ARTICLE

Received 6 Dec 2013 | Accepted 29 Apr 2014 | Published 11 Jun 2014

DOI: 10.1038/ncomms5008

OPEN

An ultralow power athermal silicon modulator

Erman Timurdogan¹, Cheryl M. Sorace-Agaskar¹, Jie Sun¹, Ehsan Shah Hosseini¹, Aleksandr Biberman¹
& Michael R. Watts¹

Silicon photonics has emerged as the leading candidate for implementing ultralow power wavelength-division-multiplexed communication networks in high-performance computers, yet current components (lasers, modulators, filters and detectors) consume too much power for the high-speed femtojoule-class links that ultimately will be required. Here we demonstrate and characterize the first modulator to achieve simultaneous high-speed (25 Gb s^{-1}), low-voltage ($0.5 V_{pp}$) and efficient 0.9 fJ per bit error-free operation. This low-energy high-speed operation is enabled by a record electro-optic response, obtained in a vertical p-n junction device that at 250 pm V^{-1} (30 GHz V^{-1}) is up to 10 times larger than prior demonstrations. In addition, this record electro-optic response is used to compensate for thermal drift over a 7.5°C temperature range with little additional energy consumption (0.24 fJ per bit for a total energy consumption below 1.03 J per bit). The combined results of highly efficient modulation and electro-optic thermal compensation represent a new paradigm in modulator development and a major step towards single-digit femtojoule-class communications.

¹Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA. Correspondence and requests for materials should be addressed to M.R.W. (email: mwatts@mit.edu).

Intra- and inter-chip interconnects, traditionally executed using parallel electrical links, are struggling to satisfy the current bandwidth, density, power consumption and cost requirements of the computing and communication industries^{1–3}. These limitations are becoming apparent in rapidly scaling massively parallel computing systems such as data centers for cloud computing, and high-performance ‘supercomputers’ used for large-scale scientific computation. While current bottlenecks within the switches and interconnects of large-scale computing systems prove an immediate need for the inclusion of inter-chip optical components, on-chip (that is, intra-chip) communication links, where the bandwidth and efficiency requirements are far more severe, provide a long-term vision. An on-chip interconnecting wire that communicates one bit over an average distance of 1 mm ($\sim 10\%$ of the die size), consumes 100 fJ per bit⁴. At the 8-nm node and beyond, over one petabit per second of data transfer will be required on large microprocessor die, leading to 100 W of communications power per chip (nearly 10 times what is acceptable in practice), and providing a clear motivation for on-chip photonic interconnects to be driven into the single-digit femtojoule per bit regime with each component within the link consuming just a fraction of that total^{2–4}.

Optical communications based on spatially division multiplexed vertical cavity surface emitting laser arrays already exist in high-performance computing systems³. Vertical cavity surface emitting laser-based links, which utilize multimode fibres, meet requirements in terms of cost and power consumption at low data rates ($\sim 100 \text{ Gb s}^{-1}$). However, they are already struggling to satisfy future high-bandwidth ($\geq 1 \text{ Tb s}^{-1}$), high-efficiency ($\sim 1 \text{ fJ}$ per bit) requirements owing to their single-channel-per-fibre nature and multimode operation. A scalable alternative using wavelength division multiplexing is required. Wavelength division multiplexing systems offer high-bandwidth operation per fibre, and can be efficiently implemented by using on-chip silicon photonic modulators to encode data on a laser comb. Silicon photonics⁵, inherently complementary metal-oxide semiconductor (CMOS) compatible, offers a high-index contrast platform that allows for single-mode high-confinement waveguides that lead to compact resonators and high-speed, low-energy consumption modulation⁶.

In the silicon platform, high-performance electro-absorptive germanium-on-silicon and electro-refractive silicon modulators have been demonstrated^{7–20}. Germanium-on-silicon modulators use the strong Franz-Keldysh and Quantum-confined Stark effects, and have enabled low-power non-resonant modulators^{7,8}. However, electro-absorption devices also act as detectors and the generated photocurrent leads to an additional, significant term in the power consumption of the modulator²¹. Further, in order to achieve a large extinction ratio (ER), all germanium-on-silicon electro-absorption modulators to date have required voltage levels that are not compatible with CMOS drive levels. Finally, since electro-absorption devices are based on a band-edge effect, the material composition inherently limits the operation to a modest wavelength range. The band-edge can be engineered to enable wideband operation, but only at the expense of additional mask layers, regrowth and fabrication complexity. In contrast, silicon modulators employ the relatively weak free-carrier plasma dispersion effect²², which induces a change in carrier concentration to alter the refractive index in the device. The index change results in a phase shift, which can be translated into an amplitude response using a Mach-Zehnder interferometer or a resonant cavity. Early on-chip demonstrations of silicon modulators used poorly confined ridge waveguides to form large Mach-Zehnder interferometers ($> 500 \mu\text{m}$ in length) and were inefficient due to large drive voltages and/or currents and high device capacitance^{9,10}. In order to enhance the electro-optic

effect, high-confinement, compact silicon resonant structures were proposed. Resonant modulators confine light in high-quality factor (Q) devices that effectively increase the optical path length and, thus, the interaction with the modulated carrier distribution. Compact devices also minimize the total device capacitance, enabling high-speed and low-power operation. Xu *et al.*¹¹ demonstrated the first silicon resonant electro-optic modulator using injection of carriers in a p-i-n junction diode. A variety of silicon resonant modulators have since been demonstrated^{13–20}. Injection-based modulators have achieved large frequency shifts, but at the expense of low-speed operation ($\sim 1.5 \text{ Gb s}^{-1}$) due to the long free-carrier lifetime ($\tau \sim 1 \text{ ns}$) within silicon diodes. Signal pre-emphasis has been used to achieve high-speed operation, but at the expense of increased power consumption and CMOS complexity¹². In order to overcome these limitations, resonant silicon modulators based on the depletion of electrons and holes have been realized and have achieved innate low-energy, low-voltage and high-bandwidth operation^{13–20}. The first depletion-based resonant modulator utilized a vertical p-n junction and achieved an error-free energy per bit of just 60 fJ per bit, paving the way for ultralow power modulators¹³. The vertical junction enabled interior p+ and n+ contacts, eliminating the need for a ridge waveguide, and preserving the hard outer wall of a step-index silicon microdisk. The hard outer wall yielded better modal confinement, allowing for more compact ($\sim 3.5 \mu\text{m}$ diameter) microdisk modulators, with larger free spectral ranges and greatly reduced total device capacitance^{13–16}. In addition, the vertical junction microdisk modulators exhibit a much larger overlap of the junction with the optical mode, a result characterized by the junction capacitance per unit volume (as discussed in the Results section), enabling electro-optic responses as large as 60 pm V^{-1} (7.5 GHz V^{-1}) (ref. 14) and 90 pm V^{-1} (11 GHz V^{-1})¹⁵. Driven in both depletion and sub-threshold forward-biased modes, these microdisk modulators achieved an error-free energy per bit of just 3 fJ at 12.5 Gb s^{-1} (ref. 14) and, an error-free $\sim 13 \text{ fJ}$ per bit modulation up to 25 Gb s^{-1} (ref. 15). These modulators have a device capacitance of 12 fF in a $1.6\text{-}\mu\text{m}^3$ junction volume¹⁴ and a device capacitance of 36 fF in a $4.6\text{-}\mu\text{m}^3$ junction volume¹⁵. The electro-optic responses were correlated with the junction capacitance per unit volume, which will be verified in the following paragraphs. When driven differentially, these modulators achieved an error-free energy per bit of just 0.9 fJ per bit at a data rate of 10 Gb s^{-1} (ref. 16). Despite the low-power operation of vertical junction modulators, ridge-based microring modulators with lateral junctions and interior and exterior contacts have persisted owing to their ease of implementation. Early demonstrations of ridge-based depletion modulators achieved operation at 10 Gb s^{-1} with an energy per bit of $\sim 50 \text{ fJ}$ per bit, an electro-optic response of 20 pm V^{-1} ($\sim 2.5 \text{ GHz V}^{-1}$) and a device capacitance of 50 fF in an $11.6\text{-}\mu\text{m}^3$ junction volume¹⁷. More recent ridge-waveguide-based microring modulators using lateral junctions have achieved an energy consumption of $\sim 7 \text{ fJ}$ per bit, an electro-optic response of 26 pm V^{-1} ($\sim 3 \text{ GHz V}^{-1}$) and a junction capacitance of 20.5 fF in a $5.4\text{-}\mu\text{m}^3$ junction volume¹⁸. However, no bit error rate information was provided. Other demonstrations at 25 Gb s^{-1} have also been made using interleaved junctions with 471 fJ per bit, an electro-optic response of 34 pm V^{-1} ($\sim 4 \text{ GHz V}^{-1}$), and a device capacitance of 184 fF in a $33.1\text{-}\mu\text{m}^3$ junction volume¹⁹ and $> 66 \text{ fJ}$ per bit without bit error rate (BER) information and with an electro-optic response of 40 pm V^{-1} ($\sim 5 \text{ GHz V}^{-1}$) and a junction capacitance of 66 fF in a $28.3\text{-}\mu\text{m}^3$ junction volume²⁰, thus highlighting the challenge of achieving low-energy error-free operation at high data rates. Yet, future communication links are expected to run at 25 Gb s^{-1} .

