Massively scalable Kerr comb-driven silicon photonic link

Anthony Rizzo [1,2,6], Asher Novick [1,6], Vignesh Gopal[1,6], Bok Young Kim [3], Xingchen Ji [1,4], Stuart Daudlin[1], Yoshitomo Okawachi [3], Qixiang Cheng [1,5], Michal Lipson [1,3], Alexander L. Gaeta[1,3] & Keren Bergman [1] ✉

The growth of computing needs for artificial intelligence and machine learning is critically challenging data communications in today's data-centre systems. Data movement, dominated by energy costs and limited 'chip-escape' bandwidth densities, is perhaps the singular factor determining the scalability of future systems. Using light to send information between compute nodes in such systems can dramatically increase the available bandwidth while simultaneously decreasing energy consumption. Through wavelength-division multiplexing with chip-based microresonator Kerr frequency combs, independent information channels can be encoded onto many distinct colours of light in the same optical fibre for massively parallel data transmission with low energy. Although previous high-bandwidth demonstrations have relied on benchtop equipment for filtering and modulating Kerr comb wavelength channels, data-centre interconnects require a compact on-chip form factor for these operations. Here we demonstrate a massively scalable chip-based silicon photonic data link using a Kerr comb source enabled by a new link architecture and experimentally show aggregate single-fibre data transmission of 512 Gb s^{-1} across 32 independent wavelength channels. The demonstrated architecture is fundamentally scalable to hundreds of wavelength channels, enabling massively parallel terabit-scale optical interconnects for future green hyperscale data centres.

With the rise of cloud-based computing, computational workloads have largely been offloaded from local machines onto the server racks of hyperscale data centres and high-performance computers. Bandwidth-hungry applications such as artificial intelligence and machine learning severely strain the current interconnects within these systems, threatening to halt further performance scaling without substantial changes to the physical hardware used to pass data between nodes[1,2]. Furthermore, the energy consumption of such data centres has become environmentally significant[3] and will be dominated by interconnect energy in the aforementioned communication-intensive workloads. Optical interconnects based on silicon photonics have been widely recognized as a promising avenue for interconnects to keep pace with these ever-growing bandwidth demands while additionally decreasing energy consumption compared to their electrical counterparts. This promise is largely due to the compact footprint of silicon-on-insulator (SOI) devices, the compatibility of SOI photonics processes with the ubiquitous complementary metal–oxide–semiconductor (CMOS) infrastructure used to fabricate microelectronic chips,

[1]Department of Electrical Engineering, Columbia University, New York, NY, USA. [2]Information Directorate, Air Force Research Laboratory, Rome, NY, USA. [3]Department of Applied Physics and Applied Mathematics, Columbia University, New York, NY, USA. [4]Present address: John Hopcroft Center for Computer Science, Shanghai Jiao Tong University, Shanghai, China. [5]Present address: Engineering Department, University of Cambridge, Cambridge, UK. [6]These authors contributed equally: Anthony Rizzo, Asher Novick, Vignesh Gopal. ✉e-mail: bergman@ee.columbia.edu

and the tremendous inherent parallelism provided by optics through dense wavelength-division multiplexing (DWDM)[4,5]. In particular, DWDM enables independent information channels to be encoded on different colours of light simultaneously in a single optical waveguide or fibre, in stark contrast to copper wire-based links, which only permit a single channel per physical connection.

The use of frequency combs for DWDM optical interconnects is an appealing prospect due to their ability to replace the currently used large arrays of lasers with a single multi-wavelength source[6,7]. For massively parallel wavelength scaling, microresonator-based Kerr frequency combs in the silicon nitride (Si_3N_4) platform show particular promise due to their compact size, CMOS compatibility and ability to generate hundreds of evenly spaced low-noise wavelength channels from a single continuous-wave (c.w.) laser source[8,9]. A recent demonstration showed nanosecond-scale optical circuit switching using a Si_3N_4 Kerr frequency comb source with chip-based arrayed waveguide gratings (AWGs) and semiconductor optical amplifiers (SOAs)[10], but a main future prospect for such sources lies in their application to massively parallel DWDM. Chip-based frequency combs have been widely demonstrated for long-reach optical communications using benchtop telecommunications equipment for filtering, modulating and receiving wavelength channels[11–14]. However, DWDM has drawn a great deal of interest for data-centre interconnects[15–17], which place stringent requirements on the energy consumption and footprint of transceivers.

Notable demonstrations with semiconductor mode-locked lasers[17,18] and distributed feedback laser (DFB) arrays[19] have shown high-bandwidth data communication with compact silicon photonic transmitter chips, but used modest single-fibre wavelength channel counts (five, eight, and eight, respectively) and lack the inherent scalability provided by Kerr combs. In particular, semiconductor mode-locked lasers are fundamentally restricted in wavelength scaling by the gain bandwidth of the active region material, and DFB arrays suffer from non-uniform channel spacing due to fabrication variations between the discrete laser cavities. Optical frequency combs provide substantial benefits when compared to other discrete parallel solutions such as vertical-cavity surface-emitting laser arrays[20], as they possess precise and stable channel spacing while naturally generating all DWDM channels on a single fibre without the need for an external multiplexer as is required in the case of vertical-cavity surface-emitting laser and DFB arrays. Previously, the only integrated Kerr comb data transmission with a chip-based transmitter used a monolithic lithium niobate chip to demonstrate comb generation, filtering and modulation on a single die, but this was restricted to a single filtered channel at megabit per second data rates due to the noisy modulational instability state of the comb[21].

In this Article, we demonstrate an integrated Kerr comb-driven silicon photonic data-communication link using a massively scalable link architecture. We design a proof-of-principle transmitter and receiver for 32 wavelength channels and observe open-eye diagrams for all channels up to 16 Gb s⁻¹, yielding an aggregate single-fibre bandwidth of 512 Gb s⁻¹. Furthermore, we confirm that the generated comb lines have a negligible power penalty compared to a tunable c.w. laser source at the same wavelength, demonstrating that each tone behaves identically to an independent c.w. carrier from an array of lasers. This Article presents a promising and realistic direction for data-centre interconnects to scale to hundreds of wavelength channels, enabling future multi-terabit-per-second chip-to-chip links operating at energies below 1 pJ bit⁻¹.

Results

Although past demonstrations have focused on coherent communications using integrated Kerr frequency combs as long-haul DWDM sources, high-performance computer and data-centre interconnects require low latency and low energy per bit, making intensity-modulated direct-detection (IM-DD) solutions more appealing than coherent solutions due to the absence of energy-intensive and latency-inducing digital signal processing (DSP). In particular, IM-DD links with native error-free signalling do not require forward error correction (FEC) and thus have no encoding overhead, easing requirements on the electronics to reconstruct the original bit stream. Additionally, coherent communication links require a local oscillator at the receiver, necessitating another light source, which consumes additional energy.

To emphasize these realistic conditions for future wide-parallel, short-reach links, we focus our analysis and experiments on modest single-channel data rates (10–16 Gb s⁻¹ λ^{-1}) using a standard non-return-to-zero on–off-keying (NRZ-OOK) modulation format. Our link analysis indicates that these modest per-channel data rates with massively parallel scaling in the number of wavelength channels per fibre lead to higher energy efficiencies while still maintaining terabit s⁻¹ per fibre aggregate bandwidths (Supplementary Notes 1 and 2). Interconnects with high bandwidth, low energy consumption and low latency are critical for future disaggregated data-centre architectures, which are uniquely enabled by the combination of these three properties[1] (Fig. 1).

Standard single-bus cascaded microresonator-based links enable DWDM by assigning each resonant modulator/filter to a particular wavelength and have sparked widespread academic and commercial interest[15,22,23]. By cascading resonators of slightly different radii in even steps along a single bus waveguide, different resonances can be used to selectively encode data onto different wavelengths with uniform channel spacing. Furthermore, the high field confinement due to the large index contrast between silicon and silicon dioxide (SiO_2) in the SOI platform permits micrometre-scale bend radii with low radiation loss, enabling resonators with up to 100× smaller footprint than comparable non-resonant devices[24]. Recent demonstrations have shown systems with 10 (ref. 25) and 20 (ref. 26) cascaded resonant modulators on a single bus, although both demonstrations were source-limited to one and four wavelength channels, respectively. Reference [26] demonstrates four simultaneously active channels at 50 Gbaud λ^{-1} with pulse amplitude modulation 4-level, but generates the four-channel wavelength source using an off-chip phase modulator at 40 GHz and a wavelength-selective switch to equalize the comb spectrum and drop every other line, achieving a final comb spacing of 80 GHz. Although these demonstrations show the potential of cascaded ring architectures to scale to high channel counts, they do not present a path to simultaneously scale the multi-wavelength source to accommodate the increase in WDM channels. As mentioned previously, chip-based Kerr frequency combs provide a compact and mass-producible platform for generating many WDM channels and thus represent a natural avenue for scaling the source capacity in cascaded resonator-based links (Fig. 2a).

The source used in this Article consists of a Si_3N_4 Kerr comb operating in the normal group velocity dispersion (GVD) regime, which demonstrates higher conversion efficiency, power per line and spectral flatness compared to soliton Kerr combs in the anomalous GVD regime[13,27–31], making normal-GVD combs better suited for data-communication applications. To enable normal-GVD comb generation, the design employs coupled Si_3N_4 ring resonators with independent integrated microheaters for tuning the coupling conditions[27,32] (Methods). The fibre-coupled output spectrum of the comb is shown in Fig. 2c, demonstrating 25 lines above −6 dBm and 33 lines above −10 dBm. We measured the output edge-coupler loss on the Si_3N_4 chip to be 3 dB, implying 25 lines with on-chip power above −3 dBm.

