

Silicon optical modulators

G. T. Reed*, G. Mashanovich, F. Y. Gardes and D. J. Thomson

Optical technology is poised to revolutionize short-reach interconnects. The leading candidate technology is silicon photonics, and the workhorse of such an interconnect is the optical modulator. Modulators have been improved dramatically in recent years, with a notable increase in bandwidth from the megahertz to the multigigahertz regime in just over half a decade. However, the demands of optical interconnects are significant, and many questions remain unanswered as to whether silicon can meet the required performance metrics. Minimizing metrics such as the device footprint and energy requirement per bit, while also maximizing bandwidth and modulation depth, is non-trivial. All of this must be achieved within an acceptable thermal tolerance and optical spectral width using CMOS-compatible fabrication processes. This Review discusses the techniques that have been (and will continue to be) used to implement silicon optical modulators, as well as providing an outlook for these devices and the candidate solutions of the future.

Network interconnects, for both conventional data networks and intra-/interchip data links, continue to scale in complexity and bandwidth year after year. The limitations of copper as an interconnect medium in terms of its loss, dispersion, crosstalk and fundamental speed are becoming increasingly obvious as interconnect densities rise. This has heralded the emergence of the optical interconnect, with silicon photonics being the leading candidate due to its unique combination of low fabrication costs, performance enhancements resulting from electronic–photonic integration, and compatibility with the world’s most successful technology for producing electronics, CMOS.

Over time, the interconnect market will see a gradual transition from electrical to optical technology, with longer links being most feasible to replace as they require larger bandwidths. Optical technology will then spread to very short distances as the technology becomes more cost effective and bandwidth requirements increase. The target application areas for optical interconnects range from high-performance computing and data centres down to mobile-to-server interconnects and desktop computers.

The dominant application for silicon photonics is in optical interconnect technology, with active optical cables already reaching the marketplace and complex intrachip interconnect technology becoming the subject of advanced research for high-performance computing. There is also a plethora of other applications that will benefit from the success of silicon photonics, including fibre-to-the-home or fibre-to-the-premises systems, environmental monitoring, biological and chemical sensing, medical and military applications, and astronomy. However, a significant amount of attention is currently focused on research into optical interconnect technology, simply because it is urgently required in many applications; recent bandwidth improvements in electronic interconnects have only been achieved at the expense of increased latency and power consumption¹. Optical interconnects could bring several major advantages over their electrical counterparts, as they enable the separation of electronic devices, and thus allow optimization of the chip layout while retaining high data rates. They can also introduce some of the more conventional advantages of optical technology, such as the reduction of electromagnetic interference, cable length and cable weight, and may also allow increased complexity in device design through optical interconnects and optimized cooling. Finally, they may save energy, retain precise clock and signal timing², and allow interconnect densities to be reduced. The potential for saving energy is important, as the amount of heat that can be removed from a chip, in a cost-effective manner, is expected to remain constant in

the foreseeable future³. Because interconnects consume 80% of a microprocessor’s power³, and server interconnect powers exceed the total power generated from solar energy⁴, optical interconnects may potentially offer substantial energy savings^{5,6}.

The ultimate physical manifestation of the silicon photonic device would be as part of an optoelectronic integrated circuit formed monolithically in silicon, which would combine photonic functionality and electronic intelligence in a seamlessly integrated design. Although components continue to be developed, interim solutions such as hybrid combinations of different materials aim to fulfil the functions that silicon is currently unable to deliver. The question remains as to whether silicon will ever deliver all of the components of such a circuit in a monolithic form, or whether a hybrid solution will prevail. Producing all components from silicon would be more elegant, but may never be realizable as a commercial success because of the ferocious performance and cost demands of the market. In this Review we discuss the progress of just one of the components that will form the optoelectronic integrated circuit of the future — the silicon optical modulator. This single component has been realized in both monolithic and hybrid forms, and it is hoped that the discussion provided here will stimulate further debate on the relative merits of each approach.

Optical modulation is one of the main required functionalities for any optical interconnect solution. An optical source can be either directly or externally modulated. External modulation offers several advantages over direct modulation: the optical source can be relatively inexpensive and its operation does not need to be compromised by direct modulation, modulation speeds can be higher, direct phase modulation is possible⁷, and optical isolation and wavelength stabilization need to be performed only once for the entire system. Furthermore, a single light source can feed multiple channels via individual modulators, thus reducing the total power budget of the system.

Modulation in silicon

An optical modulator is a device that is used to modulate (that is, vary the fundamental characteristics of) a light beam propagating either in free space or in an optical waveguide. These devices can alter different beam parameters, allowing them to be categorized as either amplitude, phase or polarization modulators. In addition, modulators can be also classified as either electro-refractive or electro-absorptive.

Applying an electric field to a material may change its real and imaginary refractive indices. A change in the real part of the refractive




Figure 1 | Cross-sections of typical device structures implementing the three different mechanisms commonly used to electrically manipulate the free-carrier concentrations in plasma-dispersion-based silicon optical modulators. a, Carrier accumulation; a thin insulating layer of SiO₂ is used to isolate two halves of the waveguide to form a capacitor structure. **b**, Carrier injection; highly doped p- and n-regions are separated by an 'intrinsic region' in which the waveguide is formed. Forward-biasing the device causes free electrons and holes to be injected into the 'intrinsic' waveguide region. **c**, Carrier depletion; lightly doped p- and n-type regions abut in the waveguide to form a p-n diode. The depletion area of the diode becomes larger with increasing reverse bias voltage.

index (Δn) with an applied electric field is known as electro-refraction, whereas a change in the imaginary part of the refractive index ($\Delta\alpha$) is known as electro-absorption. The primary electric field effects that are traditionally useful in semiconductor materials for causing either electro-absorption or electro-refraction are the Pockels effect, the Kerr effect and the Franz-Keldysh effect. However, it has been shown that these are weak in pure silicon at the telecommunications wavelengths of 1.3 μm and 1.55 μm (refs 8,9). Alternative methods are therefore required to achieve modulation in silicon. One option is thermal modulation owing to the large thermo-optic coefficient of silicon, but this is too slow for the high frequencies required by modern telecommunications applications¹⁰.

The most common method of achieving modulation in silicon devices so far has been to exploit the plasma dispersion effect, in which the concentration of free charges in silicon changes the real and imaginary parts of the refractive index⁹. Soref and Bennett⁸ evaluated changes in the refractive index Δn from experimentally produced absorption curves for a wide range of electron and hole densities, over a wide range of wavelengths. They also quantified changes in both the refractive index and absorption⁸, and produced the following expressions to evaluate changes in the carrier densities in silicon at a wavelength of 1.55 μm :

$$\Delta n = \Delta n_e + \Delta n_h = -[8.8 \times 10^{-22} \times \Delta n_e + 8.5 \times 10^{-18} \times (\Delta n_h)^{0.8}]$$

$$\Delta\alpha = \Delta\alpha_e + \Delta\alpha_h = 8.5 \times 10^{-18} \times \Delta n_e + 6.0 \times 10^{-18} \times \Delta n_h$$

where Δn_e and Δn_h are changes in refractive index resulting from changes in the free-electron and free-hole carrier concentrations, respectively, and $\Delta\alpha_e$ and $\Delta\alpha_h$ are the changes in absorption resulting from changes in the free-electron and free-hole carrier concentrations, respectively. Similarly, at a wavelength of 1.3 μm :

$$\Delta n = \Delta n_e + \Delta n_h = -[6.2 \times 10^{-22} \times \Delta n_e + 6.0 \times 10^{-18} \times (\Delta n_h)^{0.8}]$$

$$\Delta\alpha = \Delta\alpha_e + \Delta\alpha_h = 6.0 \times 10^{-18} \times \Delta n_e + 4.0 \times 10^{-18} \times \Delta n_h$$

Using these expressions, it is straightforward to calculate, for example, that a change in carrier density of the order of $5 \times 10^{17} \text{ cm}^{-3}$ results in a Δn of -1.66×10^{-3} at a wavelength of 1.55 μm . However, this is accompanied by a detrimental change in intensity due to the absorption of free carriers¹¹.