The difficulty in realizing high data rates in resonant modulators lies in the fact that, in bandwidth-limited modulators, operation at 25 Gb s^{-1} requires an electro-optic shift that is 2.5 times larger than at 10 Gb s^{-1} in order to sustain the performance (for example, insertion loss (IL) and extinction ratio (ER)). Given the square root dependence of the depletion width with applied voltage, doing so requires 6.25 times the voltage, resulting in ~ 39 times more energy per bit ($E_{\text{bit}} = CV^2/4$) than is required at 10 Gb s^{-1} . While some of the aforementioned results are impressive, none have achieved single-digit femtojoule-class communications at the high data rates that will be required in future on-chip communication links and all results thus far have been achieved in thermally stable environments that are not representative of a real application.

In this work, we present the design and demonstration of an athermal resonant silicon modulator that is compatible with a CMOS process, operates at high data rates and achieves the lowest total energy per bit performance to date. These results are made possible by combining interior circular contacts and a highly optimized vertical junction within a compact microdisk modulator. The circular contacting scheme enables a direct, short electrical path out to the vertical junction, minimizing the resistance and corresponding loss in depletion/accumulation response and maximizing the speed of the modulator. The vertical junction profile optimizes the optical mode overlap and maximizes capacitance per unit volume. As a result, we achieve a record 250 pm V^{-1} (30 GHz V^{-1}) electro-optic frequency response and, thus simultaneous error-free, high-speed (25 Gb s^{-1}), low-voltage ($0.5 V_{\text{pp}}$) and low energy per bit (0.9 fJ per bit) operation. Not only does the large electro-optic response allow for sub-1 fJ per bit operation, but it also simultaneously improves the IL and ER compared with prior demonstrations^{13–20}. Further, the record electro-optic frequency response can be used to counteract thermal shifts over a 7.5°C ($\sim 75 \text{ GHz}$) temperature (frequency) range by varying the DC bias point of the modulator. This adds little additional energy consumption (only 0.24 fJ per bit due to leakage current), while

modulation energy decreased to 0.79 fJ per bit, enabling 1.03 fJ per bit athermal operation. Over a 10°C temperature range, the modulator consumes 0.77 fJ per bit modulation energy and an additional 2 fJ per bit due to undesired leakage current, resulting in 2.77 fJ per bit total energy consumption. Finally, while these demonstrations were all performed at 25 Gb s^{-1} , the modulator demonstrates open-eye diagrams at data rates as high as 44 Gb s^{-1} when driven with higher voltages.

Results

Optimizing a depletion-mode modulator. In a resonant depletion-mode modulator, a voltage applied to a p–n junction changes the depletion width, which, through the plasma dispersion effect, results in a perturbation to the permittivity ($\Delta\epsilon$)²². This permittivity change induces a resonant frequency shift to each mode m of $\Delta\omega_m$. The magnitude of the shift can be derived from Poynting's theorem and is expressed as the ratio between the perturbed and total energy within the resonator²³.

$$\Delta\omega = \frac{-\frac{\omega_m}{4} \Delta\epsilon \int_{v_p} e_m^* \cdot e_m dv}{\frac{1}{2} \epsilon \int_{v_o} e_m^* \cdot e_m dv} \approx -\frac{\omega_m \Delta\epsilon v_p}{2 \epsilon v_o} \quad (1)$$

where, ϵ is the permittivity of silicon, ω_m the unperturbed resonant frequency, e_m the electric field distribution of mode number m , and v_p and v_o are the perturbed and total resonator volume, respectively. The approximation assumes the electric field, e_m , has a uniform flattop distribution within the silicon core. Furthermore, the total resonator volume can be approximated by $v_o = WHL$, where W , H and L are the width, height and length of the resonator core, respectively. Similarly, the perturbed volume is approximated by $v_p = \Delta w_d DL$, where Δw_d is the change in depletion region size at a given applied potential, D is a junction-profile-dependent parameter, which is shown in Fig. 1a and L is the length of the resonator. D is equal to H in lateral^{17,18} and W in vertical^{13–16} junction profiles, and to the ratio of the core area (HW) to half the length of the p–n junction period, L_{p-n} , ($D = 2HW/L_{p-n}$) for an interleaved junction profile^{19,20}.

Figure 1 | Junction profile comparison. (a) The lateral, vertical and interleaved junction profiles are illustrated for common depletion-mode modulators. The capacitance per unit length (C/L) is expressed below each junction profile. Frequency shifts based on equations (2), (3) and (6) are calculated as a function of applied potential for lateral, vertical and interleaved junction profiles, showing the large frequency shift of vertical junction profile. Calculations were performed at $\lambda = 1.55 \mu\text{m}$ and for a built-in potential of 0.7 V and a n- and p-type carrier concentration of $3 \times 10^{18} \text{ cm}^{-3}$. (b) The internal quality factors are plotted as a function of radius for an undoped microdisk and ridge resonators, showing high confinement and quality factor for the microdisk resonator. The TE radial mode profiles are overlaid with the resonator cross-sections. The thickness of the microdisk resonator was 220 nm . The ridge resonator had the cross-section of the modulator in ref. 18.