To evaluate the performance of the single-bus architecture with a Kerr comb source, we fabricated a transmitter chip with 20 cascaded microdisk modulators through the American Institute for Manufacturing (AIM) multi-project wafer (MPW) service (Methods and Extended Data Fig. 1a). We used a lensed fibre to couple the comb output to the transmitter chip input, which displayed a coupling loss of 5 dB per
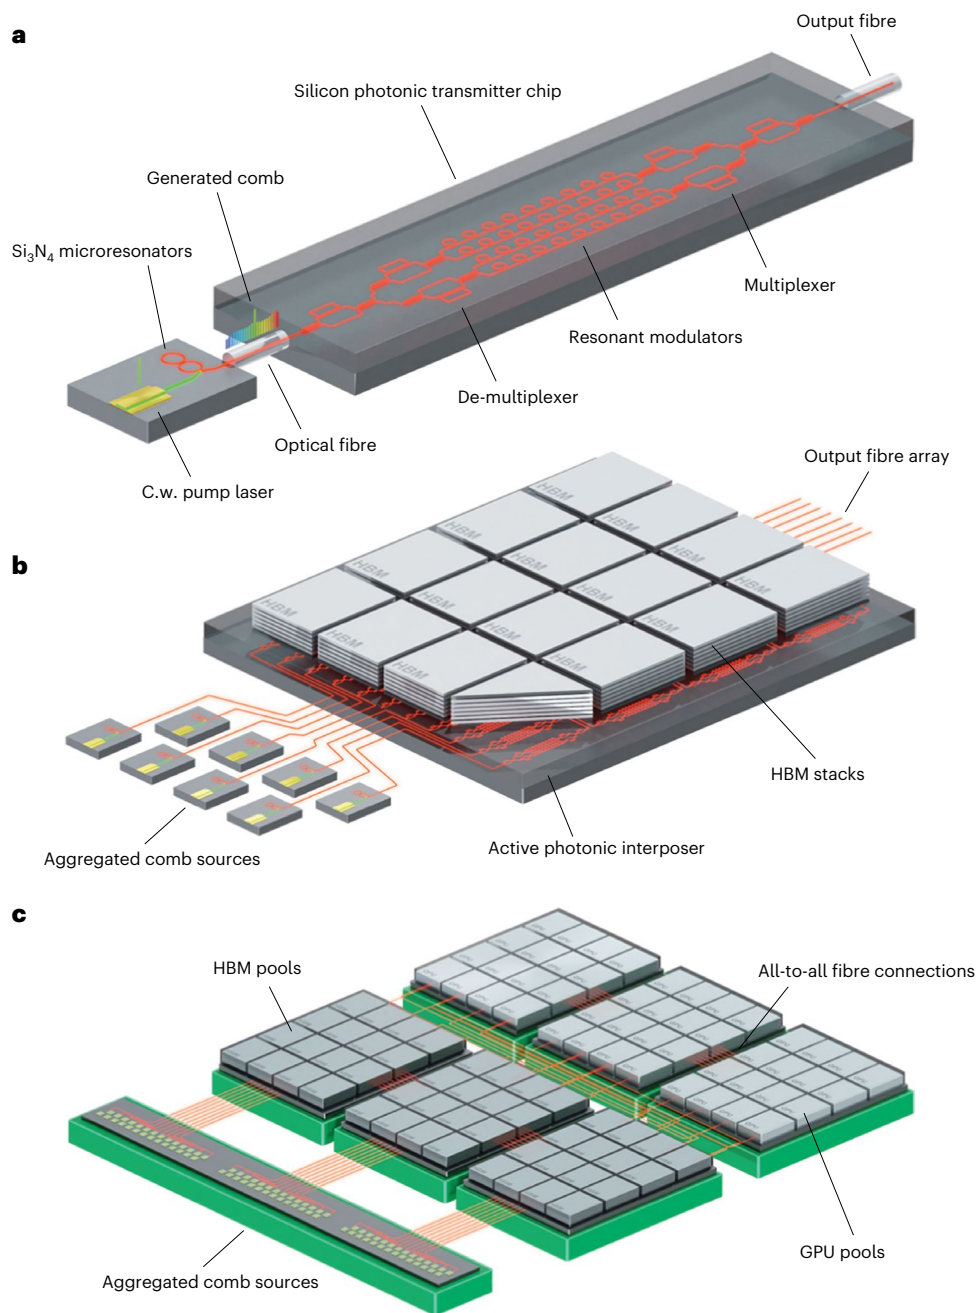
Fig. 1 | Artistic vision of a disaggregated data centre based on Kerr frequency comb-driven silicon photonic links. Hierarchical view of low-energy terabit-scale links connecting high-bandwidth memory (HBM) stacks to graphics processing units (GPUs) in a future data-centre system. The large number of distinct wavelength channels provided by integrated microresonator-based Kerr combs enable massively parallel data transmission on a single optical fibre through wavelength-division multiplexing. Resources are packaged on active photonic interposers and connected via optical fibre, exploiting the distance-agnostic nature of optical communications to pool and connect spatially distanced resources with comparable latency and energy consumption to in-package electrical signalling. **a**–**c**, This article experimentally demonstrates the arrangement in **a**, and those in **b** and **c** represent future extensions for directly integrating the photonic layer with CMOS electronics[36].

facet and thus required amplification before and after the transmitter. Although resonant modulators exhibit self-heating at modest optical input powers due to free carrier absorption, we observe that this effect does not degrade the signal integrity up to 7 dBm incident power and can be fully compensated using the device's integrated heater at timescales that are orders of magnitude below the data rate (Supplementary Note 3). We sampled multiple comb lines over 44.3 nm of spectral bandwidth in the C and L bands, with all eight lines achieving directly measured bit error rates (BERs) better than 10^{-9} for 10 Gb s^{-1} λ^{-1} and 16 Gb s^{-1} λ^{-1} without the use of FEC (Fig. 2c,d).

Despite supporting modest DWDM channel counts, single-bus cascaded resonator links run into fundamental scaling restrictions due to the limited free spectral range (FSR) of the resonators, constraining the usable bandwidth to a single FSR. To pack more channels into a single FSR, the channel spacing must be reduced. However, channel spacings below 100 GHz lead to severe cross-talk penalties[33] and thus place an upper bound on the number of allowable channels within a given FSR. To mitigate this restriction, we propose a scalable link architecture that subdivides the comb to reduce the number of channels per bus and permits channel
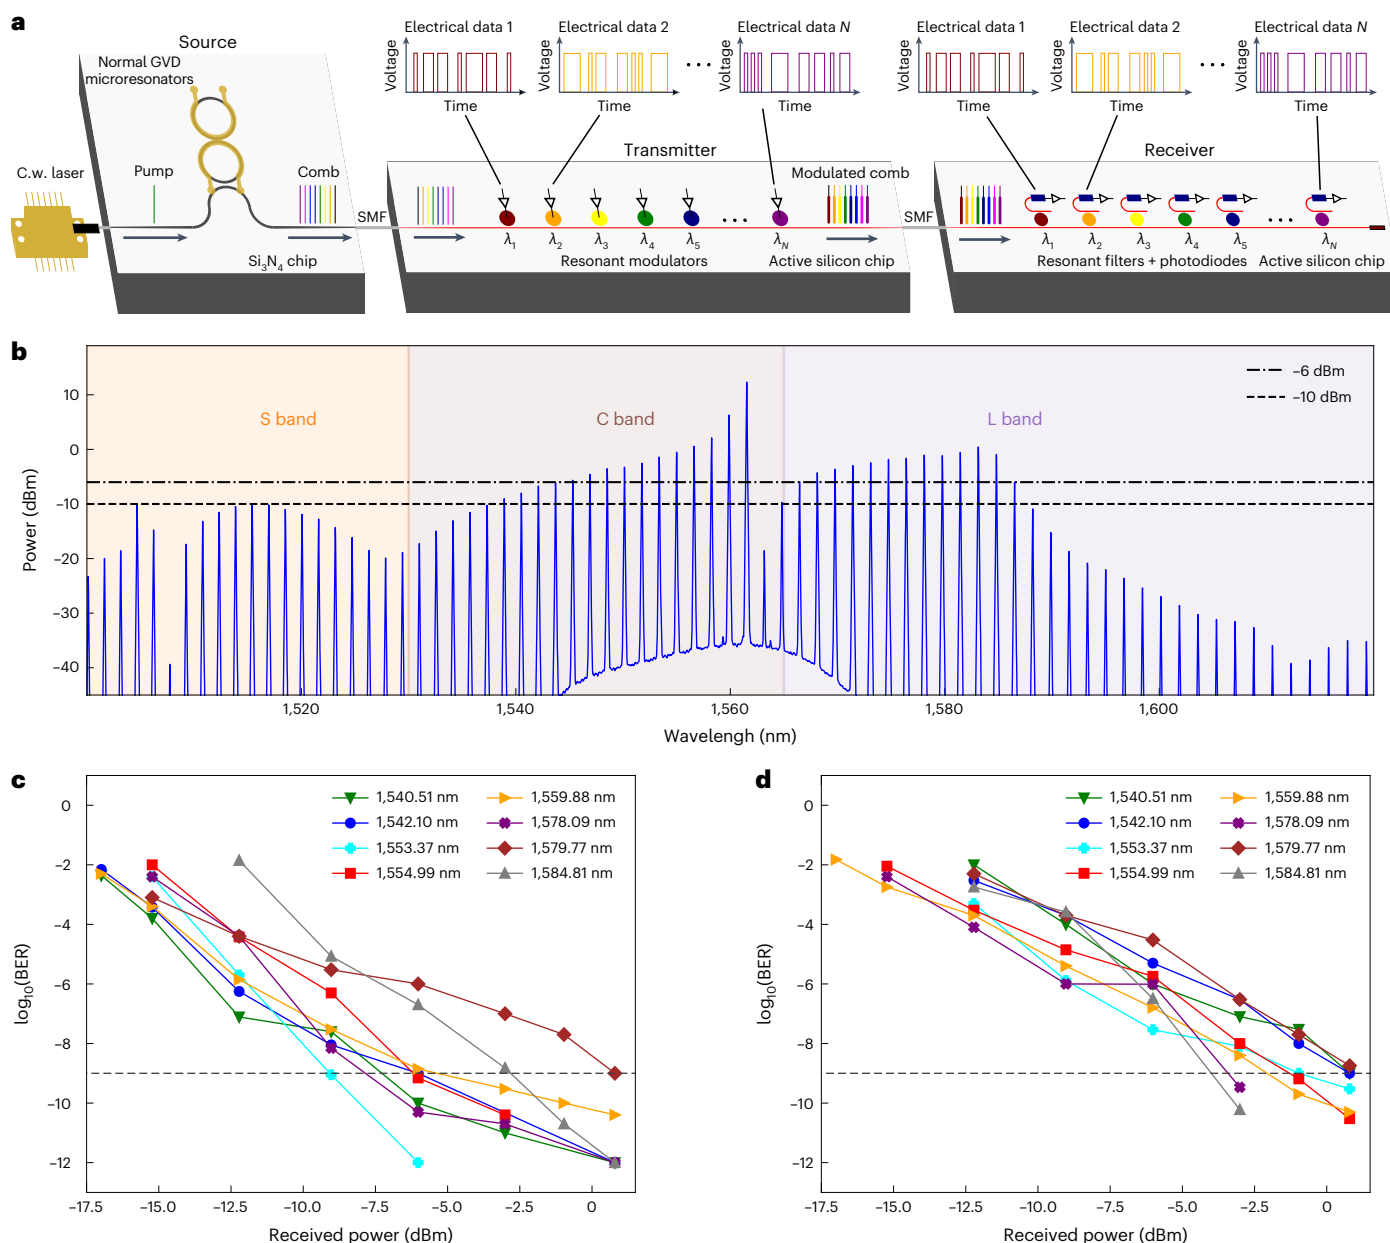
Fig. 2 | Single bus architecture and results. a, Schematic of a Kerr frequency comb-driven single bus cascaded resonator-based link. **b**, Measured normal-GVD comb spectrum with dash-dotted and dashed lines denoting −6 dBm and −10 dBm thresholds, respectively. **c,d**, BER waterfall curves for eight measured comb lines at 10 Gb s⁻¹ (**c**) and 16 Gb s⁻¹ (**d**). All lines achieve a directly measured BER better than 10⁻⁹ for both data rates without FEC.

arrangements that are not restricted to a single resonator FSR (Supplementary Note 4).

The proposed scalable link architecture is shown in Fig. 3a. By first subdividing the comb using trees of standard asymmetric Mach–Zehnder interferometers (MZIs), many of the benefits of the cascaded resonator-based architecture are preserved while providing a greater degree of modularity in channel count and channel spacing (Methods).