Electrical manipulation of the charge density interacting with the propagating light is achievable through mechanisms such as carrier injection, accumulation or depletion. These schemes are represented schematically by waveguide cross-sections in Fig. 1.

More recently, attempts have been made to investigate alternative modulation mechanisms in other materials potentially

compatible with silicon technology, such as germanium, to achieve more efficient modulation. The Franz-Keldysh effect¹² and the quantum-confined Stark effect (QCSE) are both electric-field-induced changes in optical absorption. They are related, with the Franz-Keldysh effect being the limit of the QCSE as the quantum-well layers are increased in thickness. The QCSE has more spectrally abrupt and stronger changes in absorption coefficient⁴, and it is observed in quantum well structures in which the electron and hole confinement afforded by the barriers allows for exciton enhancement of the optical absorption. However, using these effects on a silicon platform requires the introduction of a second material, which complicates device designs. In 2005, Kuo *et al.* demonstrated the QCSE in pure Ge quantum wells and Ge-rich SiGe barriers on Si (ref. 13). Similar structures exhibiting the QCSE were later integrated into modulator devices, demonstrating proof-of-concept electro-absorption modulation^{14,15}, although no waveguide-based device has yet been reported. Nevertheless this work offers significant hope that high speed, efficient QCSE modulators may emerge in the future, albeit with complex fabrication requirements.

Further modulation options are available by creating hybrid materials of silicon and other photonic materials, or through the manipulation of the crystalline material; these will be briefly summarized later.

So far we have discussed how optical modulators in silicon can cause changes in the refractive index that induce absorption. The process of inducing absorption directly modulates the intensity of a propagating mode, so it is obvious to see how modulation is achieved in such devices. However, there are essentially two options available for converting a change in refractive index into intensity modulation. First, the refractive index change can be used to shift the relative phase of two propagating waves such that they interfere either constructively or destructively. Typically, a Mach-Zehnder interferometer (MZI) is used to achieve this. Second, including a resonant structure in the device allows the refractive-index change induced in the modulator to change the resonant condition, thus allowing the device to be switched between on- and off-resonance states at any given wavelength. A discussion of the relative merits of each approach will be included in the next section, which focuses on modulator performance metrics.

Performance metrics

There are several figures of merit that are used to characterize a modulator, including its modulation speed and depth, optical bandwidth, insertion loss, area efficiency (footprint) and power consumption.

The modulation speed or bandwidth is one of the most important figures of merit for an optical modulator. Modulation bandwidth is

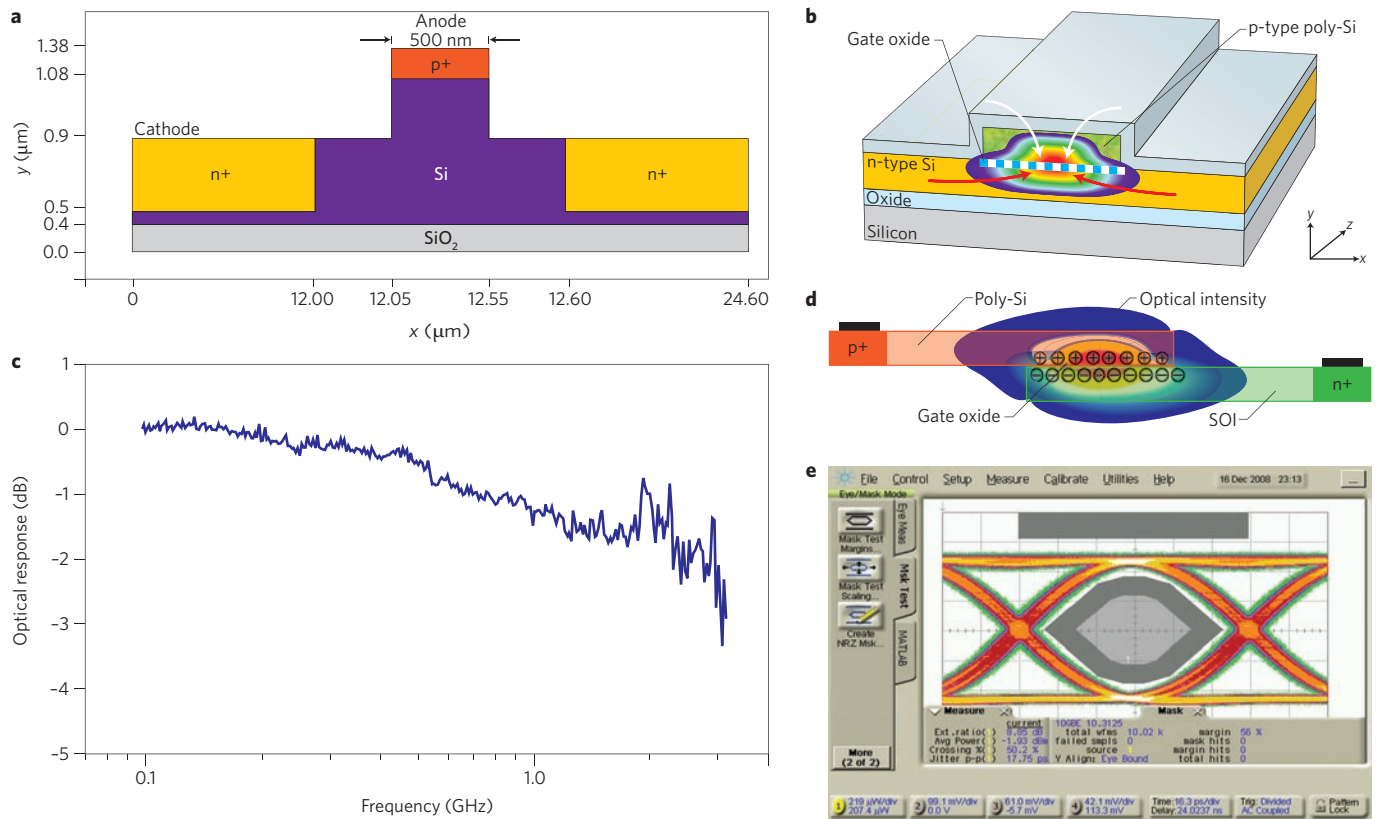


Figure 2 | The first devices to propose and realize modulation at gigahertz frequencies in silicon, together with the best present carrier-accumulation-type device. **a**, Cross-section of the first proposed gigahertz modulator, which was based on carrier injection. **b**, Cross-section of the first experimentally realized gigahertz modulator based on carrier accumulation. **c**, Electro-optic bandwidth of this device, showing a -3 dB cut-off at ~ 3 GHz. **d**, Cross-section of the state-of-the-art carrier-accumulation-type device. **e**, 10 Gbit s^{-1} eye diagram obtained from the device in **d**, showing an extinction ratio of 9 dB. Figure reproduced with permission from: **a**, ref. 58, © 2004 IEEE and ref. 75, © 2004 Univ. Surrey; **b**, ref. 98, © 2008 Univ. Surrey; **c**, ref. 60, © 2004 NPG; **d, e**, ref. 63, © 2009 OFC/NFOEC.

usually defined by the frequency at which the modulation is reduced to 50% of its maximum value. The speed of a modulator is commonly characterized by its ability to carry data at a certain rate. High modulation speeds are required for interconnect applications in which high data rates are imperative, but for certain applications (sensing, for example) only relatively moderate modulation speeds are needed.

