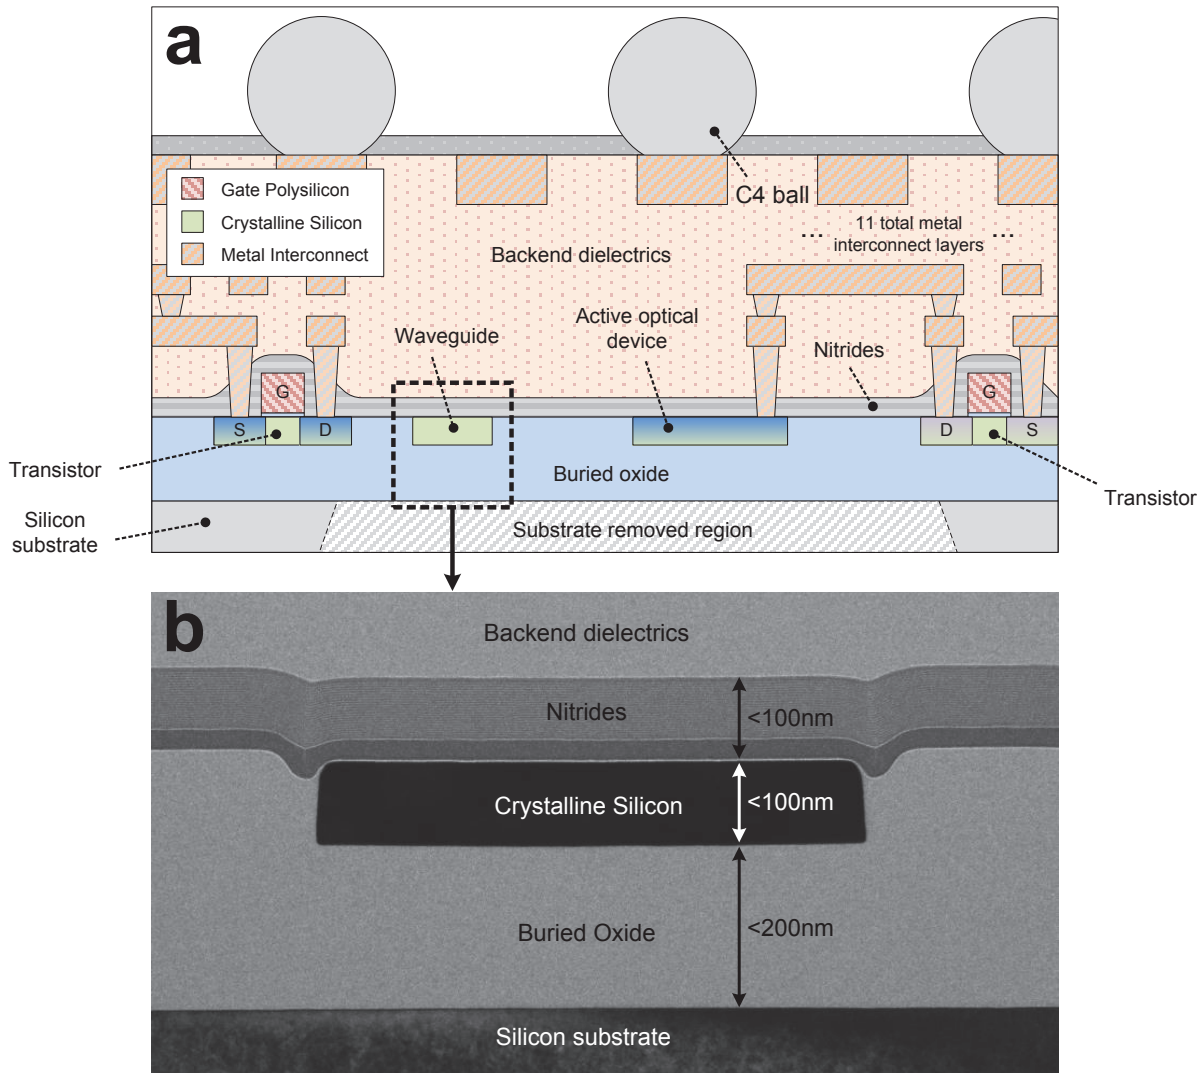

Extended Data Figure 1 | Chip cross-section. a, Full chip cross-section (not to scale) from the silicon substrate to the C4 solder balls, showing the structures of electrical transistors, waveguides, and contacted optical devices. G, S, and D mark the structures that form the gate, source, and drain, respectively, of an electrical transistor. The minimum separation between transistors and waveguides is $<1\,\mu m$, which is set only by the distance at which evanescent light from the waveguide begins to interact with the structures of the transistor. **b**, Transmission electron microscopy cross-section micrograph of an optical waveguide, before substrate removal.

