Article

An integrated CMOS–silicon photonics transmitter with a 112 gigabaud transmission and picojoule per bit energy efficiency

Check for updates

Ke Li [1,2] ✉, David J. Thomson [1] ✉, Shenghao Liu[1,3], Weiwei Zhang [1], Wei Cao [1,3], Callum G. Littlejohns[1], Xingzhao Yan[1], Martin Ebert[1], Mehdi Banakar[1], Dehn Tran[1], Fanfan Meng[1,4], Han Du[1] & Graham T. Reed [1] ✉

The widening application of advanced digital infrastructure requires the development of communications technologies with increased data transmission rates. However, ensuring that this can be achieved in an energy-efficient way is challenging. Here we report an integrated complementary metal–oxide–semiconductor/silicon-photonics-based transmitter in which a switching current is applied to the passive-equalization-network-guided silicon Mach–Zehnder modulator, rather than driving a standard Mach–Zehnder modulator with a traditional voltage swing. This approach allows the total electrical energy to be selectively distributed to different frequency components by choosing an appropriate inductance and near-end termination impedance values. With the approach, we achieve 112 gigabaud–112 gigabits per second on–off keying and 224 gigabit per second pulse-amplitude modulation with four levels–transmission with energy efficiencies below picojoules per bit, and without the need for signal-shaping functions in the data source. We also investigate the bit error rate for different electrical and optical power conditions at 100 gigabaud, including the electrical power consumption of the driver amplifier.

The development of data-intensive applications has pushed the transmission rates of optical technology within information and communication technology infrastructure to 100 gigabaud (GBd) and beyond. However, creating devices to support such speeds is challenging, with high-frequency operation often coming at the expense of power consumption for signal shaping and equalization or to mitigate problems arising from device parasitics. Thus, although increases in data capacity are essential to meet future demands, increases in power consumption are a major concern. For instance, it has been projected that information and communication technology infrastructure could use up to 21% of global electricity by 2030 (ref. 1). It is, therefore, essential to understand the practical limitations of photonics technology in terms of speed and power consumption, as well as to develop a path to power-efficient operation beyond 100 GBd.

Energy efficiency analysis of an optical transmitter is challenging and does not always provide a clear picture of a system's performance. For example, the estimations of power consumption based on device capacitance and/or the required drive voltage only correspond to the

[1]Optoelectronics Research Centre, University of Southampton, Southampton, UK. [2]Present address: Peng Cheng Laboratory, Shenzhen, China. [3]Present address: Crealights Technology Co. Ltd, Suzhou, China. [4]Present address: Semtech Semiconductor, Romsey, UK. ✉e-mail: kl@ecs.soton.ac.uk; D.Thomson@soton.ac.uk; G.Reed@soton.ac.uk

energy associated with the modulation mechanism itself and do not include the power consumption of the broadband amplifier needed to drive these devices[2–18]. Even if these small figures (femtojoules per bit) are considered as the theoretical energy efficiency limit of photonics devices, a broadband amplifier's output power efficiency (the ratio of output power to overall power consumption) could be <4% (ref. [19]) for millimetre-wave operation (Supplementary Section 1 provides a detailed analysis).

Similarly, digital signal processing (DSP) techniques are often used to push the throughput boundary of optical transmitters, but the power consumption of the DSP module is often omitted. Trade-offs between the bit error rate (BER) and electrical energy efficiency based on the required voltage swing are also sometimes made without considering the parasitic effects that arise from electro-optic integration or the nonlinearity of an actual electrical amplifier[20–22]. In some cases, amplifiers are designed to meet the needs of photonics devices without the consideration of how much electrical energy has contributed to the optical modulation itself and how much energy could be saved if photonics and electronics devices are synergistically designed.

Focusing on the relationship between the optical link budget and the power budget of electrical devices, multidimensional models of these design considerations have been built[23–26], which have tried to estimate the performance boundary for optical transceiver links. In particular, it has been shown that the power consumption from the driver amplifier can dominate the overall power breakdown[25], and models have demonstrated that the driver's energy efficiency gets worse when operating at a higher data rate[26]. These models indicate the importance of understanding the performance boundaries of an integrated optical transmitter. However, these models are only valid with certain types of circuit topology when operating within a limited speed range, and practical circuits could be dramatically different when the optical transmitter is running at a different speed with complex modulation formats (Supplementary Section 2 provides a detailed analysis).

In this Article, we report an all-silicon optical transmitter platform that is based on 28 nm bulk complementary metal–oxide–semiconductor (CMOS) and silicon photonics, and can achieve 112 GBd–112 gigabit per second on–off keying (OOK) and 224 gigabit per second pulse-amplitude modulation with four levels (PAM-4)–transmission. We examine the power consumption limits at 112 GBd and demonstrate energy efficiency below picojoule per bit without the need for conventional pre-emphasis or signal shaping in the data source. Minimizing the power consumption of the driver amplifier could directly degrade the signal-to-noise ratio (SNR) and therefore the BER. Therefore, we also experimentally investigate the trade-off between electrical energy efficiency and average received optical power for a full transceiver link.

Concept

Design considerations for optical modulators, especially silicon Mach–Zehnder modulators (MZMs), have been thoroughly analysed for decades[27,28]. Solutions and circuit topologies for electrical broadband amplifiers are even more mature, dating back to the 1940s (ref. [29]). However, to date, similarities between these two modules have not been explicitly evaluated, to the best of our knowledge.

The pre-condition for electrical amplifiers to properly operate is the appropriate d.c. biasing of its internal transistors. As indicated in Fig. 1a, the d.c. voltage difference between the gate (V_G) and source nodes (V_S) should be larger than the threshold voltage (V_{th}) of the transistor (expressed as $V_{GS} \geq V_{th}$). Similarly, as shown in Fig. 1b, for the depletion-mode p–n-junction-based silicon optical modulators, the pre-condition is that they should be reverse biased, that is, the n node is at a higher voltage than the p node. Tuning the d.c. bias considerably changes the intrinsic bandwidth of both transistor and modulator. For the 28 nm CMOS process, the n-type metal–oxide–semiconductor (NMOS) transistor transition frequency (f_t) could reach 270 GHz,

with the gate (V_G) and drain (V_D) biased at 0.7–0.8 V. Similarly, the bandwidth of a very short length of the modulator (in the depletion mode) developed in this work was found to be 70 to 160 GHz, with a reverse bias in the range between 1 and 8 V (see the 'Photonics device modelling' section).

For an NMOS-based broadband common-source amplifier, as long as the transistors are properly d.c. biased, they have a fixed gain–bandwidth product with a larger gate width providing larger voltage gain but smaller bandwidth. When the electrical amplifier sees a low-impedance load and delivers a usable voltage gain, the typical 3 dB bandwidth of the amplifier is within the range of tens of gigahertz, although the intrinsic frequency of the NMOS transistor itself could reach several hundreds of gigahertz. A similar phenomenon occurs with a silicon MZM. With a given input-voltage swing (V_{in}), the extinction ratio (ER) is proportional to the phase-shifter length, but the electro-optic (EO) bandwidth is inversely proportional. When the silicon MZM delivers a usable modulation depth with a practical input-voltage swing, the achievable EO bandwidth is within the range of tens of gigahertz, although the bandwidth for a very short length of the phase shifter can be much higher than 100 GHz. If the requirement for the voltage swing and modulation depth are omitted that allows the phase-shifter length to be short (~100 µm range), then EO bandwidths of more than 100 GHz are possible[30].

Furthermore, for a properly d.c.-biased electrical amplifier, normally increasing the transistor's width results in higher power consumption and a larger layout area. Similarly, increasing the phase-shifter length will lead to higher optical loss and a larger footprint. In summary, both electrical amplifier and silicon MZM exhibit similar performance trade-offs. Performance enhancement techniques and design philosophies that have matured for electrical amplifiers could, therefore, be utilized to develop silicon MZMs. Following this concept, an initial attempt[31] was to adopt an inductive network (T-coil peaking)[32] within the silicon MZM and driver amplifier interface (Fig. 1b, right). This enabled the demonstration that all-silicon optical transmitters can operate at 100 GBd OOK with a sufficient ER.

For energy efficiency improvements within optical transmitters, the intuitive solution is to again reference successful design approaches developed for electrical amplifiers. The figure of merit used to evaluate the power efficiency of a typical radio-frequency (RF) amplifier is normally defined as $P_E = P_{out}/P_{d.c.}$, where P_{out} is the power delivered to a relatively low-impedance load and $P_{d.c.}$ is the d.c. power drawn from the power supply. Over the decades, numerous circuit topologies[33] have been developed to enhance the power delivery of RF power amplifiers, and a similar story has occurred with the modulator's driver amplifier. Typical circuit solutions could be classified as the voltage mode and current mode (Supplementary Section 3). However, a fact that must be underlined is that when an electrical amplifier is integrated with a photonics modulator, the output of the optical transmitter is not in the electrical domain but instead is the effective optical modulation depth, which could be expressed as the ER or optical modulation amplitude (OMA) at the specific data rate in question. Therefore, for energy efficiency improvements within optical transmitters, innovation on the circuit topologies of the driver amplifier is important, but it is even more critical to ensure that the electrical energy actually contributes to optical modulation.