The relative change in permittivity ($\Delta\epsilon/\epsilon$) is related to the relative change in refractive index ($\Delta n/n$) by $\Delta\epsilon/\epsilon = ((n + \Delta n)^2 - n^2)/n^2 \approx 2\Delta n/n$. Substituting into equation (1) gives:

$$\Delta\omega = -\omega_m \frac{\Delta n \Delta w_d D}{n WH} \quad (2)$$

Equation (2) can be further simplified by relating the change in refractive index (Δn) to the acceptor- and donor-free carrier concentrations through a curve fit²² to Soref's experimental data in¹⁴.

$$\Delta n = A_{A,D} N_{A,D}^{B_{A,D}} + j C_{A,D} N_{A,D}^{D_{A,D}} \quad (3)$$

where, N_D and N_A are the donor and acceptor concentrations, respectively, and the curve fitting parameters for donors and acceptors are $A_D = -2.37 \times 10^{-23}$, $B_D = 1.08$, $C_D = 4.92 \times 10^{-26}$, $D_D = 1.2$, and $A_A = -3.93 \times 10^{-18}$, $B_A = 0.772$, $C_A = 1.96 \times 10^{-24}$, $D_A = 1.1$. Note that the imaginary part of the index change modifies the cavity loss but does not induce a real frequency shift. It is therefore the real part of the index change that matters in equation (2). Assuming, for simplicity, that $B_A \approx B_D \approx 1$, $A \approx A_A \approx A_D$ and $N_A \approx N_D \approx N$, equation (2) can be rewritten as:

$$\Delta\omega = -\omega_m \frac{AN\Delta w_d D}{n WH} \quad (4)$$

Finally, the depletion width change can be related to the device capacitance through the accumulated charge by considering $\Delta Q/L = qND\Delta w_d$ (where q is the electronic charge) and $\Delta Q \approx CV = CV_{pp}$ for $\Delta w_d \ll w_d$, where V is the applied electric potential and C is the junction capacitance, $\epsilon DL/w_d$. The final result is then:

$$\Delta\omega \approx -\omega_m \frac{A CV}{nq WHL} = -\omega_m \frac{A CV}{nq v_o} \quad (5)$$

We, thus conclude that, to first order, the frequency shift of a depletion-based modulator is directly proportional to the capacitance per unit volume of the resonator (C/v_o). Since the waveguide cross-section is generally fixed to maintain single-mode operation, it is desirable to maximize the junction capacitance per unit length to achieve the largest frequency shift. It is important to note that the depletion width and junction capacitance are a function of the applied potential. The approximate depletion width for an abrupt p-n junction as a function of voltage is given in equation (6):

$$w_d \approx 2\sqrt{\frac{\epsilon V + \phi_B}{q N}} \quad (6)$$

assuming $N_A \approx N_D \approx N$ and where ϕ_B is the built-in potential²⁴.

Figure 2 | Silicon vertical junction microdisk modulator design and characterization. (a) A three-dimensional sketch of the electro-optic silicon microdisk modulator, showing the cross-section, size, metal connections and the optical mode overlapped with the vertical p-n junction. The grey regions illustrate the depletion region in the vertical p-n junction within the microdisk and the undoped bus waveguide adjacent to the microdisk. The full doping profile, including the n+ and p+ regions as well as all of the wiring shown in the diagram was included in all calculations. The electrical equivalent diagram of the circularly contacted microdisk is also overlaid, illustrating n+, p+, p, n-doped region and vertical via resistances (R_{n+} , R_{p+} , R_p , R_n and R_{via}) as well as the junction capacitance (C_j). Owing to its cylindrical symmetry, the device has only vertical and radial resistance terms and no azimuthal resistance term, which significantly reduces the contact and device resistance. (b) A scanning electron microscopy image of the silicon microdisk modulator, revealed by dry etching the SiO_2 around the modulator to show the metal interconnect, circular contact, silicon bus waveguide and the microdisk. The signal pad, connected by short (~ 10 – $100 \mu\text{m}$) wires, is shown on the left side of the image. (c) The simulated finite element model (FEM) vertical junction carrier distribution profile is plotted as a function of applied voltage and depth. (d) The measured and simulated (green dashed curves) transmission spectra of the resonator is plotted, at $\sim 26.5^\circ\text{C}$ and with applied DC bias voltages of 0.25, 0 and -0.25 V, respectively.

Figure 3 | Energy consumption. (a) The electrical setup for the time domain reflectometry (TDR) reference measurement, configured for measuring the reflected voltage (V_{P}^-) from the open-ended on-chip pads (Z_{P}). (b) The electrical setup for the TDR reference and modulator measurement, configured for measuring the reflected voltage (V_{PM}^-), from the modulator load (Z_{M}) and identical on-chip pads (Z_{P}). The TDR input and reflected TDR output directions are depicted with + and - signs. (c) The TDR measurement of the electro-optic modulator power consumption and switching energy for a $0.5 \cdot V_{\text{PP}}$ AC-coupled drive.

The vertical junction in our modulator serves to maximize the capacitance per unit volume and thereby ensures maximal frequency shift for a given applied voltage. To illustrate the advantage of a vertical junction over the typical interleaved or lateral junctions, the resonant frequency shifts as a function of voltage are calculated using equations (2), (3) and (6), and depicted in Fig. 1a. As can be seen, the vertical junction profile provides a much larger frequency shift ($\Delta f_{\text{vertical}} \sim 30.9 \text{ GHz V}^{-1}$) at a given applied voltage than the lateral ($\Delta f_{\text{lateral}} \sim 6.9 \text{ GHz V}^{-1}$) or the interleaved ($\Delta f_{\text{interleaved}} \sim 22.7 \text{ GHz V}^{-1}$) profiles. To achieve a given frequency shift, the vertical junction profile thus requires 4.5 and 1.4 times smaller voltage drive, respectively, which translates to ~ 20 and ~ 2 times smaller energy consumption (switching energy, $E_{\text{s}} = CV^2$)²¹. Note also that the electro-optic response of the current vertical junction profile is larger than that in^{13–16,25} owing to improvements in device geometry and doping concentration, as indicated above.

However, while increasing the capacitance per unit volume in the optically active region of the resonator is helpful, reducing the overall capacitance of the device is desired to achieve the lowest switching energy. The use of an interior circular contact together with a hard outer resonator wall in the microdisk aids in doing just that by allowing for a tightly confined, compact device as shown in Figs 1b and 2a. Compared with a typical ridge-waveguide-based resonator, the overall resonator volume is minimized and the capacitance is concentrated over the optically active region of the junction as shown in Fig. 1b. The quality factor of an undoped microdisk resonator with a radius of $1.5 \mu\text{m}$ is equivalent to that of an undoped ridge resonator with a radius of $4.5 \mu\text{m}$. The sharp bending radius of microdisk resonator leads to >9 times less volume when it is compared with the typical ridge-contacted microring resonator. In addition, the interior circular contact, enabled by overlapping p+ and n+ regions within the microdisk center, also minimizes the device resistance. This is the first time such a contact has been realized within any silicon modulator. Minimizing the resistance allows for the applied voltage to be dropped directly across the optically active region of the junction, significantly improving performance as only radial- and vertical-resistive components, illustrated in Fig. 2a, are relevant, and the large azimuthal resistances, which limited the performance of prior demonstrations^{13,14,16,25}, are eliminated. Given the very short electrical paths to the vertical junction there is effectively no parasitic resistance to inhibit the electro-optic response or the bandwidth ($f_{\text{3dB}}^{\text{el}} = 1/(2\pi RC)$),

enabling much higher performance than previously possible. The low resistance and smaller overall capacitance, together with the large capacitance per unit volume, make this modulator a significant advance over prior geometries^{11–20}.

Designing a vertical junction microdisk modulator. The microdisk modulator, including the fabrication process steps (see the Methods section for details), was simulated using Synopsys's Sentaurus software suite, and the carrier concentration as a function of voltage was extracted and is shown in Fig. 2c. In this graph, the energies of the implants that form the vertical junction have been slightly adjusted to fit the experimental voltage vs spectral response curves (see the fabrication part of the Methods section for details). The depletion width can be approximated as the distance between 50% of the peak of the n and p doping concentrations. Using this approximation, the depletion width is estimated to be 124, 132 and 139 nm for applied voltages of 0.25, 0 and -0.25 V , respectively.