To verify the proposed architecture experimentally, we fabricated a 4.15 mm × 1.1 mm photonic transmitter chip and 3.1 mm × 1.1 mm receiver chip designed for 32 wavelength channels in a commercial 300-mm foundry through the AIM MPW service. We implemented (de-)interleavers as asymmetric ring-assisted MZIs (RMZIs), which exhibit flatter pass and stop bands compared to standard asymmetric MZI filters[34,35]. For the transmitter, two stages of de-interleavers subdivide the comb into four groups, which are incident on banks of eight

cascaded microdisk modulators (Fig. 3d). After modulation, the groups are recombined onto a single fibre using two stages of RMZI interleavers. For the receiver, two stages of de-interleavers are similarly used to subdivide the comb into four groups, which are incident on banks of eight cascaded microring filters with germanium photodiodes at the drop ports. Although the transmitter chip only requires a single fibre in and a single fibre out for data transmission, six ancillary optical I/O ports provide tapped light from the circuit at various points for (de-)interleaver and modulator alignment. Similarly, the receiver only requires a single fibre in, but uses seven ancillary optical I/O ports for alignment.

We optically and electrically packaged the photonic transmitter and receiver chips with eight-channel single-mode fibre (SMF) arrays at 127-µm pitch for optical I/O and 96 and 50 direct current (d.c.) wire-bonds around the periphery at 100-µm pitch for heater control for the
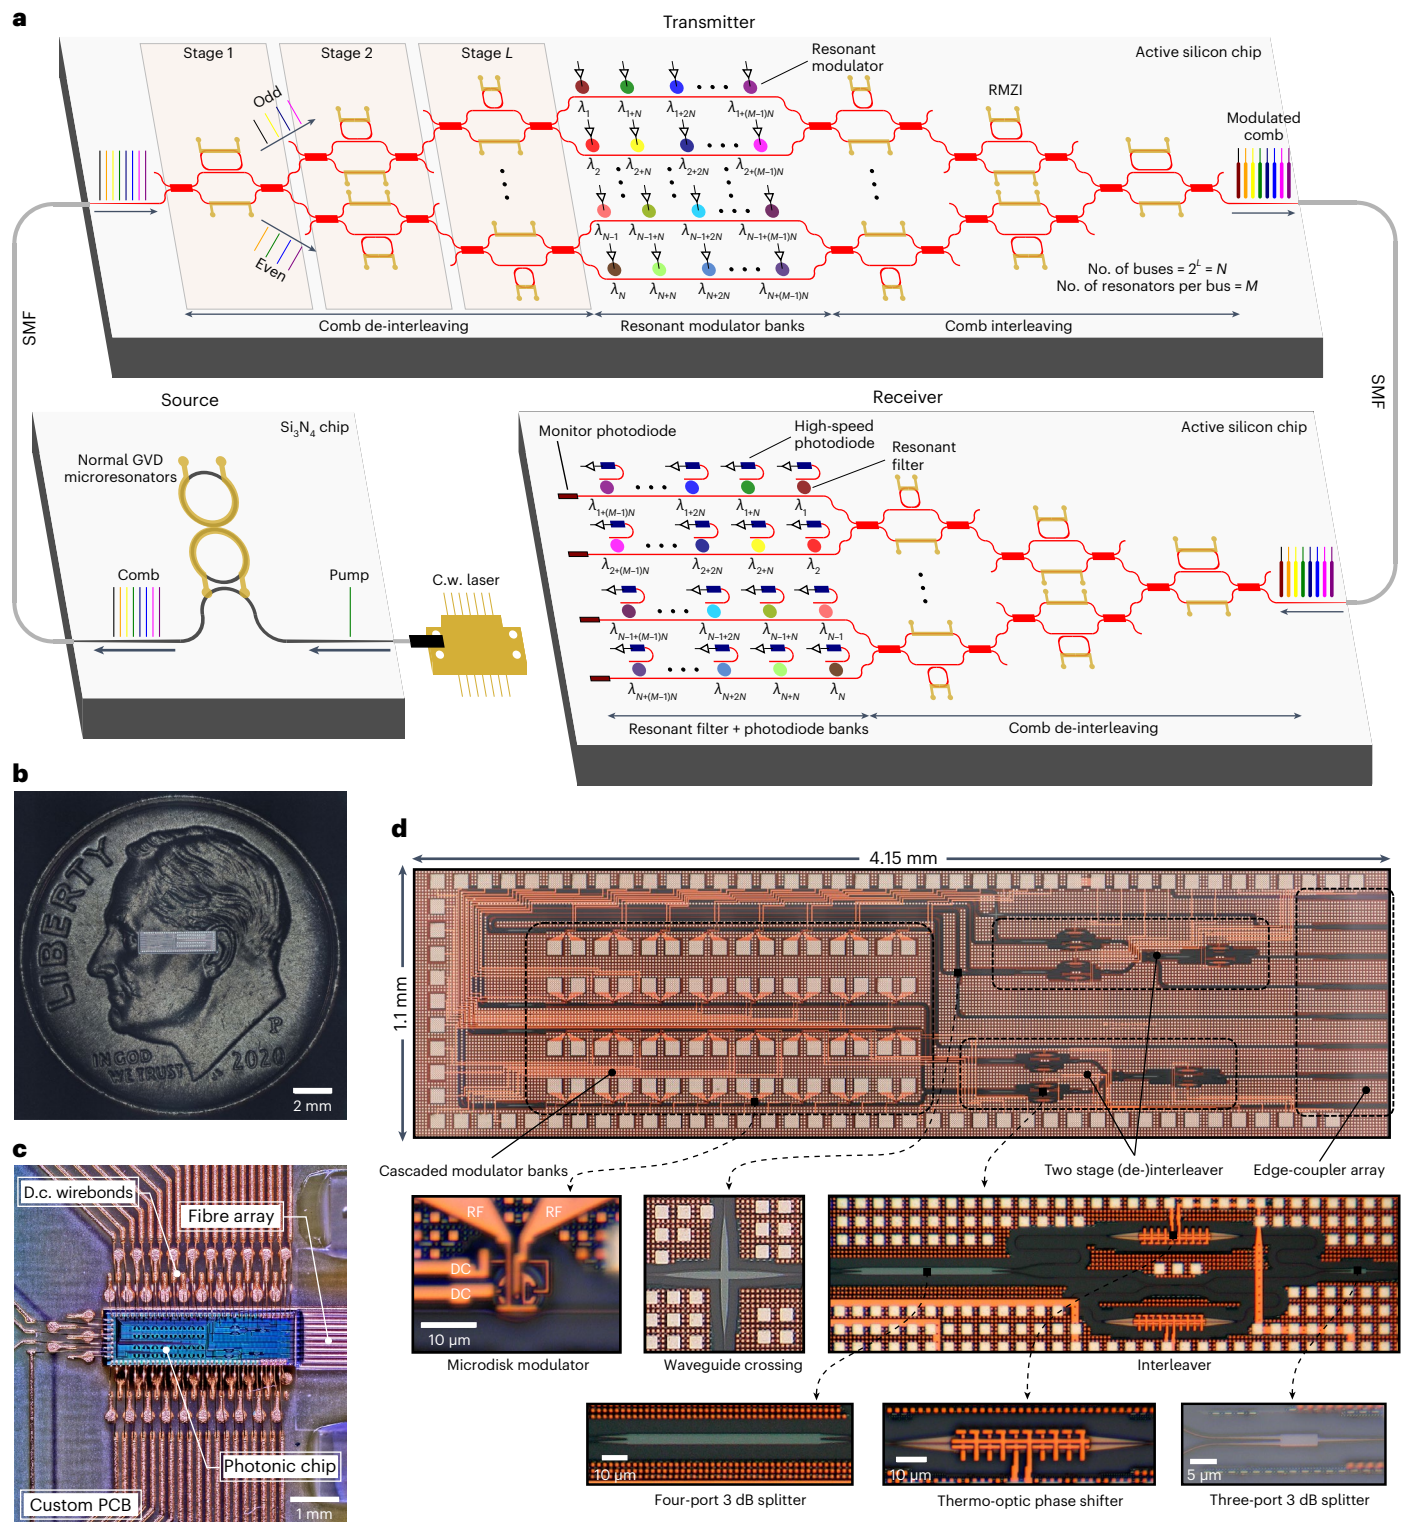
Fig. 3 | Overview of the scalable system architecture and fabricated transmitter chip. a, Schematic of the proposed scalable link architecture for a highly parallel DWDM silicon photonic interconnect using a broadband Kerr frequency comb source. **b**, Microscope image of the unpackaged chip on a US dime. **c**, Microscope image of the electrically and optically packaged system highlighting key constituent components. **d**, Microscope image of the full die before wire-bonding, with additional representative images of all active and passive devices on the chip.

transmitter and receiver, respectively (Fig. 3c and Methods). The measured optical coupling loss after ultraviolet curing with index-matching epoxy was 2.5 dB per facet, which greatly reduced the required optical power per line from the comb to close the link budget compared to the bare die case. Due to the high parasitic inductance of wirebonds, we left

the inner high-speed radiofrequency (RF) pads to the modulators and photodiodes exposed, for probing with multi-contact wedge probes. All thermal biases were addressable with uniform current–voltage characteristics (Extended Data Fig. 2), indicating 100% yield for the d.c. wirebonds and photonic devices. Furthermore, all modulators and
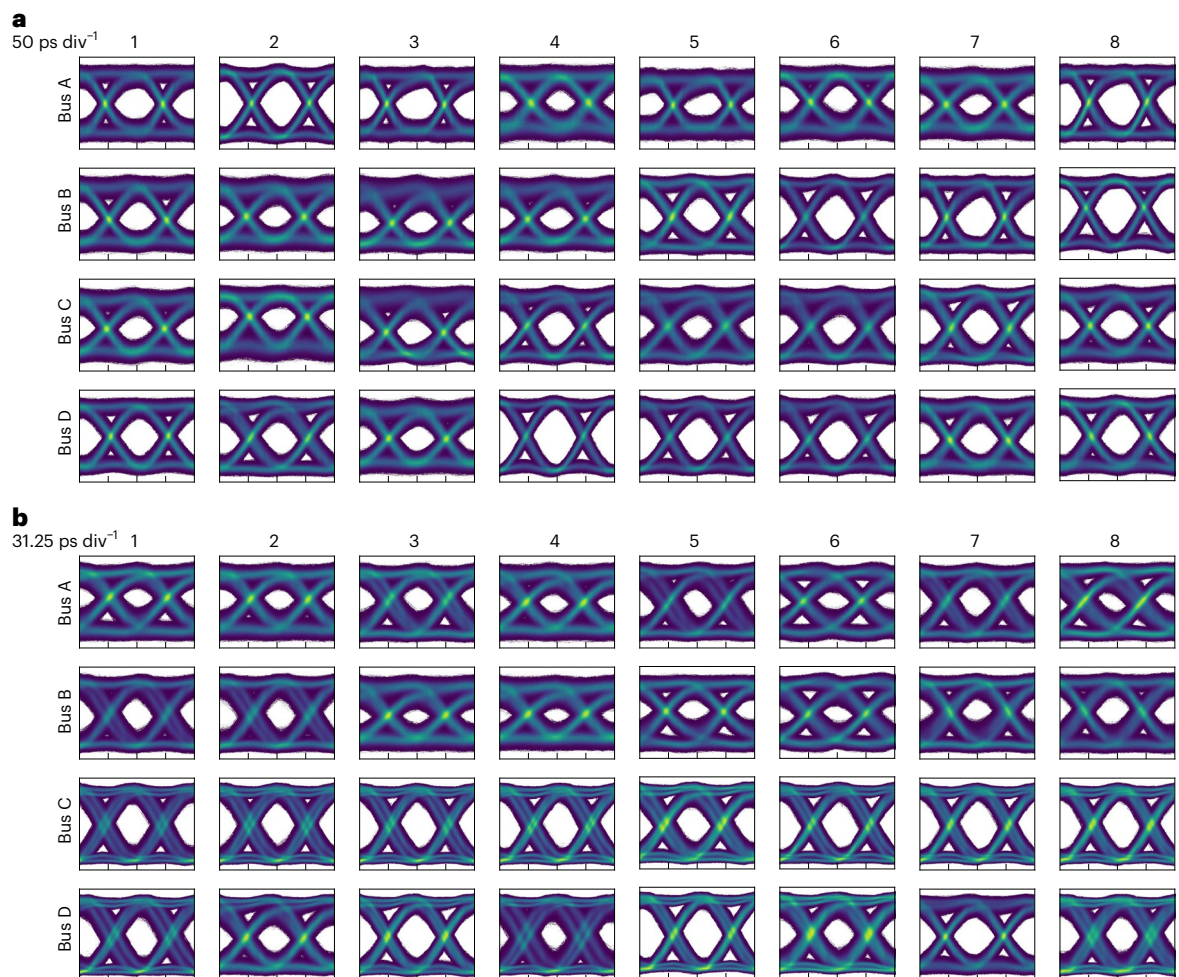
Fig. 4 | Thirty-two channel transmitter experimentally measured eye diagrams. a,b, Eye diagrams for all 32 modulators on the chip at 10 Gb s^{-1} (**a**) and 16 Gb s^{-1} (**b**). For this initial proof-of-principle demonstration, the channels were characterized one by one, rather than simultaneously.