However, optical modulator designers usually target an absolute figure, such as a small V_π or large EO bandwidth, rather than considering how the electrical energy is dissipated within the modulation procedure itself. Although many electrical integrated circuit (IC) engineers design the driver in conjunction with photonics devices, they are often guided with those top-level design parameters, often including target peak-to-peak voltage swing, output impedance, linearity and electrical bandwidth, but have not always questioned how to make the optical modulation mechanism itself more efficient. For example, the generation of a high-voltage-swing (V_{in}^+/V_{in}^-) signal at more
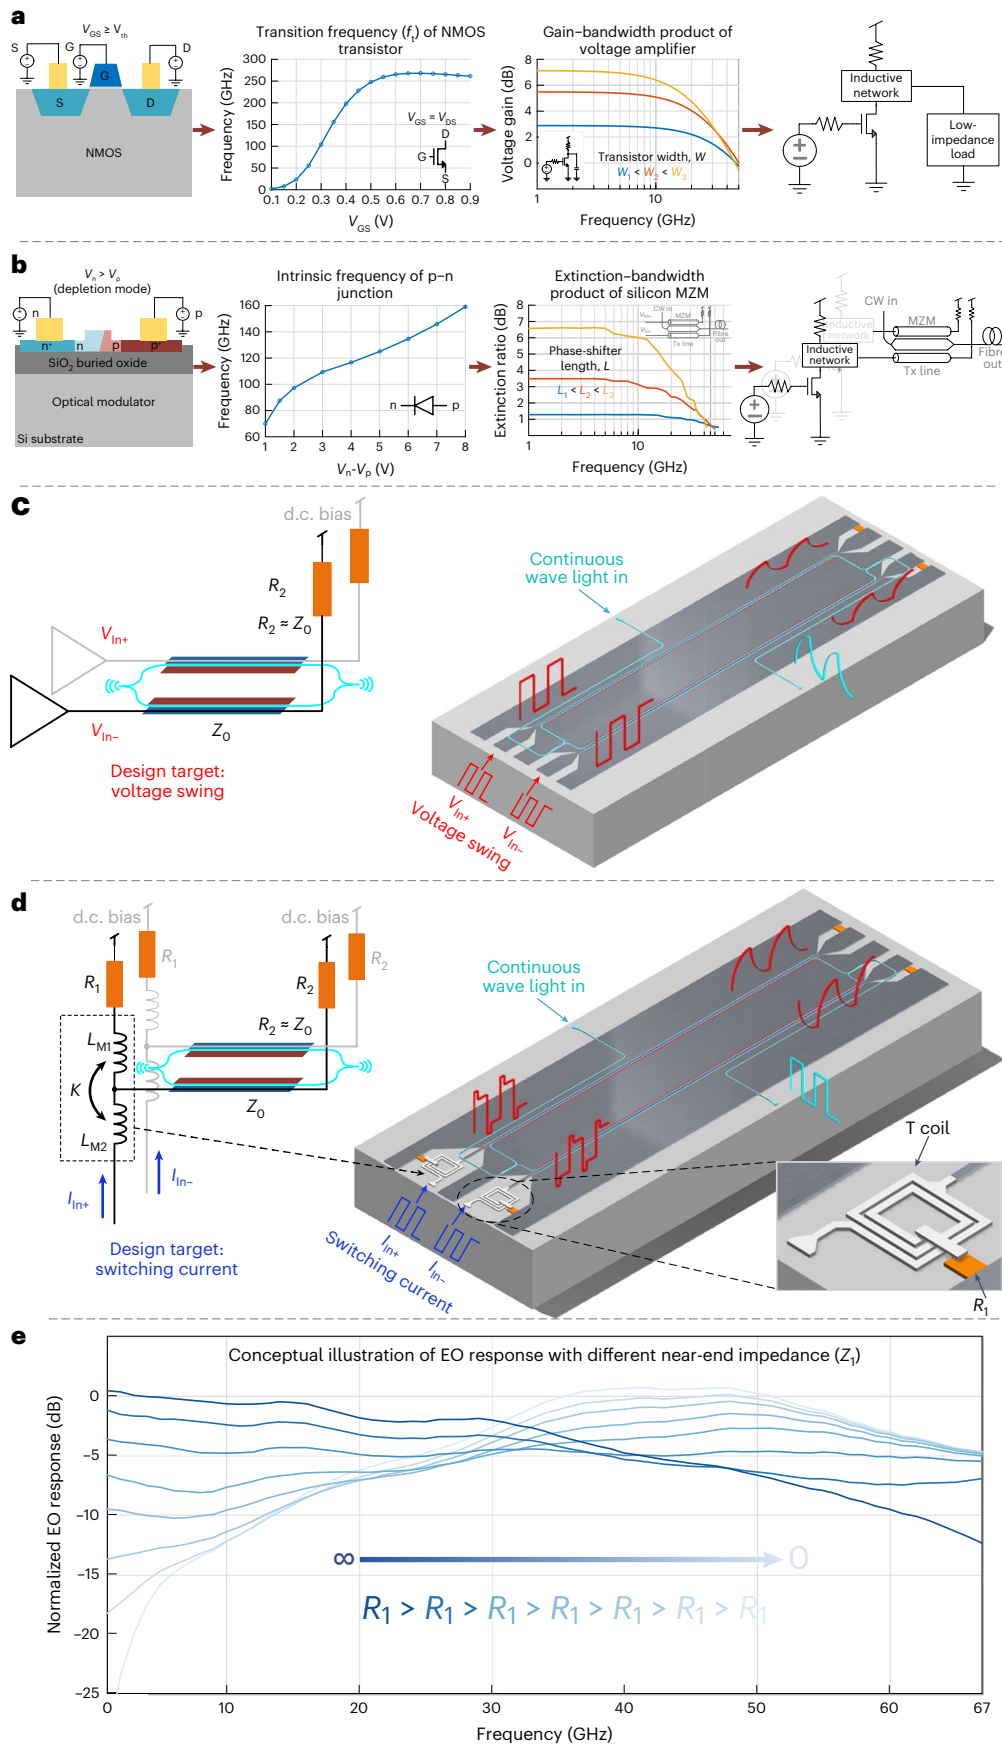
Fig. 1 | Illustration of the proposed design concept. a,b, Similarity of the design considerations among a typical broadband amplifier (**a**) and a silicon MZM (**b**). **c**, Circuit topology and behaviour of the traditional scheme. **d**, Circuit topology and behaviour of the proposed scheme. **e**, Conceptual illustration of the EO response with different near-end impedances.

than 100 GBd already requires a power-hungry driver. As depicted in Fig. 1c, this signal then suffers notable attenuation when propagating along the travelling-wave modulator electrodes with higher attenuation at higher frequencies, resulting in a bandwidth limitation. From an energy efficiency perspective, most of the energy associated with these high-frequency components is actually dissipated within the electrodes rather than contributing to optical phase change.

In contrast, perhaps a better approach is to consider the phase shifter, inductive network, far-end termination (R_2) and near-end termination (R_1) as an integrated optoelectronic device and apply a pair of switching currents (I_{in}^+/I_{in}^-). As shown in Fig. 1d, the electrical energy can then be distributed to different frequency components by changing the near-end termination impedance (R_1) and dimensions of the inductive network. If the inductive peaking network dimensions are fixed, as conceptually illustrated in Fig. 1e, the near-end termination impedance (R_1) could then be varied from ∞ to 0. When the near-end termination (R_1) is ∞, the effective circuit diagram is similar to the cases where an open-drain amplifier (open-collector amplifier for a bipolar CMOS process) is integrated into an MZM[34–37]. This is normally considered a power-efficient design since the driver amplifier only sees one termination load (R_2). However, as shown in Fig. 1e (dark-blue curve), most of the energy is distributed to low frequencies. For example, in another work[35], an additional high-frequency peaking function needed to be incorporated within the pre-driver stages, which then consumed additional power.

Some common design guidelines are still valid, including effective index matching among optics and electronics and impedance matching between the far-end termination (R_2) and characteristic impedance (Z_0) of the coplanar waveguide electrodes (CPW). The major change in the design philosophy is that we are neither targeting the absolute peak-to-peak voltage swing (V_{pp}) generated from the driver amplifier or the impedance matching between the near-end termination (R_1) and characteristic impedance (Z_0) of the CPW. Instead, within the amplifier design, the design targets are the values and signal integrities of the switching currents ($I_{in}^+ + I_{in}^-$) that are applied to the integrated optoelectronic device. Within the modulator, besides optical phase-shifter optimization itself, the value of the near-end termination (R_1) and inductive peaking network dimensions are then co-designed based on the RF loss profile of the modulator so that the majority of the energy is delivered across the modulator, causing an optical phase change over the desired frequency band.

Effectively, the combination of an inductive network, a far-end termination (R_2) and a near-end termination (R_1) could be considered as a passive equalization approach. However, the most meaningful guideline to the design is that it would be more economical and reasonable to implement the inductive networks[38] and termination resistors within the photonics chip rather than the CMOS chip, as they are compatible with standard backend processes and occupy a footprint in the micrometre range. When targeting different applications in different frequency bands, the optical modulator design involves not only optimization of the phase shifter but also a careful selection of the termination resistor values and optimization of the inductive network. The electrical driver provides the required switching current and should also be designed with the knowledge of the EO integration process and related parasitics arising from device packaging.

Experimental setup

To justify the proposed design philosophy and investigate the energy efficiency limits of all-silicon optical transmitters at more than 100 GBd, a set of silicon photonics MZMs have been fabricated (see the 'Photonics device fabrication' section) within the same wafer that not only covers a range of phase-shifter lengths (1.00, 1.27, 2.00 and 2.47 mm) for OOK but also combines two different phase-shifter lengths (1.27 mm + 2.47 mm) as a segmented modulator for PAM-4 (Fig. 2g). Regular stand-alone MZMs with phase-shifter lengths of 1.00 mm and 2.00 mm were also

fabricated to extract the fundamental performance parameters (optical insertion loss and $V_\pi \times L$) and compared with the devices (Fig. 2e,f) that are co-packaged with the driver amplifiers.