Modulation depth — also known as the extinction ratio — is defined as the ratio of I_{\max} , the intensity transmitted when the modulator is adjusted for maximum transmission, to I_{\min} , the intensity transmitted when the modulator is adjusted for minimum transmission. It is quoted in decibels and expressed as $10\log(I_{\max}/I_{\min})$. A large modulation depth is useful for long transmission distances, good bit error rates and high receiver sensitivity. Although a high modulation depth of >7 dB is always preferred for interconnect applications, 4–5 dB is often sufficient if the overall system link budget and balance of the system metrics are both taken into consideration. The performance metrics of an interconnect system are linked in a relatively simple way, making it possible to trade-off performance in one area for another. For example, a shorter modulator will result in an inferior extinction ratio, but also reduced footprint and power consumption. Consequently it is possible to optimize modulator performance for a given application, depending on which metric is most important.

Insertion loss takes into account the optical power that is lost when the modulator is added to a photonic circuit. It is a passive loss that comprises reflection, absorption and mode-coupling losses, and is significant because it contributes to the link budget

(determined by the receiver sensitivity) and to the overall end-to-end losses in the system. Although there has been a trend in silicon photonics to move to smaller, submicrometre waveguides (mostly in an attempt to reduce device size and increase performance), this move has been coupled with an increase in loss that mainly results from coupling and greater modal interaction with sidewall roughness. The performance enhancements gained by reducing the cross-section of a device should therefore be evaluated against any increase in loss that results.

In addition to the 'standard' metrics of optical modulators discussed above, power consumption — the energy expended in producing each bit of data — has become particularly important for optical interconnects. This metric is also known as the energy per bit, or more colloquially as the 'power per bit'. Some authors argue that the adoption of optical interconnects can only be justified if their power consumption does not exceed that of present electrical interconnects. Adhering to this target, future systems will need to reach power consumptions of less than 1 pJ bit^{-1} — an extremely demanding target¹⁶. Others extend this argument to limit the total (on-chip) system energy to only ~ 100 fJ bit^{-1} , such that the efficiency of the equalized on-chip electrical link at the 45 nm process node is exceeded¹⁷. Because there are several energy contributions to the power budget, this requirement implies that the energy per bit for an optical output device (a modulator in this case) must be ~ 10 fJ bit^{-1} (ref. 18). In comparison, electrical components alone currently consume 2–3 pJ bit^{-1} , which is several orders of magnitude higher than the target of 10 fJ bit^{-1} . Although




Figure 3 | Carrier-depletion-based silicon optical modulators, showing both the first proposed and present fastest devices. a, Cross-section of the first proposed device, which predicted an electro-optic bandwidth in excess of 50 GHz (ref. 64). **b**, Cross-section of the fastest realized carrier-depletion-based silicon optical modulator so far. **c**, Eye diagram obtained from the device in **b**, demonstrating data transmission at 40 Gbit s⁻¹ with a 1 dB extinction ratio. Figure **b** and **c** reproduced with permission from ref. 39, © 2007 IEEE and ref. 65, © 2007 OSA.

power consumption is certainly important, the added value that optical interconnects bring could justify their larger power consumption. Note that the typical electrical consumption of today's servers is of the order of 10–30 pJ bit⁻¹ (ref. 5) and hence one may not necessarily need an improvement of an order of magnitude to justify the added value of optical interconnects. Similarly, typical commercial 10 Gbit s⁻¹ vertical-cavity surface-emitting laser optical links have an end-to-end power consumption of the order of 10 pJ bit⁻¹ (ref. 19), which emphasizes how demanding some of the theoretical figures for power consumption are for state-of-the-art silicon modulators. For example, one of the most compact MZI modulators reported so far has an energy requirement of 5 pJ bit⁻¹ at 10 Gbit s⁻¹ for a ~200-µm-long device²⁰. Clearly this is still much larger than the target of ~10 fJ bit⁻¹ proposed above, yet it may bring significant advantages. Electro-absorption modulators could potentially offer low operating energies because they have a strong QCSE, do not rely on changes in the carrier density, and have already been demonstrated with CMOS-compatible driving voltages²¹. Electro-absorption modulators are therefore promising for both waveguide and free-space optical systems.

Another important feature of a device is its footprint. MZI-based modulators tend to require a long interaction length (millimetres) for a complete transition between a maximum and a minimum of the optical transmission. This hinders high-speed performance and results in greater insertion loss, cost and power consumption. The footprints of resonant devices, on the other hand, are typically much smaller^{22–26}. This discussion might therefore imply that modulators incorporated into resonant structures are superior to those based on MZIs. However, there are further factors to take into account that require discussion.

Optical bandwidth refers to the useful operational wavelength range of a device, with resonant-structure-based modulators tending to function over a relatively narrow band compared with MZI-based devices. Although significant footprint savings are achievable with resonant devices, the resulting bandwidth is often reduced from more than 20 nm in the case of an MZI-based device to around 100 pm in the case of ring resonators²⁴.

Narrowband devices are limited in their applications, and their main shortcoming is their high sensitivity to typical fabrication tolerances and temperature variations. High sensitivity to fabrication tolerance can have a significant effect on device operation; in ring resonators, for example, the resonant frequency shifts by approximately 0.25 nm for every nanometre increase in the average width of the ring waveguide¹⁷. Achieving the right resonant wavelength can therefore be challenging using the 130 nm CMOS generation, as three standard deviations of the fabrication tolerance is 40 nm.

This figure is significantly reduced to either 16 nm or 1.2 nm with the implementation of 45 nm or 22 nm CMOS generations, respectively²⁷. To overcome this limitation, recent studies have investigated the use of SiN rings, which are less sensitive than silicon rings to fabrication tolerance²⁸, as well as eight-channel second-order microring-resonator filter banks with accurate spectral spacing using modified scanning-electron-beam lithography²⁹.

The sensitivity of a symmetrically configured MZI to changes in temperature is theoretically very small, as it depends only on variations in the characteristics of the coupler structure. Ring resonators, on the other hand, are extremely temperature-sensitive (30–80 pm K⁻¹; refs 30,31) owing to their narrow resonant spectral width (~0.1 nm for a single ring³² and 0.5–2 nm for higher-order rings^{33,34}), and also to the large thermo-optic coefficient of silicon⁹. This means that thermo-electric coolers/controllers are required for temperature stabilization³⁵, which increases power consumption. To emphasize this issue, consider a typical server/PC system operating at a temperature of around 70 °C. Use of a ring resonator modulator in such a system would need the temperature to be held within ±1 °C to operate to specifications, thus requiring a thermo-electric cooler to be introduced. The extra power consumption that would result, together with the continuous temperature drift that would occur if the device was packaged next to a CPU, would negate the advantage of lower power consumption gained by using such a temperature-sensitive device. The modulator design must therefore mitigate temperature drift — a requirement that has spurred recent studies to minimize temperature dependence by incorporating claddings with a temperature dependence opposite to that of silicon^{30,36}. Currently, this approach is not a CMOS-compatible process. Multiple rings can also reduce temperature drift, but only over a relatively small wavelength range³⁷. Recently, passive temperature compensation has been demonstrated using a ring coupled to an MZI³⁸. Electro-absorption modulators, on the other hand, are likely to be much less temperature-sensitive than microring resonator refractive devices⁴.