photodiodes on the chip displayed open-eye diagrams up to 16 Gb s^{-1} (Fig. 4a,b), indicating uniform electro-optic bandwidth across devices and 100% yield for the high-speed photonic devices. The perfect yield and high device uniformity emphasize the advantage of fabrication in a commercial 300-mm CMOS foundry and bode well for the system's potential for future high-volume scaling.

The individual comb lines displayed a negligible power penalty compared to a tunable c.w. laser source and in some cases showed an improvement over the c.w. source (Fig. 5a) when operated in the low-noise mode-locked state (Extended Data Figs. 3 and 4). Due to the low coupling losses of the packaged transmitter chip and the low on-chip losses of the photonic devices, the comb line one channel blue-detuned from the pump (1,559.8 nm) had enough fibre-coupled optical power to close the link budget without comb amplification before the transmitter chip (Fig. 5a). This result is promising for future link iterations with reduced fibre-chip coupling losses and integrated high-sensitivity receiver modules[36,37], as it shows that the generated comb lines can natively have enough optical power per line to close the link budget. To quantify crosstalk effects between adjacent comb lines on a single bus, we simultaneously modulated two lines at 200-GHz spacing and measured the BER values for both channels as a function of received optical power (Fig. 5b). Although a slight power penalty is observed (<1 dB) compared to when only a single channel is modulated, both channels achieve simultaneous error-free operation (BER values below 10^{-12}) at 10 Gb s^{-1}.

To further verify the performance of the transmitter, we directly measured BER values for all 32 comb wavelength channels across

the C and L bands (Fig. 5c). Due to the cascaded interleavers on the transmitter chip used for subdividing the comb, no filtering was required between the Si$_3$N$_4$ comb chip and the transmitter package. All off-chip filtering was restricted to the receiver side, where we used a tunable optical bandpass filter to emulate resonant drop filters for selecting single modulated comb lines. At 10 Gb s^{-1}, all comb lines achieved natively error-free transmission, which was verified using direct BER measurements without offline processing or FEC. Direct BER measurement provides the most rigorous experimental validation of signal integrity compared to other indirect metrics such as eye quality factor[38]. At 16 Gb s^{-1}, 20 of the 32 channels were natively error-free, with the remaining channels clustered near the avoided mode-crossing region and low-power region near the blue edge of the C band. We attribute the higher BER values to the lower signal-to-noise ratio of these regions, which can be rectified in future systems by using a coherent combination of two spectrally distinct combs to selectively shape the envelope and increase the signal-to-noise ratio in problematic regions[39].

Finally, we characterized the performance of the 32-channel silicon photonic receiver (Fig. 6). The system achieves natively error-free performance for all channels at 10 Gb s^{-1} and for 6 of the 32 channels at 16 Gb s^{-1}. We attribute the deteriorated performance at 16 Gb s^{-1} to noise stemming from the long distance travelled by the generated photocurrent before electrical amplification by a commercial RF amplifier (Supplementary Note 5). In future implementations with 3D integrated photonics and CMOS electronics, the receiver sensitivity can
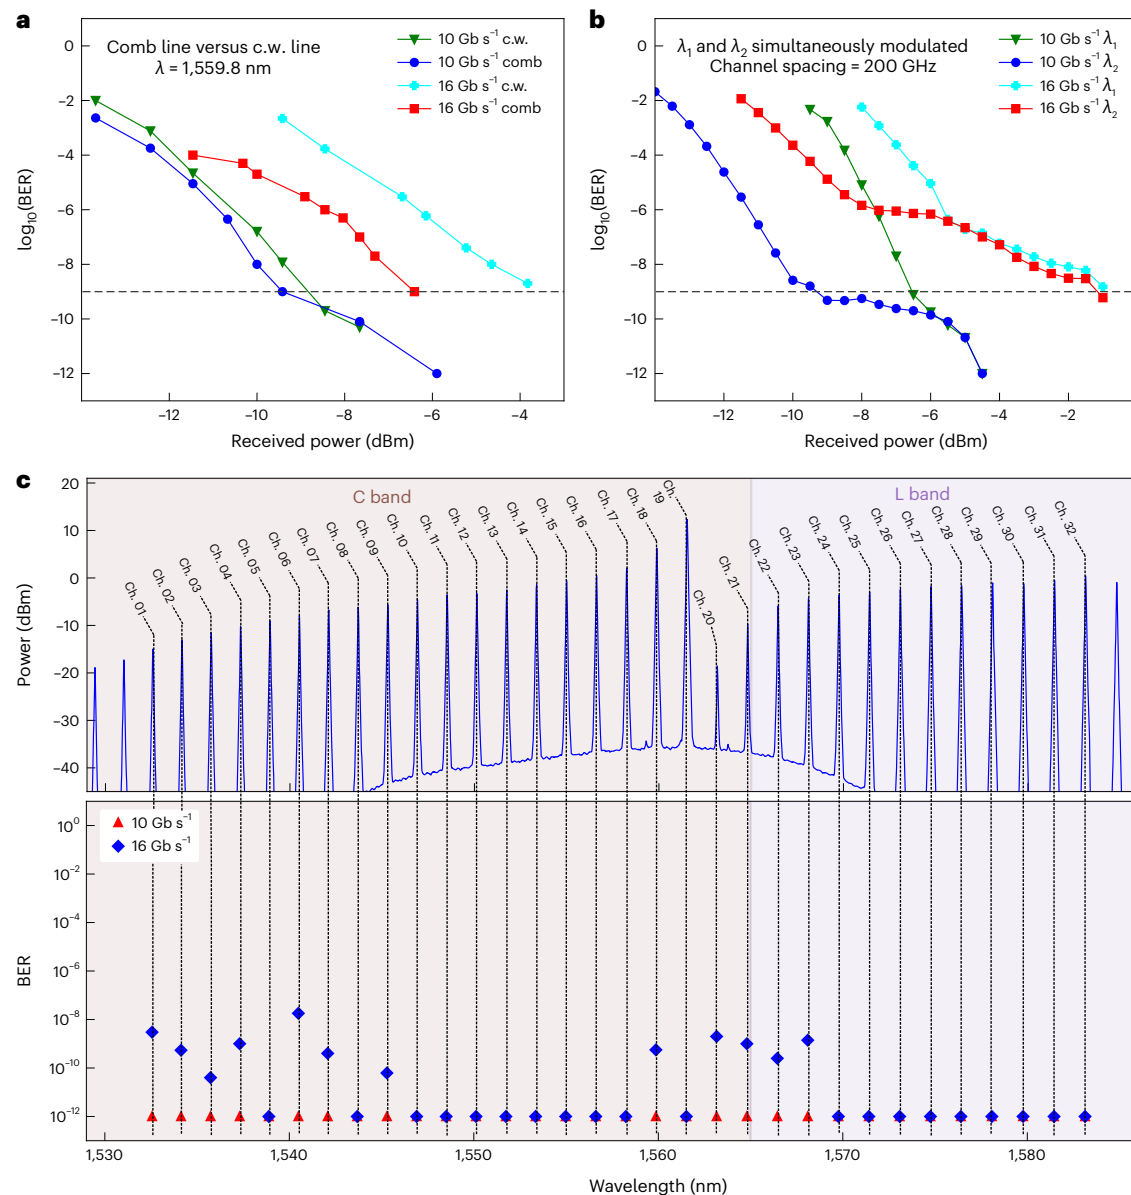
Fig. 5 | Thirty-two channel transmitter BER results. a, Measured BER versus received optical power comparing the performance of a c.w. laser to a generated comb line at 1,559.8 nm without comb amplification before the transmitter chip. **b**, Measured BER curves for two adjacent comb lines modulated simultaneously. **c**, Measured BER for all 32 comb wavelength channels spanning 50 nm of optical bandwidth across the C and L bands.

be substantially improved to −19 dBm at 16 Gb s^{-1} by packaging the transimpedance amplifiers in close spatial proximity to the photodiodes with copper pillar microbumps[37].

Although these proof-of-principle experimental demonstrations required the use of erbium-doped fibre amplifier (EDFAs) to fully characterize the system, their inclusion is not fundamental to the proposed architecture and future implementations will be able to eliminate optical amplification entirely (Supplementary Note 6). With improved photonic devices (Supplementary Note 7), improved receiver sensitivity, and the elimination of noisy broadband EDFAs (exhibiting highly wavelength-dependent noise), it is likely that every comb line across the S, C and L bands can achieve error-free performance based on their coherence and noise properties[27].