Despite the fact that we propose the electrical components of future systems to be fabricated within the photonics chip, for the purposes of this demonstration and for the ease of fabrication, the passive electrical components (Fig. 1d), including the inductive network and terminations (R_1/R_2), were realized within the electrical CMOS process (TSMC 28 nm HPC+, 1P8M5X1Z1U). To accommodate the EO integration process (see the 'Integration of electronics and photonics' section) and to enable independent d.c. signal routing for each CMOS chip, the optical path (waveguide and phase shifter) of the segmented modulator has been carefully designed (Fig. 2d). To enhance the signal integrity and minimize the slot-line mode within the CPW[39,40], gold-wire-based air-bridge bonding was deployed on all the samples. These air bridges could also be implemented within the photonics chip fabrication process if multiple metal layers are available within the backend (Supplementary Section 5). The whole electronic chip is designed via a standard analogue/RF IC design flow (see the 'Design of CMOS driver chip' section).

Device characterization

The performance of the co-packaged all-silicon transmitters (1.27 mm U-shaped and 2.47 mm U-shaped design) are analysed and compared with regular stand-alone MZMs (1.00 mm and 2.00 mm). First, the device EO response is tested with different reverse-bias voltages (Fig. 3a–d). As expected, the 3 dB EO bandwidths of the two co-packaged silicon transmitters are considerably higher than the stand-alone devices. Comparing the EO response at 65 GHz, the two co-packaged devices are at −3.0 and −4.1 dB (1.27 and 2.47 mm, respectively), whereas the two stand-alone devices are at −3.6 and −8.2 dB (1.00 and 2.00 mm, respectively). The co-packaged 2.47 mm version is just 0.5 dB lower than the 1.00 mm stand-alone device, showing that the trade-off between the phase-shifter length and EO bandwidth has been significantly alleviated within the frequency range up to 65 GHz with our approach.

The primary research target was to improve the energy efficiency, and therefore, we have characterized the properties of the eye diagrams of all the devices from 5 to 112 GBd (the maximum speed of our test equipment) and sweeping the power-supply level to the CMOS chip. The recorded ER results are summarized in Fig. 3f,g, with ten selected eye diagrams highlighted in Fig. 3h. To ensure a fair comparison, the operating point of all the devices was tuned to the quadrature point (π/2), which leads to 5–30 mW power consumption in the heating element. Since the phase error in each device is different, the required heating power in each case will also be different. To fairly compare the power consumption of each device variant, the heater power is, therefore, neglected in the power consumption results (Figs. 3–5), but should normally be considered in an overall performance evaluation. These results include the overall power consumption of the CMOS driver chip only, where the current flow through the main power supply (1.5–2.5 V) is recorded as the switching current. As shown in Fig. 3f,g, when operating at 112 GBd (OOK mode), the energy efficiencies of both devices are calculated and highlighted under different switching current conditions. The best energy efficiency is as good as 0.7 pJ bit⁻¹ with the resulting eye diagram clear and open. The modulation depths of both devices are proportional to the switching current, and longer devices provide larger modulation depths for a similar amount of switching current.