In many cases, the metrics discussed above involve trade-offs. For example, although reducing the footprint of an MZI-based modulator also reduces its power consumption and optical/radiofrequency (RF) drive signal losses, the reduction in phase shift also reduces the depth of the interferometer null as the two waves in the interferometer are not fully in antiphase, thus reducing the modulation depth³⁹. Similarly, the optical bandwidth of a ring resonator and its thermal stability can be improved by cascading multiple rings, but this method introduces the disadvantages of added complexity, a larger footprint and power consumption, and increased channel spacing³⁴.

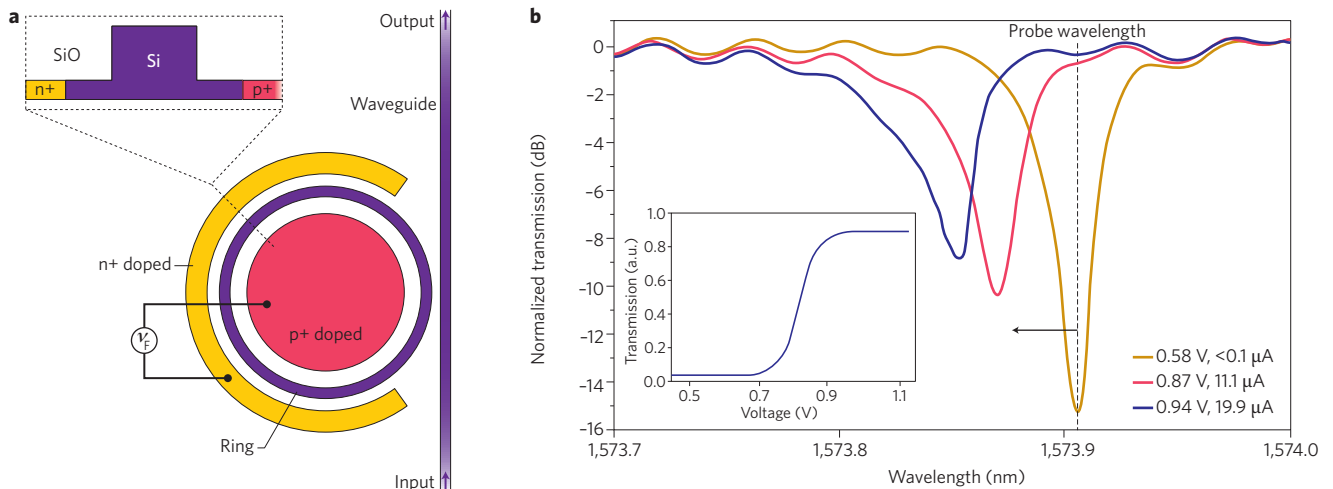


Figure 4 | The first silicon optical modulator to use a ring resonator structure to translate phase variations into intensity variations. **a**, Top-down and cross-sectional diagrams of the ring resonator structure, showing waveguide and doping positions. **b**, Typical voltage-induced spectral shift achieved from this device. Figure reproduced with permission from ref. 22, © 2005 NPG.

Another possibility for reducing device footprint is through a technique known as the slow-light approach. Slow light is a resonance effect that results when light propagates in a particular periodic structure; the light in the structure is coherently scattered and interferes with the incoming light, forming an interference pattern that moves forward slowly through the material. Such periodic structures can be found in the form of corrugated waveguides or photonic crystals. For the slow-light effect to be useful for modulation, it is necessary to operate at wavelengths away from those satisfying the Bragg condition, at which point standing waves turn the period structure into a high-reflectivity mirror. When the necessary conditions are met, slow-light modulators operate by slowing down the effective velocity of the passing light; hence the interaction with the modulating media can be enhanced, resulting in a reduction in the device length or drive voltage. However, these devices operate again through a resonant effect, and consequently optical bandwidth limitations apply.

Although this Review discusses the impact and performance of silicon-based optical modulators, it is important to also mention that the advantages of such modulators are intrinsically linked to the inherent ability of silicon photonics to facilitate optical integration at low cost. The modulator is integral to such an advantage, as a system with a high aggregate data rate can only be implemented in a high-volume application through an integrated format. Silicon photonics allows such implementation by facilitating combinations of high-data-rate modulators into a high-aggregate-bandwidth system.

State-of-the-art devices

Research into optical modulators based on silicon waveguides dates back to the mid-1980s⁴⁰. The early years of this research saw devices that were generally slow (~megahertz data rates), with seemingly little prospect for applications in high-speed systems. Most monolithic devices exploited carrier injection, using the plasma dispersion effect for modulating either the refractive index or the absorption coefficient of the material^{40–59}. In most cases p-i-n (p-type/intrinsic/n-type layers) diode structures were formed around the waveguide to electrically control the injection of electrons and holes into the path of the propagating light. To avoid excessive optical loss, the waveguides were positioned in the intrinsic region of the diode and the doped regions were positioned in such a way that the modal overlap was minimized.

Performance enhancements have been made over the years by optimizing the structure and reducing device dimensions; device bandwidths in the gigahertz regime had already been proposed by the mid-2000s. The schematic of one such device by Png *et al.*⁵⁸ is shown in Fig. 2a.

The breakthrough came in 2004 when the first monolithic silicon modulator capable of operating at speeds of >1 GHz was experimentally demonstrated by the Intel group⁶⁰. The device, shown in Fig. 2b, functioned by accumulating free carriers on either side of a dielectric layer inside a micrometre-sized waveguide (much like in a capacitor), rather than by carrier injection.

The performance of this carrier-accumulation device was later optimized, and data rates of 10 Gbit s⁻¹ were achieved by Liao *et al.* at an extinction ratio of 3.8 dB (ref. 61). Other accumulation-based modulators have been reported since then⁶², most notably the device recently demonstrated by Lightwire, which achieved 10 Gbit s⁻¹ at an extinction ratio of 9 dB (Fig. 2e)⁶³. It is also interesting to note this device demonstrated an improvement in efficiency of almost an order of magnitude over that of Liao *et al.*⁶¹, thus decreasing the active length from 3.45 mm to 480 μm . Carrier-accumulation modulators are not limited by the relatively long minority carrier lifetime in silicon, which is the case with carrier-injection modulators, but are instead dominated by the device resistance and capacitance.

Carrier depletion is another technique that manipulates free-carrier densities in a modulator to avoid the speed limitation posed by the minority carrier lifetime. Devices based on carrier depletion operate by allowing the propagating light to interact with the junction region of a p–n diode operated at reverse bias. The diode's depletion width, and therefore the free-carrier density in the waveguide, varies with the applied reverse bias. In 2005 Gardes *et al.* proposed the first modulator using this technique in a submicrometre waveguide (Fig. 3a), with theoretical models predicting an unprecedented intrinsic bandwidth of 50 GHz (ref. 64). The device featured a horizontal p–n junction across the waveguide with the highly doped regions required for the formation of resistive contacts placed at the extremities of the waveguide to avoid excessive optical loss.