The radial electric field component of the optical mode, simulated using a complex, full-vectorial and cylindrical finite difference mode solver (FDM)²⁶, is shown in Fig. 2a. As can be seen, there is a large modal overlap between the vertical depletion region and the optical mode yielding a large frequency shift. In order to ensure high-Q operation, the highly doped p+ and n+ regions, which act as low-resistance contacts, were placed $1 \mu\text{m}$ away from the outer wall of the resonator where there is only negligible overlap with the optical mode. The cylindrical FDM simulations show that the 220-nm-thick silicon microdisk with its hard outer resonator wall supports a radiation-limited internal Q of over one million for diameters as small as $4 \mu\text{m}$. Here, we chose a $4.8 \mu\text{m}$ diameter microdisk owing to contact geometry and process design rule limitations, as opposed to radiation limits, which resulted in a $2.5 \mu\text{m}^3$ junction volume. The capacitance of the modulator ($C = Q/V$), including the full doping profile, n+ and p+ regions and the wiring shown in Fig. 2a–c, was estimated to be $\sim 17 \text{ fF}$ from AC Sentaurus device simulations. The compact junction volume combined with the vertical junction and the circular contact maximize capacitance per unit volume and minimize total device capacitance and device resistance. This modulator consumes an estimated switching energy ($E_{\text{s}} = CV^2$) of $\sim 4.25 \text{ fJ}$ for an AC-coupled $0.5 V_{\text{PP}}$ drive voltage. In a non-return-to-zero (NRZ) pseudo-random bit sequence (PRBS) signal, the 0–0, 0–1, 1–0 and 1–1 transitions are equally probable ($p(0-0) = p(0-1) = p(1-1) = p(1-0) = 1/4$),

and the modulator only consumes energy from the source in the 0-1 transitions ($E_{0-1} = E_s$, $E_{0-0} = 0$, $E_{1-1} = 0$, $E_{1-0} = 0$). Therefore, the estimated energy per bit ($E_{\text{bit}} = p(0-0) \times E_{0-0} + p(0-1) \times E_{0-1} + p(1-1) \times E_{1-1} + p(1-0) \times E_{1-0} = E_s/4 = CV^2/4$) is 1.06 fJ per bit. An electrical step function from -0.25 to 0.25 V was applied in the Sentaurus device simulation to the modulator, including the full doping profile and the wiring shown in Fig. 2a. The resultant step response was fitted with an exponential rise function, revealing an estimated electrical 3 dB bandwidth of ~ 35 GHz.

Characterizing the DC optical response. A continuous-wave lightwave generated by a tunable laser source was coupled to the fundamental transverse-electric (TE)-mode of the on-chip silicon waveguide using a tapered single-mode fibre. The through port spectral response as a function of applied DC bias was measured at a temperature of 26.5°C and the results are plotted in Fig. 2d alongside the numerical simulations (see the Methods section for details). The loss within the modulator was both measured and simulated to be ~ 29 dB cm^{-1} . Losses from radiation and sidewall roughness are found to be negligible compared with the loss from free-carrier absorption. The p-n junction acts as a capacitor in the reverse-biased and sub-threshold forward-biased regimes of operation. Frequency shifts of $+8$ and -7 GHz, resulting from the charging and discharging of this capacitor, were observed under applied biases of 0.25 and -0.25 V, respectively. These shifts agree well with the theoretical predictions and reveal a record 250 pm V^{-1} (30 GHz V^{-1}) response. The optical modulation bandwidth can be extracted from the full-width-half-maximum and is equal to: $f_{3\text{dB}}^{\text{opt}} = \sqrt{\sqrt{2}-1} \times f_{\text{FWHM}}^{27}$. The $\sqrt{\sqrt{2}-1}$ factor stems from the fact that the resonator stores light (to achieve data level 0) and releases the light (to achieve data

level 1). The full-width-half-maximum was measured to be ~ 30 GHz (see Fig. 2d), which yields an optical modulation bandwidth of ~ 19.3 GHz. Combined with the simulated electrical modulation bandwidth (~ 35 GHz, see above), this yields a device bandwidth, $f_{3\text{dB}}^{\text{el-opt}}$, of ~ 17 GHz, calculated using $(1/f_{3\text{dB}}^{\text{el-opt}})^2 = (1/f_{3\text{dB}}^{\text{opt}})^2 + (1/f_{3\text{dB}}^{\text{el}})^2$.

Measuring modulation energy consumption. The electrical power and switching energy of the stand-alone modulator were measured using a time domain reflectometer (TDR) (see the Methods section for details). The measurement, shown in Fig. 3c, is performed for an AC-coupled 0.5 V_{pp} drive dropped across the modulator via a 50Ω terminated probe. The probe is terminated because without the termination, only a ~ 0.25 V_{pp} drive is required due to the large reflection that would otherwise occur at the high-impedance modulator. However, the modulator is intended to operate with an integrated CMOS circuit, which will drop the drive voltage across the modulator pins without reflection since the circuit is connected by a local deep sub-wavelength electrical interconnect and transmission line effects are not present. Thus, the 50Ω termination is required to accurately characterize the modulator performance under the intended method of operation, but it will be neither used nor required in practice. The switching energy and energy per bit of the microdisk were experimentally extracted (see the Methods section for details) to be 3.65 and 0.91 fJ per bit, respectively, which agrees well with the simulated values. The slight difference between the theoretical and experimental results can be explained by a slightly lower p+ and n+ implant activation compared with the simulations, which would reduce the overall device capacitance.

Wire and/or pad capacitance add to the device capacitance. However, with both direct integration and advanced three-dimensional integration techniques this addition can be quite low. For example, in direct integration the total via/wire capacitance has been estimated to be 5 fF (ref. 28). In wafer bonding based on three-dimensional integrated through-oxide-vias²⁹, the via capacitance has been measured to be ~ 2 fF. Therefore with through-oxide-vias, the ground and signal vias will add a total of 4 fF of capacitance, similar to the direct integration case. In either case, the parasitic capacitance will add only ~ 0.3 fJ per bit, and the modulator energy consumption will still stay close to 1 fJ per bit.

Demonstrating high-speed and error-free transmission.

High-speed optical eye diagrams, shown in Fig. 4a, were measured using a high-speed electro-optic setup (see the Methods section for details) with a digital sampling oscilloscope at data rates of 10 , 15 , 20 and 25 Gb s⁻¹. The dynamic ER and IL are shown below each eye diagram in Fig. 4a. The dynamic ER is defined as the ratio of the average data level 1, $\langle 1 \rangle$, and the average data level 0, $\langle 0 \rangle$, $\text{ER} = 10\log(\langle 1 \rangle / \langle 0 \rangle)$. The IL, illustrated in Fig. 4a, is defined as the difference between the average data level 1, $\langle 1 \rangle$, and the optical true '1' as given by the off-resonance intensity at $\lambda \sim 1,588$ nm: $\text{IL} = 10\log((\langle 1 \rangle - \langle 0 \rangle) / (1 - \langle 0 \rangle))$. Note also that the optical true '0' corresponds to the amplified spontaneous emission floor after the filter at $\lambda \sim 1,588$ nm. The modulator exhibits low IL (~ 1.0 dB) and a high ER (6.2 dB) at a data rate of 25 Gb s⁻¹ for 0.5 V_{pp} drive voltage. The increased IL at high data rates can be explained by the electro-optic bandwidth limitation.

In order to further quantify the modulator performance, bit-error-rate (BER) measurements, shown in Fig. 4b, were performed at data rates of 20 and 25 Gb s⁻¹. For data rates up to 25 Gb s⁻¹, the device achieves error-free operation ($\text{BER} < 10^{-12}$) for a PRBS pattern length of $2^{31} - 1$. We did not observe any pattern length dependence on the BER with pattern lengths ranging from $2^7 - 1$ to $2^{31} - 1$. This is owing to the use of depletion-mode operation and extremely low drive voltages, and the low power consumption of the device limiting the thermal load on the device, which normally contributes to long-pattern errors³⁰. A commercial LiNbO₃ Mach-Zehnder modulator was similarly characterized for reference. The commercial modulator is rated to a 3 dB bandwidth of 35 GHz, and driven with an AC-coupled 5.5 V_{pp} electrical signal. The power penalty was recorded as the received power difference at a BER of 10^{-12} between the silicon microdisk modulator and the commercial LiNbO₃ modulator, as depicted in Fig. 4b. Data transmission with the silicon microdisk modulator was received with a power penalty of 3.0 and 3.06 dB at data rates of 20 and 25 Gb s⁻¹, respectively. The increased power penalty at high data rates can be explained by the electro-optic bandwidth difference between the two modulators and the phase chirp that is introduced by the resonant modulator.