For comparison, we tested the stand-alone MZMs (1 and 2 mm) using a commercial 100 GBd driver (SHF 840 A), with the output swing fixed at 2.8 V_{ppd}. The two thick red curves in Fig. 3f,g show their ER performance with the baud rate, and the eye diagrams at their highest baud rates are compared with two co-packaged devices at a similar ER (Fig. 3h). When conducting these tests, the commercial driver consumes approximately 3.6 W of power, which is at least 20 times higher
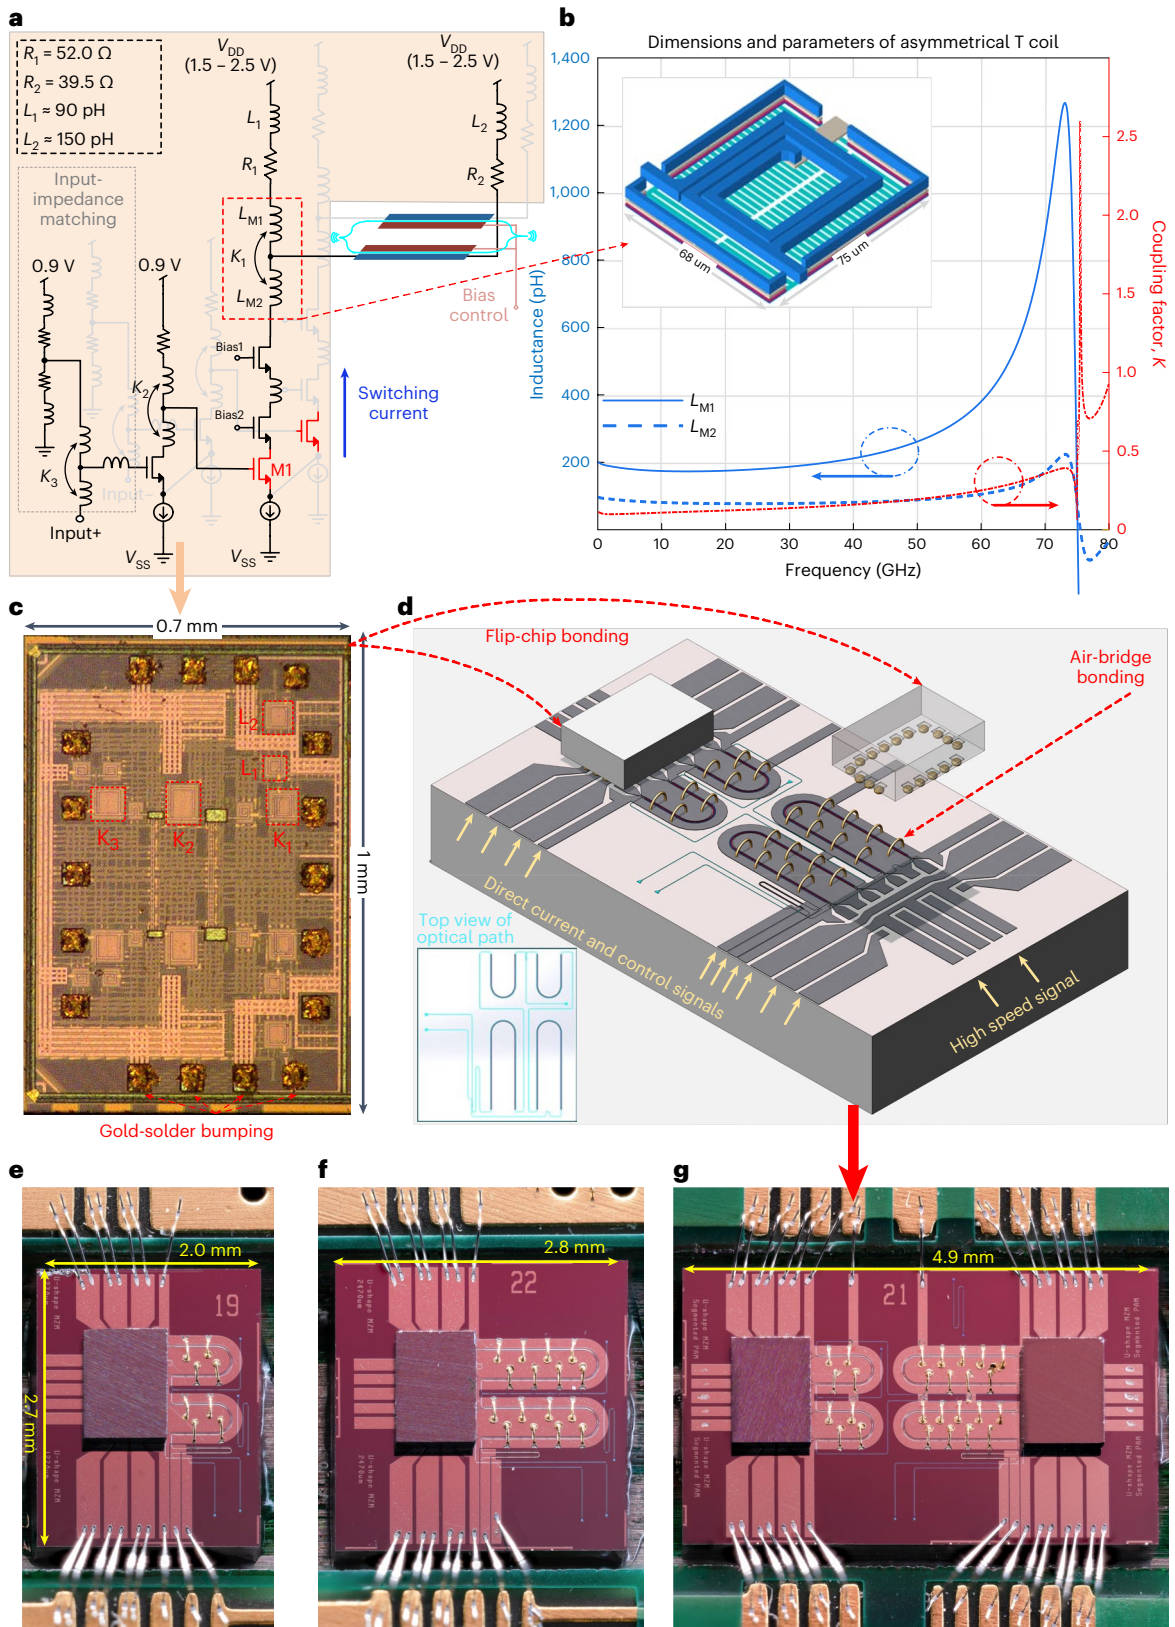
Fig. 2 | Illustration of device implementation. a, Circuit diagram and key parameters of the proposed all-silicon transmitter. **b**, Core dimensions of the inductive network. **c**, Microscopy view of the CMOS driver chip with gold-solder bumping. **d**, Conceptual illustration of the three-dimensional packaging process for the segmented U-shaped PAM-4 transmitter. **e**, Microscopy view of the co-packaged 1.27 mm U-shaped transmitter. **f**, Microscopy view of the co-packaged 2.47 mm U-shaped transmitter. **g**, Microscopy view of the co-packaged PAM-4 transmitter.
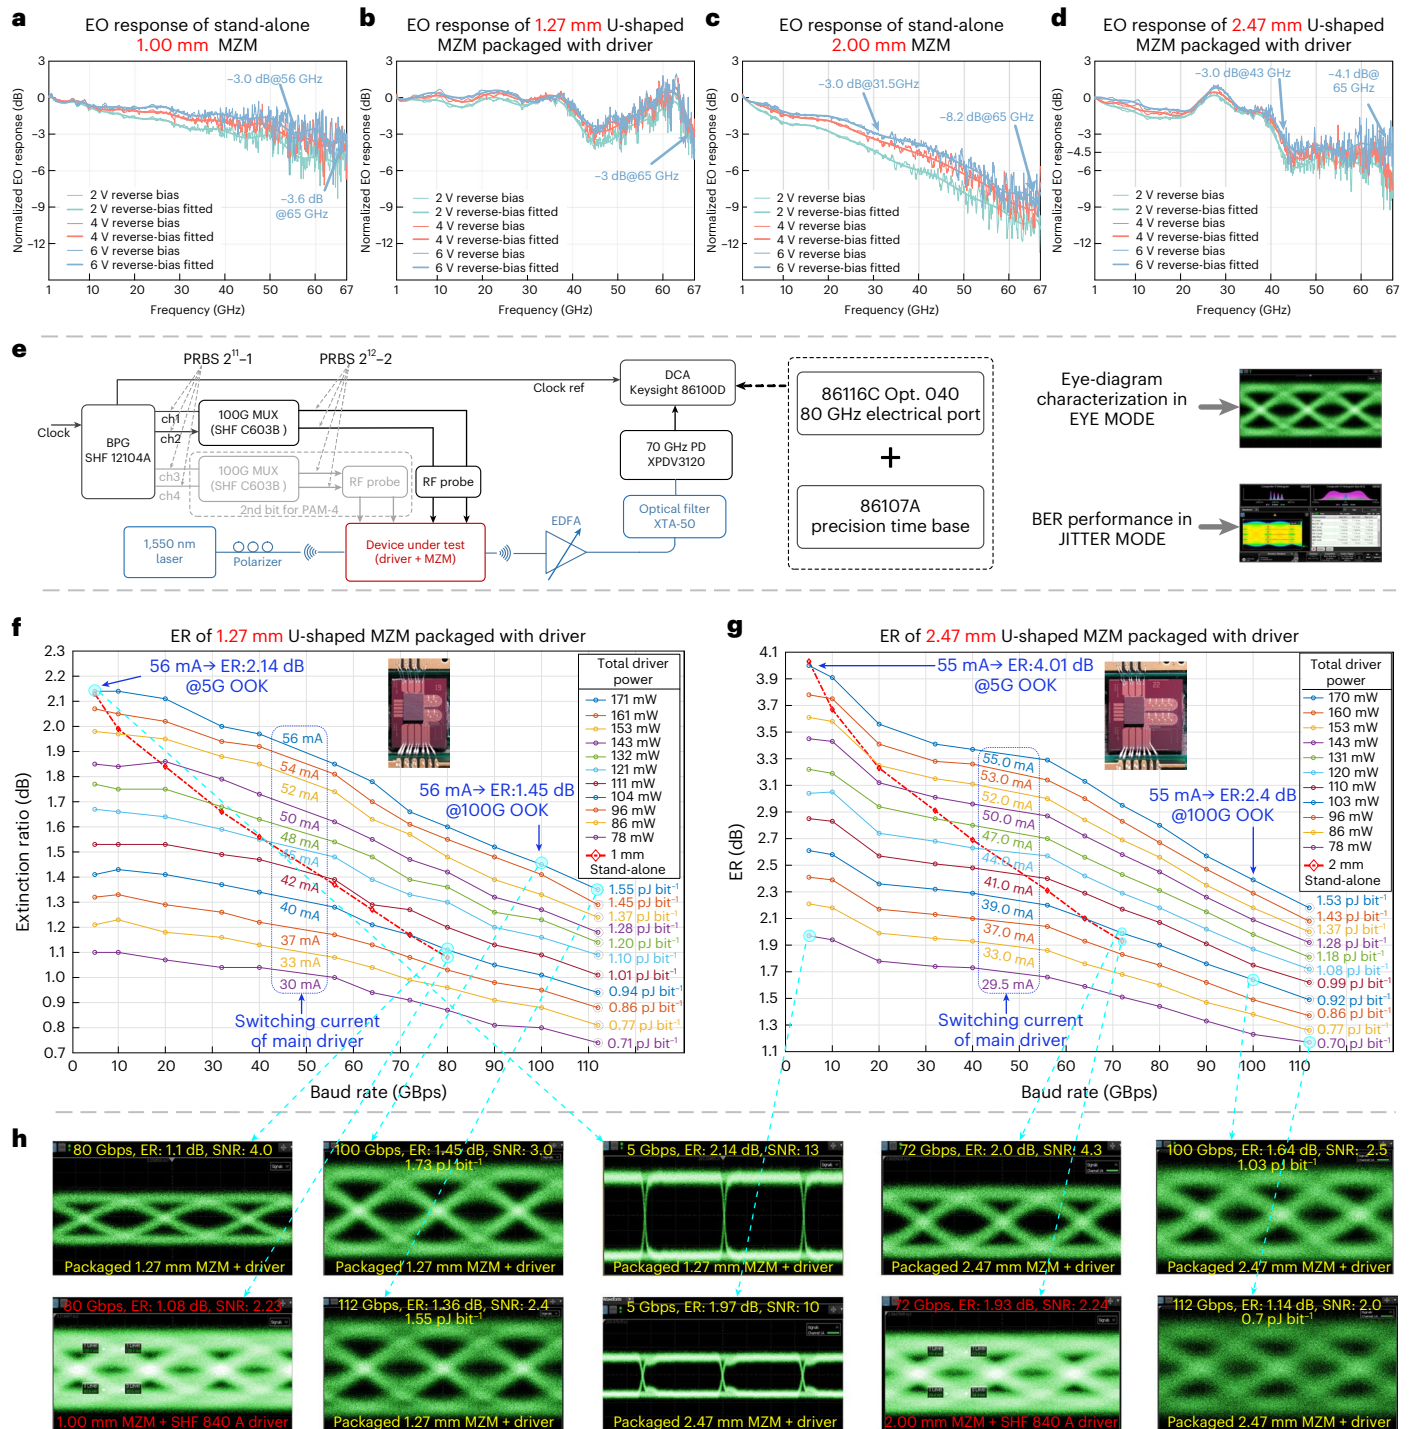
Fig. 3 | Performance of the proposed all-silicon transmitter. a, EO response of stand-alone 1.00-mm-long MZM. **b**, EO response of the co-packaged 1.27-mm-long MZM. **c**, EO response of the stand-alone 2.00-mm-long MZM. **d**, EO response of the co-packaged 2.47-mm-long MZM. **e**, Testing setup for ER and BER measurements. **f**, Comparison of the co-packaged 1.27-mm-long MZM and stand-alone 1.00-mm-long MZM under different speeds and different switching currents. **g**, Comparison of the co-packaged 2.47-mm-long MZM and stand-alone 2.00-mm-long MZM under different speeds and different switching currents. **h**, Selected eye diagrams for the co-packaged and stand-alone devices under different testing conditions.

than the driver developed in this work. It is obviously unfair to compare the absolute energy efficiency between the commercial driver and the EO co-packaged circuit developed in this work, as the commercial driver is capable of providing a higher voltage swing and incorporates additional digital control modules. However, the sharp ER roll off with speed already indicates that the standard off-shelf driver is not the optimal way to integrate with a silicon modulator, whereas the passive equalization approach adopted in this work has significantly extended the baud rate of the optical transmitter without consuming additional power.