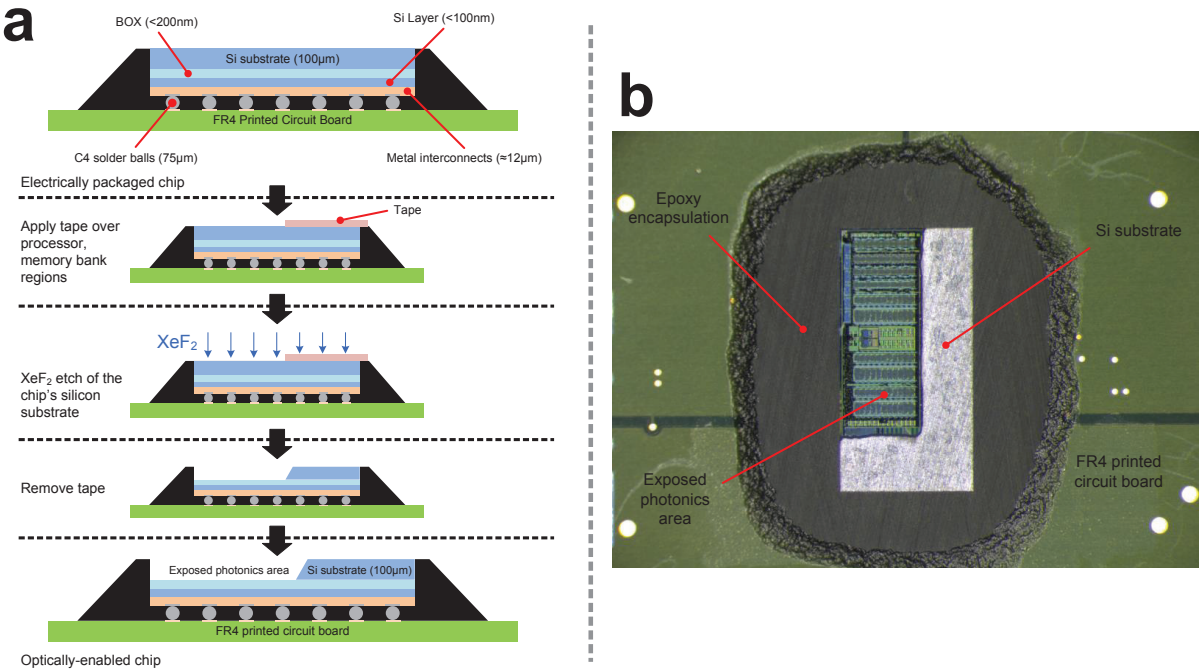

Extended Data Figure 2 | Selective substrate removal. a, Selective substrate removal steps for the flip-chip packaged chip, using tape as a coarse mask for defining areas that retain the substrate. BOX, buried oxide. **b**, Photo of a selective-substrate-removed fully electrically packaged electronic–photonic processor chip.

Extended Data Table 1 | Summary of chip characteristics

Characteristic	Value
Number of Transistors	70 Million
in Processor/Memory	60 Million
in P2M/M2P Transceivers	≈4 Million
in Standalone Transceivers	≈5 Million
Number of VGCs, Rings, PDs	851
in P2M/M2P Transceivers	324
in Standalone Sites	527
Processor Cores	2
Max Processor Frequency	1.65 GHz
L1 Instruction Cache	2× 16 KB
L1 Data Cache	2× 32 KB
L1 Vector Instruction Cache	2× 8 KB
Memory Bank	1 MB
Number of P2M/M2P Transceiver Banks	4
Transmitters/Receivers per Bank	11
Max Wavelengths	11
Theoretical throughput if all transceivers on the chip were active	550 Gb/s Tx 900 Gb/s Rx

Summary of the physical characteristics of the chip, such as the total number of electrical and optical devices, the processor parameters, and the configurations of the photonic transceiver banks connected to the processor and memory. L1 is the level 1 processor cache, Tx refers to transmit, and Rx refers to receive.

Extended Data Table 2 | Summary of transceiver performance

Property	P2M/M2P Transceivers	Standalone Transceivers
Waveguide Loss	4.3 dB/cm	4.3 dB/cm
Grating Coupler Loss	4 dB and 6 dB	1.2 dB
Tx Data Rate	2.5 Gb/s	5 Gb/s
Tx Extinction Ratio	6 dB	>6 dB
Tx Insertion Loss	3 dB	3 dB
Tx Power	0.02 pJ/bit	0.03 pJ/bit
Rx Data Rate	2.5 Gb/s	10 Gb/s
PD Responsivity	0.023 A/W	0.10 A/W
Rx OMA Sensitivity	–5 dBm@2.5 Gb/s	–7.2 dBm@10 Gb/s*
Rx Power	0.50 pJ/bit	0.30 pJ/bit
Ring Tuning Range	≈3.0 nm	≈3.0 nm
Ring Heater Tuning Efficiency	1.25 nm/mW	1.25 nm/mW
Ring Tuning Control Power	0.19 pJ/bit	0.14 pJ/bit

*Estimated OMA sensitivity using the 0.10 A/W PD.

Summary of the performance metrics of the P→M and M→P transceivers, and of the transceivers that exist in the standalone independent devices region. PD, photodetector.