Two years later a design similar to that of Gardes *et al.* was realized by Liu *et al.* from Intel⁶⁵. The authors reported initial data transmission rates of 30 Gbit s⁻¹ (ref. 65) but later improved this figure by demonstrating data transmission at 40 Gbit s⁻¹ (Fig. 3c)³⁹. Following




Figure 5 | A possible future alternative to plasma-dispersion-based silicon optical modulators is the SiGe QCSE modulator. Shown is a cross-section of the p-i-n structure on silicon with Ge/Si_{1-x}Ge_x multi-quantum-wells (MQWs) on a relaxed Si_{1-z}Ge_z buffer. Figure reproduced with permission from ref. 91, © 2010 IEEE.

this a plethora of carrier-depletion-based modulators capable of data transmission at speeds of >10 Gbit s⁻¹ were demonstrated^{23-25,66-73}, but the device of Liu *et al.*³⁹ still remains the fastest silicon modulator reported so far. The different devices described above demonstrate the possibility of achieving a high modulation efficiency using an injection device, but at the expense of modulation speed. The opposite is true for depletion devices, in which modulation speeds in excess of 30 GHz can be reached, but at reduced efficiency. The challenge remains to improve efficiency while maintaining other performance metrics.

The two principal metrics considered in this Review have been the modulation speed (bandwidth) and the area efficiency of the devices, mainly because these were the two most important metrics to the studies mentioned thus far. Over the years, however, as device speeds have increased to an acceptable level, other performance metrics such as power consumption and device footprint have grown in importance.

The use of resonant structures to decrease device footprint and therefore reduce power consumption has already been discussed. Active high-speed ring resonators were first introduced by Xu *et al.* in 2005²² (Fig. 4a). The authors introduced a p-i-n diode-based carrier injection scheme into submicrometre waveguide-based ring resonator structures, and achieved an initial data rate of 1.5 Gbit s⁻¹. The rings had diameters of 12 μm, giving a device footprint far smaller than any MZI-based modulator reported so far. The authors later reported improved results of 16 Gbit s⁻¹ at an extinction ratio of 8 dB using a pre-emphasis driving signal⁷⁴. This method of modulation was also used by Green *et al.* for a MZI-based modulator²⁰, which achieved a data rate of 10 Gbit s⁻¹ for a device length of 200 μm. The use of pre-emphasis drive signals in silicon was first proposed by Png *et al.* in 2004⁷⁵, and involves manipulation of the drive signal shape to reduce state transition times⁵⁸. Resonant-structure-based modulators have also been made using carrier-depletion structures, in which bandwidths compatible with data rates of 10 Gbit s⁻¹ have been demonstrated using a standard square-wave drive signal²³⁻²⁶. The fastest ring resonator carrier-depletion-based structure reported so far operates at a bandwidth of >35 GHz (ref. 72).

Photonic-crystal-based slow-light MZI modulators were first reported by Soljacic *et al.*⁷⁶ and replicated by other researchers⁷⁷⁻⁸⁰ using temperature modulation or through the injection or accumulation of free carriers. One of the main advantages of these devices is their ability to engineer the group velocity to allow device lengths of the order of 100 μm to be achieved, but this is often at the cost of optical bandwidth, as previously discussed⁸⁰. The dispersion engineering technique can be used⁸¹ to improve

Table 1 | A comparison of the various optical modulator metrics discussed, obtained from some key example devices (after ref. 24).

Modulation principle	Structure	Device footprint	Speed achieved	Energy per bit (fJ bit ⁻¹)	Modulation voltage	d.c. modulation depth and insertion loss	Modulation depth/speed	Working spectrum
Depletion of a horizontal p-n junction ⁶¹	MZI	~10 ⁴ μm ²	30 Gbit s ⁻¹ (more recently 40 Gbit s ⁻¹ ; ref. 39)	3 × 10 ⁴	6.5 V	>20 dB, ~7 dB	1 dB/30 Gbit s ⁻¹ , 1 dB/40 Gbit s ⁻¹	>20 nm
Forward-biased diode ²⁰	MZI	~10 ³ μm ²	10 Gbit s ⁻¹	~5 × 10 ³	7.6 V (pre-emphasis)	6-10 dB, 12 dB	—	—
Forward-biased diode ^{74,97}	Ring	~10 ² μm ²	>12.5 Gbit s ⁻¹	~300	3.5 V (pre-emphasis)	>10 dB, <0.5 dB	8 dB/16 Gbit s ⁻¹ , 3 dB/18 Gbit s ⁻¹	~0.1 nm
Reverse-biased p-n junction ²⁶	Disk	20 μm ²	10 Gbit s ⁻¹	85	3.5 V	8 dB, 1.5 dB	—	~0.1 nm
Forward-biased p-i-n diode ⁸³	Ring	~10 ² μm ²	3 Gbit s ⁻¹	86	0.5 V	7 dB, 1 dB	—	0.2 nm
Reverse-biased p-n junction ²⁴	Ring	~10 ³ μm ²	10 Gbit s ⁻¹	50	2 V	6.5 dB, 2 dB	8 dB/10 Gbit s ⁻¹	~0.1 nm

operating bandwidths to 10 nm, thus providing temperature insensitivity over a much wider range. However, so far modulation speeds of only ~ 1 GHz have been reported in these structures in silicon. Corrugated waveguide modulators that use the same technique have also been proposed⁸².

The power consumption of optical modulators has become an increasingly discussed topic in recent years. Modulators based on microring resonators detailed in recent papers by Dong *et al.*²⁴ and Chen *et al.*⁸³ have reported power consumptions of 50 fJ bit⁻¹ at 5 Gbit s⁻¹ and 86 fJ bit⁻¹ at 3 Gbit s⁻¹, respectively. However, as mentioned previously, microrings are very temperature dependent and therefore temperature stabilization may be required. For example, a power of 1 mW required for ring temperature stabilization at 10 Gbit s⁻¹ represents an additional power consumption of 100 fJ bit⁻¹. Simulations performed by Batten *et al.* have shown that the heating budget for the whole optical link should not exceed 100 fJ bit⁻¹ (ref. 84). By taking temperature stabilization into account, Stojanovic *et al.* reported power consumptions of ~ 300 fJ bit⁻¹ for their modulation scheme⁸⁵ — again significantly higher than the proposed target. Flip-chip integration and a heat sink on top of the modulator allowed the first subpicojoule-per-bit (400 fJ bit⁻¹) silicon modulator integrated with a driver circuit and embedded in a clocked digital transmitter to be achieved⁸⁶. The modulator alone consumed 120 fJ bit⁻¹. The authors stated that they were able to maintain error-free operation with a bit error rate lower than 10^{-15} at a data rate of 5 Gbit s⁻¹ with no active temperature control. A waveguide Franz–Keldysh modulator with an estimated power consumption of 50 fJ bit⁻¹ was also recently demonstrated¹².