One concern with our approach is that ringing or overshooting may exist within the waveform as an impedance mismatch exists. Therefore, in Fig. 3h (middle), two 5 Gbps eye diagrams with the maximum and minimum switching currents are shown. An almost perfect
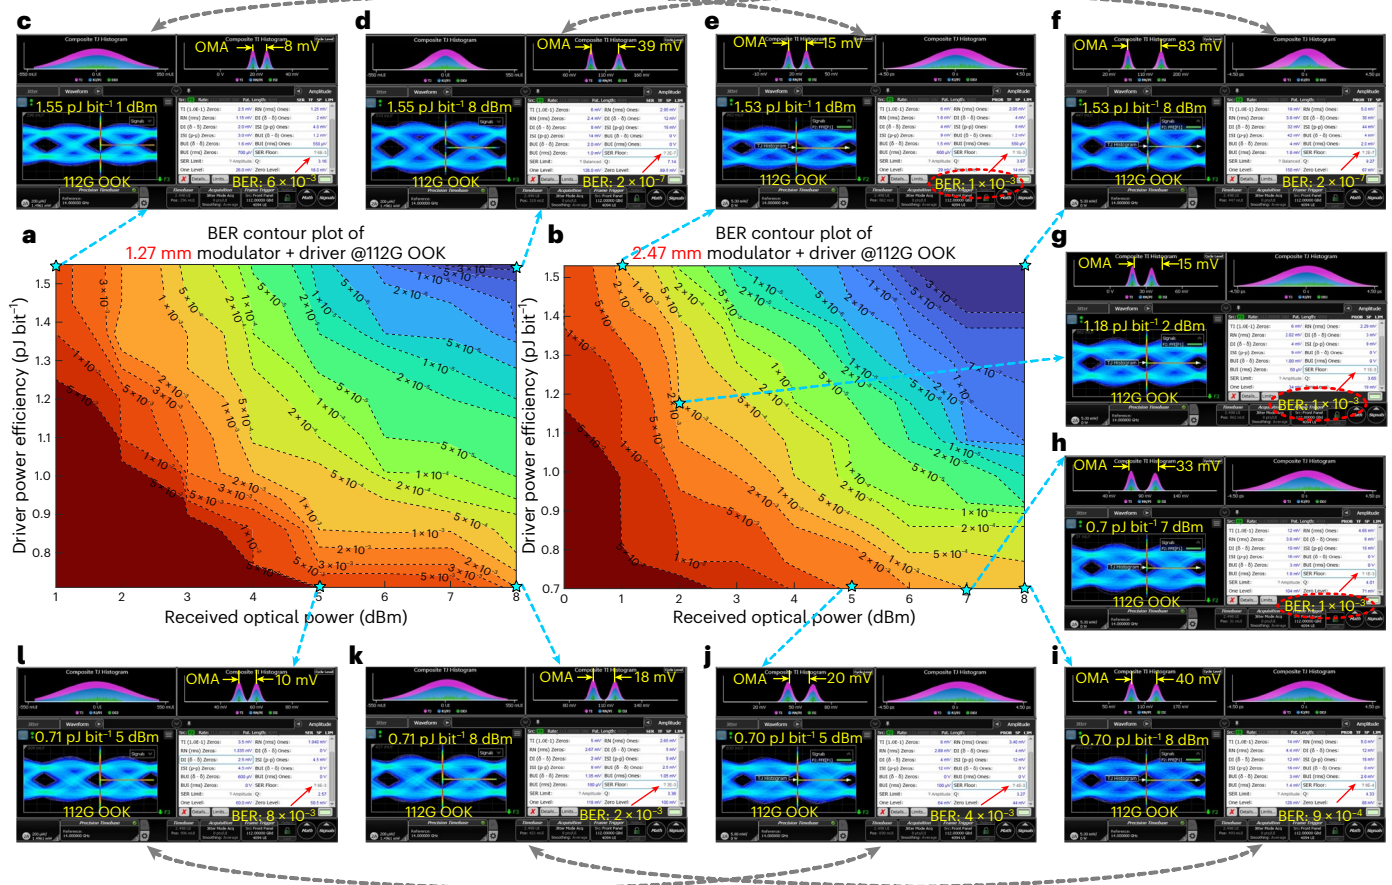
Fig. 4 | Trade-off between electrical power and optical power for 112 Gbps OOK. a, BER contour plot of the co-packaged 1.27-mm-long U-shaped MZM at 112 Gbps OOK. **b**, BER contour plot of the co-packaged 2.47-mm-long U-shaped MZM at 112 Gbps OOK. **c–l**, Eye diagrams, BER and OMA results obtained under different optical and electrical power conditions.

square-wave eye diagram demonstrates the accuracy of our analytical model and alleviates concerns regarding any near-end interface impedance mismatch.
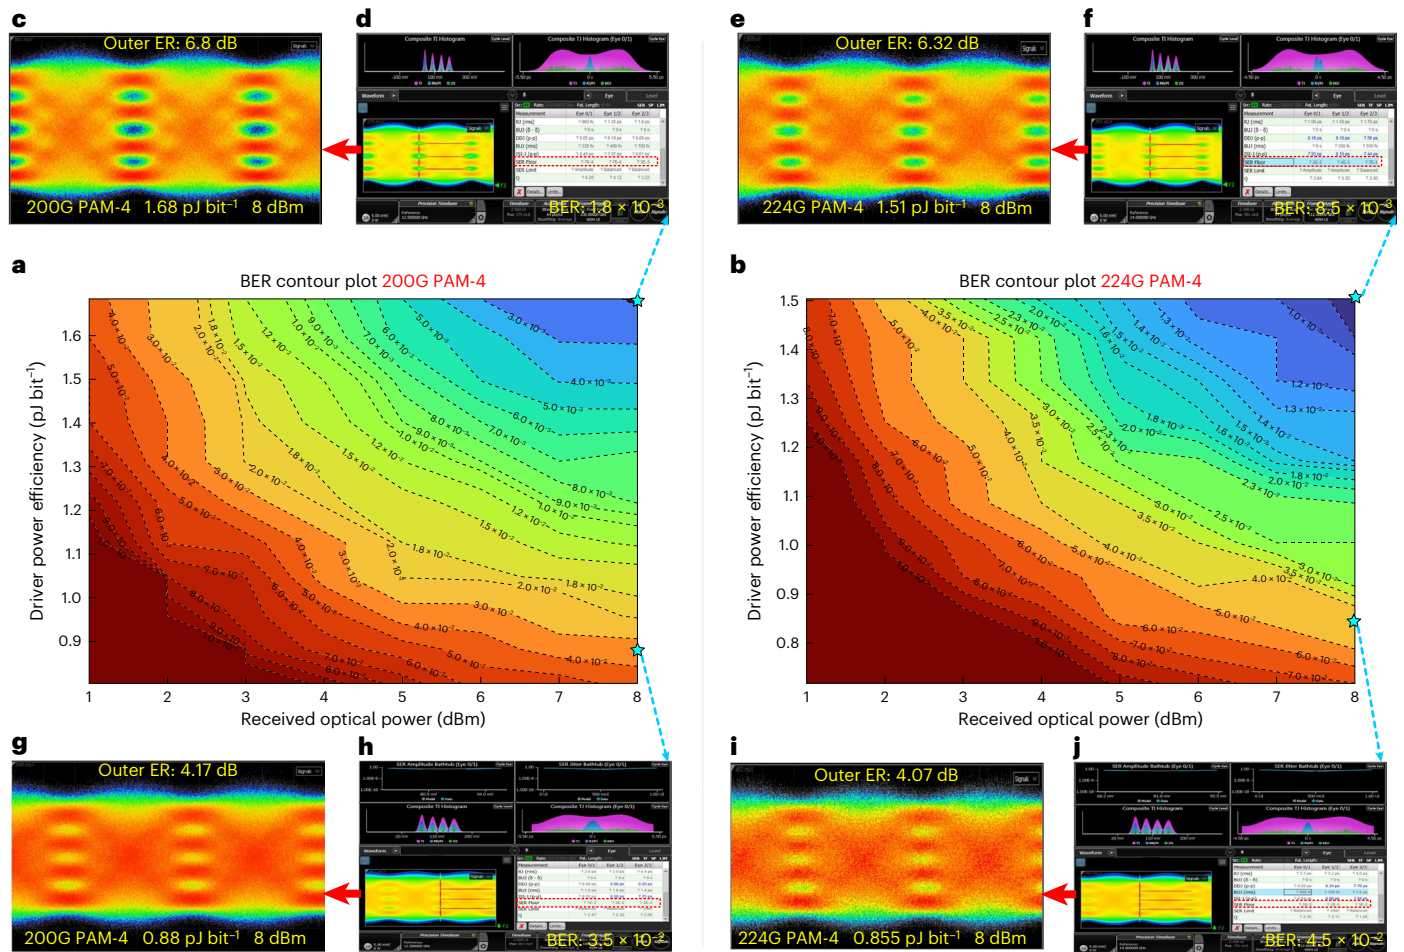
Trade-off between electrical power and optical power

The peak optical insertion loss of the co-packaged 1.27 and 2.47 mm MZMs is 3.7 and 6.9 dB, respectively, including losses from multimode interferometers, passive waveguides and phase shifters. An additional 12–13 dB of optical loss comes from the two grating couplers. Therefore, when conducting eye-diagram testing, an erbium-doped fibre amplifier (EDFA) is used to ensure that the received optical power is maintained at 8 dBm. This obviously leads to concerns about the link budget and real BER performance at the receiver side, which is directly related to the average optical power and the electrical power consumed on the transmitter side. Previous work[21,22] investigated this important relationship simply based on the estimated voltage swing applied to the modulator. However, they did not calculate the real power consumption of the driver amplifier and the obtained results also heavily relied on the DSP, a power-hungry function that is also not factored into the power calculation. In this work, we have comprehensively evaluated this relationship by sweeping the EDFA gain and the power supply of the driver to show the trade-off between the received optical power (0–8 dBm) and energy efficiency (0.70–1.54 pJ bit^{-1}) of the electrical amplifier.