Measuring electro-optic bandwidth. The electro-optic response of the modulator was characterized using an electrical signal source with a frequency range from DC to 50 GHz. The electrical signal output was calibrated to achieve a constant peak-to-peak drive voltage for the entire frequency span. The laser was then aligned to the highly linear part of the resonance at $\lambda \sim 1,578.4$ nm and a 50 mV_{pp} sinusoidal small-signal was applied to the microdisk modulator at frequencies spanning 500 MHz to 45 GHz. The through port power was measured with an external p-i-n photodiode and transimpedance amplifier receiver. The peak-to-peak AC level was recorded and plotted in Fig. 4c. The electro-optic 3 dB bandwidth was experimentally measured

Figure 4 | High-speed modulator characterization. (a) The experimentally measured high-speed optical eye diagrams of the silicon electro-optic modulator driven by a terminated probe and 0.5 V_{pp} NRZ-OOK PRBS with a pattern length of 2³¹-1 at data rates of 10-, 15-, 20- and 25-Gb s⁻¹. The extinction ratio (ER) and insertion loss (IL) are denoted below the eye diagrams at each data rate. (b) The bit error rate curves measured for the silicon microdisk and commercial LiNbO₃ modulators, for 20- and 25-Gb s⁻¹ data rates. The bit error rate curves for the commercial LiNbO₃ modulator were used as a reference to obtain the power penalty of the silicon microdisk modulator. (c) The electro-optic frequency response of the silicon modulator.

to be ~21 GHz, which agrees well with our theoretical small-signal estimate of 17 GHz. The difference is largely owing to the fact that the theoretical estimate was calculated for on-resonance operation and the modulator electro-optic bandwidth was measured slightly off-resonance to preserve linearity and to achieve a faster response, similar to 3¹.

Exhibiting thermal compensation by electro-optic tuning.

The baseline electro-optic response of this modulator, 250 pm V⁻¹ (30 GHz V⁻¹), enables 1 fJ per bit operation at 25 Gb s⁻¹, represents a new standard for resonant modulators. More importantly, the electro-optic response is sufficiently large to compensate for thermal shifts up to 10 °C. Here, we demonstrate this principle by varying the temperature of the chip from 20–30 °C (see the Methods section for details) and compensating for that variation with electro-optic tuning achieved by varying the bias voltage from -2.2 to 0.4 V. The through port spectral response as a function of applied DC bias was measured at ~26.5 °C and the results are plotted in Fig. 5a. The modulator shifts ~100 GHz over an applied voltage range from -2.2 to 0.4 V. The resonant modulator spectra without and with electro-optic thermal compensation over a temperature range of 20–30 °C are shown in Fig. 5b,c, respectively. The resonance is tuned electro-optically by utilizing the free-carrier plasma dispersion effect rather than through traditional thermo-optic tuning. The p-n junction acts as a capacitor in the reverse-biased and sub-threshold forward-biased regimes of operation. The electro-optic tuning consumes energy only if the transitions charge the capacitor. Given a microprocessor thermal time constant of $\tau \gg 1 \mu\text{s}$ ³², a worst-case electro-optic tuning power of 69.6 nW ($P_s = CV^2/\tau = 69.6 \text{ fJ } 1^{-1} \mu\text{s} = 69.6 \text{ nW}$) where C is determined from TDR measurements (see the Methods section for details), is required, resulting in a negligible energy per bit of 0.28 aJ per bit ($E_{\text{bit}} = 69.6 \text{ nW}/25 \text{ Gb s}^{-1} = 0.28 \text{ aJ}$ per bit), shown in Fig. 5g.

Up to 27.5 °C, the leakage current was sufficiently low that it contributed an additional tuning energy of only 0.24 fJ per bit. Therefore, the modulator operated over a 7.5 °C range with 1.03 fJ per bit energy consumption including both tuning (0.24 fJ per bit) and modulation energy (0.79 fJ per bit). The modulation energy is slightly decreased owing to the widened depletion region at the reverse-biased operation (equation (6)) that reduced the device capacitance, as shown in Fig. 5f,g. To achieve a full 10 °C temperature range, a -2.2 V bias is required, resulting in a leakage current of 23 μA , and consuming a static power of 50.6 μW . This corresponds to an additional energy cost of 2 fJ per bit, while modulation energy was reduced to 0.77 fJ per bit. The compensating frequency shift originates primarily from electro-optic tuning with a minor contribution from thermo-optic tuning due to the high leakage current resulting from undesired band-to-band tunnelling between the highly doped contact regions at the larger reverse biases, which can be mitigated with process optimization. The undesired leakage current also contributed to a small thermo-optic resonance frequency shift of ~7.2 GHz ($\Delta f \approx IV/7 \mu\text{W GHz}^{-1} = 7.2 \text{ GHz}$ ²⁵). While tuning the resonance appropriately for the given temperature, the modulator was driven by a terminated 0.5 V_{pp} drive with a PRBS length of 2³¹-1 at a data rate of 25 Gb s⁻¹ across the 10 °C temperature range while observing negligible degradation in eye diagrams and/or error rate. The resulting high-speed optical eye diagrams, shown in Fig. 5e, were measured using a digital sampling scope. The high ER and low IL were preserved over the 10 °C temperature range at 25 Gb s⁻¹ transmission. BER measurements, shown in Fig. 5d, were performed at a data rate of 25 Gb s⁻¹ and the modulator achieves error-free operation ($\text{BER} < 10^{-12}$) for a PRBS pattern length of 2³¹-1 over the 10 °C temperature range. The power penalty of the data transmission at a BER of 10⁻¹² due to thermal variations was within $\pm 0.14 \text{ dB}$. The total energy consumption, shown on Fig. 5f-g for each case, was determined to be 1.03 fJ per bit over a 7.5 °C temperature range and 2.77 fJ

Figure 5 | Temperature-insensitive operation. (a) The measured transmission spectra of the resonant modulator at $\sim 26.5^\circ\text{C}$ and applied DC bias voltages ranging from 0.4 to -2.2 V . (b) The measured transmission spectra of the resonant modulator over a 10°C temperature range at 0 V applied bias. The thermo-optic frequency response was extracted to be $\sim 10\text{ GHz per }^\circ\text{C}$. (c) The measured transmission spectra of the resonator over a 10°C temperature range with the modulator tuned electro-optically to compensate for thermal variations. (d) The bit error rate curves measured over a 10°C temperature range for the silicon resonant modulator, driven by a terminated probe and 0.5 V_{pp} NRZ-OOK PRBS with a pattern length of $2^{31}-1$ at a data rate of 25 Gb s^{-1} . The bit error rate curves clearly show error-free operation over the 10°C temperature range. (e) Experimentally measured high-speed optical eye diagrams of the silicon electro-optic modulator at a data rate of 25 Gb s^{-1} . The applied bias voltage (V_{DC}) and measured current (I_{DC}) are denoted above the eye diagrams at each temperature. The extinction ratio (ER) and insertion loss (IL) are denoted below the eye diagrams at each data rate and temperature. (f) The TDR measurement of the electro-optic modulator energy consumption and switching energy for a 0.5 V_{pp} drive is shown for the indicated DC bias voltages. (g) The breakdown of the thermo-optic, electro-optic and total modulator energy consumption is plotted as a function of temperature, showing the 1 fJ bit^{-1} operation over 7.5°C temperature range.

per bit over a 10°C temperature range. All prior demonstrations of resonant silicon modulators required thermo-optic tuning owing to their relatively small electro-optic frequency responses. However, implementing thermal tuning within modulators requires additional contacts, complicating and often degrading the modulator performance. When tuned over a 10°C temperature range, the energy consumption of these devices increased substantially to $\sim 82\text{ fJ per bit}$ ($7\text{ }\mu\text{W GHz}^{-1} \times 100\text{ GHz}/10^{10}\text{ bits s}^{-1} + 12\text{ fJ per bit} = 82\text{ fJ per bit}$)²⁵ and $\sim 175\text{ fJ per bit}$ ($42\text{ }\mu\text{W GHz}^{-1} \times 100\text{ GHz}/(2.5 \times 10^{10}\text{ bits s}^{-1}) + 7\text{ fJ per bit} = 175\text{ fJ per bit}$)¹⁸ at data rates of 10 and 25 Gb s^{-1} , respectively.