Discussion

In conclusion, we have demonstrated a massively scalable silicon photonic data-communication link using a normal-GVD Kerr comb source operating at a state-of-the-art aggregate bandwidth through massively parallel DWDM. In contrast with other comparable high-bandwidth solutions, which scale by increasing the single-channel baud rate and moving to higher-order modulation formats, our demonstrated approach permits modest per-channel data rates with a standard NRZ-OOK modulation format. In future system implementations without EDFAs, this approach can enable extreme energy efficiency and low latency by the elimination of power-hungry DSP and FEC while still maintaining high aggregate bandwidths per fibre. We have experimentally characterized a data transmission up to 512 Gb s^{-1} on a single fibre (465 Gb mm^{-1} shoreline bandwidth density) with continued room for scaling through a previously undemonstrated link architecture. Moreover, normal-GVD Kerr combs inherently have high conversion efficiencies (typically >30%, with 86% experimentally demonstrated[40]) and recently have been shown to synchronize to produce comb line powers otherwise unattainable[39,41], providing a clear avenue to highly efficient chip-scale multi-wavelength sources.
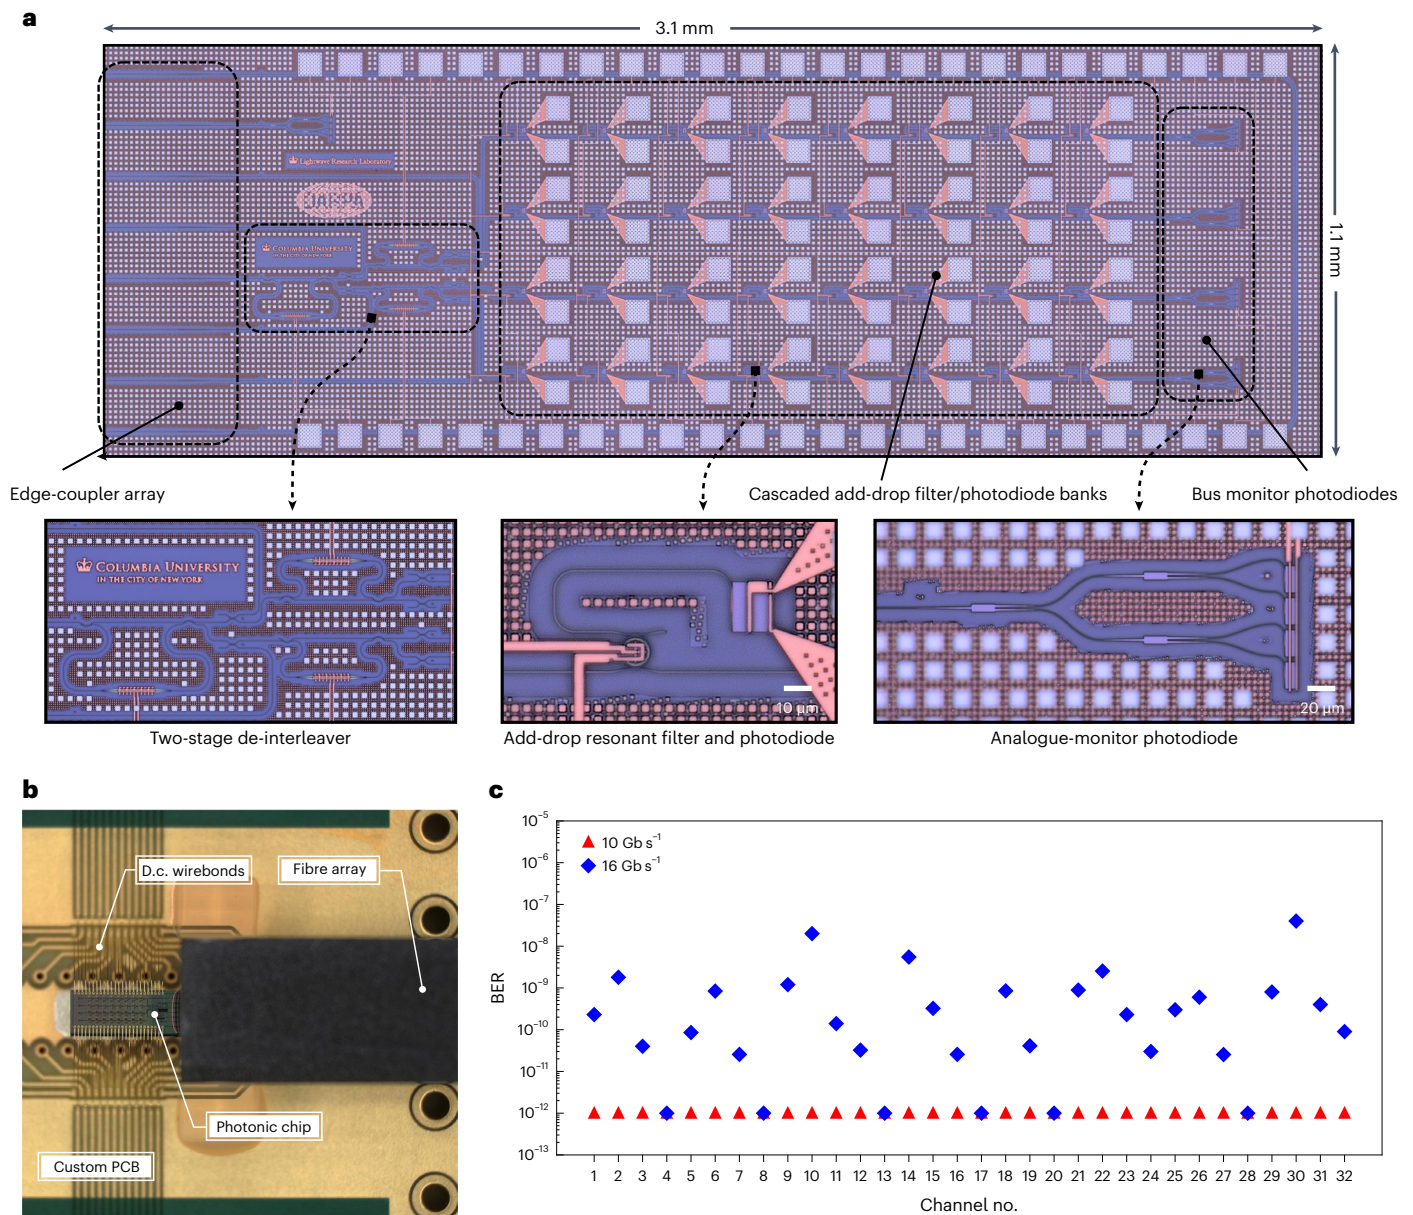
Fig. 6 | Thirty-two channel receiver results. a, Microscope image of the full die after wire-bonding, with additional representative images of all active and passive devices on the chip. **b**, Microscope image of the electrically and optically packaged system, highlighting key constituent components. **c**, BER measurements for all 32 channels on the chip at both 10 Gb s^{-1} and 16 Gb s^{-1}.

We fabricated the transmitter chip in a standard SOI process at a commercial foundry, enabling straightforward packaging with modern CMOS electronics for fully integrated transceivers either through heterogeneous[36] or monolithic[42] integration. This architecture also enables continued scaling in the number of wavelength channels as well as per-channel data rate with improved modulator designs[43] and is fully compatible with other common multiplexing techniques such as mode-division multiplexing[44,45]. Additionally, recent work has shown automated wafer-scale trimming of resonant modulators to correct for fabrication variations[46], enabling future as-designed alignment of modulators and filters to the comb grid without any thermal energy consumption.

Previous roadblocks to the use of Kerr combs for short-reach optical links have recently been overcome, with solutions for high volume manufacturability[47,48], high optical power per line with spectral flatness[41], high pump-to-comb conversion efficiency[27,31,40] and turn-key integration in standard compact laser packages[49,50].

In concert with these advancements, the work presented here represents a coming-of-age for Kerr frequency combs as optical interconnect sources, establishing them as a pragmatic contender for integration in next-generation data centres and high-performance computers.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41566-023-01244-7.

References

1. Cheng, Q., Bahadori, M., Glick, M., Rumley, S. & Bergman, K. Recent advances in optical technologies for data centers: a review. *Optica* **5**, 1354–1370 (2018).

2. Miller, D. Rationale and challenges for optical interconnects to electronic chips. *Proc. IEEE* **88**, 728–749 (2000).

3. Strubell, E., Ganesh, A. & McCallum, A. Energy and policy considerations for deep learning in NLP. In *Proc. 57th Annual Meeting of the Association for Computational Linguistics* 3645–3650 (ACL, 2019); https://aclanthology.org/P19-1355

4. Miller, D. A. B. Attojoule optoelectronics for low-energy information processing and communications. *J. Light. Technol.* **35**, 346–396 (2017).

5. Atabaki, A. H. et al. Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip. *Nature* **556**, 349–354 (2018).

6. Hu, H. & Oxenløwe, L. K. Chip-based optical frequency combs for high-capacity optical communications. *Nanophotonics* **10**, 1367–1385 (2021).

7. Gaeta, A. L., Lipson, M. & Kippenberg, T. J. Photonic-chip-based frequency combs. *Nat. Photon.* **13**, 158–169 (2019).

8. Levy, J. S. et al. CMOS-compatible multiple-wavelength oscillator for on-chip optical interconnects. *Nat. Photon.* **4**, 37–40 (2010).

9. Kippenberg, T. J., Gaeta, A. L., Lipson, M. & Gorodetsky, M. L. Dissipative Kerr solitons in optical microresonators. *Science* **361**, eaan8083 (2018).

10. Raja, A. S. et al. Ultrafast optical circuit switching for data centers using integrated soliton microcombs. *Nat. Commun.* **12**, 5867 (2021).

11. Marin-Palomo, P. et al. Microresonator-based solitons for massively parallel coherent optical communications. *Nature* **546**, 274–279 (2017).

12. Pfeifle, J. et al. Coherent terabit communications with microresonator Kerr frequency combs. *Nat. Photon.* **8**, 375–380 (2014).

13. Fülöp, A. et al. High-order coherent communications using mode-locked dark-pulse Kerr combs from microresonators. *Nat. Commun.* **9**, 1598 (2018).

14. Hu, H. et al. Single-source chip-based frequency comb enabling extreme parallel data transmission. *Nat. Photon.* **12**, 469–473 (2018).

15. Wade, M. et al. TeraPHY: a chiplet technology for low-power, high-bandwidth in-package optical I/O. *IEEE Micro* **40**, 63–71 (2020).

16. Liang, D. et al. Integrated green DWDM photonics for next-gen high-performance computing. In *Proc. Optical Fiber Communication Conference (OFC) 2020* Th1E.2 (Optical Society of America, 2020).

17. Chen, C.-H. et al. A comb laser-driven DWDM silicon photonic transmitter based on microring modulators. *Opt. Express* **23**, 21541–21548 (2015).

18. Moscoso-Mártir, A. et al. Silicon photonics transmitter with SOA and semiconductor mode-locked laser. *Sci. Rep.* **7**, 13857 (2017).

19. Wade, M. et al. An error-free 1Tbps WDM optical I/O chiplet and multi-wavelength multi-port laser. In *Proc. Optical Fiber Communication Conference (OFC) 2021* F3C.6 (eds Dong, P. et al.) (Optical Society of America, 2021).

20. Hofmann, W. et al. 1.55-μm VCSEL arrays for high-bandwidth WDM-PONs. *IEEE Photon. Technol. Lett.* **20**, 291–293 (2008).

21. Wang, C. et al. Monolithic lithium niobate photonic circuits for Kerr frequency comb generation and modulation. *Nat. Commun.* **10**, 978 (2019).

22. Xu, Q., Schmidt, B., Shakya, J. & Lipson, M. Cascaded silicon micro-ring modulators for WDM optical interconnection. *Opt. Express* **14**, 9431–9436 (2006).

23. London, Y. et al. Energy efficiency analysis of comb source carrier-injection ring-based silicon photonic link. *IEEE J. Select. Top. Quantum Electron.* **26**, 1–13 (2020).

24. Streshinsky, M. et al. Silicon parallel single mode 48×50gb/s modulator and photodetector array. *J. Light. Technol.* **32**, 3768–3775 (2014).

25. Dong, P., Lee, J., Kim, K., Chen, Y.-K. & Gui, C. Ten-channel discrete multi-tone modulation using silicon microring modulator array. In *Proc. Optical Fiber Communication Conference* W4J-4 (Optica Publishing Group, 2016).

26. Kong, D. et al. Intra-datacenter interconnects with a serialized silicon optical frequency comb modulator. *J. Light. Technol.* **38**, 4677–4682 (2020).