The results are summarized as a BER contour plot at the receiver when operating at 112 GBd (Fig. 4a,b), with ten selected eye diagrams and OMA results highlighted in Fig. 4c–l. The photodetector

(XPDV3120) used comes with a built-in 50 Ω output load and when connected to the 50 Ω electrical port of the digital communication analyser (DCA), the photodetector sees an effective load of 25 Ω, which significantly degrades the SNR. The SNR can be significantly enhanced by incorporating an integrated high-speed detector and a transimpedance amplifier[22]. Even with such an imperfect testing environment, both devices achieve a BER level down to 2×10^{-7} with an electrical energy efficiency of 1.54 pJ bit^{-1}. When the driver power is reduced to 0.7 pJ bit^{-1}, the recorded BER values are 2×10^{-3} and 9×10^{-4} for the co-packaged 1.27-mm-long and 2.47-mm-long MZMs, respectively, which are both better than hard-decision forward error correction limits (3.8×10^{-3}). Even if an additional worst-case 30 mW power (2π tuning) from the heater is included, the energy efficiency of both devices is still lower than the picojoules per bit level.

More meaningful information comes from comparing the two devices when operating under the same electrical and optical power conditions. In Fig. 4c–f,i–l, the 2.47 mm MZMs can achieve almost twice the OMA of the 1.27 mm MZM across all the power conditions and hence provides a superior BER performance in all the scenarios. This is understandable as longer devices provide a better modulation depth, and the designed inductive peaking network significantly alleviates the bandwidth drop experienced as the length is increased. However, among the different power conditions on a particular device, the operating point of the electrical amplifier has a more dominant effect on the BER performance. This phenomenon is highlighted in Fig. 4e,g,h, all of which come with the same BER level (1×10^{-3}). The minimum electrical power condition requires 33 mV OMA (Fig. 4h),

Fig. 5 | Trade-off between electrical power and optical power for 200 and 224 Gbps PAM-4. a, BER contour plot of the co-packaged segmented MZM at 200 Gbps PAM-4. **b**, BER contour plot of the co-packaged segmented MZM at 224 Gbps PAM-4. **c–j**, Eye diagrams and BER results obtained under different optical and electrical power conditions.

whereas the other two cases only need 15 mV (Fig. 4e,g). This is mainly because the transistor M1 within the electrical driver suffers a drop in transconductance (g_m) when the current flow is reduced towards the lower end (the minimum power case).

For PAM-4 testing, a second 100 G MUX supplied the data source for the second modulator segment, and a passive RF phase tuner was used to adjust its timing delay. To compensate for nonlinearity within the MZM, a slightly lower voltage swing was applied to the shorter segment. Following the similar testing procedure as with the OOK devices, the segmented PAM-4 device (Fig. 2g) was characterized over different power conditions at 200 Gbps PAM-4 and 224 Gbps PAM-4. Their BER contour plots are shown in Fig. 5a,b, with selected eye diagrams shown in Fig. 5c–j. When operating at the highest power conditions, the recorded BER is 1.8×10^{-3} (1.68 pJ bit^{-1}) for 200 Gbps PAM-4 and 8.5×10^{-3} (1.51 pJ bit^{-1}) for 224 Gbps PAM-4, which are both far better than the soft-decision forward error correction limits (2×10^{-2}). When operating with lower power consumption (<0.9 pJ bit^{-1}), the recorded BER is 3.5×10^{-2} for 200 Gbps PAM-4 and 4.5×10^{-2} for 224 Gbps PAM-4, which are both better than the 25% forward error correction limits (5.0×10^{-2}). Even if 30 mW of heater power is included, the overall energy efficiency is still better than 1 pJ bit^{-1}. The key point is that all the results shown here have been obtained without signal shaping or DSP at the transmitter side: only feed-forward equalization (FFE) has been adopted but without any other DSP within the receiver. Furthermore, if the near-end termination and inductive network had dedicated designs for the different phase-shifter lengths, we anticipate the EO bandwidth

of the longer segment (2.47 mm) could be further optimized, and the overall BER performance would be improved.

Last, Table 1 compares the performance of this work to the state-of-art results published in recent years using different types of optical transmitter. For a fair comparison, the heater power is neglected in each case. It should also be highlighted that this is the first EO-integrated optical transmitter that operates at 112 GBd with a real energy efficiency in the sub-picojoule per bit regime, and for a comparable BER level. Previous demonstrations have almost an order of magnitude high power consumption as well as require additional DSP.

Conclusions

We have reported a silicon-photonics-based modulator design approach in which the inductive networks and termination resistors are placed within the CMOS driver chip. With this approach, we achieved transmission up to 112 Gbps OOK and 224 Gbps PAM-4, with power consumption down to the sub-picojoule per bit range. We also examined the trade-off between the power consumption and BER for different optical power levels, and compared devices of different lengths. As expected, a larger power consumption leads to a superior BER performance. However, due to bandwidth compensation in the driver, longer lengths can give a better BER performance for a given power consumption. A key strength of our approach is that all the results are obtained without the need of signal shaping or DSP in the transmitter side: only FFE has been adopted within the receiver, but without any other DSP.

Table 1 | Comparison of key performance parameters for 100G+ optical transmitters

Platform and type	Data rate (bits per second format)	DSP at Tx	DSP at Rx	Driver type/ Model no.	Driver power	EO integration	Power efficiency
LiNbO$_3$ (MZI)[5]	100G OOK, 112G PAM-4	N.A.	Offline DSP, LMS equalizer, 31 taps	SHF 807	3.60 W	No	>32 pJ bit^{-1}
Polymer (MZI)[10]	200G PAM-4	Offline DSP, raised cosine filter	Offline DSP, linear equalization	SHF 804 B	2.00 W	No	>10 pJ bit^{-1}
DML[2]	256G PAM-4	No	Offline DSP, 101-tap linear, 61-tap nonlinear equalization	SHF 840 M	1.25 W	No	>4.80 pJ bit^{-1}
SOI (RRM)[16]	120G OOK, 220G PAM-4, 240G PAM-8	Offline DSP, pre-equalization	Offline DSP, NN+MLSE	SHF 807 C	3.00 W	No	>12.50 pJ bit^{-1}
GeSi EAM[13]	224G PAM-4	Offline DSP, pulse shaping/pre-emphasis	Offline DSP, FFE DFE PNLE	SHF 804 B	2.00 W	No	>8.92 pJ bit^{-1}
SOI (MZI)[43]	64G OOK, 138G PAM-4, 102G PAM-8, 408G DP-64QAM	DSP equalization	Coherent DSP	130 nm SiGe bipolar CMOS	0.66 W^a	Yes (flip-chip)	7.35 pJ bit^{-1}
SOI (RRM)[44]	112G PAM-4	No	5-tap TDECQ filter	28 nm CMOS	0.16 W^a	Yes (flip-chip)	1.43 pJ bit$^{-1\,a}$
MOSCAP (MZI)[45]	100G PAM-4	No	5-tap TDECQ filter	28 nm CMOS	0.108 W^a	Yes (flip-chip)	1.08 pJ bit$^{-1\,a}$
Plasmonic (MZI)[9]	120G OOK	No	Offline DSP, 101 tap	Bipolar	0.900 W^a	Yes (monolithic)	7.50 pJ bit$^{-1\,a}$
SOI (MZI)[31]	112G OOK	No	7-tap FFE	28 nm CMOS	0.178 W^b	Yes (flip chip)	1.59 pJ bit$^{-1\,b}$
This work SOI (MZI)	**112G OOK**	**No**	**6-tap FFE**	**28 nm CMOS**	**0.045 W^a 0.078 W^b**	**Yes (flip chip)**	**0.40 pJ bit$^{-1\,a}$, 0.70 pJ bit$^{-1\,b}$**
This work SOI (MZI)	**200G PAM-4**	**No**	**12-tap FFE**	**28 nm CMOS**	**0.107 W^a 0.177 W^b**	**Yes (flip chip)**	**0.54 pJ bit$^{-1\,a}$, 0.88 pJ bit$^{-1\,b}$**
This work SOI (MZI)	**224G PAM-4**	**No**	**12-tap FFE**	**28 nm CMOS**	**0.122 W^a 0.190 W^b**	**Yes (flip chip)**	**0.54 pJ bit$^{-1\,a}$, 0.855 pJ bit$^{-1\,b}$**

^{a}Output stage only ^{b}Power consumption includes d.c. biasing, pre-driver and output stage

Methods

Photonics device fabrication

The photonics IC was fabricated via the open-source, license-free CORNERSTONE service, which can be accessed directly (https://www.cornerstone.sotonfab.co.uk/) or via EUROPRACTICE (https://europractice-ic.com/). The devices were built following the standard multi-project wafer process on the 220 nm silicon-on-insulator (SOI) platform comprising a 220-nm-thick Si overlayer and a 2-μm-thick buried oxide layer[41]. All the lithography processes were carried out via 248 nm deep-ultraviolet projection lithography on 200 mm wafers. First, grating couplers were patterned and etched 70 nm via an inductively coupled plasma process based on SF$_6$ and C$_4$F$_8$ chemistry. Next, a low-dose p-type-doped layer was formed by ion implantation of boron. Note that to utilize the self-aligned process to ensure a repeatable p–n junction position within the waveguides[42], the entire waveguide region of the device, which was defined later in the process flow, was implanted with boron. A 200 nm SiO$_2$ layer was then deposited by plasma-enhanced chemical vapour deposition, which first acted as a hard mask for the subsequent rib waveguide etching step that left a 100-nm-thick Si slab layer; second, it was used to protect the top of the waveguides from the ensuing n-type ion implantation process using phosphorus. During this n-type ion implantation step, the wafer was tilted by 45° to move the p–n junction into the waveguide to improve the resulting device modulation efficiency. Using this so-called self-aligned process, the junction position was defined by the implantation energy, and not by a lithographic process that would introduce a variable alignment error across the wafer. Furthermore, to ensure that all the required waveguide sidewalls were implanted, irrespective of the sidewall orientation (for example, in a U-shaped modulator), six

separate implantation steps were performed, each with an additional 60° wafer rotation. Also, note that the implantation dose was increased to compensate for the p-type doping that was earlier implanted into the waveguides. Following this, high-dose p-type and n-type regions were formed by ion implantation for ohmic contacts, and the wafer was subsequently annealed to activate the doping. A 1-μm-thick SiO$_2$ top cladding layer was then deposited by plasma-enhanced chemical vapour deposition, into which vias were formed by inductively coupled plasma etching to expose the highly doped regions. A metal stack of Ti, TiN and Al was then sputtered, and finally the metal layer was etched by an inductively coupled plasma process to form the electrodes.