Integrating the modulator into a communication link. Finally, it is useful to consider the modulator performance in the context of a fully integrated communication link including the driver circuitry. A key part of this consideration is the inclusion of the circuitry that will be necessary to correct for thermal effects. Within a microprocessor, baseline temperature variations are typically within $\pm 5^\circ\text{C}$ ³², hence the reason for the comparison of the modulators across a 10°C range. Over this range, our modulator consumes a total of 2.77 fJ per bit . However, with either a slight reduction in temperature variations on a microprocessor, or with a reduction in the leakage current of our device, 1 fJ per bit operation can be maintained. Moreover, even without these improvements with the use of barrel-shifting techniques²⁸, compensation across much broader temperature ranges can be achieved while staying in the 1 fJ per bit electro-optic regime. Indeed, it is only necessary to be able to shift each ring by the channel spacing. Implementation of this technique, however, adds overhead to the power consumption, and this overhead is given by: $\sim A \times \log_2^N$, where N is the total number of resonators and A is the 2-by-1 electrical multiplexer energy consumed at the given CMOS node (Michael Georgas, Chen Sun, personal communication, 10 April 2014). For the 45 , 32 , 22 and 11 nm CMOS nodes, the estimated energy consumption of this technique for 86 resonant modulators (71 active and 15 redundant) with a channel spacing of 75 GHz will be 17.88 fJ per bit , 9.9 fJ per bit , 4.3 fJ per bit and 1.3 fJ per bit , respectively³³, which is expected to scale to $<1\text{ fJ per bit}$ below the 10 nm node³⁴. In addition to the barrel-shifting circuitry, one must also account for the modulator control circuitry. This circuitry includes the modulator driver, the feedback loop and a monitor receiver, as shown in Fig. 6a. Here, the receiver, attached to a weak drop port, constantly monitors the modulator output and

the feedback loop processes that output and locks the modulator to the laser line by adjusting the DC bias of the modulator driver as described in the preceding paragraph. The energy consumption of such a modulator driver and a monitor receiver were experimentally measured on a 45-nm CMOS node to be 20 and 69 fJ per bit , respectively, (assuming a receiver responsivity of $\sim 1\text{ AW}^{-1}$)³⁵. The energy consumption of the feedback loop was also measured on a 180-nm CMOS node to be 0.43 mW (for example, 17.2 fJ per bit at 25 Gb s^{-1} operation)³⁶. It is important, however, to note two things. First, in previous low-power modulators, which rely on thermal tuning, the lowest thermal tuning of $\sim 175\text{ fJ per bit}$ at 25 Gb s^{-1} still clearly dominates over the driving circuitry in the energy consumption budget. In contrast, for the modulator presented in this work, the circuit energy consumption now dominates the overall energy consumption, an important achievement. Second, current CMOS power consumption is technological, not fundamental, at this level as is shown in Fig. 6b. Among other things, given the CMOS energy consumption scaling³⁴, we can see that the driver will consume $<1\text{ fJ per bit}$ at or below the 10-nm CMOS node and the total circuit including feedback and the barrel shifter will consume $<1\text{ fJ per bit}$ below the 4-nm CMOS node. We therefore expect that with improvements in CMOS technology, the total modulator plus driver circuit energy consumption can be kept in the 1 fJ per bit range for this device.

Demonstrating bandwidth-limited transmission. In order to demonstrate the highest speed digital operation of the modulator, a voltage amplifier and a bias tee with a 3-dB bandwidth of 40 GHz were attached to the high-speed electro-optic setup (see the Methods section for details) and a 2.2-V_{pp} drive signal with -0.5 V DC bias was delivered to the modulator. High-speed optical eye diagrams, shown in Fig. 7, were measured with a digital sampling oscilloscope at data rates of up to 44 Gb s^{-1} . The dynamic ER and IL are shown below each eye diagram in Fig. 7. With low IL (0.9 dB) and high ER (8.0 dB), the eye diagrams are still open at 44 Gb s^{-1} . There are a couple of points to consider that enabled this result. The NRZ 44 Gb s^{-1} data stream only requires $\sim 22\text{ GHz}$ (half of the bit rate) of bandwidth since the spectrum of the NRZ signal is proportional to $\text{sinc}^2(f)$, where f is the frequency at the bit rate and the major spectral components are below half of the bit rate³⁷. Also, while the electro-optic 3 dB bandwidth was measured to be 21 GHz with a small-signal drive (50 mV_{pp}), when the modulator is driven at a reverse-biased voltage of -2.2 V , the modulator bandwidth increases to

Figure 6 | Integrated modulator components and energy scaling. (a) A diagram of the required integrated electrical circuitry to drive, and control the modulator and barrel shifter to correct for channel hopping between modulators. (b) The energy consumption scaling of the modulator subsystem is plotted as a function of CMOS node scaling, showing sub-fJ per bit driver and total electronics energy at advanced CMOS nodes.

Figure 7 | Bandwidth-limited data transmission. The experimentally measured high-speed optical eye diagrams of the silicon microdisk modulator when driven by a terminated probe and a 2.2 V_{pp} NRZ-OOK PRBS with a pattern length of $2^{31}-1$ at data rates of 25-, 34- and 44 Gb s⁻¹. The extinction ratio (ER) and insertion loss (IL) are denoted below the eye diagrams at each data rate.

22.4 GHz due to the reduction in average device capacitance, measured by TDR (see below and Fig. 5f), from ~ 16 to ~ 14 fF. The reduction in device capacitance stems from the increased depletion width at higher reverse bias, also observed in¹⁹. Finally, the modulator starts to operate in the optically limiting regime at this higher drive voltage, also observed in¹⁹. In this regime, the optical resonance limits the data levels to the optical '0' and '1', which enables relatively faster transitions. This demonstration shows that the modulator can provide on-demand bandwidth with a measured energy consumption of 17.4 fJ per bit (see the Methods section for details).

Discussion

In conclusion, we have demonstrated an athermal silicon photonic modulator with the largest electro-optic response, lowest power and lowest single-ended voltage demonstrated to date. The record 250 pm V⁻¹ (30 GHz V⁻¹) electro-optic response enables error-free 0.9 fJ per bit operation at a data rate of 25 Gb s⁻¹ with an AC-coupled drive of 0.5 V_{pp} in a 4.8-μm diameter microdisk modulator. Further, the large response leads to a small parasitic (0.24 fJ per bit) tuning energy consumption but when combined with a reduction in the modulation energy consumption to 0.79 fJ per bit at the compensated temperature, leads to a total energy consumption of 1.03 fJ per bit over a 7.5 °C temperature range. Over a 10 °C temperature range, the increased leakage current increases the tuning energy consumption to 2 fJ per bit, while sustaining error-free 0.77 fJ per bit operation at a data rate of 25 Gb s⁻¹ for a total energy consumption of 2.77 fJ per bit. Driven harder, the modulator demonstrates open-eye diagrams out to 44 Gb s⁻¹. The electro-optic response of the circularly contacted vertical junction device is up to 10 times larger than that achieved with lateral junction-based modulators, thus enabling both efficient modulation and the electro-optic control of the resonant frequency. The combined result improves upon the overall energy consumption of modulators by nearly two orders of magnitude. No other resonant modulator today is capable of achieving a similar result. Yet, these results represent only the beginning of electro-optic control of the resonant frequency. Future modulators and filters with maximized capacitance per unit volume (for example, in multiple vertical junctions such as an n-p-n junction or perhaps high spatial frequency interleaved junctions³⁸) in a smaller footprint will clear the way for direct control of even greater temperature variations with even less leakage current. While the modulator represents only one component of a silicon photonic link, as photonics becomes more closely integrated with CMOS, and as the bandwidth requirements on- and off-chip continue to rise, the power consumption of each component will be pressed to achieve new levels of efficiency in order to fit within the overall power envelope of the communication link. With their efficient 1 fJ per bit operation at the increasingly important 25 Gb s⁻¹ data rate,

and with the effective elimination of heater power consumption, these results represent an important first step towards achieving femtojoule per bit class communication links that will prove critical for future on-chip applications.