This Review has so far predominantly discussed monolithic modulators, which operate using the plasma dispersion effect, along with brief discussions of Franz–Keldysh, thermal and QCSE modulators. Some of the key devices that have been discussed are highlighted in Table 1, together with their performance metrics. Alternatives do exist, but at the expense of CMOS compatibility or fabrication complexity. Arguments for a hybrid modulator are perhaps reinforced if one accepts that other devices such as an optical source or detector can also be usefully integrated through a hybrid approach. For example, hybrid techniques have achieved the most efficient light sources in silicon by introducing III–V material onto the silicon waveguide. Debate still rages as to whether on-chip integration of the light source is required, but there is no doubt that if one hybrid component is required, then others can be included at low additional cost. Modulators formed in this way have been demonstrated^{87–89}, with the most recent study achieving a data rate of 25 Gbit s⁻¹ at an extinction ratio of 10 dB (ref. 89). Modulation through the Kerr effect was demonstrated at 10 GHz (ref. 90), although the use of an engineered optical polymer waveguide cladding was required to overcome the relatively weak nonlinearities in silicon. As previously discussed, SiGe modulators that exploit the QCSE^{13,91} (Fig. 5) and Franz–Keldysh effect⁹² have also been reported, and modulators fabricated in amorphous silicon may find use in low-temperature back-end integration schemes⁹³. Electro-absorption modulation has been demonstrated in two multistack structures based on a-Si:H/a-SiCN (ref. 94). Applying strain to crystalline silicon has also been shown to induce the otherwise absent Pockels effect⁹⁵, which could provide an alternative method for achieving modulation in the future. However, this technique requires further investigation as the theoretical coefficient is significantly smaller than in experimental demonstrations⁹⁶.

Future outlook

The development of silicon optical modulators has clearly gone through a period of rapid improvement over the past six years. Commercial products and devices, integrated with various other

photonic/electronic components, are beginning to emerge for applications such as active optical cables⁶³, which use multiplexed modulators operating at 10 Gbit s⁻¹ rather than a single 40 Gbit s⁻¹ device. This decision was surely motivated by the impractical extinction ratios achieved so far at data rates of 40 Gbit s⁻¹, as well as the ease with which suitable drive electronics can be integrated with the modulator.

Silicon optical modulators will ideally need to have high modulation speeds, large bandwidths and small footprints, as well as low losses and ultralow power consumption. They must also be CMOS-compatible. However, these requirements often contradict each other, and therefore an innovative engineering solution is necessary to achieve an optimal trade-off. This means that despite impressive results recently achieved in this field, there are still significant challenges that the silicon photonics research community must overcome before the full deployment of this technology in several important application areas becomes a reality. Nevertheless, silicon photonics offers one of the most promising approaches to achieving the high-density electronic/photonic integration necessary for developing a range of prospective applications — most notably for the optical interconnects required in high-performance computing.

The question remains as to whether the optoelectronic integrated circuit will be realized as a monolithic or hybrid circuit. In some respects this question is unimportant, because any technology develops dynamically. The overwhelming evidence is that the implementation of silicon optoelectronics is almost inevitable, regardless of the detail of the ultimate commercially successful platform, because silicon photonics has now developed enough momentum to become a sufficiently viable technological solution. As for the modulator, monolithic variants continue to improve, but it is entirely possible, given the wealth of expertise currently devoted to this subject, that a new device design could change the landscape overnight. Whether this device is monolithic or hybrid remains to be seen, and in our opinion it is largely immaterial; much like the optoelectronic integrated circuit, the final form of the technology is unimportant. What we do know is that silicon photonics will make a difference, and that it is here to stay.

References

1. <http://www.cl.cam.ac.uk/~awm22/publications/miller2009motivating.pdf>
2. Miller, D. A. B. Rationale and challenges for optical interconnects to electronic chips. *Proc. IEEE* **88**, 728–749 (2000).
3. <http://www.itrs.net/links/2007itrs/execsum2007.pdf>
4. Miller, D. Device requirements for optical interconnects to silicon chips. *Proc. IEEE* **97**, 1166–1185 (2009).
5. Pepeljugoski, P. K. *et al.* Low power and high density optical interconnects for future supercomputers. *Optical Fiber Communication Conf. paper OThX2* (2010).
6. Lee, B. G., Biberman, A., Chan, J. & Bergman, K. High-performance modulators and switches for silicon photonic networks-on-chip. *IEEE J. Sel. Top. Quant. Electron.* **16**, 6–22 (2010).
7. Pollock, C. & Lipson, M. *Integrated Photonics* Ch. 12, 301–334 (Kluwer Academic Publishers, 2003).
8. Soref, R. & Bennett, B. Electrooptical effects in silicon. *IEEE J. Quant. Electron.* **23**, 123–129 (1987).
9. Reed, G. T. & Knights, A. P. *Silicon Photonics: An Introduction* Ch. 4, 97–103 (Wiley, 2004).
10. Cocorullo, G. & Rendina, I. Thermo-optical modulation at 1.5 μm in silicon etalon. *Electron. Lett.* **28**, 83–85 (1992).
11. Reed, G. T. *Silicon Photonics: The State of the Art* Ch. 4, 95–145 (Wiley, 2008).
12. Liu, J. *et al.* Waveguide-integrated, ultralow-energy GeSi electro-absorption modulators. *Nature Photon.* **2**, 433–437 (2008).
13. Kuo, Y.-H. *et al.* Strong quantum-confined Stark effect in germanium quantum-well structures on silicon. *Nature* **437**, 1334–1336 (2005).
14. Roth, J. E. *et al.* Optical modulator on silicon employing germanium quantum wells. *Opt. Express* **15**, 5851–5859 (2007).
15. Roth, J. E. *et al.* C-band side-entry Ge quantum-well electroabsorption modulator on SOI operating at 1 V swing. *Electron. Lett.* **44**, 49–50 (2008).