Photonics device modelling

Based on the process parameters (that is, doping/implantation profile, annealing temperature and dimensions of the p–n junction and waveguide) adopted in the above photonics device, the phase shifter is modelled using SILVACO technology computer-aided design. The concentration distribution profile of electrons and holes under different d.c. bias voltages are simulated, from which the perturbation map of the refraction index under different d.c. voltages is obtained by applying plasma dispersion equations in MATLAB 2021a. The change in the effective refractive index is then calculated with a mode solver in MATLAB. For transient simulations, ideal voltage step-up and step-down functions were applied to the device and the concentration distribution profiles of electrons and holes were extracted at different time steps following the voltage change. This process was repeated for different voltage steps with the p–n junction d.c. biased from 1 to 8 V. The corresponding response of the effective refractive index change was calculated again using MATLAB and the bandwidth for a very short

length of the phase shifter calculated by using equation $f_{3dB} = 0.35/\tau_{delay}$, where τ_{delay} is the averaged time difference between 10% and 90% of the rising and falling edge of the effective refractive-index step response, respectively. Furthermore, the conductance G and capacitance C of the modelled area are extracted, which are then merged with the RF loss profiles of the electrodes (from Keysight Advanced Design System) to estimate the electrical properties (that is, characteristic impedance and effective index) of the phase shifter.

Design of CMOS driver chip

When designing the electronic driver chip, the first step is to determine the characteristic impedance (Z_0) of the CPW and the RF loss profile for different phase-shifter lengths on the photonics chip. The characteristic impedance of the phase shifter is determined to be ~35–45 Ω when the p–n junction is reverse biased from 1 to 8 V. Hence, the required far-end termination is determined to be a 39.5 Ω resistor in series with a 150 pH inductor, the latter compensates the parasitic effects associated with the resistor and input–output pads. The RF loss profile of the phase shifters at different lengths are then imported into the IC design platform, in which the dimensions of inductive peaking (K1) and the value of near-end resistors (52 Ω) are determined by sweeping a range of values. The key values of the near-end and far-end terminations are highlighted in Fig. 2a, with the dimensions of the T-coil peaking highlighted in Fig. 2b. The standard cascaded current-mode logic function is designed to generate switching currents (I_{in}^+/I_{in}^-), which can be sustained within the power supply range of 1.5–2.5 V (that is, no transistors will experience the risk of voltage breakdown). A one-stage pre-driver is designed with inductive network (K2) to ensure that transistor M1 can fully switch on/off across the whole frequency band (up to 67 GHz). A 50 Ω input-impedance-matching network is designed to interface with an external signal source and provide the necessary d.c. biasing for the pre-driver. The whole schematic is shown in Fig. 2a, with the microscopy view of the CMOS chip shown in Fig. 2c. The whole electronic chip is then designed via a standard analogue/RF IC design flow (Cadence Virtuoso and Keysight Advanced Design System), in which the parasitic effects associated with all the components, including input–output pads, solder bumps and so on are appropriately modelled (Supplementary Section 4 provides the characterization and details about the performance of the stand-alone driver amplifier).

Integration of electronics and photonics

The CMOS chip was fabricated via a shared multi-project wafer run, where the choice of bonding pad was limited to standard wire-bonding pads. Therefore, standard gold studs were first introduced on the surface of the CMOS pads, with the diameter of the gold stud controlled at ~55–60 μm. The CMOS chip was then flip-chip bonded onto the photonics chip via a thermal compression process, realized by using a flip-chip-bonding machine (Finetech lambda). The three-dimensional packaged module is then embedded within a bespoke-designed printed circuit board, which provides all the power supplies and d.c. control signals. Aluminium bonding wires were then used to electrically connect the printed circuit board and aluminium metal tracks on the photonics chip. Finally, gold-wire-based air-bridge bonding was deployed on top of the CPW tracks of the photonics chip with the loop height controlled at about 150 μm and the loop gap at about 250–300 μm.

Device characterization

Two stand-alone MZMs (1.0 mm and 2.0 mm long) were characterized to extract the fundamental process parameters (optical insertion loss and $V_\pi \times L$). The modulation efficiency ($V_\pi \times L$) is 1.9 V cm when operating at 6 V reverse bias, whereas the optical insertion loss of the phase shifter is about 2.67 dB mm^{-1}.

A Keysight four-port lightwave component analyser N4273E + N5227B was utilized to characterize the EO/electro-electro response of all the devices reported in this work. When conducting the EO response testing, the lightwave component analyser is configured in the differential mode, that is, two electrical output ports and one optical input port. Therefore, the EO response results presented in Fig. 3a–d represent the ratio of the optical response with respect to the electrical differential inputs. The power level of each electrical port is set at –12 dBm, whereas the main power supply of the driver is set at the maximum operating point of 2.5 V.

During eye-diagram testing, the reverse-bias voltage of all the MZMs is set at 6 V for consistency. The testing setup is shown in Fig. 3e. The SHF bit pattern generator provides four independent PRBS11 code signals, two of which are fed into a 100 G MUX (SHF C603 B) to generate a pattern of length $2^{12} – 2$ with a single-ended voltage swing of 0.5 V and twice the data rate of each individual signal. The voltage swing generated from the 100 G MUX is then fed into the device under test (driver + MZM) via a 67 GHz ground–signal–ground–signal–ground probe. A Keysight 81606A unit provides a tunable light source, which is fed into the device via a grating coupler. Light after the output grating coupler was amplified by an EDFA (LNA-150). The optical noise was suppressed using a 5-nm-wide band-pass filter (XTA-50/W) before being fed into a 70 GHz photodetector (XPDV3120), which is directly coupled to the 80 GHz electrical port of a Keysight DCA (Keysight Infiniium DCA-X 86100D with Agilent 86116C-040 plugin module + 86107A precision time-base plugin module). The eye-diagram characterization is done by setting the DCA without any smoothing, averaging or DSP equalization.

BER testing

BER testing uses the jitter-mode analysis in the DCA. To emulate the practical scenario at the receiver side, during OOK testing, a half-baud rate filter (56 GHz in this case) was applied together with six-tap FFE during BER testing, and the FFE coefficient was kept the same during power sweeping. Similarly, within the PAM-4 testing, a half-baud rate filter (50 and 56 GHz for 200 Gbps PAM-4 and 224 Gbps PAM-4, respectively) was applied together with 12-tap FFE within the DCA.

As suggested by the technical support team of the DCA manufacturer (Keysight), the symbol-error-rate floor represents the BER performance of the received signal when operating in the OOK mode. When it comes to the PAM-4 results, three symbol-error-rate floor results represent the BER results of the upper, middle and lower eyes. The aggregate BER of the PAM-4 waveforms is then calculated as BER = 1/2 × BER$_{upper}$ + BER$_{mid}$ + 1/2 × BER$_{low}$, which is consistent with previous work[36].

It should be highlighted that the results presented in this work are limited by the bandwidth of the available testing equipment, and hence, it was only possible to conduct testing up to 112 GBd and EO bandwidth results up to 67 GHz. Theoretically, the bandwidth limit of our device should come from the self-resonant frequency of the peaking inductors utilized in the driver design, which is about 80 GHz. Therefore, we strongly believe that it is possible for our device to demonstrate 150–160 GBd transmission if such high-speed testing equipment was available.

Data availability

The data that support the plots within this paper and other findings of this study are available via Figshare at https://doi.org/10.6084/m9.figshare.24116991. Source data are provided with this paper.

References

1. Jones, N. How to stop data centres from gobbling up the world's electricity. *Nature* **41**, 198–211 (2018).
2. Yamaoka, S. et al. Directly modulated membrane lasers with 108 GHz bandwidth on a high-thermal-conductivity silicon carbide substrate. *Nat. Photon.* **15**, 28–35 (2021).
3. Matsui, Y. et al. Low-chirp isolator-free 65-GHz-bandwidth directly modulated lasers. *Nat. Photon.* **15**, 59–63 (2020).

4. Wang, C. et al. Integrated lithium niobate electro-optic modulators operating at CMOS-compatible voltages. *Nature* **562**, 101–104 (2018).

5. He, M. et al. High-performance hybrid silicon and lithium niobate Mach–Zehnder modulators for 100 Gbit s⁻¹ and beyond. *Nat. Photon.* **13**, 359–364 (2019).