Methods

Device fabrication. The microdisk modulator was fabricated in a 300-mm CMOS foundry with a 65-nm technology node, using silicon-on-insulator wafers with a 0.225-μm top silicon layer and 2-μm buried oxide layer for optical isolation. A full silicon etch was applied to form the silicon resonant microdisk and bus waveguides. An oxidation step followed to passivate the sidewalls, which also decreased the waveguide thickness to 0.22 μm. The vertical junction, centered slightly above 110 nm thickness, was formed from arsenic (As) and boron difluoride (BF₃) implants with designed doses of 4×10^{13} and 2.77×10^{13} cm⁻², and energies of 380 and 120 keV, respectively. Through fitting the experimental spectral response of the modulator as a function of voltage, the As and BF₃ implants were determined to be formed with equivalent energies of 415 and 90 keV, respectively. The overlapping n+ and p+ -doped regions in the center of the microdisk were formed from phosphorus and BF₃ implants with doses of 5×10^{15} and 3.5×10^{15} cm⁻², and energies of 180 and 120 keV, respectively. The target p-type and n-type doping concentrations were both 2.5×10^{18} cm⁻³ and p+ and n+ -type doping concentrations were $> 2 \times 10^{20}$ cm⁻³. An oxide (SiO₂) layer with a total thickness of 3.6 μm was deposited to cover the modulator. Low-resistance contacts to p+ and n+ -doped regions were made by standard silicidation, copper-silicon circular and point contacts. Finally, the contacts were connected to on-chip ground-signal-ground (GSG)-probing pads by two metal layers.

Numerical simulation. The carrier distribution as a function of voltage, simulated by Sentaurus, was converted into an index distribution using the plasma dispersion (electro-refraction) equations in refs 14,22. The index distribution was fed into the FDM²⁶ to determine the complex propagation constant as a function of voltage. The complex propagation constant was then input to the transfer matrix of the resonant modulator³⁹ in order to theoretically predict the spectral response as a function of voltage.

Time domain reflectometer measurement. The absorbed power of the modulator was determined by subtracting two TDR measurements. First, the reflected voltage (V_P^-) from the transmission line, non-terminated GSG probe and pads (Z_P) was initially determined using the experimental setup, shown in Fig. 3a. The reflected voltage was then normalized to the emanating step voltage (V_{in}) to extract the frequency domain reflection coefficient, $S_{11}^P = V_P^- / V_{in} = (Z_P - Z_0) / (Z_P + Z_0)$, where, Z_0 , 50 Ω, is the characteristic impedance of the transmission line. The absorbed switching power in this path was determined by the following equation, $P_s^P = (1 - (S_{11}^P)^2) V_{in}^2 / Z_0$. Second, the reflected voltage (V_{PM}^-) from this path and the modulator (Z_M), connected to the identical wiring, was determined using the experimental setup, shown in Fig. 3b. The normalized frequency domain reflection coefficient and absorbed switching power was determined by the following equations, $S_{11}^{PM} = V_{PM}^- / V_{in} = (Z_{PM} - Z_0) / (Z_{PM} + Z_0)$ and $P_s^{PM} = (1 - (S_{11}^{PM})^2) V_{in}^2 / Z_0$, respectively. With close electrical and photonics integration these devices will not require long transmission lines, probing or large GSG pads. Therefore, the absorbed power of stand-alone modulator is $P_s^M = (1 - (S_{11}^M)^2) V_{in}^2 / Z_0$, which can be calculated by subtracting the previously measured power consumptions, $P_s^M = P_s^{PM} - P_s^P = ((S_{11}^P)^2 - (S_{11}^{PM})^2) V_{in}^2 / Z_0$. The switching power of the modulator (P_s^M) was integrated over time to obtain switching energy, $E_s = \int P_s^M dt$. The exact switching energy using the above equation was determined to be 3.65 fJ (0.91 fJ per bit). Alternatively, using the parallel equivalent impedance model, $1/Z_{PM} = 1/Z_P + 1/Z_M$, we can approximate the modulator reflection coefficient

with the following equation:

$$\frac{S_{11}^{pm}}{S_{11}^p} = \frac{V_{PM}^-}{V_P^-} = \frac{Z_M - Z_0 - \frac{Z_p^2}{Z_p} (Z_P + Z_0)}{Z_M + Z_0 + \frac{Z_p^2}{Z_p} (Z_P + Z_0)} \cong \frac{Z_M - Z_0}{Z_M + Z_0} = S_{11}^m$$

which was also performed in ref. 14. This approximation is accurate for $Z_P \gg Z_0^2$, and overestimates the high frequency power consumption as Z_P approaches to Z_0^2 . The approximated switching energy using the above equations was determined to be 3.84 fJ (0.96 fJ per bit), which is in good agreement with the exact switching energy.

High-speed electro-optic setup. The electrical data, encoded using a pattern generator, was a NRZ on-off-keyed signal, encoded with a $2^{31}-1$ PRBS. The output of the pattern generator was 0.5 V_{pp} with a +0.25 V DC bias, and a DC block was attached to achieve an AC-coupled 0.5 V_{pp} drive. A 50-Ω-terminated GSG probe was used to eliminate transmission line reflections and ensure that the voltage dropped across the modulator was 0.5 V_{pp}. This termination is not required when electronics are integrated on-chip since wire lengths will be on the order of 10–100 μm, where transmission line effects are negligible for wires up to a millimetre in length ($\lambda_{RF}/10$, where $\lambda_{RF} = 10$ mm) at 25 Gb s⁻¹ operation. The laser line, coupled into the microdisk, was set to $\lambda \sim 1,587.4$ nm and the through port output was then passed through an erbium-doped fibre amplifier (EDFA) to overcome fibre-to-chip coupling losses. A tunable bandpass filter with a 1-nm 3-dB bandwidth was used to filter out the amplified spontaneous emission of the EDFA. In addition, a variable optical attenuator was inserted before the high-speed p-i-n photodiode and transimpedance amplifier receiver. The received data were then fed differentially to a BER tester for evaluation. No pre-emphasis or equalization of the signal was used throughout the measurements. The total fibre-to-fibre IL was ~ 14.5 dB, and the intensity before the EDFA was ~ -9 dBm.

Temperature-controlled setup. In order to alter and record the chip temperature, a thermoelectric-cooler (TEC) was attached to the silicon substrate and a thermistor, placed adjacent to the silicon chip, was used to feedback the temperature readout. Then, a commercial TEC proportional-integral-derivative controller was used to stabilize the chip temperature from 20 to 30 °C with ± 0.05 °C accuracy. The continuous-wave laser was set to $\lambda \sim 1,587.6$ nm and fibre-to-chip coupling optimized for each temperature set point due to the thermal expansion of the TEC.