27. Kim, B. Y. et al. Turn-key, high-efficiency Kerr comb source. *Opt. Lett.* **44**, 4475–4478 (2019).

28. Jang, J. K. et al. Conversion efficiency of soliton Kerr combs. *Opt. Lett.* **46**, 3657–3660 (2021).

29. Lobanov, V. E., Lihachev, G. & Gorodetsky, M. L. Generation of platicons and frequency combs in optical microresonators with normal GVD by modulated pump. *Europhys. Lett.* **112**, 54008 (2015).

30. Liu, Y. et al. Investigation of mode coupling in normal-dispersion silicon nitride microresonators for Kerr frequency comb generation. *Optica* **1**, 137–144 (2014).

31. Xue, X. et al. Mode-locked dark pulse Kerr combs in normal-dispersion microresonators. *Nat. Photon.* **9**, 594–600 (2015).

32. Miller, S. A. et al. Tunable frequency combs based on dual microring resonators. *Opt. Express* **23**, 21527–21540 (2015).

33. Padmaraju, K., Zhu, X., Chen, L., Lipson, M. & Bergman, K. Intermodulation crosstalk characteristics of WDM silicon microring modulators. *IEEE Photon. Technol. Lett.* **26**, 1478–1481 (2014).

34. Luo, L.-W. et al. High bandwidth on-chip silicon photonic interleaver. *Opt. Express* **18**, 23079–23087 (2010).

35. Rizzo, A., Cheng, Q., Daudlin, S. & Bergman, K. Ultra-broadband interleaver for extreme wavelength scaling in silicon photonic links. *IEEE Photon. Technol. Lett.* **33**, 55–58 (2021).

36. Daudlin, S. et al. 3D-integrated multichip module transceiver for terabit-scale DWDM interconnects. In *Proc. Optical Fiber Communication Conference (OFC) 2021* Th4A.4 (eds Dong, P. et al.) (Optical Society of America, 2021).

37. Saeedi, S., Menezo, S., Pares, G. & Emami, A. A 25gb/s 3D-integrated CMOS/silicon-photonic receiver for low-power high-sensitivity optical communication. *J. Light. Technol.* **34**, 2924–2933 (2016).

38. Freude, W. et al. Quality metrics for optical signals: eye diagram, Q-factor, OSNR, EVM and BER. In *Proc. 2012 14th International Conference on Transparent Optical Networks (ICTON)* 1–4 (IEEE, 2012).

39. Kim, B. Y. et al. Coherent combining for high-power Kerr combs. In *Proc. 2022 Conference on Lasers and Electro-Optics (CLEO)* 1–2 (IEEE, 2022).

40. Zang, J., Yu, S.-P., Carlson, D. R., Briles, T. C. & Papp, S. B. Near unit efficiency in microresonator combs. In *CLEO: Science and Innovations* STh4F-3 (Optica Publishing Group, 2022).

41. Kim, B. Y. et al. Synchronization of nonsolitonic Kerr combs. *Sci. Adv.* **7**, eabi4362 (2021).

42. Giewont, K. et al. 300-mm monolithic silicon photonics foundry technology. *IEEE J. Select. Top. Quantum Electron.* **25**, 1–11 (2019).

43. Sun, J. et al. A 128gb/s PAM4 silicon microring modulator with integrated thermo-optic resonance tuning. *J. Light. Technol.* **37**, 110–115 (2019).

44. Yang, K. Y. et al. Multi-dimensional data transmission using inverse-designed silicon photonics and microcombs. *Nat. Commun.* **13**, 7862 (2022).

45. Luo, L.-W. et al. WDM-compatible mode-division multiplexing on a silicon chip. *Nat. Commun.* **5**, 3069 (2014).

46. Jayatilleka, H. et al. Post-fabrication trimming of silicon photonic ring resonators at wafer-scale. *J. Light. Technol.* **39**, 5083–5088 (2021).

47. Jin, W. et al. Hertz-linewidth semiconductor lasers using CMOS-ready ultra-high-Q microresonators.*Nat. Photon* **15**, 346–353 (2021).

48. Liu, J. et al. High-yield, wafer-scale fabrication of ultralow-loss, dispersion-engineered silicon nitride photonic circuits. *Nat. Commun.* **12**, 2236 (2021).

49. Stern, B., Ji, X., Okawachi, Y., Gaeta, A. L. & Lipson, M. Battery-operated integrated frequency comb generator. *Nature* **562**, 401–405 (2018).

50. Shen, B. et al. Integrated turnkey soliton microcombs. *Nature* **582**, 365–369 (2020).

Methods

Scalable link architecture

The source consists of an external or integrated c.w. laser that pumps a Si_3N_4 dual-microresonator system to generate a normal-GVD comb source. Standard SMF is used to connect the source to the transmitter chip. The transmitter architecture uses a binary tree of asymmetric MZIs to split the comb spectrum into even and odd groups at each stage. After the comb is divided into subgroups, the groups traverse separate buses of cascaded resonant modulators, taking advantage of the spectral selectivity of resonant cavities to encode independent data streams onto different coloured comb lines. After each channel is modulated, they are recombined using an identical binary tree of MZIs and coupled off the chip into SMF. At the receiver, the modulated comb lines traverse a final MZI binary tree and are then incident on buses with cascaded resonant filters tuned to match the respective wavelength channels. Each resonant filter has a photodetector at its drop port to convert the optical data streams back into the electrical domain. This architecture naturally enables a disaggregated laser source, which allows the pump and comb to be separately stabilized away from the harsh thermal environment of co-packaged optics and electronics.

The binary trees of MZI filters serve multiple purposes. Depending on the depth of the tree d, the channel spacing on each bus is increased by a factor of 2^d. Additionally, the number of channels per bus is reduced by a factor of 2^d, enabling channel configuration schemes that are not restricted by the FSR of the resonant modulators/filters (Supplementary Note 4). Recently, a fully automated control algorithm for aligning similar cascaded MZI binary trees was demonstrated[51], providing a straightforward path to full-system initialization and stabilization. Furthermore, this architecture can be extended for the receiver to accommodate polarization diversity while keeping the number of resonators and photodiodes constant, eliminating the need for polarization-maintaining fibre (Supplementary Note 8).

Link implementation

Our chosen transmitter implementation uses microdisk modulators with an embedded vertical p–n junction that overlaps with the optical whispering gallery mode, while the receiver implementation uses resonant add–drop filters with epitaxially grown waveguide-integrated germanium-on-silicon photodiodes at the drop port. When the modulators are driven in reverse bias, the width of the depletion region is modulated, which changes the free carrier distribution interacting with the optical mode and thus changes the effective refractive index through the plasma-dispersion effect. This change in the effective index corresponds to a shift in the resonant wavelength, which modulates the amplitude of the light to achieve levels of '0' and '1' depending on the applied bias. Due to the large overlap between the depletion region and optical mode in vertical junction devices[52], the modulators can be directly driven at CMOS-compatible voltages while maintaining a high extinction ratio and low insertion loss (Extended Data Fig. 2). Furthermore, silicon photonic modulators based on the plasma-dispersion effect have been experimentally shown to function over hundreds of nanometres[53] and thus will not restrict the wavelength scalability of the system. Similar to the resonant modulators, the add–drop microring filters also exhibit a wavelength-selective response, which is used to drop target channels from a common bus onto single photodiodes. The germanium photodiodes are reverse-biased to enhance the generated photocurrent by efficiently sweeping out the optically generated free carriers. Because the resonant wavelengths of the modulators and add–drop filters are sensitive to fabrication and temperature variations, integrated thermo-optic heaters are implemented to align the resonances to the desired comb lines.

Active silicon photonic chips

The active silicon photonic circuits were designed using a mix of custom-designed devices and verified process design kit (PDK) components provided by the foundry. Custom devices were designed using the Ansys Lumerical finite-difference time-domain solver to extract S-parameters and create compact models, which were then used to perform circuit-level simulations with the PDK devices in Ansys Lumerical INTERCONNECT. The mask layout was prepared in Cadence Virtuoso, and physical verification was performed using Mentor Graphics Calibre. Fabrication was performed as part of the AIM Photonics 300-mm MPW service. The AIM process uses a 220-nm-thick silicon waveguide layer with 2-μm-thick buried oxide. For a single-mode strip waveguide of nominal width (480 nm), the propagation loss is ~2.5 dB cm^{-1}. Additionally, the AIM process and PDK use Si_3N_4 layers for devices such as inverse-tapered edge couplers and silicon-to-Si_3N_4 escalators, indicating that monolithic integration of the comb-generating microresonators on the same die as the active silicon devices is a future possibility with improved foundry nitride processing. More information regarding the AIM process and PDK is available in ref. 54.

Si_3N_4 chip fabrication and design

Starting from a 4-inch silicon wafer, we thermally grew a 4-μm-thick oxide layer as the bottom cladding. Si_3N_4 was deposited using low-pressure chemical vapour deposition in two steps and annealed at 1,200 °C in an argon atmosphere for 3 h between steps. After Si_3N_4 deposition, we deposited a SiO_2 hard mask using plasma-enhanced chemical vapour deposition (PECVD). We patterned the devices using electron-beam lithography. Ma-N 2403 resist was used to write the pattern, and the nitride film was etched in an inductively coupled plasma reactive ion etcher using a combination of CHF_3, N_2 and O_2 gases. After stripping the oxide mask, we annealed the devices again to remove residual N–H bonds in the Si_3N_4 film. We clad the devices with 500 nm of high-temperature SiO_2 deposited at 800 °C, followed by 2 μm of SiO_2 using PECVD. A deep etched facet and an inverse taper were designed and used to minimize the edge-coupling loss. More detailed information on the device fabrication is available in refs. 55–57.

The coupled-microresonator system was designed to generate a 200-GHz-spaced comb, and the cross-section of the Si_3N_4 microrings was 730 × 1,000 nm. In this dual-resonator configuration, modes from separate mode families can spectrally overlap and couple, resulting in an avoided mode crossing with local anomalous GVD near the degeneracy point. By tuning the integrated heaters and thus altering the coupling conditions between the rings, this avoided mode crossing can be spectrally shifted and its splitting strength can be controlled. Because both of these crucial properties for comb generation are fully tunable, constraints on the device fabrication are greatly relaxed compared to soliton Kerr combs. Furthermore, the low-noise comb state can persist over a wider tuning range relative to soliton combs[58], resulting in dramatically increased resilience to environmental temperature fluctuations and thus reduced control complexity (Extended Data Fig. 3). Finally, it has been shown that such coupled-ring devices support turn-key operation with a fixed-wavelength pump laser[27], enabling fully automated Kerr comb generation and deterministic tuning to the low-noise state required for data communications. The comb used for the data-transmission experiments in this work was designed to have 200-GHz spacing to operate in the ideal regime for bandwidth and conversion efficiency while maintaining high power per line. Similar devices have shown experimentally measured pump-to-comb conversion efficiencies of up to 41%[27]. The dual resonators had platinum heaters above the Si_3N_4 waveguide layer, enabling thermo-optic tuning of an avoided mode crossing to enable arbitrary pump wavelengths for generating the low-noise comb spectrum[27,32]. To achieve the low-noise wavelength channels required for high-fidelity data transmission, the resonances were tuned to generate a low-noise comb for all experiments (Extended Data Fig. 4). More information regarding the turn-key operation and noise properties of a similar comb is provided in ref. 27.