16. Krishnamoorthy, A. V. *et al.* Potentials of group IV photonics interconnects for 'red-shift' computing applications. *Proc. 4th IEEE Int. Conf. Group IV Photonics* 1–3 (2007).
17. Barwicz, T. *et al.* Silicon photonics for compact, energy-efficient interconnects. *J. Opt. Netw.* **6**, 63–73 (2007).
18. Yoo, S. J. B. Future prospects of silicon photonics in next generation communication and computing systems. *Electron. Lett.* **45**, 584–588 (2009).
19. Alduino, A. & Paniccia, M. Wiring electronics with light. *Nature Photon.* **1**, 153–155 (2007).
20. Green, W. M., 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).
21. Helman, N. C. *et al.* Misalignment-tolerant surface-normal low-voltage modulator for optical interconnects. *IEEE J. Sel. Top. Quant. Electron.* **11**, 338–342 (2005).
22. Xu, Q., Schmidt, B., Pradhan, S. & Lipson, M. Micrometre-scale silicon electro-optic modulator. *Nature* **435**, 325–327 (2005).
23. Gardes, F. Y. *et al.* High-speed modulation of a compact silicon ring resonator based on a reverse-biased pn diode. *Opt. Express* **17**, 21986–21991 (2009).
24. Dong, P. *et al.* Low V_{pp} , ultralow-energy, compact, high-speed silicon electro-optic modulator. *Opt. Express* **17**, 22484–22490 (2009).
25. You, J.-B., Park, M., Park, J.-W. & Kim, G. 12.5 Gbps optical modulation of silicon racetrack resonator based on carrier-depletion in asymmetric p-n diode. *Opt. Express* **16**, 18340–18344 (2008).
26. Watts, M. R., Trotter, D. C., Young, R. W. & Lentine, A. L. Ultralow power silicon microdisk modulators and switches. *Proc. 5th IEEE Int. Conf. Group IV Photonics* 4–6 (2008).
27. Vlasov, Y. Silicon photonics for next generation computing systems. *Proc. 34th European Conf. Optical Communications* paper Tu.1.A.1 (2008).
28. Gondarenko, A., Levy, J. S. & Lipson, M. High confinement micron-scale silicon nitride high Q ring resonator. *Opt. Express* **17**, 11366–11370 (2009).
29. Holzwarth, C. W. *et al.* Accurate resonant frequency spacing of microring filters without postfabrication trimming. *J. Vac. Sci. Technol. B* **24**, 3244–3247 (2006).
30. Teng, J. *et al.* Athermal silicon-on-insulator ring resonators by overlaying a polymer cladding on narrowed waveguides. *Opt. Express* **17**, 14627–14633 (2009).
31. Ye, W. N., Michel, J. & Kimerling, L. C. Athermal high-index-contrast waveguide design. *IEEE Photon. Tech. Lett.* **20**, 885–887 (2008).
32. Xu, Q., Fattal, D. & Beausoleil, R. G. Silicon microring resonators with 1.5- μm radius. *Opt. Express* **16**, 4309–4315 (2008).
33. Xiao, S., Khan, M. H., Shen, H. & Qi, M. A highly compact third-order silicon microring add-drop filter with a very large free spectral range, a flat passband and a low delay dispersion. *Opt. Express* **15**, 14765–14771 (2007).
34. Xia, F., Rooks, M., Sekaric, L. & Vlasov, Y. Ultra-compact high order ring resonator filters using submicron silicon photonic wires for on-chip optical interconnects. *Opt. Express* **15**, 11934–11941 (2007).
35. Manipatruni, S. *et al.* Wide temperature range operation of micrometer-scale silicon electro-optic modulators. *Opt. Lett.* **33**, 2185–2187 (2008).
36. Lee, J.-M. *et al.* Controlling temperature dependence of silicon waveguide using slot structure. *Opt. Express* **16**, 1645–1652 (2008).
37. Vlasov, Y., Green, W. M. J. & Xia, F. High-throughput silicon nanophotonic wavelength-insensitive switch for on-chip optical networks. *Nature Photon.* **2**, 242–246 (2008).
38. Guha, B., Kyotoku, B. B. C. & Lipson, M. CMOS-compatible athermal silicon microring resonators. *Opt. Express* **18**, 3487–3493 (2010).
39. Liao, L. *et al.* 40 Gbit/s silicon optical modulator for high-speed applications. *Electron. Lett.* **43**, 1196–1197 (2007).
40. Soref, R. A. & Bennett, B. R. Kramers–Kronig analysis of electro-optical switching in silicon. *Proc. SPIE Integrated Optical Circuit Engineering IV* **704**, 32–37 (1987).
41. Lorenzo, J. P. & Soref, R. A. 1.3 μm electro-optic silicon switch. *Appl. Phys. Lett.* **51**, 6–8 (1987).
42. Friedman, L., Soref, R. A. & Lorenzo, J. P. Silicon double-injection electro-optic modulator with junction gate control. *J. Appl. Phys.* **63**, 1831–1839 (1988).
43. Treyz, G. V. Silicon Mach–Zehnder waveguide interferometers operating at 1.3 μm . *Electron. Lett.* **27**, 118–120 (1991).
44. Treyz, G. V., May, P. G. & Halbout, J.-M. Silicon optical modulators at 1.3- μm based on free-carrier absorption. *IEEE Electr. Device Lett.* **12**, 276–278 (1991).
45. Jackson, S. M. *et al.* A novel optical phase modulator design suitable for phased arrays. *J. Lightwave Technol.* **16**, 2016–2019 (1998).
46. Jackson, S. M. *et al.* Optical beamsteering using integrated optical modulators. *J. Lightwave Technol.* **15**, 2259–2263 (1997).
47. Fischer, U., Schuppert, B. & Petermann, K. Integrated optical switches in silicon based on SiGe-waveguides. *IEEE Photon. Tech. Lett.* **5**, 785–787 (1993).
48. Tang, C. K., Reed, G. T., Wilson, A. J. & Rickman, A. G. Low-loss, single-model optical phase modulator in SIMOX material. *J. Lightwave Technol.* **12**, 1394–1400 (1994).
49. Tang, C. K., Reed, G. T., Wilson, A. J. & Rickman, A. G. Simulation of a low loss optical modulator for fabrication in SIMOX material. *Proc. Symp. Mater. Res. Soc.* **298**, 247–252 (1993).
50. Tang, C. K. & Reed, G. T. Highly efficient optical phase modulator in SOI waveguides. *Electron. Lett.* **31**, 451–452 (1995).
51. Cutolo, A., Iodice, M., Spirito, P. & Zeni, L. Silicon electro-optic modulator based on a three terminal device integrated in a low-loss single-mode SOI waveguide. *J. Lightwave Technol.* **15**, 505–518 (1997).
52. Cutolo, A., Iodice, M., Irace, A., Spirito, P. & Zeni, L. An electrically controlled Bragg reflector integrated in a rib silicon on insulator waveguide. *Appl. Phys. Lett.* **71**, 199–201 (1997).
53. Sciuto, A., Libertino, S., Alessandria, A., Coffa, S. & Coppola, G. Design, fabrication, and testing of an integrated Si-based light modulator. *J. Lightwave Technol.* **21**, 228–235 (2003).
54. Hewitt, P. D. & Reed, G. T. Improving the response of optical phase modulators in SOI by computer simulation. *J. Lightwave Technol.* **18**, 443–450 (2000).
55. Hewitt, P. D. & Reed, G. T. Multi micron dimension optical p-i-n modulators in silicon-on-insulator. *Proc. SPIE* **3630**, 237–243 (1999).
56. Winney, T. *et al.* Single-chip variable optical attenuator and multiplexer subsystem integration. *Optical Fiber Communication Conf.* paper TuK4 (2002).
57. Dainesi, P. *et al.* CMOS compatible fully integrated Mach-Zehnder interferometer in SOI technology. *IEEE Photon. Tech. Lett.* **12**, 660–662 (2000).
58. Png, C. E., Chan, S. P., Lim, S. T. & Reed, G. T. Optical phase modulators for MHz and GHz modulation in silicon-on-insulator (SOI). *J. Lightwave Technol.* **22**, 1573–1582 (2004).
59. Barrios, C. A., Almeida, V. R., Panepucci, R. & Lipson, M. Electrooptic modulation of silicon-on-insulator submicrometer-size waveguide devices. *J. Lightwave Technol.* **21**, 2332–2339 (2003).
60. Liu, A. *et al.* A high-speed silicon optical modulator based on a metal-oxide-semiconductor capacitor. *Nature* **427**, 615–618 (2004).
61. Liao, L. *et al.* High speed silicon Mach-Zehnder modulator. *Opt. Express* **13**, 3129–3135 (2005).
62. Kajikawa, K., Tabei, T. & Sunami, H. An infrared silicon optical modulator of metal-oxide-semiconductor capacitor based on accumulation-carrier absorption. *Jpn. J. Appl. Phys.* **48**, 04C107 (2009).
63. http://www.ofcnfoec.org/conference_program/2009/images/09-DAndrea.pdf
64. Gardes, F. Y., Reed, G. T., Emerson, N. G. & Png, C. E. A sub-micron depletion-type photonic modulator in silicon on insulator. *Opt. Express* **13**, 8845–8854 (2005).
65. Liu, A. *et al.* High-speed optical modulation based on carrier depletion in a silicon waveguide. *Opt. Express* **15**, 660–668 (2007).
66. Marris-Morini, D. *et al.* Optical modulation by carrier depletion in a silicon PIN diode. *Opt. Express* **14**, 10838–10843, (2006).
67. Gunn, C. CMOS photonics for high-speed interconnects. *IEEE Micro* **26**, 58–66 (2006).
68. Park, J. W., You, J.-B., Kim, I. G. & Kim, G. High-modulation efficiency silicon Mach–Zehnder optical modulator based on carrier depletion in a PN Diode. *Opt. Express* **17**, 15520–15524 (2009).
69. Narasimha, A. *et al.* An ultra low power CMOS photonics technology platform for H/S optoelectronic transceivers at less than \$1 per Gbps. *Optical Fiber Communication Conf.* paper OMV4 (2010).
70. Liow, T.-Y. *et al.* Silicon modulators and germanium photodetectors on SOI: Monolithic integration, compatibility, and performance optimization. *IEEE J. Sel. Top. Quant. Electron.* **16**, 307–315 (2010).
71. Feng, N. N. *et al.* High speed carrier-depletion modulators with 1.4V-cm $V_{\pi L}$ integrated on 0.25 μm silicon-on-insulator waveguides. *Opt. Express* **18**, 7994–7999 (2010).
72. Gill, D. M. *et al.* Internal bandwidth equalization in a CMOS-compatible Si-ring modulator. *IEEE Photon. Tech. Lett.* **21**, 200–202 (2009).
73. Spector, S. J. *et al.* High-speed silicon electro-optical modulator that can be operated in carrier depletion or carrier injection mode. *Conf. Lasers and Electro-Optics* paper CFH4 (2008).
74. Manipatruni, S., Qianfan, X., Schmidt, B., Shakya, J. & Lipson, M. High speed carrier injection 18 Gb/s silicon micro-ring electro-optic modulator. *IEEE Proc. Lasers and Electro-Optics Soc.* 537–538 (2007).
75. Png, C. E. *Silicon-on-insulator phase modulators*. PhD thesis, Univ. Surrey (2004).
76. Soljacic, M. *et al.* Photonic-crystal slow-light enhancement of nonlinear phase sensitivity. *J. Opt. Soc. Am. B* **19**, 2052–2059 (2002).
77. Jiang, Y., Jiang, W., Gu, L., Chen, X. & Chen, R. T. 80-micron interaction length silicon photonic crystal waveguide modulator. *Appl. Phys. Lett.* **87**, 221105 (2005).