References

- Moore, G. E. Cramping more components onto integrated circuits. *Electronics* **38**, 114–116 (1965).
- Miller, D. A. B. Optical interconnects to silicon. *IEEE J. Sel. Top. Quant. Electron.* **6**, 1312–1317 (2000).
- Young, I. A. *et al.* Optical I/O technology for tera-scale computing. *IEEE J. Solid-State Circuits* **45**, 235–248, (2010).
- Borkar, S. & Chien, A. A. The future of microprocessors. *Commun. ACM* **54**, 67–77 (2011).
- Asghari, M. & Krishnamoorthy, A. V. Silicon photonics: energy-efficient communication. *Nat. Photon.* **5**, 268–270 (2011).
- Reed, G. T., Mashanovich, G., Gardes, F. Y. & Thomson, D. J. Silicon optical modulators. *Nat. Photon.* **4**, 518–526 (2010).
- Liu, J. *et al.* Waveguide-integrated, ultralow-energy GeSi electro-absorption modulators. *Nat. Photon.* **2**, 433–437 (2008).
- Kuo, Y.-H. *et al.* Strong quantum-confined Stark effect in germanium quantum-well structures on silicon. *Nature* **437**, 1334–1336 (2005).
- Liu, A. *et al.* A high-speed silicon optical modulator based on a metal-oxide-semiconductor capacitor. *Nature* **427**, 615–618 (2004).
- Green, W. M. J., Rooks, M. J., Sekaric, L. & Vlasov, Y. A. Ultra-compact, low RF power, 10 Gb/s silicon Mach-Zehnder modulator. *Opt. Express* **15**, 17106–17113 (2007).
- Xu, Q., Schmidt, B., Pradhan, S. & Lipson, M. Micrometre-scale silicon electrooptic modulator. *Nature* **435**, 325–327 (2005).
- Xu, Q., Maniapatruni, S., Schmidt, B., Shakyia, J. & Lipson, M. 12.5 Gbit/s carrier injection-based silicon micro-ring silicon modulators. *Opt. Express* **15**, 430–436 (2007).
- Watts, M. R., Trotter, D. C., Young, R. W. & Lentine, A. L. in: *Proceedings of the 5th IEEE International Conference on Group IV Photonics*, Sorrento, Italy, 4–6 (2008).
- Watts, M. R., Zortman, W. A., Trotter, D. C., Young, R. W. & Lentine, A. L. Vertical junction silicon microdisk modulators and switches. *Opt. Express* **19**, 21989–22003 (2011).
- Timurdogan, E., Sorace-Agaskar, C. M., Biberman, A. & Watts, M. R. in: *Optical Fiber Communication Conference* (Anaheim, CA, 2013).
- Zortman, W. A., Lentine, A. L., Trotter, D. C. & Watts, M. R. Low-voltage differentially-signaled modulators. *Opt. Express* **19**, 26017–26026 (2011).
- Dong, P. *et al.* Low V_{pp}, ultralow-energy, compact, high-speed silicon electro-optic modulator. *Opt. Express* **17**, 22484–22490 (2009).
- Li, G. *et al.* 25 Gb/s 1 V-driving CMOS ring modulator with integrated thermal tuning. *Opt. Express* **19**, 20435–20443 (2011).
- Rosenberg, J. C. *et al.* A 25 Gbps silicon microring modulator based on an interleaved junction. *Opt. Express* **20**, 26411–26423 (2012).
- Xiao, X. *et al.* 25 Gbit/s silicon microring modulator based on misalignment-tolerant interleaved PN junctions. *Opt. Express* **20**, 2507–2515 (2012).
- Miller, D. A. B. Energy consumption in optical modulators for interconnects. *Opt. Express* **20**, A293–A308 (2012).
- Soref, R. & Bennett, B. Electrooptical effects in silicon. *IEEE J. Quant. Electron.* **23**, 123–129 (1987).
- Staelin, D. H., Morgenthaler, A. W. & Kong, J. A. *Electromagnetic Waves* (Prentice Hall, 1994).
- Sze, S. M. & Ng, K. K. *Physics of Semiconductor Devices* (Wiley, 2007).
- Zortman, W. A., Lentine, A. L., Trotter, D. C. & Watts, M. R. in: *Optical Fiber Communication Conference* (Los Angeles, CA, 2012).
- Popovic, M. in: *Proceedings of Integrated Photon Research 2003* (Washington, DC, 17 June 2003).
- Gheorma, I. L. & Osgood, Jr R. M. Fundamental limitations of optical resonator based high-speed EO modulators. *IEEE Photon. Tech. Lett.* **14**, 795–797 (2002).
- Georgas, M., Leu, J., Moss, B., Sun, C. & Stojanovic, V. in *Custom Integrated Circuits Conference (CICC)* 1–8 (IEEE, 2011).
- Chen, C. L. *et al.* in: *IEEE Topical Meeting on Silicon Monolithic Integrated Circuits in RF Systems*, SiRF'09, San Diego, CA, USA. 1–4 (2009).
- Zhu, X. *et al.* in: *Conference on Optical Fiber Communication* (Anaheim, CA, 2013).
- Sacher, W. D. *et al.* Coupling modulation of microrings at rates beyond the linewidth limit. *Opt. Express* **21**, 9722–9733 (2013).
- Cher, C.-Y. & Kursun, E. Exploring the effects of on-chip thermal variation on high-performance multicore architectures. *ACM Trans. Architect. Code Optim.* **8**, 1, (2011).
- Selvaraja, S. K., Bogaerts, W., Dumon, P., Van Thourhout, D. & Baets, R. Subnanometer linewidth uniformity in silicon nanophotonic waveguide devices using CMOS fabrication technology. *IEEE J. Sel. Top. Quant. Electron.* **16**, 316–324 (2010).
- Haensch, W. *et al.* 'Silicon CMOS devices beyond scaling'. *IBM J. Res. Dev.* **50**, 339–361 (2006).
- Georgas, M. *et al.* in: *VLSI Symposium*, TAPA I, paper 6.4, Honolulu, Hawaii, USA (2014).
- Sun, C. *et al.* in: *VLSI Symposium*, TAPA I, paper 6.3 (2014).
- Gupta, M. C. & Ballato, J. *The Handbook of Photonics* 20–46 (CRC, 2007).
- Timurdogan, E., Sorace-Agaskar, C., Hosseini, E. S. & Watts, M. R. 'An interior-edge silicon microring modulator'. *J. Lightwave Technol.* **8**, 3907–3914 (2013).
- Little, B. E., Chu, S. T., Haus, H. A., Foresi, J. & Laine, J. P. Microring resonator channel dropping filters. *J. Lightwave Technol.* **15**, 998–1005 (1997).

Acknowledgements

This work was supported by the Defense Advanced Research Projects Agency (DARPA) of the United States under the E-PHI project, grant no. HR0011-12-2-0007 and DARPA POEM award no. HR0011-11-C-0100 and APIC Corporation. Cheryl Sorace-Agaskar acknowledges support from NSFGRP, No. 0645960.

Author contributions

E.T. and M.R.W. conceived the idea of the project. E.T. simulated and designed the microdisk modulator, laid out the mask, performed the experimental characterizations, analysed the data and wrote the manuscript. C.M.S.-A. did the doping and the vertical junction simulations, supported the data analysis and contributed to the writing of the manuscript. J.S. took the scanning electron microscopy image of the modulator. E.S.H. coordinated the layout. A.B. helped with the electro-optic setup. M.R.W. edited the manuscript and supervised the project.

Additional information

Competing financial interests: The authors declare no competing financial interests.

Reprints and permission information is available online at <http://npg.nature.com/reprintsandpermissions/>

How to cite this article: Timurdogan, E. *et al.* An ultralow power athermal silicon modulator. *Nat. Commun.* 5:4008 doi: 10.1038/ncomms5008 (2014).

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-sa/3.0/>