Device packaging

Electrical and optical packaging was performed by Optelligent. The active silicon photonic chip was first die-bonded to a co-designed printed circuit board (PCB). The PCB bond pads were finished with electroless nickel electroless palladium immersion gold (ENEPIG) to enable better bond reliability with gold wirebonds. All d.c. pads around the periphery of the photonic chips were wirebonded to the PCB and the internal RF pads were left exposed for probing. Eight-channel lidless V-groove fibre arrays with standard cleaved SMF-28 fibres (OZ Optics) were flipped upside down to provide clearance over the chip dicing trenches and were attached to the 127-µm-pitch edge-coupler array on the photonic chips using index-matched ultraviolet-cured epoxy after active alignment. The measured optical coupling loss after packaging was ~2.5 dB per facet.

Modulator characterization

Microdisk modulators available in the AIM PDK[54] were used for the transmitter, with four nominal wavelength designs (1,550 nm, 1,556.4 nm, 1,562.8 nm and 1,569.4 nm). The modulators each had an FSR of ~25.6 nm. Integrated thermal heaters in each modulator were used to shift the resonances to align to the comb wavelength grid, with a measured tuning efficiency of 0.5 nm mW^{-1}. Future iterations using custom modulator designs with more nominal wavelengths and high-efficiency heaters will greatly reduce the expended thermal tuning power (Supplementary Note 7). The measured loaded quality factor (Q) of a typical modulator was 5,200 ($f_{FWHM} = 39$ GHz at 1,550 nm), which indicates an optical 3 dB modulation bandwidth of $f_{3dB}^{opt} = 25$ GHz using the relation $f_{3dB}^{opt} = \sqrt{\sqrt{2}-1} \times f_{FWHM}$ (refs. 52,59). The electro-optic 3 dB bandwidth was measured to be $f_{3dB}^{el-opt} = 15$ GHz with a 0-V d.c. bias (Extended Data Fig. 2c), which is combined with the calculated optical modulation bandwidth to estimate the RC bandwidth through the relation $\left(\frac{1}{f_{3dB}^{el-opt}}\right)^2 = \left(\frac{1}{f_{3dB}^{opt}}\right)^2 + \left(\frac{1}{f_{3dB}^{el}}\right)^2$, yielding $f_{3dB}^{el} = 37.5$ GHz (ref. 52).

Interleaver characterization and alignment

The RMZI interleavers have a thermo-optic phase shifter in both the ring and bottom arm to correct for phase errors due to fabrication variations. An applied power of 25 mW to the bottom arm heater yields an FSR/2 shift of the pass and stop bands, corresponding to a π phase shift. The first de-interleaver stage has a 10% power tap in the bottom arm that is looped back to an output edge coupler to align the pass and stop bands. A tunable laser (Keysight 81608A), lightwave measurement system (Keysight 8164B) and four-channel optical power meters (Keysight N7744A) were used to sweep the (de-)interleaver spectrum while on-chip thermo-optic heaters were tuned to align the spectrum to the comb wavelength grid. Power taps were used at the end of each bus to similarly align the second-stage de-interleaver and modulator resonances. A tap before the last interleaver stage was used for final alignment to recombine the comb onto a single fibre output.

Data transmission experiments

The Si$_3$N$_4$ comb chip was pumped using a tunable c.w. laser with 7 dBm output power (Alnair Labs TLG-200), which was amplified to ~1 W using an EDFA (Amonics AEDFA-33-B-FA). The polarization of the pump was tuned with a polarization controller to ensure that the fundamental transverse-electric (TE) mode was launched into the chip edge couplers. The edge couplers on the comb chip had a measured loss of 3 dB facet^{-1}, yielding 500 mW of on-chip optical power for comb generation. The generated comb was amplified using a broadband C and L band EDFA (FiberLabs AMP-FL8021-CLB-22) and was sent to the transmitter package using SMF with a polarization controller to launch TE-polarized light into the chip. The transmitter PCB was mounted to a temperature-controlled stage (Thorlabs PTC1) to thermally stabilize the chip. D.c. biases for thermo-optic control of the interleavers and modulators were provided by a 96-channel source measurement

unit (Qontrol) connected to the PCB. The bit pattern was generated by a BER tester (BERT; Anritsu MP1900A) and was sent to the modulators using RF multi-contact wedge probes (Cascade Infinity). All data-transmission experiments used a pseudo-random bit sequence of length $2^{31} - 1$ (PRBS31) to provide the most rigorous test pattern and eliminate the possibility of pattern-dependent behaviour. After modulation, the light was coupled off the chip into SMF and amplified using a broadband C and L band EDFA (FiberLabs AMP-FL8021-CLB-22). A tunable bandpass filter (Finisar WaveShaper) was then used to select single modulated comb lines with a variable optical attenuator (VOA; Thorlabs EVOA1550A) to control the optical power to the receiver. The modulated comb lines were then converted back into the electrical domain using a photodiode and transimpedance amplifier (Thorlabs RXM40AF). The received signal was sent to a real-time oscilloscope (Keysight Infiniium Z-Series) for eye characterization or back to the BERT for BER evaluation.

For the 32-channel receiver experiments, individual comb lines were first filtered using a tunable bandpass filter (Finisar WaveShaper) and then modulated using a commercial benchtop modulator (Thorlabs MX40B). The modulated comb lines were then incident on an EDFA (FiberLabs AMP-FL8021-CLB-22) and VOA (Thorlabs EVOA1550A), and finally coupled to the packaged receiver chip. D.c. biases for tuning the resonance wavelength of the on-chip resonant filters and current readout for the on-chip monitor photodiodes were provided by a 96-channel source measurement unit (Qontrol) connected to the PCB. On-chip high-speed photodiodes were used to convert the optical signal back into the electrical domain. The converted high-speed electrical signals from the drop port photodiodes were extracted from the chip using RF multi-contact wedge probes (Cascade Infinity) and amplified with an RF amplifier (SHF Communication 826H). The amplified electrical signals were sent to a real-time oscilloscope (Keysight Infiniium Z-Series) for eye characterization or back to the BERT (Anritsu MP1900A) for BER evaluation.

Data availability

The data that support the plots in Figs. 2 and 4–6 are available at https://doi.org/10.6084/m9.figshare.22723826. All data that support the other findings of this study are available from the corresponding author upon reasonable request.

References

51. Akiyama, T. et al. Cascaded AMZ triplets: a class of demultiplexers having a monitor and control scheme enabling dense WDM on Si nano-waveguide PICs with ultralow crosstalk and high spectral efficiency. *Opt. Express* **29**, 7966–7985 (2021).

52. Timurdogan, E. et al. An ultralow power athermal silicon modulator. *Nat. Commun.* **5**, 4008 (2014).

53. Doerr, C. et al. O, E, S, C and L band silicon photonics coherent modulator/receiver. In *Proc. 2016 Optical Fiber Communications Conference and Exhibition* (*OFC*) 1–3 (IEEE, 2016).

54. Fahrenkopf, N. M. et al. The AIM photonics MPW: a highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits. *IEEE J. Select. Top. Quantum Electron.* **25**, 1–6 (2019).

55. Ji, X. et al. Ultra-low-loss on-chip resonators with sub-milliwatt parametric oscillation threshold. *Optica* **4**, 619–624 (2017).

56. Ji, X. et al. Exploiting ultralow loss multimode waveguides for broadband frequency combs. *Laser Photon. Rev.* **15**, 2000353 (2021).

57. Ji, X., Roberts, S., Corato-Zanarella, M. & Lipson, M. Methods to achieve ultra-high quality factor silicon nitride resonators. *APL Photonics* **6**, 071101 (2021).

58. Joshi, C. et al. Thermally controlled comb generation and soliton modelocking in microresonators. *Opt. Lett.* **41**, 2565–2568 (2016).

59. Gheorma, I.-L. & Osgood, R. Fundamental limitations of optical resonator based high-speed EO modulators. *IEEE Photon. Technol. Lett.* **14**, 795–797 (2002).

Author contributions

A.R. and Q.C. conceived the link architecture and performed the initial calculations. A.R. performed the circuit simulations, designed the floorplan, and drew the top level mask layout. V.G. designed the PCBs and conducted the d.c. electro-optic characterization. B.Y.K. and Y.O. designed, simulated and characterized the Kerr comb microresonators and X.J. fabricated the Si_3N_4 chip. A.N. led the high-speed testing and data-transmission experiments with assistance from A.R., S.D. and V.G. All authors helped analyse the data and A.R. prepared the paper with input from all authors. M.L., A.L.G. and K.B. supervised the project.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41566-023-01244-7.

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41566-023-01244-7.