78. Gu, L., Jiang, W., Chen, X., Wang, L. & Chen, R. T. High speed silicon photonic crystal waveguide modulator for low voltage operation. *Appl. Phys. Lett.* **90**, 071105 (2007).
79. Tanabe, T., Nishiguchi, K., Kuramochi, E. & Notomi, M. Low power and fast electro-optic silicon modulator with lateral p-i-n embedded photonic crystal nanocavity. *Opt. Express* **17**, 22505–22513 (2009).
80. Chen, X., Chen, Y. S., Zhao, Y., Jiang, W. & Chen, R. T. Capacitor-embedded 0.54 pJ/bit silicon-slot photonic crystal waveguide modulator. *Opt. Lett.* **34**, 602–604 (2009).
81. Li, J., White, T. P., O’Faolain, L., Gomez-Iglesias, A. & Krauss, T. F. Systematic design of flat band slow light in photonic crystal waveguides. *Opt. Express* **16**, 6227–6232 (2008).
82. Brimont, A., Sanchis, P. & Marti, J. Strong electro-optical modulation enhancement in a slow wave corrugated waveguide. *Opt. Express* **17**, 9204–9211 (2009).
83. Chen, L., Preston, K., Manipatruni, S. & Lipson, M. Integrated GHz silicon photonic interconnect with micrometer-scale modulators and detectors. *Opt. Express* **17**, 15248–15256 (2009).
84. Batten, C. *et al.* Building manycore processor-to-DRAM networks with monolithic silicon photonics. *16th IEEE Symp. High Performance Interconnects* 21–30 (2008).
85. Stojanovic, V., Joshi, A., Batten, C., Yong-jin, K. & Asanovic, K. Manycore processor networks with monolithic integrated CMOS photonics. *Conf. Lasers and Electro-Optics/Int. Quantum Electronics Conf.* paper CTuC3 (2009).
86. Zheng, X. *et al.* Ultra-low-energy all-CMOS modulator integrated with driver. *Opt. Express* **18**, 3059–3070 (2010).
87. Liu, L. *et al.* Carrier-injection-based electro-optic modulator on silicon-on-insulator with a heterogeneously integrated III–V microdisk cavity. *Opt. Lett.* **33**, 2518–2520 (2008).
88. Kuo, Y.-H., Chen, H.-W. & Bowers, J. E. High speed hybrid silicon evanescent electroabsorption modulator. *Opt. Express* **16**, 9936–9941 (2008).
89. Chen, H.-W., Kuo, Y. H. & Bowers, J. E. 25Gb/s hybrid silicon switch using a capacitively loaded traveling wave electrode. *Opt. Express* **18**, 1070–1075 (2010).
90. Hochberg, M. *et al.* Terahertz all-optical modulation in a silicon-polymer hybrid system. *Nature Mater.* **5**, 703–709 (2006).
91. Rong, Y. *et al.* Quantum-confined Stark effect in Ge/SiGe quantum wells on Si. *IEEE J. Sel. Top. Quant. Electron.* **16**, 85–92 (2010).
92. Liu, J. *et al.* Design of monolithically integrated GeSi electro-absorption modulators and photodetectors on an SOI platform. *Opt. Express* **15**, 623–628 (2007).
93. Preston, K., Manipatruni, S., Gondarenko, A., Poitras, C. B. & Lipson, M. Deposited silicon high-speed integrated electro-optic modulator. *Opt. Express* **17**, 5118–5124 (2009).
94. Della Corte, F. G., Rao, S., Nigro, M. A., Suriano, F. & Summonte, C. Electro-optically induced absorption in α -Si:H/ α -SiCN waveguiding multistacks. *Opt. Express* **16**, 7540–7550 (2008).
95. Jacobsen, R. S. *et al.* Strained silicon as a new electro-optic material. *Nature* **441**, 199–202 (2006).
96. Hon, N. K., Tsia, K. K., Solli, D. R., Jalali, B. & Khurgin, J. B. Stress-induced $\chi^{(2)}$ in silicon — comparison between theoretical and experimental values. *Proc. 6th IEEE Int. Conf. Group IV Photonics* 232–234 (2009).
97. Xu, Q., Manipatruni, S., Schmidt, B., Shakya, J. & Lipson, M. 12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators. *Opt. Express* **15**, 430–436 (2007).
98. Liao, L. High speed silicon-on-insulator modulators based on the free carrier plasma dispersion effect. PhD thesis, Univ. Surrey (2008).

Additional information

The authors declare no competing financial interests.

CORRIGENDUM

Silicon optical modulators

G. T. Reed, G. Mashanovich, F. Y. Gardes and D. J. Thomson

Nature Photonics **4**, 518–526 (2010); published online: 30 July 2010.

Ref. 86 was incorrectly cited in both the title and the last data row of Table 1. The correct citation is ref. 24, and this has been corrected in both the PDF and HTML versions of the article.