6. Kharel, P., Reimer, C., Luke, K., He, L. & Zhang, M. Breaking voltage–bandwidth limits in integrated lithium niobate modulators using micro-structured electrodes. *Optica* **8**, 357–363 (2021).

7. Xu, M. et al. Dual-polarization thin-film lithium niobate in-phase quadrature modulators for terabit-per-second transmission. *Optica* **9**, 61–62 (2022).

8. Haffner, C. et al. All-plasmonic Mach-Zehnder modulator enabling optical high-speed communication at the microscale. *Nat. Photon.* **9**, 525–528 (2015).

9. Koch, U. et al. A monolithic bipolar CMOS electronic–plasmonic high-speed transmitter. *Nat. Electron.* **3**, 338–345 (2020).

10. Lu, G. W. et al. High-temperature-resistant silicon-polymer hybrid modulator operating at up to 200 Gbit s⁻¹ for energy-efficient datacentres and harsh-environment applications. *Nat. Commun.* **11**, 4224 (2020).

11. Ogiso, Y. et al. Over 67 GHz bandwidth and 1.5 V Vπ InP-based optical IQ modulator with n-i-p-n heterostructure. *J. Lightwave Technol.* **35**, 1450–1455 (2017).

12. Estaran, J. M. et al. 140/180/204-Gbaud OOK transceiver for inter- and intra-data center connectivity. *J. Lightwave Technol.* **37**, 178–187 (2019).

13. Chan, D. W. U. et al. A compact 112-Gbaud PAM-4 silicon photonics transceiver for short-reach interconnects. *J. Lightwave Technol.* **40**, 2265–2273 (2022).

14. Sakib, M. et al. A high-speed micro-ring modulator for next generation energy-efficient optical networks beyond 100 Gbaud. In *2021 Conference on Lasers and Electro-Optics (CLEO)* 1–2 (IEEE, 2021).

15. Zhang, Y. et al. 200 Gbit/s optical PAM4 modulation based on silicon microring modulator. In *2020 European Conference on Optical Communications (ECOC)* 1–4 (IEEE, 2020).

16. Zhang, Y. et al. 240 Gb/s optical transmission based on an ultrafast silicon microring modulator. *Photon. Res.* **10**, 1127–1133 (2022).

17. Zhang, H. et al. 800 Gbit/s transmission over 1 km single-mode fiber using a four-channel silicon photonic transmitter. *Photon. Res.* **8**, 1176–1782 (2020).

18. Li, M., Wang, L., Li, X., Xiao, X. & Yu, S. Silicon intensity Mach–Zehnder modulator for single lane 100 Gb/s applications. *Photon. Res.* **6**, 109–116 (2018).

19. Rito, P., Líópez, I. G., Awny, A., Ulusoy, A. C. & Kissinger, D. A DC-90 GHz 4-Vpp differential linear driver in a 0.13 μm SiGe:C BiCMOS technology for optical modulators. In *2017 IEEE MTT-S International Microwave Symposium (IMS)* 439–442 (IEEE, 2017).

20. Samani, A. et al. Silicon photonic Mach–Zehnder modulator architectures for on chip PAM-4 signal generation. *J. Lightwave Technol.* **37**, 2989–2999 (2019).

21. Jacques, M. et al. 240 Gbit/s silicon photonic Mach-Zehnder modulator enabled by two 2.3-Vpp drivers. *J. Lightwave Technol.* **38**, 2877–2885 (2020).

22. Zhao, Y. et al. Silicon photonic based stacked die assembly toward 4X200-Gbit/s short-reach transmission. *J. Lightwave Technol.* **40**, 1369–1374 (2021).

23. Georgas, M. Moss, J. L. B., Sun, C. & Stojanović, V. Addressing link-level design tradeoffs for integrated photonic interconnects. In *2011 IEEE Custom Integrated Circuits Conference (CICC)* 1–8 (IEEE, 2011).

24. Zhu, K. et al. Modeling of MZM-based photonic link power budget. In *2016 IEEE Optical Interconnects Conference (OI)* 58–59 (IEEE, 2016).

25. Ahmed, A. H., Sharkia, A., Casper, B., Mirabbasi, S. & Shekhar, S. Silicon-photonics microring links for datacenters—challenges and opportunities. *IEEE J. Sel. Topics Quantum Electron.* **22**, 194–203 (2016).

26. Bahadori, M. et al. Energy-performance optimized design of silicon photonic interconnection networks for high-performance computing. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)* 326–331 (IEEE, 2017).

27. Witzens, J. High-speed silicon photonics modulators. *Proc. IEEE* **106**, 2158–2182 (2018).

28. Reed, G. T., Mashanovich, G., Gardes, F. Y. & Thomson, D. J. Silicon optical modulators. *Nat. Photon.* **4**, 518–526 (2010).

29. Ginzton, E. L., Hewlett, W. R., Jasberg, J. H. & Noe, J. D. Distributed amplification. *Proc. IRE* **36**, 956–969 (1948).

30. Han, C., Jin, M., Tao, Y., Shen, B., Shu, H. & Wang, X. Ultra-compact silicon modulator with 110 GHz bandwidth. In *Optical Fiber Communications Conference and Exhibition (OFC)* 1–3 (2022).

31. Li, K. et al. Electronic–photonic convergence for silicon photonics transmitters beyond 100 Gbps on–off keying. *Optica* **7**, 1514–1516 (2020).

32. Jun-De, J. & Shawn, H. A miniaturized 70-GHz broadband amplifier in 0.13-μm CMOS technology. *IEEE Trans. Microw. Theory Techn.* **56**, 3086–3092 (2008).

33. Johansson, T. & Fritzin, J. A review of Watt-level CMOS RF power amplifiers. *IEEE Trans. Microw. Theory Techn.* **62**, 111–124 (2014).

34. Jyo, T., Nagatani, M., Ozaki, J., Ishikawa M. & Nosaka, H. 12.3 A 48GHz BW 225mW/ch linear driver IC with stacked current-reuse architecture in 65nm CMOS for beyond-400Gb/s coherent optical transmitters. In *2020 IEEE International Solid-State Circuits Conference—(ISSCC)* 212–214 (IEEE, 2020).

35. Jyo, T., Nagatani, M., Ogiso, Y., Yamanaka, S. & Nosaka, H. An over 67-GHz bandwidth 21-dB gain 4.5-Vppd linear modulator driver for 100-GBd coherent optical transmitter. *IEEE Microw. Wireless Compon. Lett.* **31**, 705–708 (2021).

36. Xiong, C. et al. Monolithic 56Gb/s silicon photonic pulse-amplitude modulation transmitter. *Optica* **3**, 1060–1065 (2016).

37. Li, K. et al. Co-Design of electronics and photonics components for silicon photonics transmitters. In *2018 European Conference on Optical Communication (ECOC)* 1–3 (IEEE, 2018).

38. Gill, D. M. et al. Distributed electrode Mach-Zehnder modulator with double-pass phase shifters and integrated inductors. *Opt. Express* **23**, 16857–16865 (2015).

39. Hao, X. et al. Demonstration and characterization of high-speed silicon depletion-mode Mach–Zehnder modulators. *IEEE J. Sel. Topics Quantum Electron.* **20**, 23–32 (2014).

40. Tu, X. et al. Silicon optical modulator with shield coplanar waveguide electrodes. *Opt. Express* **22**, 23724–23731 (2014).

41. Littlejohns, C. G. et al. CORNERSTONE's silicon photonics rapid prototyping platforms: current status and future outlook. *Appl. Sci.* **10**, 8201 (2020).

42. Thomson, D. J. et al. High contrast 40Gbit/s optical modulation in silicon. *Opt. Express* **19**, 11507–11516 (2011).

43. Ahmed, A. H. et al. A dual-polarization silicon-photonic coherent transmitter supporting 552Gb/s/wavelength. *IEEE J. Solid-State Circuits* **55**, 2597–2608 (2020).

44. Li, H., Hsu, C.-M., Sharma, J., Jaussi, J. & Balamurugan, G. A 100-Gb/s PAM-4 optical receiver with 2-tap FFE and 2-tap direct-feedback DFE in 28-nm CMOS. *IEEE J. Solid-State Circuits* **57**, 44–53 (2022).

45. Talkhooncheh, A. H. et al. A 2.4pJ/b 100Gb/s 3D-integrated PAM-4 optical transmitter with segmented SiP MOSCAP modulators and a 2-channel 28nm CMOS driver. In *2022 IEEE International Solid-State Circuits Conference (ISSCC)* 284–286 (IEEE, 2022).

Author contributions

K.L. conceived the idea and led the research. G.T.R. supervised the project and managed the silicon photonics group. K.L. modelled and designed the electronic devices. D.J.T., S.L., W.C. and K.L. modelled and designed the photonics device. D.J.T., X.Y., C.G.L., M.B., W.Z., M.E., D.T. and H.D. conducted the photonics device fabrication. K.L. conducted the EO integration process. K.L., W.Z. and W.C. conducted the device characterization and data transmission testing. K.L. and F.M. designed the printed circuit board. W.C. took photos of the fully packaged devices and K.L. composed the three-dimensional images. K.L. analysed the data and co-wrote the manuscript with D.J.T., C.G.L. and G.T.R. All authors commented on the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41928-023-01048-1.

Correspondence and requests for materials should be addressed to Ke Li, David J. Thomson or Graham T. Reed.

Peer review information *Nature Electronics* thanks Shiyoshi Yokoyama 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.