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

Peer review information *Nature Photonics* thanks David Moss and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.
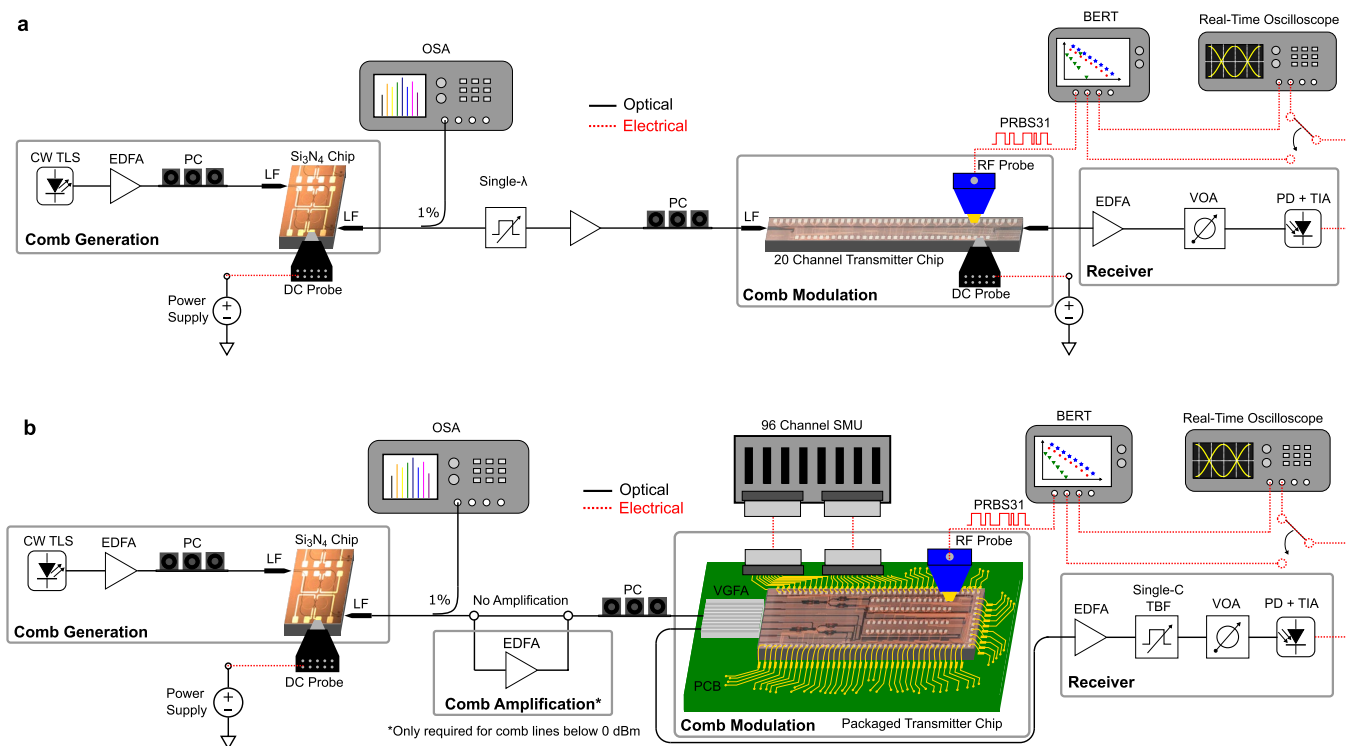
Extended Data Fig. 1 | Data transmission experimental setups. a, Experimental setup used for the single bus data transmission demonstrations. b, Experimental setup used for the packaged 32-channel integrated transmitter data transmission experiments. CW TLS: continuous-wave tunable laser source, EDFA: erbium-doped fibre amplifier, PC: polarisation controller, LF: lensed fibre, Si$_3$N$_4$: silicon nitride, DC: direct current, OSA: optical spectrum analyser, VGFA: v-groove fibre array, SMU: source measurement unit, PCB: printed circuit board, RF: radio frequency, TBF: tunable bandpass filter, VOA: variable optical attenuator, PD: photodetector, TIA: trans-impedance amplifier, BERT: bit-error rate tester, PRBS: pseudo-random bit sequence.
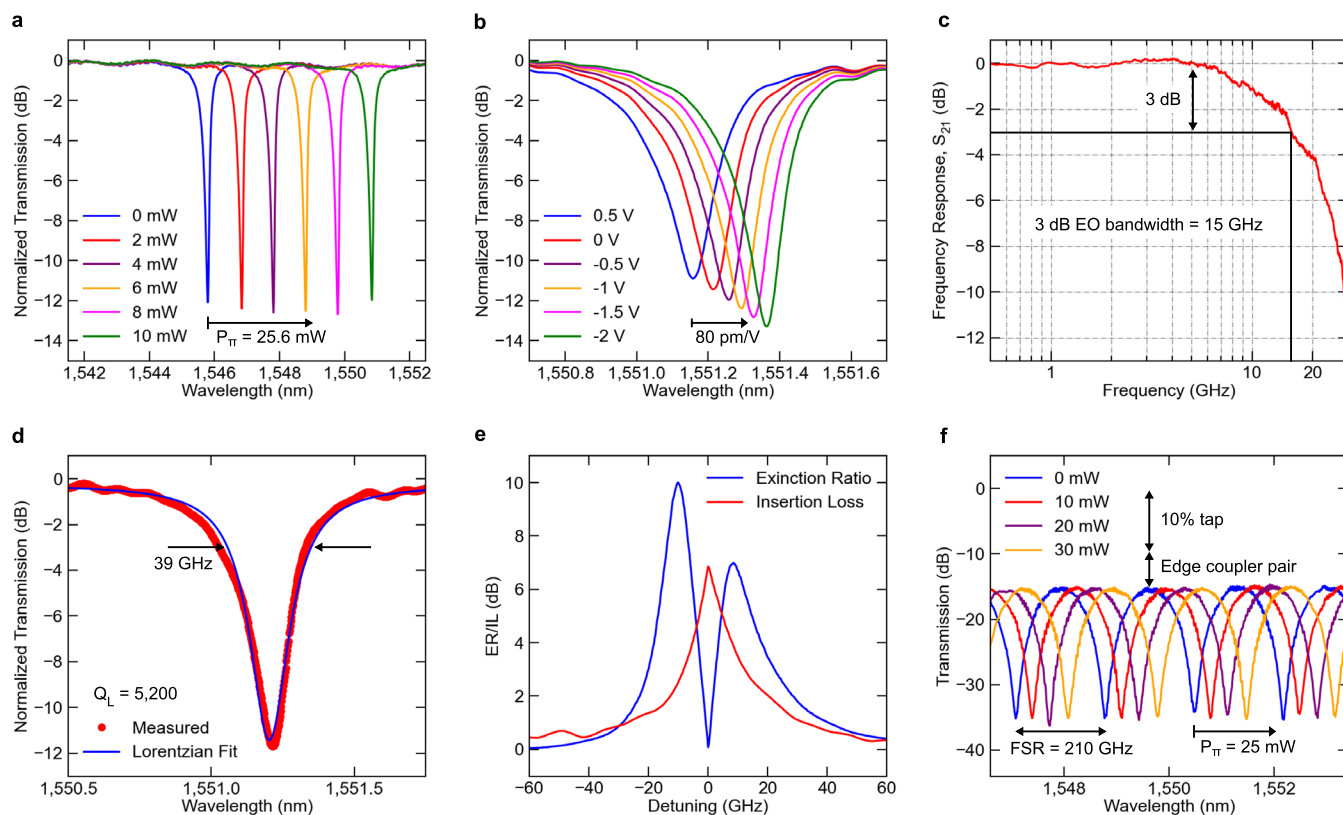
Extended Data Fig. 2 | Experimental characterisation of photonic devices.
a, Thermo-optic response of the microdisk heaters as a function of dissipated power showing a tuning efficiency of 0.5 nm/mW (P_π = 25.6 mW). b, Modulator depletion response displaying an ultra-high modulation efficiency of 80 pm/V. c, Measured S_{21} electro-optic response of the microdisk modulator with a 0 V DC bias demonstrating a 3 dB bandwidth of 15 GHz. d, Measured resonance of the modulator with a Lorentzian fit indicating a loaded quality factor (Q_L) of 5,200. e, Extinction ratio (ER) and insertion loss (IL) tradeoff for different biasing points of the modulator with a 1.3 V peak-to-peak swing. The asymmetry is due to the increasing Q_L with reverse bias, which narrows the resonance linewidth and increases the static extinction ratio as it is driven closer to critical coupling. f, Interleaver response with applied power to the bottom arm phase shifter showing a tuning efficiency of P_π = 25 mW and FSR = 210 GHz. The vertical lines indicate a 10 dB loss from the 10% tap and a 5 dB loss between the edge coupler pair for an inferred coupling loss of 2.5 dB/facet.

a

b
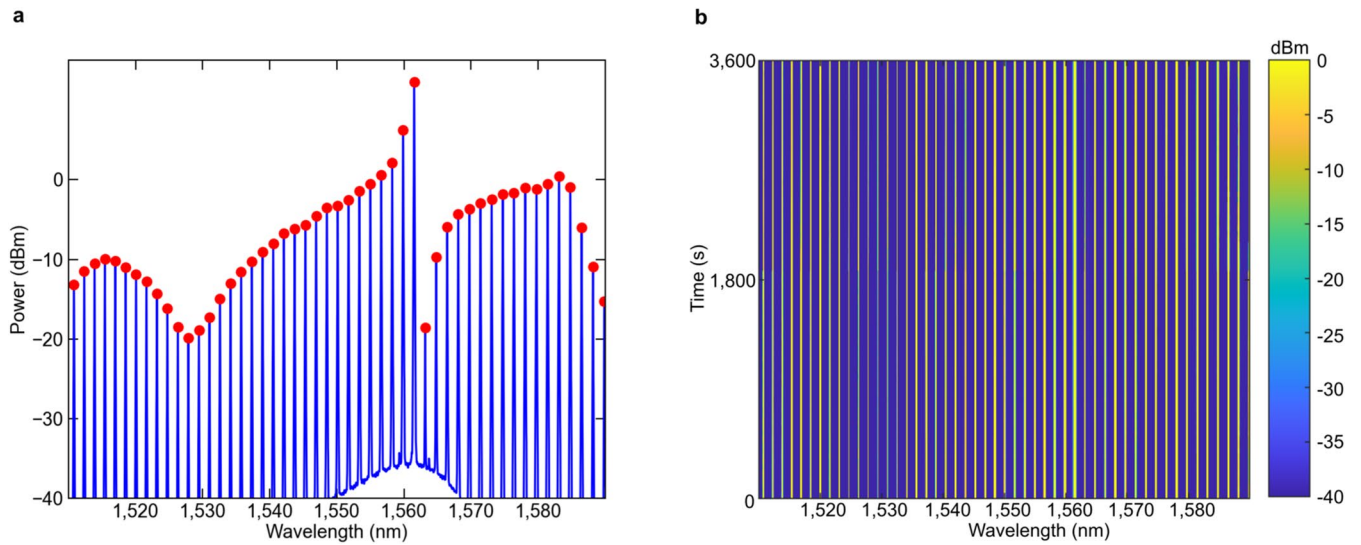
Extended Data Fig. 3 | Comb stability over time. a, Measured comb spectrum in the low noise state with peaks annotated. b, Measured power per line of comb in a, measured every 30 seconds for 1 hour with an optical power scale saturation of 0 dBm for visual clarity. No closed loop thermal control is used, demonstrating long term line stability despite environmental fluctuations.

a

b
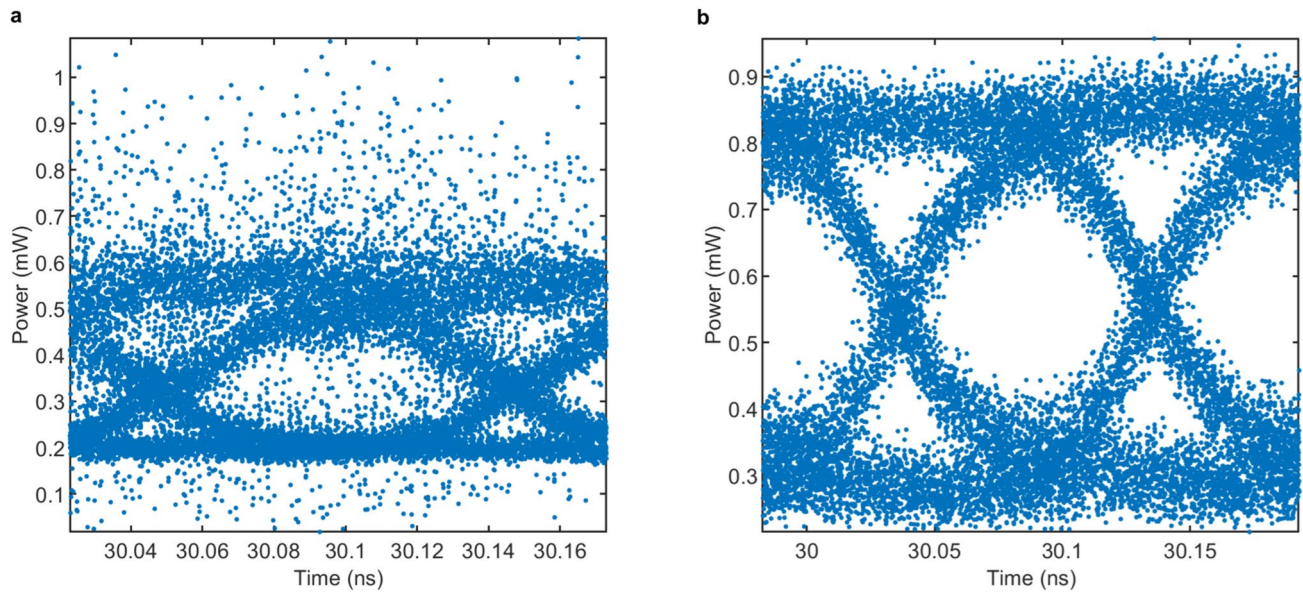
Extended Data Fig. 4 | Impact of comb noise state on signal quality.
a, Measured eye diagram of a comb line (λ = 1,559.8 nm) at 10 Gb/s when the comb is in the high-noise state, with a quality factor (Q) of ≈ 2.5 and an estimated bit

error rate (BER) of 6×10^{-3}. b, Measured eye diagram of the same comb line at 10 Gb/s when the comb is in the low-noise state, with a Q of ≈ 5.8 and an estimated BER of 4×10^{-9}.